

Kitchen Math

Some Fly by Night TCS and Stuffs

Hang Su

Last updated: October 20, 2026

Contents

1	Probability and Computing	4
1.1	Events and Probability	4
1.2	Discrete Random Variable and Expectation	6
1.3	Moments and Deviations	8
1.3.1	Weak Law of Large Numbers (WLLN)	10
1.4	Chernoff and Hoeffding Bounds	11
1.4.1	Hoeffding's Bound	17
1.4.2	Quicksort Runtime Analysis	19
1.4.3	Network Routing Problems	20
1.5	Balls, Bins, and Random Graphs	24
1.5.1	A few things I should have noted about e	24
1.5.2	Birthday Paradox and Balls-and-Bins model	25
1.5.3	Poisson Distribution	28
1.5.4	Poisson Approximation	29
1.5.5	Recap on Coupon Collector's Problem	32
1.5.6	Hashing and Random Graphs	33
1.5.7	Hamiltonian Cycles in Random Graphs	38
1.6	The Probabilistic Method	41
1.6.1	A few things I should have noted about factorial	41
1.6.2	Basic Counting Argument	43
1.6.3	The Expectation Argument and Derandomization by Conditional Expectation	44
1.6.4	Sample and Modify	47
1.6.5	The Second Moment Method and the Conditional Expectation Inequality	48
1.6.6	Lovász Local Lemma	53
1.6.7	Algorithmic Lovász Local Lemma	58
1.7	Markov Chains and Random Walks	64
1.7.1	Stirling's Formula	64
1.7.2	Markov Chains and Classification of States	66
1.7.3	Strong Law of Large Numbers (SLLN)	75
1.7.4	Stationary Distributions	83
1.7.5	Random Walks on Undirected Graphs	95
1.7.6	Parrondo's Paradox	101
1.8	Continuous Distributions and the Poisson Process	102

1.9	The Normal Distribution	103
1.10	Entropy, Randomness, and Information	104
1.10.1	The Entropy Function	104
1.10.2	Entropy and Binomial Coefficients	107
1.10.3	Entropy: A Measure of Randomness	108
1.10.4	Entropy: A Limit of Data Compression	117
1.10.5	Shannon's Theorem	121
1.11	The Monte Carlo Method	130
1.11.1	The Monte Carlo Method	130
1.11.2	DNF Counting	130
1.11.3	From Approximate Sampling to Approximate Counting	130
1.11.4	The Markov Chain Monte Carlo Method (MCMC)	130
1.12	Coupling of Markov Chains	131
1.12.1	Variation Distance and Mixing Time	131
1.12.2	Coupling	132
1.12.3	Variation Distance is Nonincreasing	139
1.12.4	Geometric Convergence	143
1.12.5	Application in Approximately Sampling Proper Colorings	145
1.12.6	Path Coupling	146
1.13	Martingales	154
1.13.1	Martingales	154
1.13.2	Stopping Times	157
1.13.3	Wald's Equation	160
1.13.4	Azuma-Hoeffding's Inequality	165
1.14	Sample Complexity, VC Dimension, and Rademacher Complexity	174
1.15	Pairwise Independence and Universal Hash Function	175
1.15.1	Pairwise Independence	175
1.15.2	Chebyshev's Inequality for Pairwise Independent Random Variables	175
1.15.3	Universal Families of Hash Functions	175
1.15.4	Application in Finding Heavy Hitters in Data Streams	175
1.16	Power Laws and Related Distributions	176
1.17	Balanced Allocations and Cuckoo Hashing	177
2	Pseudorandomness	178
3	Probabilistic Method	179
4	Various Papers	180
4.1	Leftover Hash Lemma and HILL Entropy	180
4.2	Discrete Gaussian Leftover Hash Lemma	181
4.3	Schwartz-Zippel by Dana Moshkovitz	182
4.4	Ligero Interleaved Linear Code Testing	183
4.5	Interleaved Linear Code Testing with Logarithmic Randomness	186
4.6	Basefold Random Linear Foldable Code and the Distance Proof	187
4.7	Bulletproof	188
4.8	Sublinear Lattice ZKP for Arithmetic Circuits	189
4.9	A Non-PCP Approach to Succinct Quantum-Safe ZK	193
4.10	Short Invertible Polynomial Ring Elements	203
4.11	Lattice-Based Exact Argument of Knowledge	207
4.12	Practical Product Proof from Lattice Commitments	218
4.13	Practical Exact Proofs from Lattices	231
4.14	Lattice Proofs for Exact Euclidean Norm Bounds and Quadratic Relations	239

4.15 Practical Sublinear Proofs for R1CS from Lattices	254
4.16 LaBRADOR	255
4.17 Greyhound	256
4.18 Lattice-Based Succinct Arguments for NP with Polylog Verification	257



1 Probability and Computing

1.1 Events and Probability

Exercise 1.1.1 (Exercise 1.6 [MU17]). Consider the following balls-and-bin game. We start with one black ball and one white ball in a bin. We repeatedly do the following: choose one ball from the bin uniformly at random, then put the ball back in the bin with another ball of the same color. We repeat until there are n balls in the bin. Show that the number of white balls is equally likely to be any number between 1 and $n - 1$.

Proof. We rethink the behavior as: choose a color from the current balls in bin, and add one more ball with same color to the bin. Then we introduce a random variable X_n representing the number of white balls in the bin, where there are n balls in bin in total. The probability of k white balls among n balls can be represented by

$$\begin{aligned} \Pr[X_n = k] &= \Pr[X_n = k \mid X_{n-1} = k - 1] \cdot \Pr[X_{n-1} = k - 1] \\ &\quad + \Pr[X_n = k \mid X_{n-1} = k] \cdot \Pr[X_{n-1} = k] \\ &= \frac{k-1}{n-1} \cdot \Pr[X_{n-1} = k - 1] + \left(1 - \frac{k}{n-1}\right) \cdot \Pr[X_{n-1} = k]. \end{aligned}$$

By induction we can have that, the probability of X_n being k is $1/(n-1)$. □

Exercise 1.1.2 (Exercise 1.14 [MU17]). Suppose I am playing in a tournament against a player never played before. Consider 3 probabilities for my prior model: we are equally talented, and each of us is equally likely to win each game; I am slightly better, and I win each game independently with probability 0.6; or he is slightly better, and he wins each game independently with probability 0.6. Before we play, I think that each of these 3 probabilities is equally likely.

In the match we play until one player wins three games. I win the second game, but he wins the first, third, and fourth. After this match, with what probability should I believe that my opponent is slightly better than I am?

Proof. This is the Bayesian prior to posterior update.

Let E be the event that “I won the second and he won the first/third/fourth”.

Let C_0 be the event that he is better, C_1 be the event that I am better, C_2 be the event that we are equally talented. Then C_0, C_1, C_2 are disjoint events, with prior model on C_i being $\Pr[C_i] = 1/3$. The posterior model on C_0 is derived by the Bayes’ law

$$\begin{aligned} \Pr[C_0 \mid E] &= \frac{\Pr[E \mid C_0] \cdot \Pr[C_0]}{\sum_{i \in [0,2]} \Pr[E \mid C_i] \cdot \Pr[C_i]} \\ &= \frac{\Pr[E \mid C_0]}{\sum_{i \in [0,2]} \Pr[E \mid C_i]}, \end{aligned}$$

and each $\Pr[E \mid C_i] = p_i \cdot (1 - p_i)^3$, where p_i is the probability of me winning a game preconditioned in C_i .

Therefore, the exact probability of he is better preconditioned that E happened is

$$\Pr[C_0 \mid E] = \frac{0.6^3 \cdot 0.4}{0.6^3 \cdot 0.4 + 0.5^4 + 0.4^3 \cdot 0.6} \approx 0.46. \quad \square$$

Exercise 1.1.3 (Exercise 1.18 [MU17]). We have function $F : [0, n - 1] \rightarrow [0, m - 1]$, and we know $F((x + y) \bmod n) = (F(x) + F(y)) \bmod m$ for $x, y \in [0, n - 1]$. The only way we have for evaluating F is to use a lookup table that stores the values of F . An adversary changed the value of $1/5$ of the table entries.

Describe a simple randomized algorithm that, given an input z , outputs a value that equals $F(z)$ with probability at least $1/2$. The algorithm should be working for every $z \in [0, n - 1]$, regardless of what values the adversary changed. The algorithm should use as few lookups and as little computation as possible.

Suppose I allow you to repeat the initial algorithm three times, what is the probability of the enhanced version of algorithm returning the correct answer?

Proof. This looks very much like BLR linearity testing [BLR93]. We construct a random function as follows:

- Sample $\alpha \xleftarrow{R} [0, n - 1]$, and query $F(\alpha)$ and $F((z + \alpha) \bmod n)$.
- Output $(F((z + \alpha) \bmod n) - F(\alpha)) \bmod m$.

Since the error probability at a single query is $1/5$, the probability of this algorithm giving right answer is lower bounded by $3/5 \geq 1/2$.

Now that the error probability of our algorithm is upper bounded by $2/5$, if we repeat the algorithm 3 times, the algorithm returning none of them being correct has probability upper bounded by $(2/5)^3$.

- The probability of existing a correct answer is lower bounded by $1 - (2/5)^3$.
- If the majority are correct, then $(3/5)^3 + 3 \times (2/5) \times (3/5)^2 \geq 81/125$.

Try majority vote from 3 repetition as an enhanced version, then the correctness is lower bounded by $81/125$. $\square$

1.2 Discrete Random Variable and Expectation

Exercise 1.2.1 (Exercise 2.2 [MU17]). A monkey types on a 26-letter keyboard that has lowercase letters only. Each letter is chosen independently and uniformly at random from the alphabet. If the monkey types 1,000,000 letters, what is the expected number of times the sequence “proof” appears?

Proof. This is exactly a brain teaser in linearity of expectation.

Say X_i are 0-1 random variables indicating if the sequence “proof” appears at i^{th} alphabet. Then $X = \sum X_i$ is the random variable for the number of appearances of the sequence “proof”.

By linearity of expectation,

$$\mathbb{E}[X] = \mathbb{E}\left[\sum_{i \in [1, 999996]} X_i\right] = \sum_{i \in [1, 999996]} \mathbb{E}[X_i] = 999996 \times \left(\frac{1}{26}\right)^5.$$

□

Exercise 1.2.2 (Exercise 2.12 [MU17]). We draw cards uniformly at random with replacement from a deck of n cards. What is the expected number of cards we must draw until we have seen all n cards in the deck? If we draw $2n$ cards, what is the expected number of cards in the deck that are not chosen at all? Chosen exactly once?

Proof. The problem is a combination of “coupon collection” problem and “indicator” problem.

First one is the “coupon collection” problem, then let X_i be the number of draws to draw i^{th} different card, which is a geometrically distributed random variable with parameter $\frac{n-i+1}{n}$, and $X = \sum X_i$ be the total number of draws to draw all n different cards. $\mathbb{E}[X] = n \ln n + \Theta(n)$.

Later, we introduce 0-1 random variables $\{Y_i\}_{i \in [1, n]}$ indicating if the i^{th} cards is not drawn once in the $2n$ draws, then $\Pr[Y_i = 1] = \left(\frac{n-1}{n}\right)^{2n}$. Let $Y = \sum Y_i$ for the number of cards not drawn at all in $2n$ draws, then

$$\mathbb{E}[Y] = \mathbb{E}\left[\sum_{i \in [1, n]} Y_i\right] = \sum_{i \in [1, n]} \mathbb{E}[Y_i] = n \cdot \left(\frac{n-1}{n}\right)^{2n}.$$

We also introduce 0-1 random variables $\{Z_i\}_{i \in [1, n]}$ indicating if the i^{th} card is drawn only once in the $2n$ draws, then $\Pr[Z_i = 1] = 2n \cdot \frac{1}{n} \cdot \left(\frac{n-1}{n}\right)^{2n-1}$. Let $Z = \sum Z_i$ for the number of cards drawn only once in $2n$ draws, then

$$\mathbb{E}[Z] = \mathbb{E}\left[\sum_{i \in [1, n]} Z_i\right] = \sum_{i \in [1, n]} \mathbb{E}[Z_i] = 2n \cdot \left(\frac{n-1}{n}\right)^{2n-1}.$$

□

Exercise 1.2.3 (Exercise 2.16 [MU17]). Suppose we flip a coin n times to obtain a sequence of flips $\{X_i\}_{i \in [1, n]}$. A streak of flips is a consecutive subsequence of flips that are all the same.

- Let n be a power of 2. Show that the expected number of streaks of length $\log_2 n + 1$ is $1 - o(1)$.
- Show that, for sufficiently large n , the probability that there is no streak of length at least $\lfloor \log_2 n - 2 \log_2 \log_2 n \rfloor$ is less than $1/n$.

Proof. We first consider the expected number of streaks of length at least $\log_2 n + 1$. This is solved by the linearity of expectation, by introducing $\{X_i\}_{i \in [1, k]}$ 0-1 random variables, where $k = n - \log_2 n$, indicating the i^{th} flip starts a streak of length at least $\log_2 n + 1$. The probability for $X_i = 1$ can be upper bounded as follows

$$\Pr[X_i = 1] = 2 \cdot \frac{1}{2^{\log_2 n + 1}} = \frac{1}{n},$$

as streak on either face works. Therefore, number of streaks is $X = \sum_{i \in [1, k]} X_i$, and the expectation is

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}\left[\sum_{i \in [1, k]} X_i\right] = \sum_{i \in [1, k]} \mathbb{E}[X_i] \\ &= (n - \log_2 n) \mathbb{E}[X_i] \\ &= \frac{n - \log_2 n}{n} = 1 - o(1). \end{aligned}$$

We now show that the probability of no streaks of length at least $\ell = \lfloor \log_2 n - 2 \log_2 \log_2 n \rfloor$ is less than $1/n$. We consider n/ℓ segments of ℓ length flips, and thus introduce n/ℓ independent 0-1 random variables indicating if the i^{th} segment starts a streak of length at least ℓ . Therefore $\Pr[Y_i = 1] = 2^{-\ell}$.

Though n/ℓ segments is less than the total number of potential segments among n flips, it suffices for giving a sufficiently low probability. We now bound the probability of all n/ℓ segments has no streaks of length ℓ by

$$\Pr\left[\bigcap_{i \in [1, n/\ell]} Y_i = 0\right] = \left(1 - \frac{2}{2^\ell}\right)^{n/\ell} = \left(1 - \frac{1}{2^{\ell-1}}\right)^{2^{\ell-1} \cdot n/(\ell \cdot 2^{\ell-1})} \leq \left(\frac{1}{e}\right)^{n/(\ell \cdot 2^{\ell-1})}.$$

Let $\eta = n/(\ell \cdot 2^{\ell-1})$. By $\lfloor x \rfloor \leq x$ for any x , we lower bound η by

$$\begin{aligned} \eta &= \frac{n}{\ell \cdot 2^{\ell-1}} = \frac{2n}{\lfloor \log_2 n - 2 \log_2 \log_2 n \rfloor} \cdot \frac{1}{2^{\lfloor \log_2 n - 2 \log_2 \log_2 n \rfloor}} \\ &\geq \frac{2n}{\log_2 n - 2 \log_2 \log_2 n} \cdot \frac{\log_2^2 n}{n} = \frac{2 \log_2^2 n}{\log_2 n - 2 \log_2 \log_2 n} \\ &\geq 2 \log_2 n. \end{aligned}$$

Therefore, $\exp(-\eta)$ is at most $1/n$, and thus completes the proof. $\square$

Exercise 1.2.4 (Exercise 2.20 [MU17]). A permutation on the numbers $[1, n]$ can be represented as a function $\pi : [1, n] \mapsto [1, n]$. A fixed point of a permutation π is a value for which $\pi(x) = x$. Find the expected number of fixed points for a permutation chosen uniformly at random for all permutations.

Proof. This is another indicator counting problem.

We introduce n 0-1 random variables $\{X_i\}_{i \in [1, n]}$ indicating if i^{th} value is a fixed point for the permutation. Then $\Pr[X_i = 1] = 1/n$, and thus the expectation of $X = \sum X_i$, which is the number of fixed points in π , is

$$\mathbb{E}[X] = \mathbb{E}\left[\sum_{i \in [1, n]} X_i\right] = \sum_{i \in [1, n]} \mathbb{E}[X_i] = 1.$$

$\square$

Theorem 1.2.5 (Jensen's Inequality). *If f is a convex function, then*

$$\mathbb{E}[f(X)] \geq f(\mathbb{E}[X]).$$

If f is concave, then

$$\mathbb{E}[f(X)] \leq f(\mathbb{E}[X]).$$

Definition 1.2.6. The expression $\mathbb{E}[Y | X]$ is a random variable $f(X)$ that takes on value $\mathbb{E}[Y | X = x]$ when $X = x$.

Lemma 1.2.7.

$$\mathbb{E}[Y] = \mathbb{E}[\mathbb{E}[Y | X]].$$

1.3 Moments and Deviations

Exercise 1.3.1 (Exercise 3.9 [MU17]). Let $\{X_i\}_{i \in [1, m]}$ be Bernoulli random variables that are not necessarily mutually independent, and $X = \sum_i X_i$. Then

$$\mathbb{E}[X^2] = \sum_{i \in [1, m]} \Pr[X_i = 1] \mathbb{E}[X \mid X_i = 1].$$

Proof. We can derive the proof by

$$\begin{aligned} \mathbb{E}[X^2] &= \mathbb{E}\left[\sum_{i \in [1, m]} X_i X\right] = \sum_{i \in [1, m]} \mathbb{E}[X_i X] \\ &= \sum_{i \in [1, m]} \Pr[X_i = 1] \mathbb{E}[X \mid X_i = 1], \end{aligned}$$

where the first two equalities hold from linearity of expectation, and the last one holds by conditional expectation. $\square$

Lemma 1.3.2. Let $\{X_i\}_{i \in [1, m]}$ be Bernoulli random variables that are not necessarily mutually independent, and $X = \sum_i X_i$. Then

$$\begin{aligned} \text{Var}[X] &= \sum_{i \in [1, m]} \text{Var}[X_i] + \sum_{\substack{i, j \in [1, m] \\ i \neq j}} \text{Cov}[X_i, X_j] \\ &= \sum_{i \in [1, m]} \text{Var}[X_i] + 2 \cdot \sum_{\substack{i, j \in [1, m] \\ i < j}} \text{Cov}[X_i, X_j]. \end{aligned}$$

Exercise 1.3.3 (Exercise 3.11 [MU17]). Recall the Bubblesort algorithm. Determine the variance of the number of inversions that need to be corrected by Bubblesort.

Proof. We introduce $\{X_{i,j}\}$ 0-1 random variables, where $j \in [1, n]$ and $i \in [1, j-1]$. $X_{i,j}$ is 1 if and only if the i^{th} and the j^{th} element are inverted during the sorting. The variance of $X = \sum_{i,j} X_{i,j}$ can be expressed by lemma 1.3.2

$$\text{Var}[X] = \sum_{\substack{j \in [1, n] \\ i \in [1, j-1]}} \text{Var}[X_{i,j}] + 2 \cdot \sum_{\substack{j_0, j_1 \in [1, n] \\ i_0 \in [1, j_0-1], i_1 \in [1, j_1-1] \\ (i_0, j_0) < (i_1, j_1)}} \text{Cov}[X_{i_0, j_0}, X_{i_1, j_1}].$$

$X_{i,j} = 1$ if and only if the i^{th} element is greater than j^{th} element, and thus $\Pr[X_{i,j} = 1] = 1/2$. Since $\mathbb{E}[X_{i,j}] = 1/2$, therefore $\text{Var}[X_{i,j}] = 1/4$.

For $\text{Cov}[X_{i_0, j_0}, X_{i_1, j_1}]$, if the two pairs are disjoint, that $\{i_0, j_0\} \cap \{i_1, j_1\} = \emptyset$, then X_{i_0, j_0} and X_{i_1, j_1} are independent, therefore $\text{Cov}[X_{i_0, j_0}, X_{i_1, j_1}] = 0$.

On the other hand, if i_0, i_1, j_0, j_1 have index overlap, there are $\binom{n}{3}$ choices, each breaking into 3 categories:

- if $i_0 = i_1$, then $X_{i_0, j_0} \cdot X_{i_1, j_1} = 1$ if and only if i_0^{th} element is the largest among 3, and thus $\mathbb{E}[X_{i_0, j_0} \cdot X_{i_1, j_1}] = 1/3$.
- if $j_0 = i_1$, then $X_{i_0, j_0} \cdot X_{i_1, j_1} = 1$ if and only if j_0^{th} element is in the middle, and i_0^{th} element is the largest, which halves the probability, and therefore $\mathbb{E}[X_{i_0, j_0} \cdot X_{i_1, j_1}] = 1/6$.
- if $j_0 = j_1$, then $X_{i_0, j_0} \cdot X_{i_1, j_1} = 1$ if and only if j_0^{th} element is the smallest among 3, and thus $\mathbb{E}[X_{i_0, j_0} \cdot X_{i_1, j_1}] = 1/3$.

We conclude that

$$\text{Var}[X] = \sum_{j \in [1, n]} \sum_{i \in [1, j-1]} \frac{1}{4} + 2 \cdot \binom{n}{3} \cdot \left(\frac{1}{3} + \frac{1}{3} + \frac{1}{6} - \frac{3}{4}\right) = \frac{n(n-1)}{8} + 2 \cdot \frac{1}{12} \cdot \binom{n}{3} = \frac{1}{72} n(n-1)(2n+5).$$

$\square$

Exercise 1.3.4 (Exercise 3.13 [MU17]). Find an example of random variable with finite j^{th} moment for $j \in [1, k]$, but an unbounded $k + 1^{\text{th}}$ moment.

Proof. Let $X \in \mathbb{N}^+$ have $\Pr[X = x] = c \cdot x^{-k-2}$ with $c = (\sum_{x \in \mathbb{N}^+} x^{-k-2})^{-1}$. $\mathbb{E}[X^k]$ converges, $\mathbb{E}[X^{k+1}]$ doesn't. $\square$

Remark 1.3.5. Ever heard of Pareto tails?

Theorem 1.3.6. For any random variable X and finite expectation $\mathbb{E}[X]$, finite $\mathbb{E}[X^2]$, and finite median m ,

- $\mathbb{E}[X]$ is the value c that minimizes the expression $\mathbb{E}[(X - c)^2]$.
- m is the value c that minimizes the expression $\mathbb{E}[|X - c|]$.

Proof. The first one is immediate, as $\mathbb{E}[(X - c)^2]$ minimizes at $c = \mathbb{E}[X]$ by expanding the expression.

For second case, we assume that there exists c such that $\mathbb{E}[|X - c|] < \mathbb{E}[|X - m|]$, or $\mathbb{E}[|X - m| - |X - c|] > 0$. Since $c > m$ or $c < m$ works the same by symmetry, WLOG we let $c > m$, then

$$|X - m| - |X - c| = \begin{cases} m - c & X \leq m \\ 2X - m - c & m \leq X \leq c \\ c - m & X \geq c \end{cases}.$$

For $m < X < c$, we have $2X - m - c \leq c - m$, and for $X \geq c$ the expression equals $c - m$, so

$$\mathbb{E}[|X - m| - |X - c|] \leq (m - c) \Pr[X \leq m] + (c - m) \Pr[X > m].$$

Since for a median m , $\Pr[X \leq m] \geq 1/2$ and $\Pr[X > m] \leq 1/2$, the right-hand side is at most 0, contradicting the prior assumption of existence of c , and proves that m minimizes the expression $\mathbb{E}[|X - c|]$. $\square$

Lemma 1.3.7. Let X be a random variable, with mean μ , median m , and finite standard deviation σ . Then $|\mu - m| \leq \sigma$.

Proof.

$$\begin{aligned} |\mathbb{E}[X] - m| &\leq \mathbb{E}[|X - m|] \\ &\leq \mathbb{E}[|X - \mathbb{E}[X]|] \\ &\leq \sqrt{\mathbb{E}[(X - \mathbb{E}[X])^2]} = \sigma, \end{aligned}$$

where the first and the third inequality follow from theorem 1.2.5, and the second one follows from theorem 1.3.6. $\square$

Exercise 1.3.8 (Exercise 3.18 [MU17]). Show that, for a random variable X with standard deviation $\sigma[X]$ and $t \in \mathbb{R}^+$:

- $\Pr[X - \mathbb{E}[X] \geq t\sigma[X]] \leq 1/(1 + t^2)$;
- $\Pr[|X - \mathbb{E}[X]| \geq t\sigma[X]] \leq 2/(1 + t^2)$.

Proof. Cantelli's inequality (or one-sided Chebyshev's inequality) can be proved by:

$$\begin{aligned} \Pr[X - \mathbb{E}[X] \geq t\sigma[X]] &= \Pr[X - \mathbb{E}[X] + u \geq t\sigma[X] + u] \\ &\leq \Pr[(X - \mathbb{E}[X] + u)^2 \geq (t\sigma[X] + u)^2] \\ &\leq \frac{\mathbb{E}[(X - \mathbb{E}[X] + u)^2]}{(t\sigma[X] + u)^2} = \frac{\mathbb{E}[(X - \mathbb{E}[X])^2] + u^2}{(t\sigma[X] + u)^2} \\ &= \frac{\sigma[X]^2 + u^2}{(t\sigma[X] + u)^2} = \varphi(u). \end{aligned}$$

The last inequality was established by Chebyshev's inequality.

$\varphi(u)$ is minimized at $u = \sigma[X]/t$, and thus the first inequality is upper bounded by $1/(1 + t^2)$.

By symmetry, the second inequality is proven. $\square$

Exercise 1.3.9 (Exercise 3.19 [MU17]). Using exercise 1.3.8, show lemma 1.3.7.

Proof. By exercise 1.3.8, we have $\Pr[X \leq \sigma + \mu] \geq 1/2$ and $\Pr[X \geq \mu - \sigma] \geq 1/2$.

By definition of median m , $\Pr[X \geq m] \geq 1/2$ and $\Pr[X \leq m] \geq 1/2$.

Since $\sigma > 0$ is finite, μ is finite, we have $\mu - \sigma \leq m \leq \mu + \sigma$, and thus completes the proof for $|m - \mu| \leq \sigma$. $\square$

Exercise 1.3.10 (Exercise 3.20 [MU17]). Let Y be a nonnegative integer-valued random variable with positive expectation. Prove

$$\frac{\mathbb{E}[Y]^2}{\mathbb{E}[Y^2]} \leq \Pr[Y \neq 0] \leq \mathbb{E}[Y].$$

Proof. By conditional expectation, we have

$$\begin{aligned} \mathbb{E}[Y^2] &= \mathbb{E}[Y^2 \mid Y = 0] \cdot \Pr[Y = 0] + \mathbb{E}[Y^2 \mid Y \neq 0] \cdot \Pr[Y \neq 0] \\ &= \mathbb{E}[Y^2 \mid Y \neq 0] \cdot \Pr[Y \neq 0] \\ &\geq \mathbb{E}[Y \mid Y \neq 0]^2 \cdot \Pr[Y \neq 0], \end{aligned}$$

where the last inequality follows from Jensen's inequality.

Again by conditional expectation, we have $\mathbb{E}[Y] = \mathbb{E}[Y \mid Y \neq 0] \cdot \Pr[Y \neq 0]$, and thus

$$\frac{\mathbb{E}[Y]^2}{\mathbb{E}[Y^2]} \leq \Pr[Y \neq 0].$$

Since Y is an integer-valued random variable,

$$\begin{aligned} \Pr[Y \neq 0] &= 0 \cdot \Pr[Y = 0] + \sum_{y \neq 0} 1 \cdot \Pr[Y = y] \\ &\leq \sum_{y \in \mathbb{N}} y \cdot \Pr[Y = y] = \mathbb{E}[Y]. \end{aligned}$$

Another way of seeing it is by a direct application of Markov's inequality, that $\Pr[Y \geq 1] \leq \mathbb{E}[Y]$. $\square$

1.3.1 Weak Law of Large Numbers (WLLN)

Theorem 1.3.11 (WLLN). If $\{X_i\}_{i \in [1, n]}$ are independent, identically distributed random variables with mean μ and standard deviation σ , then for any constant $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \Pr \left[\left| \frac{1}{n} \sum_i X_i - \mu \right| > \varepsilon \right] = 0.$$

1.4 Chernoff and Hoeffding Bounds

Exercise 1.4.1 (Exercise 4.9 [MU17]). Suppose that we can obtain independent samples $\{X_i\}_{i \in [1, n]}$ of a random variable X and that we want to use these samples to estimate $\mathbb{E}[X]$. Using t samples, we use $(\sum_{i \in [1, t]} X_i)/t$ for our estimate of $\mathbb{E}[X]$. We want the estimate to be within $\varepsilon \mathbb{E}[X]$ from the true value of $\mathbb{E}[X]$ with probability at least $1 - \delta$. We may not be able to use Chernoff's bound directly to bound how good our estimate is if X is not a 0-1 random variable, and we do not know its moment generating function. We develop an alternative approach that requires only having a bound on the variance of X . Let $r = \sqrt{\text{Var}[X]}/\mathbb{E}[X]$.

- Show using Chebyshev's inequality that $O(r^2/\varepsilon^2\delta)$ samples are sufficient to solve the problem.
- Suppose that we need only a weak estimate that is within $\varepsilon \mathbb{E}[X]$ of $\mathbb{E}[X]$ with probability at least $3/4$. Argue that $O(r^2/\varepsilon^2)$ samples are enough for this weak estimate.
- Show that, by taking the median of $O(\log(1/\delta))$ independent weak estimates, we can obtain an estimate within $\varepsilon \mathbb{E}[X]$ of $\mathbb{E}[X]$ with probability at least $1 - \delta$. Conclude that we need only $O((r^2 \log(1/\delta))/\varepsilon^2)$ samples.

Proof. Let $Y = \sum_{i \in [1, t]} X_i$, then $\mathbb{E}[Y] = t \cdot \mathbb{E}[X]$ and $\text{Var}[Y] = \text{Var}[\sum_{i \in [1, t]} X_i] = t \cdot \text{Var}[X]$.

By Chebyshev's inequality,

$$\begin{aligned} \Pr \left[\left| \frac{1}{t} \cdot Y - \mathbb{E}[X] \right| \geq \varepsilon \mathbb{E}[X] \right] &= \Pr [|Y - \mathbb{E}[Y]| \geq \varepsilon \mathbb{E}[Y]] \\ &\leq \frac{\text{Var}[Y]}{\varepsilon^2 \mathbb{E}[Y]^2} = \frac{\text{Var}[X]}{t \varepsilon^2 \mathbb{E}[X]^2} = \frac{r^2}{t \varepsilon^2} \leq \delta. \end{aligned}$$

Now that $t \geq r^2/\delta\varepsilon^2$, thus $t = O(r^2/\delta\varepsilon^2)$.

Suppose $\delta \leq 1/4$, then $t \geq 4r^2/\varepsilon^2$, and thus $t = O(r^2/\varepsilon^2)$.

When it comes to using a median of s weak estimates to boost the confidence in estimation, we have a trick called *Median of Mean (MoM)* [JV86]. The key idea is to take advantage of the “structure” of median value, that half of the weak estimates are smaller, while the other half are greater.

Since we hope with probability at least $1 - \delta$, the median of weak estimates is within $\varepsilon \mathbb{E}[X]$ range, then we say the bad event is the median of weak estimates is out of $\varepsilon \mathbb{E}[X]$ range. Equivalently, a bad event occurs when more than half of the weak estimates are above $(1 + \varepsilon) \mathbb{E}[X]$, or more than half of them are below $(1 - \varepsilon) \mathbb{E}[X]$.

Form each weak estimate from an independent batch of $O(r^2/\varepsilon^2)$ samples, so the weak estimates are independent. In this sense, we introduce independent 0-1 random variables (indicators) $\{1_{L,i}\}_{i \in [1, s]}$ and $\{1_{R,i}\}_{i \in [1, s]}$ where

- $1_{L,i} = 1$ only when i^{th} weak estimate is smaller than $(1 - \varepsilon) \mathbb{E}[X]$,
- $1_{R,i} = 1$ only when i^{th} weak estimate is greater than $(1 + \varepsilon) \mathbb{E}[X]$,

and $1_L = \sum_{i \in [1, s]} 1_{L,i}$, $1_R = \sum_{i \in [1, s]} 1_{R,i}$. Since a weak estimate is within $\varepsilon \mathbb{E}[X]$ range with probability at least $3/4$, we have both $\Pr[1_{L,i} = 1]$ and $\Pr[1_{R,i} = 1]$ upper bounded by $1/4$, and thus both $\mathbb{E}[1_L]$ and $\mathbb{E}[1_R]$ are at most $s/4$.

We first bound 1_L , and derive 1_R by symmetry. By Chernoff's bound,

$$\Pr \left[1_L \geq \frac{s}{2} \right] \leq \mathbb{E}[\exp(t \cdot 1_L)] / \exp\left(\frac{s}{2} \cdot t\right) \leq \exp\left(\frac{s}{4} \cdot (e^t - 1) - \frac{s}{2} \cdot t\right) = \exp(\varphi(t)).$$

Then we have $\varphi'(t) = \frac{s}{4} \cdot e^t - \frac{s}{2}$, and derive that $\varphi(t)$ min by $t = \ln 2$ for a positive t . Now that

$$\exp\left(\left(\frac{1}{4} - \frac{\ln 2}{2}\right) \cdot s\right) \leq \frac{\delta}{2},$$

we have $s \geq ((\ln 2)/2 - 1/4)^{-1} \cdot \ln(2/\delta)$, thus $s = O(\log(1/\delta))$.

We thus conclude that we need $O(r^2 \log(1/\delta)/\varepsilon^2)$ samples in total. $\square$

Exercise 1.4.2 (Exercise 4.12 [MU17]). Consider a collection $\{X_i\}_{i \in [1, n]}$ of n independent geometrically distributed random variables with mean 2. Let $X = \sum_{i \in [1, n]} X_i$ and $\delta > 0$.

- Derive a bound on $\Pr[X \geq (1 + \delta) \cdot 2n]$ by applying Chernoff's bound to a sequence of $(1 + \delta) \cdot 2n$ fair coin tosses.
- Directly derive a Chernoff's bound on $\Pr[X \geq (1 + \delta) \cdot 2n]$ using the MGF for geometric random variables.
- Compare these 2 bounds derived above.

Proof. A way to model the “Coupon Collection” alike problem for bounding the sum of geometrically distributed random variables is transforming to an upper bound to the sum of a sequence of Bernoulli trials.

Since a geometrically distributed random variable can be seen as the number of Bernoulli trials until one success, then for event $X \geq (1 + \delta) \cdot 2n$, there are at most n successes among $(1 + \delta) \cdot 2n$ independent Bernoulli trials (in the worst case, the last trial must be success among n successes in total $(1 + \delta) \cdot 2n$ trials).

We write Y_i as i^{th} independent Bernoulli trial with success probability being $1/2$ among $(1 + \delta) \cdot 2n$ trials, we write $Y = \sum_{i \in [1, (1+\delta) \cdot 2n]} Y_i$, and $\mathbb{E}[Y] = (1 + \delta) \cdot n$. We thus can upper bound $\Pr[X \geq (1 + \delta) \cdot 2n]$ by $\Pr[Y \leq n]$, as we can loosen the boundary condition on the worst case, where last trial must be successful among n successes.

By Chernoff's bound for bounding the deviation below the mean, we have

$$\begin{aligned} \Pr[X \geq (1 + \delta) \cdot 2n] &\leq \Pr[Y \leq n] \\ &\leq \frac{\mathbb{E}[\exp(tY)]}{\exp(n \cdot t)} \\ &\leq \frac{\exp((1 + \delta) \cdot n \cdot (e^t - 1))}{\exp(n \cdot t)} \\ &= \exp((1 + \delta) \cdot n \cdot (e^t - 1) - n \cdot t) = \exp(\varphi(t)) \end{aligned}$$

with $t < 0$, and min by $t = -\ln(1 + \delta)$ deriving an upper bound $((1 + \delta) \cdot e^{-\delta})^n$.

If we brute force for a bound on $X = \sum_{i \in [1, n]} X_i$, the MGF of X_i is $M_{X_i}(t) = \sum_{k \geq 1} \exp(kt) \cdot 2^{-k} = \sum_{k \geq 1} (e^t/2)^k$. When $0 < t < \ln 2$, $M_{X_i}(t) = e^t/(2 - e^t)$. Again by Chernoff's bound for bounding the deviation above the mean,

$$\begin{aligned} \Pr[X \geq (1 + \delta) \cdot 2n] &\leq \left(\frac{M_{X_i}(t)}{\exp((1 + \delta) \cdot 2t)} \right)^n \\ &= ((2 - e^t) \exp((2\delta + 1) \cdot t))^{-n} \\ &= (2 \exp((2\delta + 1) \cdot t) - \exp((2\delta + 2) \cdot t))^{-n} = \varphi(t)^{-n}, \end{aligned}$$

and $\varphi'(t) = (4\delta + 2) \exp((2\delta + 1) \cdot t) - (2\delta + 2) \exp((2\delta + 2) \cdot t)$, and thus $\varphi(t)$ max by $t = \ln(2\delta + 1) - \ln(\delta + 1)$,

$$\Pr[X \geq (1 + \delta) \cdot 2n] \leq \left(\frac{1}{\delta + 1} \cdot \left(\frac{2\delta + 1}{\delta + 1} \right)^{2\delta + 1} \right)^{-n} = \left((1 + \delta) \cdot \left(\frac{\delta + 1}{2\delta + 1} \right)^{2\delta + 1} \right)^{-n}.$$

By section 1.5.1, we know the first one is looser. □

Exercise 1.4.3 (Exercise 4.8 [MU17]). We show how to construct a random permutation π on $[1, n]$, given a black box that outputs numbers independently and uniformly at random from $[1, k]$ where $k \geq n$. If we compute a function $f : [1, n] \rightarrow [1, k]$ with $f(i) \neq f(j)$ for $i \neq j$, this yields a permutation, by outputting $[1, n]$ according to the order of $\{f(i)\}_{i \in [1, n]}$. To construct such a function f , do the following for $i \in [1, n]$: choose $f(i)$ by repeatedly drawing number from the black box and set $f(i)$ to the first number that is $f(i) \neq f(j)$ for $j < i$.

Prove that this approach gives a permutation chosen uniformly at random for all permutations. Find the expected number of calls to the black box that are needed when $k = n$ and $k = 2n$. For the case $k = 2n$, argue that the probability that each call to the black box assigns a value of $f(j)$ to some j is at least $1/2$. Based on this, use a Chernoff bound to bound the probability that the number of calls to the black box is at least $4n$.

Proof. At the first sight, this is a “Coupon Collection” problem, where we can introduce n independent geometrically distributed random variables $\{X_i\}_{i \in [1, n]}$ each with success probability $p_i = (k+1-i)/k$, and we write $X = \sum_{i \in [1, n]} X_i$. When $k = 2n$, each $p_i \geq 1/2$ as $p_n = (n+1)/2n$.

For expectation of X , $\mathbb{E}[X] = \sum_{i \in [1, n]} \mathbb{E}[X_i]$ by linearity of expectation. For $\mathbb{E}[X_i]$, by conditional expectation and memoryless of geometric distribution,

$$\begin{aligned}\mathbb{E}[X_i] &= \mathbb{E}[X_i \mid X_{i,0} = 1] \cdot \Pr[X_{i,0} = 1] + \mathbb{E}[X_i \mid X_{i,0} = 0] \cdot \Pr[X_{i,0} = 0] \\ &= 1 \cdot p_i + (1 + \mathbb{E}[X_i]) \cdot (1 - p_i),\end{aligned}$$

$\mathbb{E}[X_i] = p_i^{-1} = k/(k-i+1)$. Thus $\mathbb{E}[X] = n \ln n + \Theta(n)$ when $k = n$, $\mathbb{E}[X] = 2n(\ln 2n - \ln n) + \Theta(n)$ when $k = 2n$.

For upper bounding $\Pr[X \geq 4n]$, we first upper bound by $\Pr[Y \geq 4n]$, where $\{Y_i\}_{i \in [1, n]}$ are n independent geometrically distributed random variables, each with success probability $1/2$, and we write $Y = \sum_{i \in [1, n]} Y_i$. Then we upper bound $\Pr[Y \geq 4n]$ by $\Pr[Z \leq n]$ through the previous trick of turning sum of geometrically distributed random variables into sum of a sequence of Bernoulli trials. We introduce $4n$ independent Bernoulli distributed random variables $\{Z_i\}_{i \in [1, 4n]}$ with $\Pr[Z_i = 1] = 1/2$, and we write $Z = \sum_{i \in [1, 4n]} Z_i$.

By Chernoff’s bound, we have

$$\begin{aligned}\Pr[X \geq 4n] &\leq \Pr[Y \geq 4n] \\ &\leq \Pr[Z \leq n] \\ &\leq \exp(2n(e^t - 1) - nt) = \exp(\varphi(t)),\end{aligned}$$

with $\varphi(t)$ min at $t = -\ln 2$, thus derives an upper bound by $(2/e)^n$. $\square$

Exercise 1.4.4 (Exercise 4.13 [MU17]). Let $\{X_i\}_{i \in [1, n]}$ be independent Poisson trials such that $\Pr[X_i = 1] = p$. Let $X = \sum_{i \in [1, n]} X_i$ so that $\mathbb{E}[X] = pn$. Let

$$F(x, p) = x \ln \frac{x}{p} + (1-x) \ln \frac{1-x}{1-p}.$$

- Show that, for $1 \geq x > p$,

$$\Pr[X \geq xn] \leq e^{-nF(x, p)}.$$

- Show that, when $0 < x, p < 1$, we have $F(x, p) - 2(x-p)^2 \geq 0$.
- Argue that

$$\begin{aligned}\Pr[X \geq (p + \varepsilon)n] &\leq e^{-2n\varepsilon^2}, \\ \Pr[X \leq (p - \varepsilon)n] &\leq e^{-2n\varepsilon^2}, \\ \Pr[|X - pn| \geq \varepsilon n] &\leq 2e^{-2n\varepsilon^2}.\end{aligned}$$

Proof. For Chernoff’s bound on $\Pr[X_i \geq x]$, we typically turn to $1 + kx \leq e^{kx}$ for $k > 0$, then

$$\begin{aligned}\frac{\mathbb{E}[\exp(tX_i)]}{\exp(xt)} &= \frac{p \cdot e^t + 1 - p}{\exp(xt)} = \frac{1 + p \cdot (e^t - 1)}{\exp(xt)} \\ &\leq \exp(p \cdot (e^t - 1) - xt) = \exp(\varphi(t)),\end{aligned}$$

to derive an upper bound by choosing t . A different bound can be derived by avoiding the trick we often apply,

$$\frac{\mathbb{E}[\exp(tX_i)]}{\exp(xt)} = p \cdot \exp((1-x) \cdot t) + (1-p) \cdot \exp(-xt) = \varphi(t),$$

then we have

$$\varphi'(t) = p(1-x) \cdot \exp((1-x) \cdot t) - x(1-p) \cdot \exp(-xt).$$

When $t = \ln(1-p) + \ln x - \ln(1-x) - \ln p > 0$, $\varphi(t)$ min by

$$\begin{aligned}\min_t \varphi(t) &= p \cdot \left(\frac{(1-p)x}{(1-x)p} \right)^{1-x} + (1-p) \cdot \left(\frac{(1-p)x}{(1-x)p} \right)^{-x} \\ &= \left(p \cdot \frac{(1-p)x}{(1-x)p} + 1-p \right) \cdot \left(\frac{(1-p)x}{(1-x)p} \right)^{-x} \\ &= \frac{1-p}{1-x} \cdot \left(\frac{(1-p)x}{(1-x)p} \right)^{-x}.\end{aligned}$$

By taking $\ln$ both sides, we have

$$\begin{aligned}\ln \min_t \varphi(t) &= \ln \frac{1-p}{1-x} - x \ln \frac{(1-p)x}{(1-x)p} \\ &= (1-x) \ln \frac{1-p}{1-x} + x \ln \frac{p}{x} = -F(x, p),\end{aligned}$$

thus deriving a bound for $\Pr[X_i \geq x] \leq \exp(-F(x, p))$. Finally, when bounding $\Pr[X \geq xn]$,

$$\begin{aligned}\Pr[X \geq xn] &\leq \frac{\mathbb{E}[\exp(tX)]}{\exp(xnt)} \\ &= \frac{\prod_{i \in [1, n]} \mathbb{E}[\exp(tX_i)]}{\exp(xnt)} = \left(\frac{\mathbb{E}[\exp(tX_i)]}{\exp(xt)} \right)^n \\ &\leq \exp(-nF(x, p)).\end{aligned}$$

Taking partial derivative against x in $F(x, p)$, we have

$$\frac{\partial}{\partial x} F(x, p) = \ln \frac{x}{p} + x \cdot \frac{p}{x} \cdot \frac{1}{p} - \ln \frac{1-x}{1-p} + (1-x) \cdot \frac{1-p}{1-x} \cdot \frac{-1}{1-p} = \ln \frac{x}{p} - \ln \frac{1-x}{1-p},$$

and it is 0 when $x = p$, taking second order partial derivative against x , we have

$$\frac{\partial^2}{\partial x^2} F(x, p) = \frac{1}{x} - \frac{1-p}{1-x} \cdot \frac{-1}{1-p} = \frac{1}{x} + \frac{1}{1-x} \geq 4,$$

and therefore $F(x, p) \geq F(p, p) + 2(x-p)^2 = 2(x-p)^2$ by Taylor expansion.

We can plug in the result $F(x, p) \geq 2(x-p)^2$, and derive

$$\Pr[X \geq (p + \varepsilon)n] \leq \exp(-nF(p + \varepsilon, p)) \leq \exp(-2n\varepsilon^2).$$

When $x < p$ and we upper bound $\Pr[X_i \leq x]$, $t = \ln(1-p) + \ln x - \ln(1-x) - \ln p < 0$ gives $\ln \min_t \varphi(t) = -F(x, p)$, then $\Pr[X_i \leq x] \leq \exp(-F(x, p))$. Thus,

$$\Pr[X \leq (p - \varepsilon)n] \leq \exp(-nF(p - \varepsilon, p)) \leq \exp(-2n\varepsilon^2).$$

Eventually, we derive bounds on both sides by

$$\Pr[|X - pn| \geq \varepsilon n] \leq 2 \exp(-2n\varepsilon^2).$$

Remark 1.4.5. $F(x, p)$ is exactly the KL divergence [KL51] between Bernoulli distributions with parameters x and p .

□

Exercise 1.4.6 (Exercise 4.14 [MU17]). Let $\{X_i\}_{i \in [1, n]}$ be independent Poisson trials such that $\Pr[X_i] = p_i$, and let $\{a_i\}_{i \in [1, n]}$ be real numbers in $[0, 1]$. Let $X = \sum_{i \in [1, n]} a_i X_i$ and $\mu = \mathbb{E}[X]$. Then the following Chernoff bound holds: for any $\delta > 0$,

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu.$$

Prove a similar bound for the probability that $X \leq (1 - \delta)\mu$ for $0 < \delta < 1$.

Proof. By Jensen's inequality over concave functions,

$$\mathbb{E}[\exp(t a_i X_i)] \leq \mathbb{E}[\exp(t X_i)]^{a_i}$$

for all t , as $f(x) = x^{a_i}$ is concave if $0 \leq a_i \leq 1$.

By Chernoff's upper tail and lower tail bound,

$$\begin{aligned} \Pr[X_i \geq (1 + \delta)p_i] &\leq \frac{\mathbb{E}[\exp(t X_i)]}{\exp((1 + \delta)p_i \cdot t)} \leq \frac{\exp(p_i \cdot (e^t - 1))}{\exp((1 + \delta)p_i \cdot t)}, \text{ min by } t = \ln(1 + \delta) > 0 \text{ by } \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{p_i}, \\ \Pr[X_i \leq (1 - \delta)p_i] &\leq \frac{\mathbb{E}[\exp(t X_i)]}{\exp((1 - \delta)p_i \cdot t)} \leq \frac{\exp(p_i \cdot (e^t - 1))}{\exp((1 - \delta)p_i \cdot t)}, \text{ min by } t = \ln(1 - \delta) < 0 \text{ by } \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{p_i}. \end{aligned}$$

By Chernoff's bound and Jensen's inequality, for upper tail bound, given $t > 0$ and $\delta > 0$,

$$\begin{aligned} \Pr[X \geq (1 + \delta)\mu] &\leq \frac{\mathbb{E}[\exp(t X)]}{\exp((1 + \delta)\mu \cdot t)} = \prod_{i \in [1, n]} \left(\frac{\mathbb{E}[\exp(t a_i X_i)]}{\exp((1 + \delta)a_i p_i \cdot t)} \right) \\ &\leq \prod_{i \in [1, n]} \left(\frac{\mathbb{E}[\exp(t X_i)]}{\exp((1 + \delta)p_i \cdot t)} \right)^{a_i} \\ &\leq \prod_{i \in [1, n]} \left(\frac{\exp(p_i \cdot (e^t - 1))}{\exp((1 + \delta)p_i \cdot t)} \right)^{a_i} = \left(\frac{\exp(e^t - 1)}{\exp((1 + \delta) \cdot t)} \right)^\mu, \end{aligned}$$

and right hand side is minimized by

$$\min_{t > 0} \left(\frac{\exp(e^t - 1)}{\exp((1 + \delta) \cdot t)} \right)^\mu = \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu.$$

Symmetrically, for lower tail bound, given $t < 0$ and $0 < \delta < 1$,

$$\begin{aligned} \Pr[X \leq (1 - \delta)\mu] &\leq \frac{\mathbb{E}[\exp(t X)]}{\exp((1 - \delta)\mu \cdot t)} = \prod_{i \in [1, n]} \left(\frac{\mathbb{E}[\exp(t a_i X_i)]}{\exp((1 - \delta)a_i p_i \cdot t)} \right) \\ &\leq \prod_{i \in [1, n]} \left(\frac{\mathbb{E}[\exp(t X_i)]}{\exp((1 - \delta)p_i \cdot t)} \right)^{a_i} \\ &\leq \prod_{i \in [1, n]} \left(\frac{\exp(p_i \cdot (e^t - 1))}{\exp((1 - \delta)p_i \cdot t)} \right)^{a_i} = \left(\frac{\exp(e^t - 1)}{\exp((1 - \delta) \cdot t)} \right)^\mu, \end{aligned}$$

with right hand side minimized by

$$\min_{t < 0} \left(\frac{\exp(e^t - 1)}{\exp((1 - \delta) \cdot t)} \right)^\mu = \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu.$$

□

Exercise 1.4.7 (Exercise 4.7 [MU17]). Let $X = \sum_{i \in [1, n]} X_i$, where $\{X_i\}_{i \in [1, n]}$ are independent 0-1 random variables. Let $\mu = \mathbb{E}[X]$. Choose any μ_L and μ_H such that $\mu_L \leq \mu \leq \mu_H$. Then, for any $\delta > 0$,

$$\Pr[X \geq (1 + \delta)\mu_H] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu_H}.$$

Similarly, for any $0 < \delta < 1$,

$$\Pr[X \leq (1 - \delta)\mu_L] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\mu_L}.$$

Proof. The MGF of X can be relaxed in following manner

$$\mathbb{E}[\exp(tX)] = \prod_{i \in [1, n]} \mathbb{E}[\exp(tX_i)] \leq \prod_{i \in [1, n]} \exp(p_i \cdot (e^t - 1)) = \exp(\mu \cdot (e^t - 1)).$$

Given $\mu_L \leq \mu \leq \mu_H$,

- When $t > 0$, $\exp(\mu \cdot (e^t - 1)) \leq \exp(\mu_H \cdot (e^t - 1))$.
- When $t < 0$, $\exp(\mu \cdot (e^t - 1)) \leq \exp(\mu_L \cdot (e^t - 1))$.

Then for upper tail with $t > 0$ and $\delta > 0$,

$$\Pr[X \geq (1 + \delta)\mu_H] \leq \frac{\exp(\mu \cdot (e^t - 1))}{\exp((1 + \delta)\mu_H \cdot t)} \leq \frac{\exp(\mu_H \cdot (e^t - 1))}{\exp((1 + \delta)\mu_H \cdot t)} = \left(\frac{\exp(e^t - 1)}{\exp((1 + \delta) \cdot t)} \right)^{\mu_H},$$

the right hand side is minimized by

$$\min_{t > 0} \left(\frac{\exp(e^t - 1)}{\exp((1 + \delta) \cdot t)} \right)^{\mu_H} = \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu_H}.$$

For lower tail with $t < 0$ and $0 < \delta < 1$,

$$\Pr[X \leq (1 - \delta)\mu_L] \leq \frac{\exp(\mu \cdot (e^t - 1))}{\exp((1 - \delta)\mu_L \cdot t)} \leq \frac{\exp(\mu_L \cdot (e^t - 1))}{\exp((1 - \delta)\mu_L \cdot t)} = \left(\frac{\exp(e^t - 1)}{\exp((1 - \delta) \cdot t)} \right)^{\mu_L},$$

the right hand side is minimized by

$$\min_{t < 0} \left(\frac{\exp(e^t - 1)}{\exp((1 - \delta) \cdot t)} \right)^{\mu_L} = \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\mu_L}.$$

Moreover, for lower tail bound,

$$\Pr[X \leq (1 - \delta)\mu_L] \leq \Pr[X \leq (1 - \delta)\mu] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\mu_L}.$$

The first inequality holds from $\mu_L \leq \mu$, the second one holds from Chernoff's bound, and the third one holds from the fact that for $0 < \delta < 1$,

$$\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \leq \exp\left(-\frac{\delta^2}{2}\right) \leq 1.$$

This chain of inequality comes in handy when we cannot compute the exact expectation, yet we can still derive a lower bound for the expectation, and give a looser version of tail bound. The sandwiched inequality chain can even be applied with previous weighted version of Chernoff bound with weights $\{a_i\}_{i \in [1, n]}$ in $(0, 1)$. $\square$

1.4.1 Hoeffding's Bound

Hoeffding's bound says that a sum of independent bounded random variables concentrates around the sum of their expectations. The probability of being far decays exponentially fast in the squared deviation, namely $\exp(-O(\varepsilon^2))$.

Lemma 1.4.8 (Hoeffding's Lemma [Hoe63]). *Let X be a random variable with $X \in [a, b]$ and $\mathbb{E}[X] = 0$, then for any λ ,*

$$\mathbb{E}[\exp(\lambda X)] \leq \exp\left(\frac{\lambda^2(b-a)^2}{8}\right).$$

Proof. Since $f(x) = \exp(\lambda x)$ is convex, then

$$f(\alpha a + (1-\alpha)b) \leq \alpha f(a) + (1-\alpha)f(b).$$

For $x \in [a, b]$, let $x = \alpha a + (1-\alpha)b$, then $\alpha = (b-x)/(b-a)$.

Since $\mathbb{E}[X] = 0$, then by linearity of expectation,

$$\mathbb{E}[\exp(\lambda X)] \leq \mathbb{E}\left[\frac{b-X}{b-a} \exp(\lambda a) + \frac{X-a}{b-a} \exp(\lambda b)\right] = \frac{b}{b-a} \exp(\lambda a) - \frac{a}{b-a} \exp(\lambda b).$$

Let $\theta = -a/(b-a) \in [0, 1]$, and

$$\varphi(t) = -\theta t + \ln(1 - \theta + \theta \exp(t)),$$

then when $t = \lambda(b-a)$,

$$\begin{aligned} \exp(\varphi(\lambda(b-a))) &= \exp(-\theta \lambda(b-a)) \cdot (1 - \theta + \theta \exp(\lambda(b-a))) \\ &= (1 - \theta) \cdot \exp(-\theta \lambda(b-a)) + \theta \cdot \exp((1 - \theta)\lambda(b-a)) \\ &= (1 - \theta) \cdot \exp(\lambda a) + \theta \cdot \exp(\lambda b) \\ &= \frac{b}{b-a} \exp(\lambda a) - \frac{a}{b-a} \exp(\lambda b). \end{aligned}$$

Since $\varphi(0) = 0$,

$$\varphi'(t) = -\theta + \frac{\theta \exp(t)}{1 - \theta + \theta \exp(t)}$$

means $\varphi'(0) = 0$, and

$$\varphi''(t) = \frac{\theta \exp(t)}{1 - \theta + \theta \exp(t)} - \left(\frac{\theta \exp(t)}{1 - \theta + \theta \exp(t)}\right)^2 \leq \frac{1}{4}.$$

By the Taylor expansion and the Taylor remainder theorem in Lagrange form, there is t' between 0 and t such that

$$\varphi(t) = \varphi(0) + \varphi'(0) \cdot t + \frac{\varphi''(t')}{2} \cdot t^2 \leq \frac{t^2}{8}.$$

Therefore,

$$\mathbb{E}[\exp(\lambda X)] \leq \exp(\varphi(\lambda(b-a))) \leq \exp\left(\frac{\lambda^2(b-a)^2}{8}\right).$$

□

Theorem 1.4.9 (Hoeffding's Bound [Hoe63]). *Let $X_1, \dots, X_n$ be independent random variables such that for all $1 \leq i \leq n$, $\mathbb{E}[X_i] = \mu$ and $X_i \in [a, b]$. Then*

$$\Pr\left[\left|\frac{1}{n} \sum_{i \in [1, n]} X_i - \mu\right| \geq \varepsilon\right] \leq 2 \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right).$$

Proof. Let $Z_i = (X_i - \mu)/n$, then $\mathbb{E}[Z_i] = 0$ and $Z_i \in [(a - \mu)/n, (b - \mu)/n]$ for $i \in [1, n]$. Let $Z = \sum_{i \in [1, n]} Z_i$. For $\lambda \geq 0$,

$$\begin{aligned} \Pr[Z \geq \varepsilon] &= \Pr[\exp(\lambda Z) \geq \exp(\lambda \varepsilon)] \\ &\leq \frac{1}{\exp(\lambda \varepsilon)} \cdot \mathbb{E}[\exp(\lambda Z)] = \frac{1}{\exp(\lambda \varepsilon)} \cdot \prod_{i \in [1, n]} \mathbb{E}[\exp(\lambda Z_i)] \\ &\leq \frac{1}{\exp(\lambda \varepsilon)} \cdot \left(\exp\left(\frac{\lambda^2(b-a)^2}{8n^2}\right) \right)^n = \exp\left(\frac{(b-a)^2}{8n} \cdot \lambda^2 - \varepsilon \cdot \lambda\right) \\ &\leq \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right), \end{aligned}$$

where the second inequality is from Markov inequality over MGF, the third equality is from $X_1, \dots$ being independent, the fourth inequality is from lemma 1.4.8, and the last inequality is by minimizing at $\lambda = 4n\varepsilon/(b-a)^2$.

The lower bound argument is similar, for $\lambda < 0$,

$$\begin{aligned} \Pr[Z \leq -\varepsilon] &= \Pr[\exp(\lambda Z) \geq \exp(-\lambda \varepsilon)] \\ &\leq \frac{1}{\exp(-\lambda \varepsilon)} \cdot \mathbb{E}[\exp(\lambda Z)] = \exp(\lambda \varepsilon) \cdot \prod_{i \in [1, n]} \mathbb{E}[\exp(\lambda Z_i)] \\ &\leq \exp\left(\frac{(b-a)^2}{8n} \cdot \lambda^2 + \varepsilon \cdot \lambda\right) \leq \exp\left(-\frac{2n\varepsilon^2}{(b-a)^2}\right), \end{aligned}$$

where the last inequality is by minimizing at $\lambda = -4n\varepsilon/(b-a)^2$. $\square$

Theorem 1.4.10 (Generalized Hoeffding's Bound [Hoe63]). *Let $X_1, \dots, X_n$ be independent random variables such that for all $1 \leq i \leq n$, $\mathbb{E}[X_i] = \mu_i$ and $X_i \in [a_i, b_i]$, then*

$$\Pr\left[\left|\sum_{i \in [1, n]} X_i - \sum_{i \in [1, n]} \mu_i\right| \geq \varepsilon\right] \leq 2 \exp\left(-\frac{2\varepsilon^2}{\sum_{i \in [1, n]} (b_i - a_i)^2}\right).$$

Proof. Let $Z_i = X_i - \mu_i$, then $\mathbb{E}[Z_i] = 0$ and $Z_i \in [a_i - \mu_i, b_i - \mu_i]$ for $i \in [1, n]$. Let $Z = \sum_{i \in [1, n]} Z_i$. For $\lambda \geq 0$,

$$\begin{aligned} \Pr[Z \geq \varepsilon] &= \Pr[\exp(\lambda Z) \geq \exp(\lambda \varepsilon)] \\ &\leq \frac{1}{\exp(\lambda \varepsilon)} \cdot \mathbb{E}[\exp(\lambda Z)] = \frac{1}{\exp(\lambda \varepsilon)} \cdot \prod_{i \in [1, n]} \mathbb{E}[\exp(\lambda Z_i)] \\ &\leq \frac{1}{\exp(\lambda \varepsilon)} \cdot \prod_{i \in [1, n]} \exp\left(\frac{\lambda^2(b_i - a_i)^2}{8}\right) = \exp\left(\sum_{i \in [1, n]} \frac{(b_i - a_i)^2}{8} \cdot \lambda^2 - \varepsilon \cdot \lambda\right) \\ &\leq \exp\left(-\frac{2\varepsilon^2}{\sum_{i \in [1, n]} (b_i - a_i)^2}\right), \end{aligned}$$

where the last inequality is by minimizing at $4\varepsilon/\sum_{i \in [1, n]} (b_i - a_i)^2$.

The lower bound argument is similar, for $\lambda < 0$,

$$\begin{aligned} \Pr[Z \leq -\varepsilon] &= \Pr[\exp(\lambda Z) \geq \exp(-\lambda \varepsilon)] \\ &\leq \frac{1}{\exp(-\lambda \varepsilon)} \cdot \mathbb{E}[\exp(\lambda Z)] = \exp(\lambda \varepsilon) \cdot \prod_{i \in [1, n]} \mathbb{E}[\exp(\lambda Z_i)] \\ &\leq \exp(\lambda \varepsilon) \cdot \prod_{i \in [1, n]} \exp\left(\frac{\lambda^2(b_i - a_i)^2}{8}\right) \leq \exp\left(-\frac{2\varepsilon^2}{\sum_{i \in [1, n]} (b_i - a_i)^2}\right), \end{aligned}$$

where the last inequality is by minimizing at $-4\varepsilon/\sum_{i \in [1, n]} (b_i - a_i)^2$. $\square$

1.4.2 Quicksort Runtime Analysis

Exercise 1.4.11 (Exercise 4.21 [MU17]). We prove that the Randomized Quicksort algorithm sorts a set of n numbers in time $O(n \log n)$ with high probability. Consider the following view of Randomized Quicksort. Every point in the algorithm where it decides on a pivot element is called a *node*. Suppose the size of the set to be sorted at a particular node is s . The node is called *good* if the pivot element divides the set into two parts, each of size not exceeds $2s/3$. Otherwise the node is called *bad*. The nodes can be thought of as forming a tree in which the root node has the whole set to be sorted and its children have the two sets formed after the first pivot step and so on.

- Show that the number of good nodes in any path from the root to a leaf in this tree is not greater than $c \log_2 n$, where c is some positive constant.
- Show that, w.h.p. (greater than $1 - 1/n^2$), the number of nodes in a given root to leaf path of the tree is not greater than $c' \log_2 n$, where c' is another constant.
- Show that, w.h.p. (greater than $1 - 1/n$), the number of nodes in the longest root to leaf path is not greater than $c' \log_2 n$.
- Show that the running time of Quicksort is $O(n \log n)$ with probability at least $1 - 1/n$.

Proof. Suppose a given node with all the internal nodes being good nodes, then the longest such path gives the upper bound on the number of good nodes in any path of the tree. This can be proven by contradiction.

Suppose the path with the largest number of good nodes is not with all good nodes, looking like $\mathbf{u} = \mathbf{v}_0 \parallel c \parallel \mathbf{v}_1$, where c is a not good node. Then the suffix path $c \parallel \mathbf{v}_1$ has to be the path with the largest number of good nodes among the subtree with root c ; otherwise there exists a path $\mathbf{s}$ in subtree with root c with more good nodes, then $\mathbf{u}$ has less good nodes than $\mathbf{v}_0 \parallel \mathbf{s}$, contradicting with $\mathbf{u}$ being the path with the most good nodes among the tree.

We assume $c \parallel \mathbf{v}_1$ being the path with the largest number of good nodes among the subtree rooted by c , then the largest number of good nodes under subtree rooted by c is $|\mathbf{v}_1|$, so is the subtree rooted by $v_{1,0}$. Since the subtree rooted by c as strictly more nodes than the subtree rooted by $v_{1,0}$, then the largest number of good nodes of subtree rooted by c is lower bounded by $|\mathbf{v}_1|$. Thus we complete the proof that the largest number of good nodes from a tree is from the longest path with all nodes being good nodes.

The longest path full of good nodes has length $\log_{3/2} n = \log_2 n \cdot (\log_2 3 - 1)^{-1} = c \log_2 n$, by partition on boundary of good node condition, and choose the big piece side.

One observation: a path has at most $c \log_2 n$ good nodes.

Let Y be the length of a path, and we want to upper bound $\Pr[Y \geq c' \log_2 n]$. By prior observation, the path can have at most $c \log_2 n$ good nodes, and we can turn the bound on path length into the bound on the number of good nodes. If the path survives ℓ levels, then among those first ℓ levels we have seen at most $c \log_2 n$ good nodes (since too many of good nodes stops a path from growing). Therefore $\Pr[Y \geq c' \log_2 n] \leq \Pr[X \leq c \log_2 n]$, where $\{X_i\}_{i \in [1, c' \log_2 n]}$ are independent Bernoulli trials with parameter $1/3$, and $X = \sum_{i \in [1, c' \log_2 n]} X_i$.

We want $c < c'/3$ for lower tail bounding, and c' sufficiently large. By Chernoff's bound,

$$\begin{aligned} \Pr[Y \geq c' \log_2 n] &\leq \Pr[X \leq c \log_2 n] \\ &\leq \exp\left(-\frac{1}{2} \cdot \mu \cdot \delta^2\right) = n^{-\Theta(1)}, \end{aligned}$$

where $\mu = \frac{1}{3} c' \log_2 n$ and $\delta = 1 - \frac{3c}{c'}$.

For bounding the longest path is no longer than $c' \log_2 n$, we union bound against all leaves (at most n), that the probability of existing a path longer than $c' \log_2 n$ is at most $n \cdot 1/n^2$, hence the path is at most $c' \log_2 n$ long with probability $1 - 1/n^2 \cdot n = 1 - 1/n$.

Each row of partition is dominated by runtime in $O(n)$, and the tree height is at most $c' \log_2 n$ with probability at least $1 - 1/n$, hence that the runtime of the algorithm is $O(n \log n)$ with probability greater than $1 - 1/n$. $\square$

1.4.3 Network Routing Problems

Exercise 1.4.12 (Exercise 4.22 [MU17]). Consider the bit-fixing routing algorithm for routing a permutation on the n -cube. Suppose that n is even. Write each source node as $a_s \parallel b_s$, with a_s and b_s of length $n/2$. Let the destination of s 's packet be $b_s \parallel a_s$. Show that this permutation causes the bit-fixing routing algorithm to take $\Omega(\sqrt{N})$ steps.

Proof. Since the permutation is $a_s \parallel b_s \mapsto b_s \parallel a_s$, we can view the algorithm as 2 phases by fixing a_s and b_s . WLOG we analyze the property from phase 1, as the 2 phases are symmetric.

We write address format $a_s \parallel b_s$ in $a_{s,0} \parallel a_{s,1} \parallel b_{s,0} \parallel b_{s,1}$, where $a_s = a_{s,0} \parallel a_{s,1}$ and $b_s = b_{s,0} \parallel b_{s,1}$. During the bit fixing, the packet from $a_s \parallel b_s$ traverses to $b_{s,0} \parallel a_{s,1} \parallel b_{s,0} \parallel b_{s,1}$. Eventually, when all a_s bits are fixed to b_s , the packet arrives $b_s \parallel b_s$.

We notice that there are up to $\sqrt{N}$ such intermediate addresses $b_s \parallel b_s$, yet in the n -cube there are $N = 2^n$ packets. Thus, nodes with addresses $b_s \parallel b_s$ will be traversed by $\sqrt{N}$ packets in total.

Another interesting observation is: the packet routing in such permutation is not “load-balancing”.

View each bit being fixed as a packet traversing from one half-hypercube to another half-hypercube. By prior discussion, we noticed that the first phase bit fixing is “localized”, namely all $\sqrt{N}$ packets from $* \parallel b_s$ go to $b_s \parallel b_s$, and it takes no other packets. To be more generalized, all packets from $* \parallel a_{s,1} \parallel b_{s,0} \parallel b_{s,1}$ go to $b_{s,0} \parallel a_{s,1} \parallel b_{s,0} \parallel b_{s,1}$. The earlier the bit being fixed, the less packets traversing through half-hypercubes. With more bits fixed in the front ($b_{s,0}$) and less bits flexible in the back ($a_{s,1}$), more packets aggregated from sub-hypercubes are forced to take 1 path to merge sub-hypercubes (Think of 8 packets over 3 bit hypercube, the 110 – 111 edge will take 4 packages). Therefore, right before n^{th} bit being fixed, where $n = |a_s|$, there are $\sqrt{N}/2$ packets aggregated for 1 bit fix to $b_s \parallel b_s$, that the edge will be used by $\sqrt{N}/2$ packets, and thus the edge has an $\Omega(\sqrt{N})$ delay.

We conclude that the deterministic bit-fixing algorithm takes $\Omega(\sqrt{N})$ steps to route in this permutation. $\square$

Exercise 1.4.13 (Exercise 4.23 [MU17]). Consider the following modification to the bit-fixing routing algorithm for routing a permutation on the n -cube. Instead of fixing the bits in order from 1 to n , each packet chooses a random order (independent of other packet's choices) and fixes the bits in that order. Show that there is a permutation for which this algorithm requires $2^{\Omega(n)}$ steps with high probability.

Proof. We use the previous permutation $a_s \parallel b_s \mapsto b_s \parallel a_s$, and $2\sqrt{N}$ special cases, $a_s \parallel 0$ and $0 \parallel b_s$, are interesting. For these permutations, we would like to lower bound the number of packets that traverse $0 \parallel 0$.

To reach $0 \parallel 0$, the non-zero side should be fixed to zero first, before any bit on zero side being flipped to non-zero. WLOG we discuss the case that non-zero on the left side, the non-zero on right side follows from symmetry. We introduce 0-1 independent random variables $\{X_{i,j}\}$ where $i \in [1, \frac{n}{2}]$ and $j \in [1, \binom{n/2}{i}]$, indicating if j^{th} packet with i ones traverses $0 \parallel 0$. The routing should fix i one bits first, thus $\Pr[X_{i,j} = 1] = \binom{2i}{i}^{-1}$.

Before moving to the lower bound the packet numbers, we need to derive the expectation of $X = \sum X_{i,j}$,

$$\mathbb{E}[X] = \sum_{i \in [1, \frac{n}{2}]} \sum_{j \in [1, \binom{n/2}{i}]} \mathbb{E}[X_{i,j}] = \sum_{i \in [1, \frac{n}{2}]} \frac{\binom{n/2}{i}}{\binom{2i}{i}} \geq \sum_{i \in [1, \frac{n}{2}]} \frac{1}{4^i} \cdot \binom{n/2}{i} = \left(\frac{5}{4}\right)^{n/2} - 1 = \mu_L,$$

the relaxation on $\binom{2i}{i}$ follows

$$\binom{2i}{i} \leq \sum_{k \in [0, 2i]} \binom{2i}{k} = 2^{2i} = 4^i.$$

The previous lower tail bound inequality chain comes in handy, that

$$\Pr[X \leq (1 - \delta) \cdot \mu] \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}}\right)^\mu \leq \left(\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}}\right)^{\mu_L}.$$

Since $\mu_L = 2^{\Omega(n)}$, we conclude that the algorithm on this permutation has runtime $2^{\Omega(n)}$ time steps. $\square$

Exercise 1.4.14 (Exercise 4.24 [MU17]). Assume we use the randomized routing algorithm for n -cube network to route a total of up to $p2^n$ packets, where each node is the source of no more than p packets, and each node is the destination of no more than p packets.

- Give a high probability bound on the runtime of the algorithm.
- Give a high probability bound on the maximum number of packets at any nodes at any step of the execution of the routing algorithm.

Proof. The analysis follows from the runtime analysis for randomized bit-fixing routing [VB81], but the number of packets increase from 2^n to $p \cdot 2^n$. There are 2 phases, by first routing $p \cdot 2^n$ packets to random intermediate nodes, then route these nodes to their destinations. WLOG we analyze the first phase's runtime, as they are symmetric.

Follow from the proof structure, we first bound the number of “active packets” over certain path, then upper bound the probability of the path with not many “active packets” yet runs slow, finally we derive a probability upper bound by union bounding all paths over the n -cube.

We recall the notion of “active packets”, that a packet is “active” at a node v_{i-1} on the path. If v_{i-1} and v_i are adjacent on path, and diff by j^{th} bit, in order for a packet to be “active”, its bits should be fixed before j^{th} bit when it reaches v_{i-1} . With that said, there are $p2^{j-1}$ packets sharing same suffix as v_{i-1} from j^{th} bit, while the probability of sampling a random intermediate address with prefixed $j-1$ bits matching v_{i-1} is $2^{-(j-1)}$. Thus, the expected number of “active packets” at a node on an edge is p .

Given a path of length m , the path is at most n as there are n bits to fix. Let X be the number of “active packets” on this path, then the expected number of “active packets” is upper bounded by $\mathbb{E}[X] = mp \leq np$. By Chernoff's bound, the probability of having too many “active packets” is upper bounded as follows

$$\Pr[X \geq 6np \geq 6mp] \leq 2^{-6np}.$$

By conditional probability, we derive the following upper bound by sums of (conditional) probabilities

$$\begin{aligned} \Pr[A] &= \Pr[A \mid B] \cdot \Pr[B] + \Pr[A \mid \neg B] \cdot \Pr[\neg B] \\ &\leq \Pr[B] + \Pr[A \mid \neg B]. \end{aligned}$$

We upper bound the probability of the runtime of the path being at least $30np$ time steps by either “too many active packets” or “not many active packets but running slow” as follows

$$\begin{aligned} \Pr[T \geq 30np] &= \Pr[T \geq 30np \mid X \geq 6np] \cdot \Pr[X \geq 6np] + \Pr[T \geq 30np \mid X < 6np] \cdot \Pr[X < 6np] \\ &\leq \Pr[X \geq 6np] + \Pr[T \geq 30np \mid X < 6np]. \end{aligned}$$

For “not many active packets but running slow” case, the runtime of the path is measuring how many packets stay on the path, turns into the time steps used to route all the packets. Only “active packets” can stay on the path, with probability $1/2$ to branch into the next node on path. An “active packet”'s stay on the path can be modeled by the number of failures in a geometrically distributed random variable, and it leaves the path on a successful trial with probability $1/2$. Upper bound on the sum of geometrically distributed random variables can be relaxed by upper bounding the sum of a serial of Bernoulli trials with parameter $1/2$ here

$$\Pr[T \geq 30np \mid X < 6np] \leq \Pr\left[\sum_{i \in [1, 36np]} Y_i < 6np\right] \leq \exp\left(-\frac{1}{2} \cdot 18np \cdot \left(\frac{2}{3}\right)^2\right) = \exp(-4np) \leq 2^{-3np-1}.$$

Therefore, the probability for the path's runtime being at least $30np$ is upper bounded by 2^{-3np}

$$\Pr[T \geq 30np] \leq 2^{-6np} + 2^{-3np-1} \leq 2^{-3np}.$$

Union bounding over 2^{2n} paths, the probability of the first phase runtime exceeding $30np$ is at most 2^{2n-3np} .

On the number of packets at any nodes at any time step, we can upper bound by the number of packets that will traverse the node. Let X_i be the number of packets arriving the node on i^{th} bit, $X = \sum_{i \in [1, n]} X_i$ be the number of packets using the node. Since $\mathbb{E}[X_i] = p$ by $p2^i$ source packets and 2^{-i} probability in random address prefix matching the node's address's prefix, by linearity of expectation, $\mathbb{E}[X] = np$.

Once we have an expectation, by Chernoff's bound we have $\Pr[X \geq 6np] \leq 2^{-6np}$ for this node at a certain step. By union bound over all nodes and time steps, the queue exceeding $6np$ has probability at most $2^{-6np} \cdot 2^n \cdot O(np)$. $\square$

Exercise 1.4.15 (Exercise 4.26 [MU17]). Given a network that is an undirected graph G , where nodes represent processors and the edges between the nodes represent wires. We are also given a set of N packets to route. For each packet we are given a source node, a destination node, and the exact route that the packet should take from source to destination. In each time step, at most one packet can traverse an edge. A packet can wait at any node during any time step, and we assume unbounded queue sizes at each node.

A schedule for a set of packets specifies the timing for the movement of packets along their respective routes. That is, it specifies which packet should move and which should wait at each time step. Our goal is to produce a schedule for the packets that tries to minimize the total time and the maximum queue size needed to route all the packets to their destination.

- The dilation d is the maximum distance traveled by any packet. The congestion c is the maximum number of packets that must traverses a single edge during the entire course of the routing. Argue that the time required for any schedule should be at least $\Omega(c + d)$.
- Consider the following unconstrained schedule, where many packets may traverse an edge during a single time step. Assign each packet an integral delay x , chosen randomly, independently, and uniformly from the interval $[1, \lceil \alpha c / \log(Nd) \rceil]$, where α is a constant. A packet waits in its source node for x time steps, then it moves on to its final destination through its specified route without ever stopping. Give an upper bound on the probability that more than $O(\log(Nd))$ packets use a particular edge e at a particular time step t .
- Again using the unconstrained schedule, show that the probability that more than $O(\log(Nd))$ packets pass through any edge at any time step is at most $1/(Nd)$ for a positive constant α .
- Use the unconstrained schedule to devise a simple randomized algorithm that, with high probability, produces a schedule of length $O(c + d \log(Nd))$ using queues of size $O(\log(Nd))$ and following the constraint that at most one packet crosses an edge per time step.

Proof. The length of the schedule is the runtime of the routing algorithm, and it is lower bounded by: the maximum distance a packet can traverse in the network, and the maximum congestion that delays the package traversal. Therefore, let T denotes the time steps of the routing algorithm, and $T \geq \max(c, d)$, and we conclude that $T = \Omega(c + d)$.

Considering how many packets traverse edge e at the t^{th} time step, we introduce $\{X_i\}_{i \in [1, N]}$ independent 0-1 random variables, where $X_i = 1$ stands for i^{th} packet traverses e at t^{th} time step. Since there are at most $c \leq N$ packets traverse e , and if they traverse e at the t^{th} time step, they had to be starting by one of the $\lceil \alpha c / \log Nd \rceil$ delays, we here upper bound the expectation of $X = \sum_{i \in [1, N]} X_i$ by

$$\mathbb{E}[X] \leq c \cdot \frac{\log(Nd)}{\alpha c} = \frac{\log(Nd)}{\alpha}.$$

By exercise 1.4.7, given $\mu_H \geq \mu$,

$$\Pr[X \geq (1 + \delta)\mu_H] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu_H} \leq \exp(-\mu_H \delta^2 / 3),$$

we have $X \geq (1 + \delta) \log(Nd) / \alpha$ probability upper bounded by $(Nd)^{-\delta^2 / 3\alpha}$.

Bounding the number of packets traverses any edge at any time step is equivalent to bounding the congestion of the schedule, which can be achieved by union bounding through all the time steps and all the edges. Each packet can go as far as d edges, and the packet arrives destination with at most $d + \lceil \alpha c / \log(Nd) \rceil$ time steps (by assumption that it does not stop once). If $\delta = 1$ and $\alpha < 1/9$, the upper bounded probability is

$$\begin{aligned} Nd \cdot \left(d + \left\lceil \frac{\alpha c}{\log Nd} \right\rceil \right) \cdot (Nd)^{-\frac{\delta^2}{3\alpha}} &\leq Nd \cdot Nd \cdot (Nd)^{-\frac{\delta^2}{3\alpha}} \\ &\leq (Nd)^2 \cdot (Nd)^{-3} = (Nd)^{-1}. \end{aligned}$$

If the tuned algorithm has schedule of length $O(d \log(Nd) + c)$ with queue size $O(\log(Nd))$, try scale the schedule by $2 \log(Nd)/\alpha$. By previous result, with the assumption that multiple packets crossing an edge at a time with no stop, the probability of more than $2 \log(Nd)/\alpha$ packets passing through any edge at any time step is at most $(Nd)^{-1}$. In the prior model, each time step the edge let packets crossing in a batch, and the batch is at most $2 \log(Nd)/\alpha$.

Now removing the assumption of multiple packets crossing an edge at a time, by scaling $2 \log(Nd)/\alpha$ to the schedule and queue of size $O(\log(Nd))$, given an edge allowing only 1 packet at a time, the batched moving can be simulated with the devised algorithm. $\square$

Remark 1.4.16. The intuition is: we first let everyone run non-stop with random starts, count how bad the pileups get, then spread those pileups thin by stretching time proportionally, and leading to an average case non stop routing. This is the Leighton-Maggs-Rao trick [LMR94].

1.5 Balls, Bins, and Random Graphs

1.5.1 A few things I should have noted about e

We begin with the definition of e and e^{-1}

$$\lim_{n \rightarrow \infty} \left(1 + \frac{1}{n}\right)^n = e, \quad \lim_{n \rightarrow \infty} \left(1 - \frac{1}{n}\right)^n = \frac{1}{e}.$$

Both $f(n) = (1 + 1/n)^n$ and $g(n) = (1 - 1/n)^n$ are monotone, proven as follows:

- Let $\varphi_f(n) = n \ln(1 + 1/n)$, then the first order derivative

$$\varphi'_f(n) = \ln\left(1 + \frac{1}{n}\right) + n \cdot \frac{1}{1 + 1/n} \cdot \left(-\frac{1}{n^2}\right) = \ln\left(1 + \frac{1}{n}\right) - \frac{1}{1 + n}.$$

Let $t = 1 + 1/n > 1$, then

$$\varphi'_f(n) = \ln\left(1 + \frac{1}{n}\right) - \frac{1}{1 + n} = \ln\left(1 + \frac{1}{n}\right) + \frac{n}{1 + n} - 1 = \ln t + \frac{1}{t} - 1.$$

As $t > 1$ and by integral,

$$\ln t = \int_1^t \frac{1}{x} dx > (t - 1) \cdot \frac{1}{t},$$

we have $\varphi'_f(n) > 0$, and thus $\varphi_f(n)$ is monotone increasing.

- Similarly, let $\varphi_g(n) = n \ln(1 - 1/n)$, then the first order derivative

$$\varphi'_g(n) = \ln\left(1 - \frac{1}{n}\right) + n \cdot \frac{1}{1 - 1/n} \cdot \left(-\frac{-1}{n^2}\right) = \ln\left(1 - \frac{1}{n}\right) + \frac{1}{n - 1}.$$

Let $t = 1 - 1/n$ that $0 < t < 1$, then

$$\varphi'_g(n) = \ln\left(1 - \frac{1}{n}\right) + \frac{1}{n - 1} = \ln\left(1 - \frac{1}{n}\right) + \frac{n}{n - 1} - 1 = \ln t + \frac{1}{t} - 1.$$

As $0 < t < 1$ and by integral again,

$$\ln t = \int_1^t \frac{1}{x} dx = - \int_t^1 \frac{1}{x} dx > -(1 - t) \cdot \frac{1}{t},$$

we have $\varphi'_g(n) > 0$, and thus $\varphi_g(n)$ is monotone increasing.

Moreover, by the Taylor expansion

$$e^x = \sum_{i \geq 0} \frac{x^i}{i!},$$

then $e^x \geq 1 + x$ for all x . We derive following handy inequality from the Taylor expansion.

Lemma 1.5.1. For integer $x \geq 0$ and $n \geq 0$, we have

$$\frac{1}{x!} \leq \frac{e^n}{n^x}.$$

Proof. By Taylor expansion, since

$$\frac{n^x}{x!} \leq \sum_{m \geq 0} \frac{n^m}{m!} = e^n,$$

therefore

$$\frac{1}{x!} \leq \frac{e^n}{n^x}.$$

□

Exercise 1.5.2 (Exercise 5.7 [MU17]). Use the Taylor expansion

$$\ln(1+x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \dots$$

to prove that, for any $|x| \leq 1$, $e^x(1-x^2) \leq 1+x \leq e^x$.

Proof. We prove the left side $e^x(1-x^2) \leq 1+x$ first. Take $\ln$ on both sides, we have left side $x + \ln(1+x) + \ln(1-x)$.

Now by the Taylor expansion, for $|x| \leq 1$, we have

$$\begin{aligned} x + \ln(1+x) + \ln(1-x) &= x + x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} \dots - x - \frac{x^2}{2} - \frac{x^3}{3} - \frac{x^4}{4} \dots \\ &= x - 2 \cdot \frac{x^2}{2} - 2 \cdot \frac{x^4}{4} \dots \\ &\leq x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \dots = \ln(1+x). \end{aligned}$$

The right side holds from the Taylor expansion of e^x , and we conclude that for $|x| \leq 1$, $e^x(1-x^2) \leq 1+x \leq e^x$. $\square$

Lemma 1.5.3. For any integer $k \geq 1$, by concavity of $\ln x$,

$$\frac{1}{2} (\ln k + \ln(k+1)) \leq \int_k^{k+1} \ln x dx.$$

Lemma 1.5.4 (Lemma 5.8 [MU17]).

$$n! \leq e\sqrt{n} \left(\frac{n}{e}\right)^n.$$

Proof. Since $\ln n! = \sum_{i \in [1, n]} \ln i$, and we have the following from lemma 1.5.3

$$\sum_{i \in [1, n]} \ln i \leq \int_1^n \ln x dx + \frac{1}{2} \ln n = n \ln n - n + 1 + \frac{1}{2} \ln n,$$

which completes the proof by exponentiating over e on both sides. $\square$

Corollary 1.5.5. By lemma 1.5.1 and lemma 1.5.4, we have

$$\left(\frac{n}{e}\right)^n \leq n! \leq e\sqrt{n} \cdot \left(\frac{n}{e}\right)^n \text{ and } \frac{1}{e\sqrt{n}} \cdot \left(\frac{e}{n}\right)^n \leq \frac{1}{n!} \leq \left(\frac{e}{n}\right)^n.$$

1.5.2 Birthday Paradox and Balls-and-Bins model

Suppose there are n possible birthdays, and there are m people in a room. Then, the probability that all m are having different birthdays is

$$\prod_{j \in [0, m-1]} \left(1 - \frac{j}{n}\right) \leq \prod_{j \in [0, m-1]} \exp\left(-\frac{j}{n}\right) \leq \exp\left(-\frac{m^2}{2n}\right).$$

Hence, when $m = \sqrt{2n \ln 2}$, the probability goes to $1/2$.

But through union bound, we can also give a loose bound as follows. Consider event E_k being the k^{th} people failed to have distinct birthday from all previous $k-1$ people. Then $\Pr[E_k] = (k-1)/n$. Through union bound, the probability of the first k people failed to have distinct birthdays is upper bounded by

$$\Pr\left[\bigcup_{i \in [1, k]} E_i\right] \leq \sum_{i \in [1, k]} \Pr[E_i] = \frac{k(k-1)}{2n}.$$

The upper bound indicates that when $k = \lceil \sqrt{n} \rceil$, the probability of having distinct birthdays is at most $1/2$.

Moreover, assuming that first $\lceil \sqrt{n} \rceil$ birthdays are distinct, then the next $\lceil \sqrt{n} \rceil$ birthdays are different from the first $\lceil \sqrt{n} \rceil$ birthdays with probability upper bounded by

$$\left(1 - \frac{\lceil \sqrt{n} \rceil}{n}\right)^{\lceil \sqrt{n} \rceil} \leq \frac{1}{e}.$$

Therefore, once there are $2\lceil \sqrt{n} \rceil$ people out there, the probability of having all distinct birthdays is at most $1/e < 1/2$.

Remark 1.5.6. The idea of union bound is used throughout the Poisson approximation and balls-and-bins model later on, after introduction of theorem 1.5.15.

Exercise 1.5.7 (Exercise 5.4 [MU17]). Consider the extended birthday paradox, but with no 3 having same birthdays.

Proof. Let there be k people, and there be n days in a year.

Suppose there are j pairs of same birthdays, then we first consider the number of pair choices out of k people. We can choose 2 people continuously from k people, then rule out all the duplicate choices, as we don't care about the ordering, and thus we have number of choices in total

$$\binom{k}{2} \binom{k-2}{2} \cdots \binom{k-2j+2}{2} \cdot \frac{1}{j!} = \frac{k!}{2^j(k-2j)!} \cdot \frac{1}{j!} = \binom{k}{2; \dots; 2} \cdot \frac{1}{j!}.$$

For each choice, there are $k-j$ days to be distinct, the rest j people should have same birthdays as their pairs, leading to probability

$$\frac{1}{n^j} \cdot \prod_{i \in [1, k-j]} \left(1 - \frac{i-1}{n}\right).$$

Now that the probability p_j with j pairs of same birthdays is

$$p_j = \frac{k!}{2^j(k-2j)!} \cdot \frac{1}{j!} \cdot \frac{1}{n^j} \cdot \prod_{i \in [1, k-j]} \left(1 - \frac{i-1}{n}\right),$$

summing over p_j over $j \in [0, \lfloor \frac{k}{2} \rfloor]$ is the exact probability. $\square$

Exercise 1.5.8 (Exercise 5.8 [MU17]). n balls are thrown independently and uniformly at random into n bins.

- Find the conditional probability that bin 1 has one ball given that exactly one ball fell into the first three bins.
- Find the conditional expectation of the number of balls in bin 1 under the condition that bin 2 received no balls.
- Write an expression for the probability that bin 1 receives more balls than bin 2.

Proof. If conditioned that there are only 1 ball in first 3 bins, then the probability that a ball is in bin 1 is $1/3$.

We introduce independent Bernoulli trials $\{X_i\}_{i \in [1, n]}$ with probability $1/n$ indicating if the i^{th} ball falls in bin 1. Then the conditional probability of $X_i = 1$ given that the i^{th} ball does not fall into bin 2 is

$$\Pr[X_i = 1 \mid i^{\text{th}} \text{ ball does not fall in bin 2}] = \frac{1}{n-1}.$$

Therefore, the conditional expectation of $X = \sum_{i \in [1, n]} X_i$ under the condition that bin 2 has no balls is

$$\begin{aligned} \mathbb{E}[X \mid \text{bin 2 has no balls}] &= \sum_{i \in [1, n]} \mathbb{E}[X_i \mid \text{bin 2 has no balls}] \\ &= \sum_{i \in [1, n]} 1 \cdot \Pr[X_i = 1 \mid i^{\text{th}} \text{ ball does not fall in bin 2}] = \frac{n}{n-1}. \end{aligned}$$

There are only 3 cases related to balls in bin 1 and bin 2, one of more, less, or equal. By symmetry, former 2 cases are equal. Therefore, we need only the equal case to derive the exact result.

We choose $k \in [0, \lfloor \frac{n}{2} \rfloor]$ balls each for bin 1 and bin 2, and the probability is given by

$$p_{e,k} = \binom{n}{k; k} \left(\frac{1}{n}\right)^k \left(\frac{1}{n}\right)^k \left(1 - \frac{2}{n}\right)^{n-2k} = \binom{n}{k; k} \left(\frac{1}{n}\right)^{2k} \left(1 - \frac{2}{n}\right)^{n-2k}.$$

Therefore, the exact probability for the equal case is $p_e = \sum_k p_{e,k}$, and the desired probability is $(1 - p_e)/2$. $\square$

Lemma 1.5.9 (Lemma 5.1 [MU17]). *When n balls are thrown independently and uniformly at random into n bins, the probability that the maximum load is more than $3 \ln n / \ln \ln n$ is at most $1/n$ for n sufficiently large.*

Proof. Suppose for the i^{th} bin receives at least M balls, let it be event E_i , then the probability is

$$\Pr[E_i] = \binom{n}{M} \frac{1}{n^M}.$$

If any bins are receiving with at least M balls, by union bound, the probability is upper bounded by

$$\Pr\left[\bigcup_{i \in [1, n]} E_i\right] \leq \sum_{i \in [1, n]} \Pr[E_i] = n \binom{n}{M} \frac{1}{n^M}.$$

Since

$$\binom{n}{M} \frac{1}{n^M} \leq \frac{1}{M!} \leq \left(\frac{e}{M}\right)^M,$$

where second inequality follows from lemma 1.5.1, the probability is then upper bounded by $n (e/M)^M$.

Choose $M = 3 \ln n / \ln \ln n$, then when n is sufficiently large, the probability upper bound is at most $1/n$. $\square$

Exercise 1.5.10 (Exercise 5.12 [MU17]). The following problem models a simple distributed system wherein agents contend for resources but “back off” in the face of contention. Balls represent agents, and bins represent resources.

The system evolves over rounds. Every round, balls are thrown independently and uniformly at random into n bins. Any ball that lands in a bin by itself is served and removed from consideration. The remaining balls are thrown again in the next round. We begin with n balls in the first round, and we wish when every ball is served.

- If there are b balls at the start of a round, what is the expected number of balls at the start of the next round?
- Suppose that every round the number of balls served was exactly the expected number of balls to be served. Show that all the balls would be served in $O(\log \log n)$ rounds. (Hint: If x_j is the expected number of balls left after j rounds, show and use that $x_{j+1} \leq x_j^2/n$.)

Proof. We first consider the expected number of remaining balls. Only the bin with one and only one ball can serve the ball, and thus we can disregard the ball onwards. We introduce n independent 0-1 random variables $\{X_i\}_{i \in [1, n]}$ for i^{th} bin has one and only one ball to serve. Therefore, the probability for a bin with only one ball is

$$\Pr[X_i = 1] = \binom{b}{1} \frac{1}{n} \left(1 - \frac{1}{n}\right)^{b-1} = \frac{b}{n} \cdot \left(1 - \frac{1}{n}\right)^{b-1},$$

where b is the number of balls remaining by now.

Let $X = \sum_{i \in [1, n]} X_i$ be the number of bins can serve (or balls to disregard), by linearity of expectation, we have

$$\mathbb{E}[X] = n \cdot \frac{b}{n} \cdot \left(1 - \frac{1}{n}\right)^{b-1} = b \cdot \left(1 - \frac{1}{n}\right)^{b-1}.$$

Therefore, the remaining balls next round is $b - b \cdot (1 - 1/n)^{b-1}$.

Let b be the remaining balls to serve after round j , then $x_{j+1} = b - b \cdot (1 - 1/n)^{b-1}$. Since

$$\left(1 - \frac{1}{n}\right)^{b-1} \geq 1 - \frac{b-1}{n},$$

we have $x_{j+1} \leq b(b-1)/n$, and therefore $x_{j+1} \leq x_j^2/n$.

Now we write $y_j = x_j/n$, and therefore $y_{j+1} \leq y_j^2$, and we want to know the number of rounds to reach a $y = 1/n$. Suppose $y_j = 1/2$ for some constant number of rounds j , we count the number of additional rounds r until we have

$$y_j^{2^r} = \left(\frac{1}{2}\right)^{2^r} \leq \frac{1}{n}.$$

Therefore $r = O(\log \log n)$, that all balls will be served in $O(\log \log n)$ rounds. $\square$

1.5.3 Poisson Distribution

Exercise 1.5.11 (Exercise 5.6 [MU17]). Let X be a Poisson random variable with mean μ , representing the number of a page of this book. Each error is independently a grammatical error with probability p and a spelling error with probability $1 - p$. If Y and Z are random variables representing the number of grammatical and spelling error (respectively) on a page of this book, prove that Y and Z are Poisson random variables with means μp and $\mu(1 - p)$, respectively. Also, prove that Y and Z are independent.

Proof. The sum of a finite number of independent Poisson random variables is a Poisson random variable, by looking at the MGFs. Now think in a reverse way, that a Poisson random variable can be thinned down to a finite number of independent Poisson random variables, which is the Poisson Thinning (Splitting) Property.

Let a Poisson variable X with mean μ represent the errors of a page, then Y be the random variable represent the grammatical errors with probability p in these errors. The probability of $Y = k$ is the summation of all conditional probability conditioned over X by

$$\begin{aligned} \Pr[Y = k] &= \sum_{x \geq k} \Pr[Y = k \cap X = x] \\ &= \sum_{x \geq k} \Pr[Y = k \mid X = x] \cdot \Pr[X = x] \\ &= \sum_{x \geq k} \binom{x}{k} p^k (1-p)^{x-k} \cdot \frac{e^{-\mu} \mu^x}{x!} = \frac{e^{-\mu} p^k \mu^k}{k!} \sum_{x \geq k} (1-p)^{x-k} \cdot \frac{\mu^{x-k}}{(x-k)!} \\ &= \frac{e^{-\mu} p^k \mu^k}{k!} \cdot e^{(1-p)\mu} = \frac{e^{-p\mu} (p\mu)^k}{k!}, \end{aligned}$$

which is the probability of a Poisson random variable with mean $p\mu$ being k . By symmetry, Z is a Poisson random variable with mean $(1-p)\mu$ by substituting p with $1-p$.

For Y and Z being independent, we prove by the property of independence as follows:

$$\begin{aligned} \Pr[Y = y \cap Z = z] &= \Pr[Y = y \cap Z = z \mid X = y+z] \cdot \Pr[X = y+z] \\ &= \binom{y+z}{y} p^y (1-p)^z \cdot \frac{e^{-\mu} \mu^{y+z}}{(y+z)!} = p^y (1-p)^z \cdot \frac{e^{-\mu} \mu^{y+z}}{y! z!} \\ &= \frac{e^{-p\mu} p^y \mu^y}{y!} \cdot \frac{e^{-(1-p)\mu} (1-p)^z \mu^z}{z!} \\ &= \Pr[Y = y] \cdot \Pr[Z = z]. \end{aligned}$$

$\square$

Lemma 1.5.12 (Exercise 5.14 [MU17]). *If Z is a Poisson random variable of mean $\mu \geq 1$ integer, then $\Pr[Z \geq \mu] \geq 1/2$.*

- Show that $\Pr[Z = \mu + h] \geq \Pr[Z = \mu - h - 1]$ for $0 \leq h \leq \mu - 1$.
- Argue that $\Pr[Z \geq \mu] \geq 1/2$.

Proof. We write $\Pr[Z = \mu + h] = \delta_h \cdot \Pr[Z = \mu - h - 1]$, where $h \in [0, \mu - 1]$. Therefore,

$$\delta_h = \prod_{j \in [\mu-h, \mu-1]} \frac{\mu}{j} \cdot \prod_{j \in [\mu+1, \mu+h]} \frac{\mu}{j} = \prod_{k \in [1, h]} \frac{\mu^2}{(\mu-k)(\mu+k)} = \prod_{k \in [1, h]} \alpha_k.$$

As $\alpha_k \geq 1$ for $k \in [1, h]$, therefore $\delta_h \geq 1$. Thus $\Pr[Z \leq \mu - 1] \leq \Pr[Z \geq \mu]$, and $\Pr[Z \geq \mu] \geq 1/2$.

Moreover, since $\delta_0 = 1$, therefore $\Pr[Z = \mu - 1] = \Pr[Z = \mu]$. □

Theorem 1.5.13 (Poisson Median Range Theorem [Cho94]). *Let Z be a Poisson random variable of mean $\mu > 0$, and m be the median of Z . Then*

$$\mu - \ln 2 \leq m \leq \mu + \frac{1}{3}.$$

If μ is an integer such that $\mu \geq 1$, then $\mu = m$, and thus $\Pr[Z \geq \mu] \geq 1/2$ and $\Pr[Z \leq \mu] \geq 1/2$.

1.5.4 Poisson Approximation

Theorem 1.5.14 (Theorem 5.6 [MU17]). *Let $\{X_i^{(m)}\}_{i \in [1, n]}$ be the number of balls in each bin of the m balls and n bins model, and $\{Y_i^{(m)}\}_{i \in [1, n]}$ be independent Poisson random variables of mean m/n . The distribution of $(X_1^{(k)}, \dots, X_n^{(k)})$ is the same as the distribution of $(Y_1^{(m)}, \dots, Y_n^{(m)})$ conditioned on $Y^{(m)} = \sum_{i \in [1, n]} Y_i^{(m)} = k$, regardless of the value of m .*

Theorem 1.5.15 (Theorem 5.7 [MU17]). *Let $f(x_1, \dots, x_n)$ be a nonnegative function. Then*

$$\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \leq e\sqrt{m} \cdot \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right].$$

Proof. By conditional expectation and theorem 1.5.14, we have

$$\begin{aligned} \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right] &= \sum_{k \geq 0} \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \mid Y^{(m)} = k \right] \Pr[Y^{(m)} = k] \\ &= \sum_{k \geq 0} \mathbb{E} \left[f \left(X_1^{(k)}, \dots, X_n^{(k)} \right) \right] \Pr[Y^{(m)} = k] \\ &\geq \mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \Pr[Y^{(m)} = m]. \end{aligned}$$

By definition of the Poisson random variable, $Y^{(m)}$ is a Poisson random variable of mean $\mu = m$, thus

$$\Pr[Y^{(m)} = m] = e^{-m} \cdot \frac{m^m}{m!}.$$

By lemma 1.5.4, $\Pr[Y^{(m)} = m] \geq (e\sqrt{m})^{-1}$. Therefore, $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \leq e\sqrt{m} \cdot \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right]$. □

Corollary 1.5.16. *Let E be an event. If E has probability p in the Poisson case, where there are n independent Poisson random variables of mean m/n , then E has probability at most $e\sqrt{m} \cdot p$ in the exact m balls and n bins case.*

Proof. Immediately from theorem 1.5.15. □

Theorem 1.5.17 (Theorem 5.10 [MU17]). *Let $f(x_1, \dots, x_n)$ be a nonnegative function such that $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ is either monotonically increasing or monotonically decreasing in m . Then*

$$\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \leq 2 \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right].$$

Exercise 1.5.18 (Exercise 5.15 [MU17]). *Prove that if $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ is monotonically increasing in m , then*

$$\mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right] \geq \mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \Pr \left[Y^{(m)} \geq m \right],$$

again under the condition that f is nonnegative.

Make a similar statement for the case when $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ is monotonically decreasing in m .

Moreover, prove theorem 1.5.17 for the case that $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ is monotonically increasing in m .

Proof. We derive the chain of inequalities as follows

$$\begin{aligned} \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right] &= \sum_{k \geq 0} \mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \mid Y^{(m)} = k \right] \Pr \left[Y^{(m)} = k \right] \\ &= \sum_{k \geq 0} \mathbb{E} \left[f \left(X_1^{(k)}, \dots, X_n^{(k)} \right) \right] \Pr \left[Y^{(m)} = k \right] \\ &\geq \sum_{k \geq m} \mathbb{E} \left[f \left(X_1^{(k)}, \dots, X_n^{(k)} \right) \right] \Pr \left[Y^{(m)} = k \right] \\ &\geq \sum_{k \geq m} \mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \Pr \left[Y^{(m)} = k \right] \\ &= \mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \Pr \left[Y^{(m)} \geq m \right]. \end{aligned}$$

The first equality holds from conditional expectation, the second equality holds from theorem 1.5.14, the third holds from nonnegativity of f , the fourth holds from $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ being monotonically increasing in m .

Similarly, if f is nonnegative and $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ is monotonically decreasing in m , the following holds

$$\mathbb{E} \left[f \left(Y_1^{(m)}, \dots, Y_n^{(m)} \right) \right] \geq \mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right] \Pr \left[Y^{(m)} \leq m \right].$$

By lemma 1.5.12, theorem 1.5.13, and $Y^{(m)}$ being a Poisson random variable of mean $\mu = m \geq 1$ integer, therefore we prove theorem 1.5.17 for the $\mathbb{E} \left[f \left(X_1^{(m)}, \dots, X_n^{(m)} \right) \right]$ being monotonically increasing/decreasing in m . $\square$

Corollary 1.5.19 (Corollary 5.11 [MU17]). *Let E be an event whose probability is monotonically increasing or decreasing in the number of balls. If E has probability p in the Poisson case, then E has probability at most $2p$ in the exact case.*

Proof. Immediately follow from theorem 1.5.17. $\square$

Lemma 1.5.20 (Lemma 5.12 [MU17]). *When n balls are thrown independently and uniformly at random into n bins, the max load is at least $\ln n / \ln \ln n$ with probability at least $1 - 1/n$ for n sufficiently large.*

Proof. In the Poisson case, there are n independent Poisson random variables $\{Z_i\}_{i \in [1, n]}$ of mean 1.

If $Z_i \geq M$ for some load M , then

$$\Pr[Z_i < M] < 1 - \Pr[Z_i = M] = 1 - \frac{1}{eM!}.$$

Considering all n Poisson random variables, then the probability of all loads being lower than M is

$$\Pr \left[\bigcap_{i \in [1, n]} Z_i < M \right] < \left(1 - \frac{1}{eM!} \right)^n \leq \exp \left(-\frac{n}{eM!} \right).$$

From corollary 1.5.16 or corollary 1.5.19, we want the probability upper bound to be less than $1/n^2$, as $e\sqrt{n}/n^2 = o(1)$. By lemma 1.5.4 and relax by sufficiently large n and M by

$$M! \leq e\sqrt{M} \left(\frac{M}{e} \right)^M \leq M \left(\frac{M}{e} \right)^M,$$

therefore let $M = \ln n / \ln \ln n$, with less than $1/n$ probability in exact case the max load is no greater than M . $\square$

Remark 1.5.21. For n balls and n bins, the maximum load is $\Theta(\log n / \log \log n)$ from lemma 1.5.9 and lemma 1.5.20 with probability approaches 1 as $n \rightarrow \infty$.

Exercise 1.5.22 (Exercise 5.13 [MU17]). Suppose that we vary the balls-and-bins process as follows. For convenience let the bins be numbered from 0 to $n-1$. There are $\log_2 n$ players. Each player choose a starting position $\ell \xleftarrow{\mathcal{R}} [0, n-1]$, then places one ball in each of the bins numbered $\ell \bmod n$ to $\ell + n/\log_2 n - 1 \bmod n$. Argue that the maximum load in this case is only $O(\log \log n / \log \log \log n)$ with probability that approaches 1 as $n \rightarrow \infty$.

Proof. For the j^{th} player and the i^{th} bin, we introduce an independent 0-1 random variable $X_{i,j}$, indicating if the j^{th} player puts ball in the i^{th} bin. $\Pr[X_{i,j} = 1] = 1/\log_2 n$, as the j^{th} player starts at most $n/\log_2 n$ bins before the i^{th} bin. Therefore, $X_i = \sum_{j \in [1, \log_2 n]} X_{i,j}$ is the load of the i^{th} bin, and by linearity of expectation, $\mathbb{E}[X_i] = \sum_j \mathbb{E}[X_{i,j}] = 1$. We can apply Chernoff's tail bound to upper bound the probability for $X_i \geq 1 + \delta$ by

$$\Pr[X_i \geq 1 + \delta] \leq \frac{e^\delta}{(1 + \delta)^{1+\delta}}.$$

By union bound, if there exists bins with load exceeding $1 + \delta$, the probability is upper bounded by $n \cdot e^\delta / (1 + \delta)^{1+\delta}$.

Let $t = 1 + \delta$ be the load bound, then the probability upper bound is $n/e \cdot (e/t)^t$. We want to have

$$\lim_{n \rightarrow \infty} \frac{n}{e} \cdot \left(\frac{e}{t} \right)^t = 0$$

to bound all of the bin loads. Taking $\ln$ on both sides, $t \ln t - t = \omega(\ln n)$, giving only $t = O(\log n / \log \log n)$.

But there are at most $\log_2 n$ players and n balls, such result follow from vanilla remark 1.5.21. The problem lies in that we did not utilize the structure of the player model, giving a loose bound.

Looking at the balls-and-bins structure, we can divide the circle into $\log_2 n$ intervals, each has length $n/\log_2 n$. Let $\{Y_i\}_{i \in [1, \log_2 n]}$ be independent random variables for number of players in the i^{th} interval, following $\log_2 n$ balls and $\log_2 n$ bins model. An observation is that, for a bin in the i^{th} interval, its load is at most $Y_i + Y_{i-1}$, relaxed by

$$Y_i + Y_{i-1} \leq 2 \max(Y_i, Y_{i-1}).$$

Therefore, the max load of a bin is $O(Y_i)$.

We continue with Chernoff tail bound for Y_i , then union bound over $\log_2 n$ intervals, rather than over all n bins. Since $\mathbb{E}[Y_i] = 1$, the probability for existing $Y_i \geq 1 + \delta$ is $\log_2 n \cdot e^\delta / (1 + \delta)^{1+\delta}$.

Again we write $t = 1 + \delta$ to be the number of players in an interval, and we want

$$\lim_{n \rightarrow \infty} \frac{\log_2 n}{e} \cdot \left(\frac{e}{t} \right)^t = 0,$$

and therefore $t \ln t - t = \omega(\log \log n)$. We conclude that $t = O(\log \log n / \log \log \log n)$, and the load is $O(t)$.

This $\log n$ balls and $\log n$ bins max load result matches remark 1.5.21. $\square$

1.5.5 Recap on Coupon Collector's Problem

Theorem 1.5.23 (Theorem 5.13 [MU17]). *Let X be the number of coupons observed before obtaining one of each of n types of coupons. Then, for a constant c ,*

$$\lim_{n \rightarrow \infty} \Pr[X > n \ln n + cn] = 1 - e^{-e^{-c}}.$$

Proof. We can view this problem in balls-and-bins model: If balls are thrown independently and uniformly at random into bins, how many balls are thrown until all bins have at least one ball? We begin with a Poisson approximation, and later demonstrate that the Poisson approximation gives the right answer in the limit.

We begin by Poisson approximation, supposing $\{X_i\}_{i \in [1, n]}$ independent Poisson random variable of mean $\ln n + c$, then the expected number of balls is $m = n \ln n + cn$. Let E be the event that no bin is empty, and let $X = \sum_{i \in [1, n]} X_i$ be the number of balls in bin. Since $\Pr[X_i = 0] = (ne^c)^{-1}$, then

$$\lim_{n \rightarrow \infty} \Pr[E] = \lim_{n \rightarrow \infty} \left(1 - \frac{1}{ne^c}\right)^n = e^{-e^{-c}}.$$

Let $b = \sqrt{2m \ln m}$. By conditional probability, we have

$$\begin{aligned} \Pr[E] &= \Pr[E \mid |X - m| > b] \Pr[|X - m| > b] \\ &\quad + \Pr[E \mid |X - m| \leq b] \Pr[|X - m| \leq b]. \end{aligned}$$

Lemma 1.5.24. $\Pr[|X - m| > b] = o(1)$.

Proof. We can use the Chernoff bound such that $\Pr[|X - m| > b] < 2 \exp(-2 \ln m / 3) = 2m^{-2/3} = o(1)$. □

Lemma 1.5.25. $|\Pr[E \mid |X - m| \leq b] - \Pr[E \mid X = m]| = o(1)$.

Proof. Since $\Pr[E \mid X = k]$ is monotonically increasing in k , by conditional probability

$$\begin{aligned} \Pr[E \mid |X - m| \leq b] &= \sum_{k \in [-b, b]} \Pr[E \mid X = m + k] \cdot \Pr[X = m + k \mid |X - m| \leq b] \\ &\leq \sum_{k \in [-b, b]} \Pr[E \mid X = m + b] \cdot \Pr[X = m + k \mid |X - m| \leq b] \\ &= \Pr[E \mid X = m + b]. \end{aligned}$$

Therefore, we relax the term by

$$|\Pr[E \mid |X - m| \leq b] - \Pr[E \mid X = m]| \leq \Pr[E \mid X = m + b] - \Pr[E \mid X = m - b],$$

which is interpreted as: when $m - b$ balls are thrown, there exists empty bins; after $m + b$ balls, all bins are filled.

Let A be the event that $m - b$ balls thrown with empty bins, S be the set for all possible empty bins $B \subseteq [1, n]$ after $m - b$ initial throws, and F_B be the event that B empty bins are filled after another $2b$ balls. The exact term follows

$$\begin{aligned} \Pr[E \mid X = m + b] - \Pr[E \mid X = m - b] &= \Pr[A] \cdot \sum_{B \in S} (\Pr[\text{empty bins } B \subseteq [1, n] \mid A] \cdot \Pr[F_B]) \\ &\leq \sum_{B \in S} (\Pr[\text{empty bins } B \subseteq [1, n] \mid A] \cdot \Pr[F_B]) \\ &\leq \max_{B \in S} \Pr[F_B]. \end{aligned}$$

Since $\Pr[F_B] \leq \Pr[F_r]$ for some $r \in B$ by $F_B \subseteq F_r$, the probability for all empty bins being filled is upper bounded by the probability of one of the empty bins in B being filled, we upper bound $\Pr[F_r] \leq 2b/n = o(1)$ by union bound.

Therefore, we conclude the resulting difference by $o(1)$. □

By lemma 1.5.24 and lemma 1.5.25, we conclude that

$$\begin{aligned}\lim_{n \rightarrow \infty} \Pr[E] &= \lim_{n \rightarrow \infty} o(1) + (\Pr[E \mid X = m] + o(1)) \cdot (1 - o(1)) \\ &= \lim_{n \rightarrow \infty} \Pr[E \mid X = m],\end{aligned}$$

indicating as $n \rightarrow \infty$, for m balls n bins exact case, all n bins are filled with probability $e^{-e^{-c}}$. $\square$

1.5.6 Hashing and Random Graphs

Exercise 1.5.26 (Exercise 5.16 [MU17]). We consider another way to obtain Chernoff-like bounds in the settings of balls and bins without using theorem 1.5.15. Consider n balls and n bins model. Let $X_i = 1$ iff the i^{th} bin is empty, and $X = \sum_i X_i$. Let $\{Y_i\}_{i \in [1, n]}$ be independent Bernoulli random variables with probability $(1 - 1/n)^n$, and $Y = \sum_i Y_i$.

- Show that $\mathbb{E} \left[\prod_{i \in [1, k]} X_i \right] \leq \mathbb{E} \left[\prod_{i \in [1, k]} Y_i \right]$ for any $k \geq 1$.
- Show that $\mathbb{E} [\exp(tX)] \leq \mathbb{E} [\exp(tY)]$ for any $t \geq 0$.
- Derive a Chernoff bound for $\Pr[X \geq (1 + \delta) \mathbb{E}[X]]$.

Proof. Considering the event that first k bins are empty in n balls and n bins model, the probability is $(n - k)^n / n^n$. On the other hand, since each Y_i are independent, then $\prod_{i \in [1, k]} Y_i = 1$ has probability $(1 - 1/n)^{nk}$.

Since $(1 - 1/n)^k \geq 1 - k/n$, then $\mathbb{E} \left[\prod_{i \in [1, k]} X_i \right] \leq \mathbb{E} \left[\prod_{i \in [1, k]} Y_i \right]$ for any $k \geq 1$.

Expanding the prior claim, for any subsets of indices $S \subseteq [1, n]$, $\Pr \left[\prod_{i \in S} X_i = 1 \right] \leq \Pr \left[\prod_{i \in S} Y_i = 1 \right]$. We expand MGF function in its Taylor expansion, then we derive the following

$$\begin{aligned}\mathbb{E} [\exp(tX)] &= \mathbb{E} \left[\sum_{k \geq 0} \frac{(tX)^k}{k!} \right] = \sum_{k \geq 0} \frac{t^k}{k!} \cdot \mathbb{E} [X^k] = \sum_{k \geq 0} \frac{t^k}{k!} \cdot \mathbb{E} \left[\left(\sum_{i \in [1, n]} X_i \right)^k \right] \\ &= \sum_{k \geq 0} \frac{t^k}{k!} \cdot \sum_{S \subseteq [1, n], |S| \leq k} \Pr \left[\prod_{i \in S} X_i = 1 \right] \\ &\leq \sum_{k \geq 0} \frac{t^k}{k!} \cdot \sum_{S \subseteq [1, n], |S| \leq k} \Pr \left[\prod_{i \in S} Y_i = 1 \right] = \mathbb{E} [\exp(tY)].\end{aligned}$$

Since each Y_i are independent, let $p = (1 - 1/n)^n$, and therefore

$$\mathbb{E} [\exp(tY)] = \prod_{i \in [1, n]} \mathbb{E} [\exp(tY_i)] = (p \cdot (e^t - 1) + 1)^n.$$

By linearity of expectation, we have $\mathbb{E}[X] = \sum_{i \in [1, n]} \mathbb{E}[X_i] = np$. Therefore, the Chernoff bound is derived as follows

$$\Pr[X \geq (1 + \delta) np] \leq \frac{\mathbb{E} [\exp(tX)]}{\exp((1 + \delta) np t)} \leq \frac{\mathbb{E} [\exp(tY)]}{\exp((1 + \delta) np t)} \leq \left(\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right)^{np}.$$

$\square$

Remark 1.5.27. This trick in balls-and-bins model is analogous to the Poisson approximation. The spirit is the same: both are replacing dependent counts with independent random variables whose marginal behaviors matches that of the exact model. Poisson approximation captures more than just the expected load per bin, it captures the correct distributional shape and small-count probabilities, and approximate the joint distribution via independence, giving a limit theorem for the exact case. The Bernoulli trick preserves the expected emptiness per bin via independent random variables, allowing Chernoff to be applied, but it is more of a bounding device than a limit theorem.

Remark 1.5.28. This trick can be extended to m balls and n bins, and be used in Bloom Filter analysis on the number of zeroes (k hash functions, m disallowed passwords, and n bits to store, yielding km balls and n bins model).

Exercise 1.5.29 (Exercise 5.23 [MU17]). Bloom filter can be used to estimate set differences. Suppose there are sets X and Y both with n elements. Create Bloom filters for X and Y , using the same number of bits m and the same k hash functions. Determine the expected number of bits where Bloom filters differ as a function of m, n, k , and $|X \cap Y|$.

Proof. Let $A = X \setminus Z$, $B = Y \setminus Z$, $Z = X \cap Y$, and $|Z| = c$. We want to know the probability of some bit being set by only one of A or B , and not by Z . Therefore, we define following events

- $E_{i,0}$ be the event that i^{th} bit set by A not by B ,
- $E_{i,1}$ be the event that i^{th} bit set by B not by A ,
- $E_{i,2}$ be the event that i^{th} bit not set by Z ,

and let $\{W_i\}_{i \in [1,m]}$ be 0-1 random variables such that $W_i = 1$ if and only if $(E_{i,0} \cup E_{i,1}) \cap E_{i,2}$, and $W = \sum_i W_i$.

Since $E_{i,0}$ and $E_{i,1}$ are symmetric, WLOG we analyze $E_{i,0}$. Assuming hash function set a random bit to 1, then

$$\Pr[E_{i,0}] = \left(1 - \frac{1}{m}\right)^{k(n-c)} \cdot \left(1 - \left(1 - \frac{1}{m}\right)^{k(n-c)}\right).$$

On the other hand, not being set by Z for i^{th} bit is event $E_{i,2}$, which follow $\Pr[E_{i,2}] = (1 - 1/m)^{kc}$.

Therefore, $\Pr[W_i = 1] = 2 \Pr[E_{i,0}] \Pr[E_{i,2}]$, and by linearity of expectation,

$$\mathbb{E}\left[W = \sum_{i \in [1,m]} W_i\right] = 2m \left(1 - \frac{1}{m}\right)^{kn} \cdot \left(1 - \left(1 - \frac{1}{m}\right)^{k(n-c)}\right).$$

□

Definition 1.5.30. In $G_{n,p}$ model we consider all undirected graphs on n distinct vertices, an edge is connected between 2 distinct vertices with probability p , thus a graph with a given set of m edges has probability $p^m(1-p)^{\binom{n}{2}-m}$. The expected number of edges in the graph is $p\binom{n}{2}$, and each vertex has expected degree $(n-1)p$.

In $G_{n,N}$ model we consider all undirected graphs on n distinct vertices with exact N edges. There are $\binom{n}{2}$ distinct graphs to select with equal probability.

Remark 1.5.31. The relation between $G_{n,p}$ and $G_{n,N}$ has following properties: Let $p = N/\binom{n}{2}$, the number of edges in $G \xleftarrow{R} G_{n,p}$ is concentrated around N . Conditioned on $G \xleftarrow{R} G_{n,p}$ having N edges, then G is uniform over $G_{n,N}$.

Such relation is similar to the one between Poisson approximation and the balls-and-bins model.

Lemma 1.5.32. The distribution of $G_{n,p}$ conditioned on the graph sampled has N edges is the same as $G_{n,N}$ regardless of N .

Lemma 1.5.33. Let E be an event. If $\Pr[E] = t$ in the $G_{n,p}$ model, then $\Pr[E] = O(Nt)$ in the $G_{n,N}$ model where $N = p\binom{n}{2}$.

Proof. By the probability lower bounding from conditional probability (trick used in theorem 1.5.15), we have

$$\begin{aligned} \Pr[E \text{ in } G_{n,p}] &= \sum_{k \geq 0} \Pr[E \text{ in } G_{n,p} \mid \text{graph sampled in } G_{n,p} \text{ has } k \text{ edges}] \cdot \Pr[\text{graph sampled in } G_{n,p} \text{ has } k \text{ edges}] \\ &\geq \Pr[E \text{ in } G_{n,N}] \cdot \Pr[\text{graph sampled in } G_{n,p} \text{ has } N \text{ edges}]. \end{aligned}$$

Let $M = (1-p)\binom{n}{2}$. To lower bound the probability of graph sampled in $G_{n,p}$ with N edges, the probability follows

$$\Pr[\text{graph sampled in } G_{n,p} \text{ has } N \text{ edges}] = \binom{\binom{n}{2}}{N} \cdot p^N (1-p)^M = (N+M)! \cdot \frac{p^N}{N!} \cdot \frac{(1-p)^M}{M!}.$$

By result from corollary 1.5.5, we can lower bound components as follows:

$$\begin{aligned} (N+M)! &\geq \left(\frac{N+M}{e}\right)^{N+M} \\ \frac{p^N}{N!} &\geq \frac{p^N}{e\sqrt{N}} \cdot \left(\frac{e}{N}\right)^N = \frac{e^N}{e\sqrt{N}} \cdot (N+M)^{-N} \\ \frac{(1-p)^M}{M!} &\geq \frac{(1-p)^M}{e\sqrt{M}} \cdot \left(\frac{e}{M}\right)^M = \frac{e^M}{e\sqrt{M}} \cdot (N+M)^{-M}, \end{aligned}$$

The probability is lower bounded by $O(N^{-1})$, therefore the probability for the event happening in $G_{n,N}$ is $O(Nt)$. $\square$

Remark 1.5.34. Using Stirling bound, then the probability for the event to happen in $G_{n,N}$ can be improved to $O(nt)$.

Definition 1.5.35. A graph property is a property that holds for a graph, and all the isomorphisms of the graph. We say a graph property is monotone increasing if whenever the property holds for $G = (V, E)$, it holds for any graph $G' = (V', E')$ with $E \subseteq E'$. Monotone decreasing property is defined similarly, that whenever the property holds for $G = (V, E)$, it holds for any graph $G' = (V', E')$ with $E' \subseteq E$.

Lemma 1.5.36 (Lemma 5.14 [MU17]). *For a given monotone increasing graph property, let $P_{n,N}$ be the probability that the property holds for a graph in $G_{n,N}$ and $P_{n,p}$ be the probability that the property holds for a graph in $G_{n,p}$. Let $p^+ = (1+\varepsilon)N/\binom{n}{2}$ and $p^- = (1-\varepsilon)N/\binom{n}{2}$ for a constant $\varepsilon \in (0, 1)$. Then*

$$P_{n,p^-} - e^{-O(N)} \leq P_{n,N} \leq P_{n,p^+} + e^{-O(N)}.$$

Proof. The main theme of the proof is by relaxation from conditional probability, then apply Chernoff bound.

Let X be the random variable for the number of edges of a graph that is sampled from G_{n,p^-} . We relax P_{n,p^-} by

$$\begin{aligned} P_{n,p^-} &= \sum_{k \leq N} P_{n,k} \cdot \Pr[X = k] + \sum_{k > N} P_{n,k} \cdot \Pr[X = k] \\ &\leq P_{n,N} \cdot \Pr[X \leq N] + \sum_{k > N} P_{n,k} \cdot \Pr[X = k] \\ &\leq P_{n,N} + \Pr[X > N] \end{aligned}$$

The first equality follow from conditional probability, the second inequality holds from monotone increasing property has $P_{n,N} \geq P_{n,k}$ for all $k \leq N$, and the last inequality holds from bounding valid probability by 1.

Now we bound $\Pr[X > N]$ by Chernoff bound, as X is the sum of $\binom{n}{2}$ independent Bernoulli random variables with probability p^- , then

$$\Pr[X > N] = \Pr\left[X > \frac{1}{1-\varepsilon} p^- \binom{n}{2}\right] < \Pr\left[X > (1+\varepsilon) p^- \binom{n}{2}\right] \leq \exp\left(-\frac{1}{3} \varepsilon^2 p^- \binom{n}{2}\right) = \exp\left(-\frac{\varepsilon^2(1-\varepsilon)}{3} N\right),$$

as $1/(1-\varepsilon) > 1+\varepsilon$ for $0 < \varepsilon < 1$, which proves the $P_{n,N} \geq P_{n,p^-} - e^{-O(N)}$.

Similarly, let X be the random variable for the number of edges of a graph that is sampled from G_{n,p^+} . Then

$$\begin{aligned} P_{n,p^+} &= \sum_{k < N} P_{n,k} \cdot \Pr[X = k] + \sum_{k \geq N} P_{n,k} \cdot \Pr[X = k] \\ &\geq P_{n,N} \cdot \Pr[X \geq N] \\ &\geq P_{n,N} - \Pr[X < N]. \end{aligned}$$

Again by Chernoff bound,

$$\Pr[X < N] = \Pr\left[X < \frac{1}{1+\varepsilon} p^+ \binom{n}{2}\right] < \Pr\left[X < \left(1-\frac{\varepsilon}{2}\right) p^+ \binom{n}{2}\right] \leq \exp\left(-\frac{\varepsilon^2}{8} p^+ \binom{n}{2}\right) = \exp\left(-\frac{\varepsilon^2(1-\varepsilon)}{8} N\right),$$

as $1/(1+\varepsilon) < 1-\varepsilon/2$ for $0 < \varepsilon < 1$, which proves the $P_{n,N} \leq P_{n,p^+} + e^{-O(N)}$, and we complete the proof. $\square$

Remark 1.5.37. The spirit is similar to the Poisson approximation theorem 1.5.17. In the n balls m bins model, we replace dependent bins with independent m Poisson random variables of mean n/m . In the random graph case, we replace dependent N edges among all possible $\binom{n}{2}$ edges with each lined up with probability $N/\binom{n}{2}$ independently.

Lemma 1.5.38. *For a given monotone increasing graph property, let $P_{n,N}$ and $P_{n,p}$ be notions follow from lemma 1.5.36. Let $N = p\binom{n}{2}$, $N^+ = (1 + \varepsilon)N$, and $N^- = (1 - \varepsilon)N$ for a constant $\varepsilon \in (0, 1)$. Then*

$$P_{n,N^-} - e^{-O(N)} \leq P_{n,p} \leq P_{n,N^+} + e^{-O(N)}.$$

Proof. We prove with a similar strategy to lemma 1.5.36.

Let X be the random variable for the number of edges of a graph sampled from $G_{n,p}$. X can be seen as the sum of $\binom{n}{2}$ independent Bernoulli random variables with probability p . By conditional probability,

$$\begin{aligned} P_{n,p} &= \sum_{k \leq N^+} P_{n,k} \cdot \Pr[X = k] + \sum_{k > N^+} P_{n,k} \cdot \Pr[X = k] \\ &\leq P_{n,N^+} \cdot \Pr[X \leq N^+] + \sum_{k > N^+} P_{n,k} \cdot \Pr[X = k] \\ &\leq P_{n,N^+} + \Pr[X > N^+]. \end{aligned}$$

By Chernoff bound, the upper tail of X can be bounded as follows

$$\Pr[X > N^+] \leq \exp\left(-\frac{1}{3}\varepsilon^2 N\right) = \exp(-O(N)).$$

Similarly, again by conditional probability,

$$\begin{aligned} P_{n,p} &= \sum_{k < N^-} P_{n,k} \cdot \Pr[X = k] + \sum_{k \geq N^-} P_{n,k} \cdot \Pr[X = k] \\ &\geq P_{n,N^-} \cdot \Pr[X \geq N^-] \\ &\geq P_{n,N^-} - \Pr[X < N^-]. \end{aligned}$$

By Chernoff bound, the lower tail bound of X can be bounded as follows

$$\Pr[X < N^-] \leq \exp\left(-\frac{1}{2}\varepsilon^2 N\right) = \exp(-O(N)).$$

Therefore, we complete the proof on both sides. $\square$

Exercise 1.5.39 (Exercise 5.19 [MU17]). An undirected graph on n vertices is disconnected if there exists a set of $k < n$ vertices such that there is no edge between this set and the rest of the graph. Otherwise, the graph is connected. Show that there exists a constant c such that if $N \geq cn \log n$ then, with probability $1 - o(1)$, $G \xrightarrow{R} G_{n,N}$ is connected.

Proof. Let $P_{n,N}$ be the probability of $G \xrightarrow{R} G_{n,N}$ being connected, $P_{n,p}$ be the probability of $G \xrightarrow{R} G_{n,p}$ being connected, where $p = N/\binom{n}{2}$. The proof strategy follows: Since connectedness is a monotone increasing graph property, we apply lemma 1.5.36 such that we bound $P_{n,N}$ by

$$P_{n,p^-} - \exp(-O(N)) \leq P_{n,N} \leq P_{n,p^+} + \exp(-O(N)),$$

where $p^+ = (1 + \varepsilon)p$ and $p^- = (1 - \varepsilon)p$. Suppose P_{n,p^+} and P_{n,p^-} are both $1 - o(1)$, then we complete the proof.

We continue by analyzing $G_{n,p}$ model. For disconnected subset, the vertex size ranges in $[1, \lfloor \frac{n}{2} \rfloor]$ by symmetry. Let E be the event that $G \xrightarrow{R} G_{n,p}$ being disconnected. Then $\Pr[E]$ is union bounded by

$$\Pr[E] \leq \sum_{k \in [1, \lfloor \frac{n}{2} \rfloor]} \binom{n}{k} (1-p)^{k(n-k)} = \sum_{k \in [1, \lfloor \frac{n}{2} \rfloor]} a_k.$$

We want to show that $\Pr[E] = o(1)$ as $n \rightarrow \infty$. Observing that

$$\frac{a_{k+1}}{a_k} = \frac{n-k}{k+1} \cdot (1-p)^{n-2k-1},$$

then for a fixed k , as $n \rightarrow \infty$, $a_{k+1}/a_k \rightarrow 0$. Suppose for some K , we have $\frac{n}{2} \cdot (1-p)^{n-2K-1} < 1$, then $\{a_k\}_{k \in [1, K]}$ are monotone decreasing. We thus separate into 2 cases:

- $\sum_{k \leq K} a_k \leq Ka_1 = Kn(1-p)^{n-1}$, and the upper bound approaches $Kn \exp(-p(n-1))$ as $n \rightarrow \infty$.
- $\sum_{k > K} a_k < 2^n (1-p)^{\alpha n^2}$ by binomial coefficients for some α , and the upper bound approaches 0 when $n \rightarrow \infty$.

Let $p = O(\log n/n)$ with a sufficiently large constant, then the former bound approaches 0 when $n \rightarrow \infty$.

The rest follow from both $P_{n,p^+} = P_{n,p^-} = 1 - o(1)$ when ε is sufficiently small. Bounding by lemma 1.5.36, then when $N \geq cn \log n$ we complete the proof. $\square$

Theorem 1.5.40 (Theorem 5.15 [MU17]). *Let $N = \frac{1}{2}(n \ln n + cn)$. Then the probability that there are no isolated vertices (vertices with degree 0) in $G_{n,N}$ converges to $e^{-e^{-c}}$ as $n \rightarrow \infty$.*

Proof. The main idea is utilizing the balls-and-bins model to model adding edges in the $G_{n,N}$ model.

We introduce edges to the graph $G = (V, E)$, with $|V| = n$ and $E = \emptyset$ on initialization.

Repeat $\ell = N + \sqrt{n}$ times, each time toss 2 balls into n bins at random.

- If they get into the i^{th} and j^{th} bin, with $i \neq j$, edge $(i, j) \notin E$, and $|E| \leq N$, introduce the edge.
- Otherwise, take back the 2 balls from the bins.

Conditioning on having N edges, the distribution of the resulting G is equivalent to sampling from $G_{n,N}$.

Lemma 1.5.41. *After repeating ℓ times, as $n \rightarrow \infty$, $\Pr[|E| = N] = 1 - o(1)$.*

Proof. There are 2 mutually exclusive cases that an edge introduction trial get rejected: Either the 2 balls are tossed into the same bin, with probability $r_0 = n / \binom{n}{2}$; or rejected by prior introduction, upper bounded by $r_1 = N / \binom{n}{2}$.

Let $\{X_i\}_{i \in [1, \ell]}$ be independent 0-1 random variables of probability $r = r_0 + r_1$, and $X = \sum_i X_i$. Then X is a random variable upwards relaxing the number of rejections in the edge introduction trials.

Since by linearity of expectation $\mathbb{E}[X] = \ell r = O(\log^2 n)$, then by Markov/Chernoff concentration inequality, we have $\Pr[X > \sqrt{n}] = o(1)$, which completes the proof. $\square$

Let E_i be the event that G has no isolated vertices after the i^{th} edge introduction trial. Next corollary is immediate.

Corollary 1.5.42. *As $n \rightarrow \infty$, $\Pr[E_\ell] = \Pr[E_\ell \mid |E| = N]$.*

Let $M_{n,m}$ be the event that n balls and m bins model having no empty bin.

On one hand, $\Pr[E_\ell]$ is upper bounded by the 2ℓ balls and n bins model without empty bins, as:

- Repeated edge introduction does not matter in bin emptiness.
- 2 balls into the same bin means bin filled in the balls-bins model, rather than being rejected in the $G_{n,N}$ model.

A direct application of theorem 1.5.23 implies that as $n \rightarrow \infty$, $\Pr[E_\ell] \leq \Pr[M_{2\ell, n}] = e^{-e^{-c}}$.

On the other hand, $\Pr[E_\ell \mid |E| = N]$ is lower bounded by the $2N$ balls and n bins model with no empty bin, as

- Repeated edge introduction is allowed in the balls-and-bins model.
- 2 balls tossed into the same bin fills 1 bin rather than 2.

A direct application of corollary 1.5.42 and theorem 1.5.23 implies that as $n \rightarrow \infty$, $\Pr[E_\ell] \geq \Pr[M_{2N,n}] = e^{-e^{-c}}$.

Since we have sandwiched $\Pr[E_\ell] = e^{-e^{-c}}$ as $n \rightarrow \infty$, therefore we conclude the proof ¹. $\square$

Theorem 1.5.43. Let $p = (\ln n + c)/n$. The probability that there is no isolated vertices in $G_{n,p}$ converges to $e^{-e^{-c}}$ as $n \rightarrow \infty$.

Proof. This is immediate after theorem 1.5.40 and lemma 1.5.38. $\square$

1.5.7 Hamiltonian Cycles in Random Graphs

Definition 1.5.44. Let G be an undirected graph. Suppose that a path $P = v_1, \dots, v_k$ is a simple path in G , and that (v_k, v_i) is an edge in G . Then

$$P = v_1, \dots, v_i, v_k, v_{k-1}, \dots, v_{i+1}$$

is also a simple path, which we refer to as the rotation of P with the rotation edge (v_k, v_i) .

On input a graph $G = (V, E)$ with associated “used/unused-edges” lists for each vertex, the modified Hamiltonian cycle finding algorithm follows:

- Start with a random vertex as the head of the path.
- Repeat the following until a Hamiltonian cycle is closed, or the “unused-edges” of the head vertex is empty.
 - Let the current path be $P = v_1, \dots, v_k$ with v_k being the head.
 - Execute one of the following 3 cases with probabilities specified as follows.
 - * With probability $1/n$, reverse the path and make v_1 the new head.
 - * With probability $|\text{used-edges}(v_k)|/n$, sample $(v_k, v_i) \xleftarrow{R} \text{used-edges}(v_k)$, rotate P with edge (v_k, v_i) , and let v_{i+1} be new head (if the edge is (v_k, v_{k-1}) , take no action).
 - * Otherwise, sample $(v_k, u) \xleftarrow{R} \text{unused-edges}(v_k)$. If u is not on the path, make u the new head by $v_{k+1} = u$; or $u = v_i$, then rotate P by (v_k, v_i) and make v_{i+1} the new head (if the edge is (v_k, v_{k-1}) , take no action).
 - Update “unused-edges” and “used-edges” list accordingly.

We construct a special random graph model similar to $G_{n,p}$. Assuming each of the $n - 1$ possible edges connected to a vertex v is initially on the “unused-edges” list for vertex v independently with probability q . We also assume the order of these edges are in random order.

Remark 1.5.45. One way of looking at it is, before beginning the Hamiltonian cycle finding algorithm, we create the “unused-edges” list for each vertex v by inserting each possible edge (u, v) with probability q . The corresponding graph G is the graph including all edges that were inserted into some “unused-edges” list. But notice that an edge (u, v) could initially be on the “unused-edges” list for v but not for u . The independence among each “unused-edges” list is helpful for analysis onwards.

The following lemma says, regardless of the viewed vertices, the next vertex to view is still uniformly random, if the “unused-edges” list is not empty in the head vertex.

Lemma 1.5.46. Suppose the modified Hamiltonian cycle algorithm runs on the random graph model as described. Let V_i be the head vertex after the i^{th} time step. Suppose for any vertex u , as long as the i^{th} time step the unused edges list is not empty for the head vertex,

$$\Pr[V_{i+1} = u \mid V_i = v_i, V_{i-1} = v_{i-1}, \dots, V_1 = v_1] = \frac{1}{n}.$$

¹The proof is inspired by the [math stackexchange link](#). More specifically, I was thinking of sandwiching by $G_{n,p}$, but this solution suggested a way of sampling $G_{n,N}$ graph from the balls-and-bins model, and inspired me sandwiching by different number of balls.

Proof. The first 2 cases are both with probability $1/n$ to make u new head of path.

For the last case, there are at most $p = n - |\text{used-edges}(v_i)| - 1$ options for v_i 's unused edges, at the i^{th} time step. To make u the new head, u has to be in the previously sampled unused-edges(v_k), and the probability follows

$$\Pr \left[V_{i+1} = u \mid \bigcap_{j \in [1, i]} V_j = v_j \cap \text{third case} \right] = q \cdot \sum_{k \in [0, p-1]} \left(\binom{p-1}{k} q^k (1-q)^{p-1-k} \cdot \frac{1}{k+1} \right) - (1-q)^{p-1} = \frac{1}{n}.$$

(or actually we can prove by symmetry, that for any possible option, the probability is the same, and it splits up the whole sample space.) Thus we complete the proof. $\square$

Now the problem looks exactly like the coupon collector's problem, that the probability of finding a new vertex to add to the path, when there are k vertices are left to be added, is k/n .

Theorem 1.5.47. *Suppose the input to the modified Hamiltonian cycle algorithm initially has the probability of adding edges to the unused edges list $q \geq 20 \ln n/n$. Then the algorithm managed to find a Hamiltonian cycle in $O(n \ln n)$ time steps with probability $1 - O(n^{-1})$.*

Proof. We start by upper bounding the event $\mathcal{E}$ that the algorithm failed to find such cycle, splitting by:

- The algorithm ran for $3n \ln n$ steps with no empty unused edges list, and failed to close a Hamiltonian cycle.
- The algorithm drained some unused edges list in the first $3n \ln n$ steps, and failed to close a Hamiltonian cycle.

Let the first event be $\mathcal{E}_1$. We relax the second event by lifting the constraint of “failed to close a Hamiltonian cycle”, and we let the relaxed event be $\mathcal{E}_2$. We thus have $\Pr[\mathcal{E}] \leq \Pr[\mathcal{E}_1] + \Pr[\mathcal{E}_2]$.

We start with $\mathcal{E}_1$ bounding first. Failing to close a Hamiltonian cycle can be considered as either failing to traverse all vertices, or failing to close by traversing back to starting vertex. We model the vertices traversal conditioned on no unused edges list empty by balls-and-bins model in coupon collector's problem, and let $\mathcal{E}_{1,a}$ be the event that all coupons are collected before $2n \ln n$ time steps. By conditional probability, we have

$$\begin{aligned} \Pr[\mathcal{E}_1] &= \Pr[\mathcal{E}_1 \mid \mathcal{E}_{1,a}] \Pr[\mathcal{E}_{1,a}] + \Pr[\mathcal{E}_1 \mid \neg \mathcal{E}_{1,a}] \Pr[\neg \mathcal{E}_{1,a}] \\ &\leq \Pr[\mathcal{E}_1 \mid \mathcal{E}_{1,a}] + \Pr[\neg \mathcal{E}_{1,a}], \end{aligned}$$

where the first term means not closing the Hamiltonian cycle in at least $n \ln n$ steps after traversing all vertices, and the second term means the coupon collector's problem does not stop before $2n \ln n$ steps.

The first term is upper bounded by

$$\Pr[\mathcal{E}_1 \mid \mathcal{E}_{1,a}] \leq \left(1 - \frac{1}{n}\right)^{n \ln n} \leq \frac{1}{n}.$$

The second term is upper bounded by union bounding over all n vertices not being picked after $2n \ln n$ trials

$$\Pr[\neg \mathcal{E}_{1,a}] \leq n \left(1 - \frac{1}{n}\right)^{2n \ln n} \leq \frac{1}{n}.$$

Therefore, $\Pr[\mathcal{E}_1] \leq 2/n = O(n^{-1})$.

Now we bound $\Pr[\mathcal{E}_2]$. In this case, either all vertices have at least $10 \ln n$ unused edges initially, but some gets hit hard; or some vertices have too few unused edges on initialization. Let $\mathcal{E}_{2,a}$ be the event that some vertices have too few unused edges. By conditional probability

$$\begin{aligned} \Pr[\mathcal{E}_2] &= \Pr[\mathcal{E}_2 \mid \mathcal{E}_{2,a}] \Pr[\mathcal{E}_{2,a}] + \Pr[\mathcal{E}_2 \mid \neg \mathcal{E}_{2,a}] \Pr[\neg \mathcal{E}_{2,a}] \\ &\leq \Pr[\mathcal{E}_2 \mid \neg \mathcal{E}_{2,a}] + \Pr[\mathcal{E}_{2,a}]. \end{aligned}$$

The first term can be upper bounded by some vertex being visited too often. An unused edges list can be emptied after at least $10 \ln n$ visits. Let X be the number of visits, which is upper bounded by a sum of $3n \ln n$ independent Bernoulli trials of probability $1/n$. By Chernoff bound,

$$\Pr[X \geq 10 \ln n] \leq \frac{1}{n^2}.$$

By union bounding over all n vertices, the $\Pr[\mathcal{E}_2 \mid \neg \mathcal{E}_{2,a}] \leq 1/n$.

The second term can be bounded similarly. Let Y be the initial number of unused edges. Then

$$\mathbb{E}[Y] = (n-1)q \geq 19 \ln n$$

for a sufficiently large n . By Chernoff bound for $\Pr[Y \leq 10 \ln n]$, then union bounding over n vertices, $\Pr[\mathcal{E}_{2,a}] \leq 1/n$.

Therefore, $\Pr[\mathcal{E}] \leq 4/n = O(n^{-1})$, and we complete the proof. $\square$

Remark 1.5.48. The operation of rotating in definition 1.5.44 is derived from [Pós76], which is often referred to as the Pósa's rotation-extension technique.

1.6 The Probabilistic Method

1.6.1 A few things I should have noted about factorial

We bounded around factorial in corollary 1.5.5. We can derive a tighter lower bound from following lemmas.

Lemma 1.6.1. *For any integer $k \geq 1$, by concavity of $\ln x$,*

$$\ln k \geq \int_{k-\frac{1}{2}}^{k+\frac{1}{2}} \ln x dx.$$

Lemma 1.6.2.

$$n! \geq \sqrt{2n} \left(\frac{n}{e}\right)^n.$$

Proof. Since $\ln n! = \sum_{i \in [1, n]} \ln i$, and we have the following from lemma 1.6.1

$$\begin{aligned} \sum_{i \in [1, n]} \ln i &\geq \int_{\frac{1}{2}}^{n+\frac{1}{2}} \ln x dx = (x \ln x - x) \Big|_{\frac{1}{2}}^{n+\frac{1}{2}} \\ &= \left(\left(n + \frac{1}{2}\right) \ln \left(n + \frac{1}{2}\right) - \left(n + \frac{1}{2}\right) \right) - \left(\frac{1}{2} \ln \frac{1}{2} - \frac{1}{2} \right) \\ &\geq \left(n + \frac{1}{2}\right) \ln n - n + \frac{1}{2} \ln 2 = n \ln n - n + \frac{1}{2} (\ln(2n)), \end{aligned}$$

which completes the proof by exponentiating over e on both sides. $\square$

Corollary 1.6.3. *By lemma 1.5.4 and lemma 1.6.2, we have*

$$\sqrt{2n} \cdot \left(\frac{n}{e}\right)^n \leq n! \leq e\sqrt{n} \cdot \left(\frac{n}{e}\right)^n \text{ and } \frac{1}{e\sqrt{n}} \cdot \left(\frac{e}{n}\right)^n \leq \frac{1}{n!} \leq \frac{1}{\sqrt{2n}} \cdot \left(\frac{e}{n}\right)^n.$$

Lemma 1.6.4. *By corollary 1.6.3, we have*

$$\binom{2n}{n} = \Theta\left(\frac{1}{\sqrt{n}} \cdot 2^{2n}\right).$$

Proof. Since $\binom{2n}{n} = (2n)!/(n!)^2$, then we have

$$\begin{aligned} \binom{2n}{n} &\leq e\sqrt{2n} \cdot \left(\frac{2n}{e}\right)^{2n} \cdot \left(\frac{1}{\sqrt{2n}} \cdot \left(\frac{e}{n}\right)^n\right)^2 = \frac{e}{\sqrt{2n}} \cdot 2^{2n} \\ \binom{2n}{n} &\geq 2\sqrt{n} \cdot \left(\frac{2n}{e}\right)^{2n} \cdot \left(\frac{1}{e\sqrt{n}} \cdot \left(\frac{e}{n}\right)^n\right)^2 = \frac{2}{e^2\sqrt{n}} \cdot 2^{2n}, \end{aligned}$$

which completes the proof. $\square$

Lemma 1.6.5. *For $s = O(\sqrt{n})$, $(n-s)! \cdot (n+s)! = \Theta(1) \cdot (n!)^2$.*

Proof.

$$\frac{(n-s)!}{n!} \cdot \frac{(n+s)!}{n!} = \prod_{i \in [1, s]} \frac{n-s+i}{n+i} = \prod_{i \in [1, s]} \left(1 - \frac{s}{n+i}\right) = \Theta\left(\exp\left(-\frac{s^2}{n}\right)\right) = \exp(-\Theta(1)).$$

$\square$

Corollary 1.6.6. *Immediate from lemma 1.6.5, for $s = O(\sqrt{n})$,*

$$\binom{2n}{n+s} = \Theta(1) \cdot \binom{2n}{n}.$$

Corollary 1.6.7. Immediate from lemma 1.6.5, for $s = O(\sqrt{n})$,

$$\binom{2n+s}{n} \binom{2n-s}{n} = \Theta(1) \cdot \binom{2n}{n}^2.$$

Exercise 1.6.8 (Exercise 6.16.b [MU17]). Let $\{Y_i\}_{i \in [1, 2n]}$ be independent random variables uniform over $\{0, 1\}$. Show that there exists a positive constant c such that, for sufficiently large n ,

$$\Pr \left[\left| \sum_{i \in [1, n]} Y_i - \sum_{i \in [n+1, 2n]} Y_i \right| > c\sqrt{n} \right] > \frac{1}{2}.$$

Proof. Let $Y = \sum_i Y_i$, $A = \sum_{i \in [1, n]} Y_i$, and $B = \sum_{i \in [n+1, 2n]} Y_i$. Let E_t be the event $|A - B| > c\sqrt{n}$ when $Y = t$.

Let k be a positive constant. By conditional probability (similar to lemma 1.5.38),

$$\Pr[|A - B| > c\sqrt{n}] \geq \Pr[|A - B| > c\sqrt{n} \mid |Y - n| \leq k\sqrt{n}] - \Pr[|Y - n| > k\sqrt{n}],$$

where $\Pr[|Y - n| > k\sqrt{n}] < 2 \exp(-k^2/3)$ by Chernoff bound.

We continue by conditional probability, that

$$\begin{aligned} \Pr[|A - B| > c\sqrt{n} \mid |Y - n| \leq k\sqrt{n}] &= \sum_t \Pr[E_t \mid Y = t] \cdot \Pr[Y = t \mid |Y - n| \leq k\sqrt{n}] \\ &= 1 - \sum_t \Pr[\neg E_t \mid Y = t] \cdot \Pr[Y = t \mid |Y - n| \leq k\sqrt{n}]. \end{aligned}$$

Event $\neg E_t$ is equivalent to $A \in [(t - c\sqrt{n})/2, (t + c\sqrt{n})/2] = R_t$ when $Y = t$. Therefore,

$$\begin{aligned} \Pr[\neg E_t \mid Y = t] &= \sum_{a \in R_t} \Pr[A = a \mid Y = t] \\ &\leq (c\sqrt{n} + 1) \cdot \max_a \Pr[A = a \mid Y = t] \\ &= (c\sqrt{n} + 1) \cdot \max_a \binom{n}{a} \binom{n}{t-a} \binom{2n}{t}^{-1} = (c\sqrt{n} + 1) \cdot p(a). \end{aligned}$$

By $p(a+1)/p(a)$,

$$\frac{p(a+1)}{p(a)} = \frac{(a+1)(n-t+a+1)}{(n-a)(t-a)},$$

$p(a)$ peaks at $a = \lfloor \frac{t}{2} \rfloor$ giving $\left(\frac{n}{\frac{t}{2}}\right)^2 \binom{2n}{t}^{-1}$. On the other hand, by lemma 1.6.4, the following is immediate

$$p\left(\frac{n}{2}\right) = \left(\frac{n}{\frac{n}{2}}\right)^2 \binom{2n}{n}^{-1} = \Theta\left(\frac{1}{\sqrt{n}}\right).$$

By corollary 1.6.6, for $t \in [n - k\sqrt{n}, n + k\sqrt{n}]$, we have $p(a) = \Theta\left(\frac{1}{\sqrt{n}}\right)$.

Choose c sufficiently small such that $\Pr[\neg E_t \mid Y = t] \leq 1/2$ for all $t \in [n - k\sqrt{n}, n + k\sqrt{n}]$, then

$$\sum_t \Pr[\neg E_t \mid Y = t] \cdot \Pr[Y = t \mid |Y - n| \leq k\sqrt{n}] \leq \frac{1}{2},$$

which completes the proof. $\square$

Exercise 1.6.9 (Exercise 6.16.c [MU17]). Let $\{Y_i\}_{i \in [1, 2n]}$ be independent random variables uniform over $\{0, 1\}$, and let $\{b_i\}_{i \in [1, 2n]}$ each be either $+1$ or -1 . Show that there exists a positive constant c such that, for sufficiently large n ,

$$\Pr \left[\left| \sum_{i \in [1, 2n]} b_i Y_i \right| > c\sqrt{n} \right] > \frac{1}{2}.$$

Proof. We claim that either the $\{b_i\}$ are imbalanced, that $|B| = |\sum_i b_i| > \ell\sqrt{n}$; or they are balanced. WLOG let $B > 0$.

For the imbalanced case, let $Z = \sum_i b_i Y_i$. By linearity of expectation, $\mathbb{E}[Z] = B/2$. Since $\{Y_i\}$ are independent, by Hoeffding's bound

$$\Pr \left[Z \leq \mathbb{E}[Z] - \sqrt{2n} \right] \leq \exp \left(-\frac{2(\sqrt{2n})^2}{\sum_i b_i^2} \right) = \exp(-2),$$

as $\sum_i b_i^2 = 2n$. Since $\Pr \left[Z \leq \mathbb{E}[Z] - \sqrt{2n} \right] \leq e^{-2}$, then for $\ell \geq 2(c + \sqrt{2})$,

$$\Pr \left[Z \geq c\sqrt{n} \right] \geq \Pr \left[Z \geq \mathbb{E}[Z] - \sqrt{2n} \right] > 1 - e^{-2} > 1/2$$

the imbalanced case is proved.

Now for the balanced case, we reuse the structure and notion in exercise 1.6.8 (the same window argument) that

$$\begin{aligned} \Pr[\neg E_t \mid Y = t] &\leq (c\sqrt{n} + 1) \cdot \max_a \Pr[A = a \mid Y = t] \\ &= (c\sqrt{n} + 1) \cdot \max_a \binom{n+B/2}{a} \binom{n-B/2}{t-a} \binom{2n}{t}^{-1} = (c\sqrt{n} + 1) \cdot p(a). \end{aligned}$$

By $p(a+1)/p(a)$,

$$\frac{p(a+1)}{p(a)} = \frac{(n+B/2-a)(t-a)}{(a+1)(n-B/2-t+a+1)},$$

where $p(a)$ peaks at $\lfloor t \cdot \frac{2n+B}{4n} \rfloor$.

On the other hand, by corollary 1.6.6 and corollary 1.6.7, the max of $p(a)$ is still $\Theta\left(\frac{1}{\sqrt{n}}\right)$, completing the proof. $\square$

1.6.2 Basic Counting Argument

Exercise 1.6.10 (Exercise 6.10 [MU17]). A family of subsets $\mathcal{F}$ of $[1, n]$ is called an antichain if there is no pair of sets A and B in $\mathcal{F}$ satisfying $A \subset B$.

- Let f_k be the number of sets in $\mathcal{F}$ with size k . Show that

$$\sum_{k \in [0, n]} f_k \cdot \binom{n}{k}^{-1} \leq 1.$$

- Argue that $|\mathcal{F}| \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}$ for any antichain $\mathcal{F}$.

Proof. We start choosing a random permutation of $[1, n]$. Let $X_k = 1$ if the first k numbers in the permutation yield a set in $\mathcal{F}$, and let $X = \sum_{i \in [1, n]} X_i$. Then X is at most 1 by antichain property, as no larger subset containing all prior elements, and X can be 0 as no element in $\mathcal{F}$ exists on the random enumeration starting from 0.

Therefore, for a $\mathcal{F}$, when we are iterating over all permutations, if a k sized subset is in $\mathcal{F}$, we will encounter $k!(n-k)!$ permutations containing such subset, and we derive

$$\sum_{k \in [0, n]} f_k \cdot k!(n-k)! \leq n!,$$

as a permutation can have at most 1 element in $\mathcal{F}$, and the first result is immediate.

Since $\binom{n}{k} \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}$ for all $k \in [0, n]$, then

$$|\mathcal{F}| = \sum_{k \in [0, n]} f_k \leq \sum_{k \in [0, n]} \binom{n}{\lfloor \frac{n}{2} \rfloor} \cdot f_k \cdot \binom{n}{k}^{-1} \leq \binom{n}{\lfloor \frac{n}{2} \rfloor}.$$

$\square$

Remark 1.6.11. This is a result from [Spe28] with a Lubell [Lub66] flavored proof.

Another way of seeing it is, a k sized subset in $\mathcal{F}$ appears in a random permutation over $[1, n]$ with probability $k!(n-k)!/n! = \binom{n}{k}^{-1}$. By linearity of expectation,

$$\mathbb{E}[X] = \sum_{k \in [0, n]} \mathbb{E}[X_k] = \sum_{k \in [0, n]} \frac{f_k}{\binom{n}{k}}.$$

On the other hand, X is a 0-1 random variable, hence $\mathbb{E}[X] \leq 1$, which completes the proof.

1.6.3 The Expectation Argument and Derandomization by Conditional Expectation

Lemma 1.6.12 (The Expectation Argument). *For random variable X with $\mathbb{E}[X] = \mu$, $\Pr[X \geq \mu] > 0$ and $\Pr[X \leq \mu] > 0$.*

Lemma 1.6.13. *Let $\{X_i\}_{i \in [1, n]}$ be 0-1 random variables, $X = \sum_i X_i$ with $\mathbb{E}[X] = m \cdot 2^{-k}$ for some integer $m \geq 1$ and $k \geq 0$. Then for all $X < \mathbb{E}[X]$, $X \leq \mathbb{E}[X] - 2^{-k}$; for all $X > \mathbb{E}[X]$, $X \geq \mathbb{E}[X] + 2^{-k}$.*

Lemma 1.6.14. *Let $\{X_i\}_{i \in [1, n]}$ be 0-1 random variables, $X = \sum_i X_i$ with $\mathbb{E}[X] = m \cdot 2^{-k}$ for some integer $m \geq 1$ and $k \geq 0$. Then*

$$\begin{aligned} \Pr[X \geq \mathbb{E}[X]] &\geq \frac{1}{2^k(n - \mathbb{E}[X] + 1)}, \\ \Pr[X \leq \mathbb{E}[X]] &\geq \frac{1}{2^k \mathbb{E}[X] + 1}. \end{aligned}$$

Proof. Let $p = \Pr[X \geq \mathbb{E}[X]]$. By conditional expectation and lemma 1.6.13,

$$\begin{aligned} \mathbb{E}[X] &= \sum_{i < \mathbb{E}[X]} i \Pr[X = i] + \sum_{i \geq \mathbb{E}[X]} i \Pr[X = i] \\ &\leq (\mathbb{E}[X] - 2^{-k})(1 - p) + np, \end{aligned}$$

yielding $p \geq 1/(2^k(n - \mathbb{E}[X] + 1))$.

Let $p = \Pr[X \leq \mathbb{E}[X]]$. By conditional expectation and lemma 1.6.13,

$$\begin{aligned} \mathbb{E}[X] &= \sum_{i \leq \mathbb{E}[X]} i \Pr[X = i] + \sum_{i > \mathbb{E}[X]} i \Pr[X = i] \\ &\geq (\mathbb{E}[X] + 2^{-k})(1 - p), \end{aligned}$$

yielding $p \geq 1/(2^k \mathbb{E}[X] + 1)$. □

Exercise 1.6.15 (Exercise 6.1 [MU17]). Consider an SAT instance with m clauses and each clause has exactly k literals. Give a Las Vegas algorithm that finds an assignment satisfying at least $m(1 - 2^{-k})$ clauses, and its expected runtime. Give a derandomization of the randomized algorithm using the method of conditional expectation.

Proof. For a k -literal SAT clause, it is satisfied with probability $1 - 2^{-k}$. Let X_i be the 0-1 random variable indicating if the i^{th} clause is satisfied, and $X = \sum_i X_i$ be the random number of satisfied clauses. By linearity of expectation, $\mathbb{E}[X] = m(1 - 2^{-k})$.

The Las Vegas algorithm searching for assignment satisfying at least $m(1 - 2^{-k})$ clauses works by keep sampling assignments in $O(km)$ time, and checks if the assignment satisfies sufficient number of clauses in $O(km)$ time.

By lemma 1.6.14, $\Pr[X \geq \mathbb{E}[X]] \geq (2^k(m - \mathbb{E}[X]) + 1)^{-1}$.

The expected number of sampling is at most $2^k(m - \mathbb{E}[X]) + 1$, and the expected runtime is $O(km2^k(m - \mathbb{E}[X]))$.

Let there be $n = O(km)$ variables for the SAT instance, σ be a random permutation of x_i , and x_i be the choice of the i^{th} variable. The derandomized algorithm from conditional expectation wants to prove

$$\mathbb{E}[X] \leq \mathbb{E}[X \mid x_1, \dots, x_n],$$

the right hand side is the number of satisfied clauses after all x_i are decided. We prove by induction as follows.

The base case $\mathbb{E}[X] = \mathbb{E}[X \mid x_1]$ is immediate by symmetry.

The intermediate step $\mathbb{E}[X \mid x_1, \dots, x_i] \leq \mathbb{E}[X \mid x_1, \dots, x_{i+1}]$ can be shown by conditional expectation, as

$$\mathbb{E}[X \mid x_1, \dots, x_i] = \frac{1}{2} \sum_{x_{i+1} \in \{0,1\}} \mathbb{E}[X \mid x_1, \dots, x_{i+1}] \leq \max_{x_{i+1} \in \{0,1\}} \mathbb{E}[X \mid x_1, \dots, x_{i+1}].$$

We can greedily choose the next x_{i+1} that satisfies more clauses.

The n step greedy algorithm suffices to find an assignment satisfying at least $m(1-2^{-k})$ clauses by our analysis. $\square$

Exercise 1.6.16 (Exercise 6.2 [MU17]). Prove that for every integer n , there exists a 2-coloring of the edges of K_n so that the total number of monochromatic K_4 is at most $\binom{n}{4}2^{-5}$, and show the following

- A randomized algorithm to find a coloring with at most $\binom{n}{4}2^{-5}$ monochromatic K_4 in expected poly time in n .
- A deterministic polynomial time algorithm to find such a coloring by derandomization.

Proof. There are $\binom{n}{4}$ copies of K_4 in a K_n , and a K_4 is 2-colored monochromatic with probability 2^{-5} .

Let X_i be the 0-1 random variable indicating the i^{th} K_4 being monochromatic, and $X = \sum_i X_i$ be the number of monochromatic K_4 in K_n . By linearity of expectation, $\mathbb{E}[X] = \binom{n}{4}2^{-5}$. Hence, the existence of 2-coloring with at most $\binom{n}{4}2^{-5}$ monochromatic K_4 is immediate by the expectation argument.

We can keep randomly sampling 2-coloring K_n until we see at most $\mathbb{E}[X]$ copies of monochromatic K_4 .

By lemma 1.6.14, $\Pr[X \leq \mathbb{E}[X]] \geq (2^5 \mathbb{E}[X] + 1)^{-1}$.

The expected sampling time is at most $2^5 \mathbb{E}[X] + 1$, and expected runtime is $O\left(\binom{n}{4}^2\right) = O(n^8)$.

Let x_i be the i^{th} edge's color choice among all $N = \binom{n}{2}$ edges. The derandomized algorithm wants to show

$$\mathbb{E}[X] \geq \mathbb{E}[X \mid x_1, \dots, x_N],$$

right hand side is the number of monochromatic K_4 after fixing all vertices' colors. We prove by induction as follows.

The base case $\mathbb{E}[X] = \mathbb{E}[X \mid x_1]$ is immediate by symmetry.

The intermediate step $\mathbb{E}[X \mid x_1, \dots, x_i] \geq \mathbb{E}[X \mid x_1, \dots, x_{i+1}]$ can be shown by conditional expectation, as

$$\mathbb{E}[X \mid x_1, \dots, x_i] = \frac{1}{2} \sum_{x_{i+1} \in \{0,1\}} \mathbb{E}[X \mid x_1, \dots, x_{i+1}] \geq \min_{x_{i+1} \in \{0,1\}} \mathbb{E}[X \mid x_1, \dots, x_{i+1}].$$

We can greedily choose the next x_{i+1} that bring less monochromatic K_4 .

The N step greedy algorithm suffices to 2-color K_n with at most $\binom{n}{4}2^{-5}$ monochromatic K_4 by our analysis. $\square$

Remark 1.6.17. This appears in The Probabilistic Method [AS16], and uses the “pessimistic estimator” idea [Rag88].

Exercise 1.6.18 (Exercise 6.3 [MU17]). Given a n -vertex undirected graph, consider the following method of generating an independent set. Given a permutation σ of V , define $S(\sigma) \subset V$ as follows: for each vertex i , $i \in S(\sigma)$ if and only if no neighbor j of i precedes i in the permutation σ .

- Show $\mathbb{E}[|S(\sigma)|] = \sum_i 1/(d_i + 1)$, where $\{d_i\}_{i \in [1,n]}$ denotes the degree of vertices.
- Prove that G has an independent set of size at least $\sum_i 1/(1 + d_i)$.

Proof. Given a random σ , the resulting $S(\sigma)$ has each vertex with no neighbors, which is an independent set.

Let $\{X_i\}_{i \in [1, n]}$ be 0-1 random variables indicating if the i^{th} vertex is the first among all its neighbors, and $X = \sum_i X_i$ is the random variable for $|S(\sigma)|$. The i^{th} vertex is the first among all d_i neighbors and itself with probability $1/(1 + d_i)$. Therefore, by linearity of expectation, $\mathbb{E}[X] = \sum_i 1/(1 + d_i)$.

The existence of an independent set at least $\sum_i 1/(1 + d_i)$ large is immediate by the expectation argument. $\square$

Exercise 1.6.19 (Exercise 6.5 [MU17]). We have shown using the probabilistic method that, if a graph G has n nodes and m edges, then there exists a partition of the n nodes into 2 sets such that at least $m/2$ edges cross the partition. Improve this result slightly: show that there exists a partition such that at least $mn/(2n - 1)$ edges cross the partition.

Proof. The $m/2$ result derive from 2-coloring the vertices each with probability $1/2$, and same color vertices goes to the same set. In this way, edges with two ends colored differently are included in the cut, with a probability of $1/2$. Existence of cut with at least $m/2$ edges is immediate by linearity of expectation and the expectation argument.

On the other hand, n vertices can be partitioned into sets with $\lfloor \frac{n}{2} \rfloor$ and $\lceil \frac{n}{2} \rceil$ vertices each, and improve by the expectation argument. Let the event of choosing 2 vertices with different colors be E , the probability is

$$\Pr[E] = \frac{2 \lfloor \frac{n}{2} \rfloor \lceil \frac{n}{2} \rceil}{n(n-1)}.$$

- When n even, the probability is $n/(2n - 2) > n/(2n - 1)$.
- When n odd, the probability is $(n + 1)/2n > n/(2n - 1)$.

By linearity of expectation, the expected size of the cut is $m \Pr[E] \geq mn/(2n - 1)$. Hence, the existence of a cut with at least $mn/(2n - 1)$ is immediate by the expectation argument. $\square$

Exercise 1.6.20 (Exercise 6.6 [MU17]). A k -cut is a partition of the vertices into k disjoint sets, and the value of a cut is the weight of all edges crossing from one of the k sets to another. Show that any graph G with m edges has a k -cut with value at least $(k - 1)m/k$. Show how to use derandomization to deterministically find such a cut.

Proof. We assign one out of the k colors to each of the n vertices with probability $1/k$. The probability of 2 vertices having same color is $1/k$, and therefore an edge being included into a k -cut has probability $(k - 1)/k$.

By linearity of expectation, the expected k -cut size is $m(k - 1)/k$. Hence, the existence of k -cut with size at least $m(k - 1)/k$ is immediate by the expectation argument.

When it comes to the derandomized deterministic algorithm for finding a k -cut that is larger than $m(k - 1)/k$, we let $\{S_i\}_{i \in [1, k]}$ be the sets of partitioned vertices, and we write $C(\{S_i\}_{i \in [1, k]})$ for size of the k -cut. We pick an arbitrary emulation of all vertices $\{v_i\}_{i \in [1, n]}$, and write $\{x_i\}_{i \in [1, n]}$ for each vertex's color choice. We prove by induction for

$$\mathbb{E}[C(\{S_i\}_{i \in [1, k]})] \leq \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_n],$$

where the right hand side is the k -cut size determined by the derandomized algorithm, where the color choices are $\{x_i\}_{i \in [1, n]}$, hence the algorithm gives a cut whose size is at least $\mathbb{E}[C(\{S_i\}_{i \in [1, k]})] = m(k - 1)/k$.

The base case $\mathbb{E}[C(\{S_i\}_{i \in [1, k]})] \leq \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1]$ is immediate by symmetry.

The intermediate step

$$\mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_i] \leq \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_{i+1}]$$

can be shown by conditional expectation, that

$$\begin{aligned} \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_i] &= \frac{1}{k} \sum_{x_{i+1} \in [1, k]} \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_{i+1}] \\ &\leq \max_{x_{i+1} \in [1, k]} \mathbb{E}[C(\{S_i\}_{i \in [1, k]}) \mid x_1, \dots, x_{i+1}]. \end{aligned}$$

In fact, we greedily find the next x_{i+1} without computing all conditional expectations on x_{i+1} 's choice. We choose the x_{i+1} such that v_{i+1} contribute most edges to the k -cut. Equivalently, we put v_{i+1} to the partition with the fewest v_{i+1} 's colored neighbors, such that v_{i+1} introduces most k -cut edges with colored neighbors.

This n steps greedy searching algorithm suffices to determine a k -cut at least $m(k-1)/k$ large by our analysis. $\square$

Exercise 1.6.21 (Exercise 6.9 [MU17]). A tournament is a graph on n vertices with exactly one directed edge between each pair of vertices. If vertices represent players, then each edge can be thought of as the result of a match between the two players: the edge points to the winner. A ranking is an ordering of the n players from best to worst (ties are not allowed). Given the outcome of a tournament, one might wish to determine a ranking of the players. A ranking is said to disagree with a directed edge from y to x if y is ahead of x in the ranking.

- Prove that, for every tournament, there exists a ranking that disagrees with at most 50% of the edges.
- Prove that, for sufficiently large n , there exists a tournament such that every ranking disagrees with at least 49% of the edges in the tournament.

Proof. Fix a tournament, and let X_i be the 0-1 random variable indicating if a random ranking disagrees with the i^{th} of all $\binom{n}{2}$ edges of the tournament, then $X = \sum_i X_i$ is the total number of disagreement.

For any unordered pair (x, y) , each ordering is equally likely in a random ranking. Therefore, $\Pr[X_i = 1] = 1/2$. By linearity of expectation,

$$\mathbb{E}[X] = \sum_i \mathbb{E}[X_i] = \frac{1}{2} \binom{n}{2}.$$

By the expectation argument, there exists a ranking disagreeing at most 50% of the edges, proving the first part.

Fix a ranking. Let $Y = \sum_j Y_j$ be the random variable for the number of disagreement with a random tournament, and Y_j be the 0-1 random variable indicating if the ranking disagrees with the j^{th} edge of the random tournament.

Let the fixed ranking be the k^{th} among all $n!$ ranking, and let E_k be the event that the fixed k^{th} ranking has at most 49% of disagreement against a random tournament, namely $Y \leq (1 - \delta) \mathbb{E}[Y]$, where $\delta = 1/50$. Since $\{Y_j\}$ are independent, we can use concentration inequality like Chernoff bound to upper bound the lower tail of Y as follows

$$\Pr[Y \leq (1 - \delta) \mathbb{E}[Y]] \leq \exp(-\delta^2 \mu / 2) = \exp\left(-\frac{1}{50^2} \cdot \frac{1}{2} \binom{n}{2} \cdot \frac{1}{2}\right).$$

Union bounding over all rankings yields

$$\Pr\left[\bigcup_{k \in [1, n!]} E_k\right] \leq \sum_{k \in [1, n!]} \Pr[E_k] = n! \cdot \exp\left(-\frac{1}{100^2} \binom{n}{2}\right) \leq n^n \cdot \exp\left(-\frac{1}{100^2} \binom{n}{2}\right) = \exp\left(n \ln n - \frac{1}{100^2} \binom{n}{2}\right),$$

and the bound is less than 1 if n is sufficiently large. Thus, for a tournament, the probability of existing a ranking permutation with at most 49% disagreement is strictly less than 1. Hence,

$$\Pr\left[\bigcap_{k \in [1, n!]} \neg E_k\right] = 1 - \Pr\left[\bigcup_{k \in [1, n!]} E_k\right] > 0.$$

$\square$

1.6.4 Sample and Modify

Exercise 1.6.22 (Exercise 6.7 [MU17]). A Hypergraph $H = (V, E)$, where V is the set of vertices and E is the set of hyperedges, and every hyperedge in E is a subset of V . In particular, r -uniform hypergraph is one where the size of each hyperedge is r . A standard graph is a 2-uniform hypergraph. A dominating set in a hypergraph H is a set of vertices $S \subset V$ such that $e \cap S \neq \emptyset$ for every $e \in E$, namely S hits every hyperedge in H .

Let $|V| = n$ and $|E| = m$, show that there is a dominating set S with size at most $np + (1 - p)^r m$ for $0 \leq p \leq 1$.

Proof. Given a hypergraph $H = (V, E)$ with $|V| = n$ and $|E| = m$, we first sample across V with each vertices to be chosen with probability p , then check if any hyperedges are not covered, choose a vertex from each of the hyperedge. The sampling stage has expected chosen vertices np , leaving expected unchosen hyperedges $m(1-p)^r$ to be fixed. $\square$

Exercise 1.6.23 (Exercise 6.8 [MU17]). Prove that, for every integer n , there is a way to 2-color the edges of K_n such that there is no monochromatic cliques of size k when

$$x = n - \binom{n}{k} 2^{1-\binom{k}{2}}.$$

Proof. We start 2-coloring K_n , keep removing vertices and attached edges until there is no monochromatic K_k .

Let each edge in K_n be 2-colored with probability $1/2$, and let X_i be the 0-1 random variable indicating the i^{th} K_k being monochromatic. There are $\binom{n}{k}$ K_k 's in a K_n , hence $X = \sum_i X_i$ is the number of monochromatic K_k 's in the K_n .

For a K_k , the probability of being monochromatic is $2^{1-\binom{k}{2}}$ as it can be one of 2 colors. By linearity of expectation,

$$\mathbb{E}[X] = \sum_i \mathbb{E}[X_i] = \binom{n}{k} 2^{1-\binom{k}{2}}.$$

By expectation argument, there exists a 2-coloring such that $X \leq \mathbb{E}[X]$.

To modify up to no monochromatic K_k , we can remove a vertex and the attached edges in each K_k , therefore we need to remove at most X vertices from K_n . Since there exists a 2-coloring for K_n containing at most $\mathbb{E}[X]$ K_k 's, then there exists a 2-coloring with no monochromatic K_k for K_m that $m \geq n - \mathbb{E}[X]$.

“no monochromatic K_k ” is a monotonely decreasing graph property, therefore any subgraph of such 2-colored K_m has this property. Therefore, a K_x in this K_m has no monochromatic K_k , which completes the proof. $\square$

1.6.5 The Second Moment Method and the Conditional Expectation Inequality

Theorem 1.6.24 (The Second Moment Method (Chebyshev's version)). Let X be a random variable, then

$$\Pr[X = 0] \leq \frac{\text{Var}[X]}{\mathbb{E}[X]^2}.$$

Proof. Since

$$\Pr[X = 0] \leq \Pr[|X - \mathbb{E}[X]| \geq \mathbb{E}[X]] \leq \frac{\text{Var}[X]}{\mathbb{E}[X]^2},$$

which is immediate by Chebyshev's inequality. $\square$

Lemma 1.6.25. For integer valued non-negative random variable X , $\Pr[X > 0] \leq \mathbb{E}[X]$ by Markov's inequality.

Theorem 1.6.26 (Cauchy-Schwarz). For random variables X and Y (that may not be independent),

$$\mathbb{E}[XY]^2 \leq \mathbb{E}[X^2] \cdot \mathbb{E}[Y^2].$$

Proof. Let $Z = X - t \cdot Y$, then the second moment of Z is

$$\mathbb{E}[Z^2] = \mathbb{E}[X^2] + t^2 \cdot \mathbb{E}[Y^2] - 2t \cdot \mathbb{E}[XY] \geq 0,$$

which holds for any t . Minimizing over t , and we have $\mathbb{E}[X^2] \cdot \mathbb{E}[Y^2] \geq \mathbb{E}[XY]^2$. $\square$

Corollary 1.6.27. For random variable X with $\mathbb{E}[X] \geq 0$, $0 \leq t \leq \mathbb{E}[X]$, and $1_t = 1$ if and only if $X \geq t$, by theorem 1.6.26

$$\Pr[X \geq t] \geq \frac{\mathbb{E}[X \cdot 1_t]^2}{\mathbb{E}[X^2]}.$$

Theorem 1.6.28 (The Second Moment Method (Cauchy-Schwarz's version)). *For random variable $X \geq 0$,*

$$\mathbb{E}[X] \geq \Pr[X > 0] \geq \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}.$$

Proof. Let $1_+ = 1$ if and only if $X > 0$, then by theorem 1.6.26 and corollary 1.6.27,

$$\mathbb{E}[X^2] \mathbb{E}[1_+]^2 \geq \mathbb{E}[X^2] \Pr[X > 0] \geq \mathbb{E}[X \cdot 1_+]^2 \geq \mathbb{E}[X]^2,$$

and the last inequality holds by non-negativity of X . □

Remark 1.6.29. A weaker version of theorem 1.6.28 is in exercise 1.3.10. We can relax non-negative X to $\mathbb{E}[X] \geq 0$.

Theorem 1.6.30 (Paley-Zygmund). *For random variable $X \geq 0$ with finite $\mathbb{E}[X^2]$, and $0 \leq \alpha < 1$. Then*

$$\Pr[X \geq \alpha \mathbb{E}[X]] \geq (1 - \alpha)^2 \cdot \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}.$$

Proof. Let $1_\alpha = 1$ if and only if $X \geq \alpha \mathbb{E}[X]$, and $1_{\bar{\alpha}} = 1$ if and only if $X < \alpha \mathbb{E}[X]$, then

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}[X \cdot 1_\alpha] + \mathbb{E}[X \cdot 1_{\bar{\alpha}}] \\ &\leq \mathbb{E}[X \cdot 1_\alpha] + \alpha \cdot \mathbb{E}[X], \end{aligned}$$

and therefore $\mathbb{E}[X \cdot 1_\alpha] \geq (1 - \alpha) \cdot \mathbb{E}[X]$. By theorem 1.6.26 and corollary 1.6.27, we have

$$\mathbb{E}[X \cdot 1_\alpha]^2 \leq \mathbb{E}[X^2] \cdot \mathbb{E}[1_\alpha] = \mathbb{E}[X^2] \cdot \Pr[X \geq \alpha \mathbb{E}[X]].$$

Hence, the result is immediate by

$$\Pr[X \geq \alpha \mathbb{E}[X]] \geq \frac{\mathbb{E}[X \cdot 1_\alpha]^2}{\mathbb{E}[X^2]} \geq (1 - \alpha)^2 \cdot \frac{\mathbb{E}[X]^2}{\mathbb{E}[X^2]}.$$

□

Remark 1.6.31. The theorem 1.6.30 can be relaxed from non-negative X to $\mathbb{E}[X] \geq 0$.

Exercise 1.6.32 (Exercise 6.11 [MU17]). Prove a threshold for the existence of triangles in the $G_{n,p}$. Let $\{X_i\}$ each be the random value for the i^{th} triangle appearing in the graph among the $\binom{n}{3}$ triplets of vertices, and let $X = \sum_i X_i$.

- Show $\mathbb{E}[X]$, and show if $pn \rightarrow 0$, then $\Pr[X > 0] \rightarrow 0$.
- Show $\text{Var}[X_i] \leq p^3$.
- Show that $\text{Cov}[X_i, X_j] = p^5 - p^6$ for $O(n^4)$ pairs of triangle triplets $i \neq j$, otherwise 0.
- Show $\text{Var}[X] = O(n^3 p^3 + n^4(p^5 - p^6))$.
- Show if $pn \rightarrow \infty$ then $\Pr[X = 0] \rightarrow 0$.

Proof. By linearity of expectation, $\mathbb{E}[X] = \binom{n}{3} p^3 = O(n^3 p^3)$. If $pn \rightarrow 0$, then $\mathbb{E}[X] \rightarrow 0$.

For variance, we have $\text{Var}[X_i] = \mathbb{E}[X_i^2] - \mathbb{E}[X_i]^2 = p^3 - p^6 \leq p^3$.

For covariance, when the triplets overlap up to 1 vertex, no edges are shared, and $\mathbb{E}[X_i X_j] = \mathbb{E}[X_i] \mathbb{E}[X_j]$, hence $\text{Cov}[X_i, X_j] = 0$. Otherwise, when 2 vertices are shared, $\mathbb{E}[X_i X_j] = p^5$, and there are in total $\binom{n}{2;1;1}$ csaes.

Therefore, the variance for X is derived as

$$\begin{aligned} \text{Var}[X] &= \sum_i \text{Var}[X_i] + \sum_{i \neq j} \text{Cov}[X_i, X_j] = \binom{n}{3} (p^3 - p^6) + \binom{n}{2;1;1} (p^5 - p^6) \\ &= O(n^3 p^3 + n^4(p^5 - p^6)). \end{aligned}$$

When $pn \rightarrow \infty$, by theorem 1.6.24, $\Pr[X = 0] \leq O((pn)^{-1}) \rightarrow 0$. □

Exercise 1.6.33 (Exercise 6.12 [MU17]). Use exercise 1.3.1 to derive the variance of the number of K_4 in $G_{n,p}$.

Proof. Let $\{X_i\}$ be the 0-1 random variable for the existence of the i^{th} K_4 among all possible $\binom{n}{4}$ K_4 , and let $X = \sum_i X_i$. By exercise 1.3.1,

$$\mathbb{E}[X^2] = \sum_i \mathbb{E}[X | X_i = 1] \Pr[X_i = 1].$$

We know $\Pr[X_i = 1] = p^6$, and among all $\binom{n}{4}$ possible K_4 , if $X_i = 1$,

- If there are no vertex overlap, then $\binom{n-4}{4}$ K_4 with probability p^6 .
- If there is 1 vertex overlapped, then $\binom{n-4}{3} \binom{4}{1}$ K_4 with probability p^6 .
- If there are 2 vertices overlapped, then $\binom{n-4}{2} \binom{4}{2}$ K_4 with probability p^5 .
- If there are 3 vertices overlapped, then $\binom{n-4}{1} \binom{4}{3}$ K_4 with probability p^3 .
- If there are 4 vertices overlapped, then it has to be i^{th} K_4 with probability 1.

By summation,

$$\mathbb{E}[X | X_i = 1] = \binom{n-4}{4} p^6 + \binom{n-4}{3} \binom{4}{1} p^6 + \binom{n-4}{2} \binom{4}{2} p^5 + \binom{n-4}{1} \binom{4}{3} p^3 + 1,$$

therefore,

$$\begin{aligned} \mathbb{E}[X^2] &= \binom{n}{4} p^6 \left(\binom{n-4}{4} p^6 + \binom{n-4}{3} \binom{4}{1} p^6 + \binom{n-4}{2} \binom{4}{2} p^5 + \binom{n-4}{1} \binom{4}{3} p^3 + 1 \right) \\ &= \binom{n}{4;4} p^{12} + \binom{n}{3;3;1} p^{12} + \binom{n}{2;2;2} p^{11} + \binom{n}{3;1;1} p^9 + \binom{n}{4} p^6, \end{aligned}$$

hence

$$\begin{aligned} \text{Var}[X] &= \left(\binom{n}{4;4} + \binom{n}{3;3;1} - \binom{n}{4}^2 \right) p^{12} + \binom{n}{2;2;2} p^{11} + \binom{n}{3;1;1} p^9 + \binom{n}{4} p^6 \\ &= O(n^7 p^{12}) + O(n^6 p^{11}) + O(n^5 p^9) + O(n^4 p^6). \end{aligned}$$

When $p = o(n^{-2/3})$, by lemma 1.6.25, that $\Pr[X \geq 1] \leq \mathbb{E}[X] = o(1)$.

When $p = \omega(n^{-2/3})$, by theorem 1.6.24 and $\mathbb{E}[X]^2 = O(n^8 p^{12}) = \omega(1)$, we have

$$\Pr[X = 0] \leq \frac{\text{Var}[X]}{\mathbb{E}[X]^2} = o(1).$$

□

Theorem 1.6.34 (Conditional Expectation Inequality). Let $\{X_i\}_{i \in [1,n]}$ be 0-1 random variables, and $X = \sum_i X_i$. Then

$$\Pr[X > 0] \geq \sum_{i \in [1,n]} \frac{\Pr[X_i = 1]}{\mathbb{E}[X | X_i = 1]}.$$

Proof. Let $Y = 1/X$ when $X \neq 0$, and 0 otherwise. Then $\mathbb{E}[XY] = \Pr[X > 0]$, and

$$\begin{aligned} \mathbb{E}[XY] &= \sum_{i \in [1,n]} \mathbb{E}[Y | X_i = 1] \Pr[X_i = 1] \\ &\geq \sum_{i \in [1,n]} \frac{\Pr[X_i = 1]}{\mathbb{E}[X | X_i = 1]}, \end{aligned}$$

where the first equality holds from exercise 1.3.1, and the last inequality holds from theorem 1.2.5.

□

Remark 1.6.35. Previous exercise 1.6.33 $p = \omega(n^{-2/3})$ case can be solved with theorem 1.6.34 by

$$\begin{aligned} \Pr[X > 0] &\geq \sum_i \frac{\Pr[X_i = 1]}{\mathbb{E}[X \mid X_i = 1]} \\ &= \frac{\binom{n}{4} p^6}{\binom{n-4}{4} p^6 + \binom{n-4}{3} \binom{4}{1} p^6 + \binom{n-4}{2} \binom{4}{2} p^5 + \binom{n-4}{1} \binom{4}{3} p^3 + 1}. \end{aligned}$$

When $p = \omega(n^{-2/3})$, both $\binom{n}{4} p^6$ and $\binom{n-4}{4} p^6$ are $\omega(1)$, while the others are $o(1)$, proving the second part of the proof.

Exercise 1.6.36 (Exercise 6.13 [MU17]). Consider the problem of whether graphs in $G_{n,p}$ have K_k of constant size k . Suggest a threshold function and generalize the argument for K_4 using either theorem 1.6.24 or theorem 1.6.34, to prove that your threshold function is correct for K_5 .

Proof. We reuse notions in exercise 1.6.33, let X_i be the random variable for the existence of the i^{th} K_k in $G_{n,p}$, and let $X = \sum_i X_i$. We guess the threshold $p = n^{-2/(k-1)}$, which balances both the dominating numerator and denominator terms in theorem 1.6.34, and convert $\mathbb{E}[X] = \binom{n}{k} p^{k(k-1)/2} = O(n^k p^{k(k-1)/2})$ into constant for lemma 1.6.25.

When $p = o(n^{-1/2})$, by linearity of expectation, $\mathbb{E}[X] = \binom{n}{5} p^{10} = o(1)$. By lemma 1.6.25, $\Pr[X > 0] \leq \mathbb{E}[X] = o(1)$.

When $p = \omega(n^{-1/2})$, we apply the theorem 1.6.34. Among all $\binom{n}{5}$ possible K_5 , conditioned that $X_i = 1$,

- If there are no vertex overlap, then $\binom{n-5}{5} K_5$ with probability p^{10} .
- If there is 1 vertex overlapped, then $\binom{n-5}{4} \binom{5}{1} K_5$ with probability p^{10} .
- If there are 2 vertices overlapped, then $\binom{n-5}{3} \binom{5}{2} K_5$ with probability p^9 .
- If there are 3 vertices overlapped, then $\binom{n-5}{2} \binom{5}{3} K_5$ with probability p^7 .
- If there are 4 vertices overlapped, then $\binom{n-5}{1} \binom{5}{4} K_5$ with probability p^4 .
- If there are 5 vertices overlapped, then it has to be the i^{th} K_5 with probability 1.

By summation,

$$\mathbb{E}[X \mid X_i = 1] = \binom{n-5}{5} p^{10} + \binom{n-5}{4} \binom{5}{1} p^{10} + \binom{n-5}{3} \binom{5}{2} p^9 + \binom{n-5}{2} \binom{5}{3} p^7 + \binom{n-5}{1} \binom{5}{4} p^4 + 1.$$

Since $\binom{n-5}{5} p^{10}$ and $\binom{n}{5} p^{10}$ are $\omega(1)$, while all other terms are $o(1)$, therefore by theorem 1.6.34, $\Pr[X > 0] \rightarrow 1$. $\square$

Exercise 1.6.37 (Exercise 6.14 [MU17]). Consider $G \in G_{n,p}$ with $p = c \ln n / n$. Use theorem 1.6.24 or theorem 1.6.34 to prove that if $c < 1$ then, for any constant $\varepsilon > 0$ and for n sufficiently large, the graph has isolated vertices with probability at least $1 - \varepsilon$.

Proof. Let X_i be random variable for the i^{th} vertex being isolated in the graph, and $X = \sum_i X_i$.

By theorem 1.6.34, among n vertices, conditioned that $X_i = 1$,

- If a vertex is not the i^{th} vertex, it has probability $(1 - p)^{n-2}$ to be isolated.
- Otherwise, the i^{th} vertex is isolated with probability 1 by conditioning.

By summing,

$$\lim_{n \rightarrow \infty} \mathbb{E}[X \mid X_i = 1] = \lim_{n \rightarrow \infty} 1 + (n-1)(1-p)^{n-2} = 1 + \frac{n}{n^c}.$$

Since $\Pr[X_i = 1] = (1-p)^{n-1}$, which is $1/n^c$ as $n \rightarrow \infty$, therefore by theorem 1.6.34,

$$\lim_{n \rightarrow \infty} \Pr[X > 0] \geq \frac{n/n^c}{1 + n/n^c} = \frac{n}{n^c + n},$$

which approaches 1 if $c < 1$. $\square$

Exercise 1.6.38 (Exercise 6.15 [MU17]). Consider a graph in $G_{n,p}$ where $p = 1/n$. Let X be the number of triangles in the graph. Show that

$$\Pr[X \geq 1] \leq \frac{1}{6},$$

and that

$$\lim_{n \rightarrow \infty} \Pr[X \geq 1] \geq \frac{1}{7}.$$

Proof. There are $\binom{n}{3}$ triangles in K_n , while a triangle exists with probability p^3 . Let $\{X_i\}$ be the set of 0-1 random variables indicating if the i^{th} triangle exists in the graph, then $X = \sum_i X_i$. By linearity of expectation,

$$\mathbb{E}[X] = \sum_i \mathbb{E}[X_i] = \binom{n}{3} p^3 = \frac{(n-1)(n-2)}{6n^2} \leq \frac{1}{6}.$$

On the other hand, by lemma 1.6.25

$$\mathbb{E}[X] = \sum_{x \in [0, \binom{n}{3}]} x \Pr[X = x] \geq \sum_{x \in [1, \binom{n}{3}]} \Pr[X = x] = \Pr[X \geq 1],$$

which completes the proof for the first part.

On the other hand, by theorem 1.6.34 we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \Pr[X \geq 1] &\geq \lim_{n \rightarrow \infty} \sum_{i \in [1, \binom{n}{3}]} \frac{\Pr[X_i = 1]}{\mathbb{E}[X \mid X_i = 1]} \\ &= \lim_{n \rightarrow \infty} \frac{\binom{n}{3} p^3}{\mathbb{E}[X \mid X_i = 1]}. \end{aligned}$$

For all $\binom{n}{3}$ possible triangles,

- If there are no vertex overlap, then $\binom{n-3}{3}$ triangles with probability p^3 .
- If there are 1 vertex overlap, then $\binom{n-3}{2} \binom{3}{1}$ triangles with probability p^3 .
- If there are 2 vertex overlap, then $\binom{n-3}{1} \binom{3}{2}$ triangles with probability p^2 .
- If there are 3 vertex overlap, then 1 triangle with probability 1.

Summing, we have

$$\begin{aligned} \mathbb{E}[X \mid X_i = 1] &= \sum_{j \in [1, \binom{n}{3}]} \mathbb{E}[X_j \mid X_i = 1] \\ &= 1 + \binom{n-3}{1} \binom{3}{2} p^2 + \binom{n-3}{2} \binom{3}{1} p^3 + \binom{n-3}{3} p^3. \end{aligned}$$

When $n \rightarrow \infty$, $\binom{n}{3} p^3 \rightarrow 1/6$ and $\mathbb{E}[X \mid X_i = 1] \rightarrow 7/6$, which proves the second part. $\square$

Remark 1.6.39. We attempted exercise 1.6.8 from theorem 1.6.30, which is inspired by [AS16] to use the second/fourth moment method, just to avoid the bounding over binomial coefficients, but it did not bound as good.

Let $\{Y_i\}_{i \in [2n]}$ be independent Bernoulli trials from even coins, $\{Z_i\}_{i \in [n]}$ each be $Y_i - Y_{i+n}$, and $Z = \sum_i Z_i$. Since Z is symmetric over 0, then event $|Z| > c\sqrt{n}$ is equivalent to $Z^2 > c'n$.

For random variable Z^2 , we want to prove $\Pr[Z^2 > c'n] \geq 1/2$ for some constant $c' > 0$, which fits theorem 1.6.30:

- The second moment is $\mathbb{E}[Z^2] = \sum_i \mathbb{E}[Z_i^2] = n/2$.
- The fourth moment is $\mathbb{E}[Z^4] = \sum_i \mathbb{E}[Z_i^4] + \sum_{i \neq j} \mathbb{E}[Z_i^2 Z_j^2] = \sum_i \mathbb{E}[Z_i^4] + \sum_{i \neq j} \mathbb{E}[Z_i^2] \mathbb{E}[Z_j^2] = n/2 + \binom{n}{2} (4)/4$.

Therefore, even when $\Pr[Z^2 > 0] < 1/3$ by Paley-Zygmund theorem 1.6.30, which is too loose.

1.6.6 Lovász Local Lemma

The symmetric Lovász Local Lemma requires that the bad events are all upper bounded by a probability p , and each bad event depends on at most d others.

By definition, an event E_{n+1} is mutually independent of $\{E_i\}_{i \in [1, n]}$, if for any subset $S \subseteq [1, n]$,

$$\Pr \left[E_{n+1} \mid \bigcap_{i \in S} E_i \right] = \Pr [E_{n+1}].$$

Onwards, we denote E_S for $\bigcap_{i \in S} E_i$, and $\bar{E}_S$ for $\bigcap_{i \in S} \neg E_i$.

Definition 1.6.40 (Dependency Graph). A dependency graph for events $\{E_i\}$ is a graph $G = (V, E)$, such that E_i is independent of the events $\{E_j \mid (i, j) \notin E\}$. The degree of the dependency graph is the maximum degree of any vertex in the graph.

Lemma 1.6.41 (Symmetric Lovász Local Lemma [EL75]). Let $\{E_i\}_{i \in [1, n]}$ be a set of events, and the following holds:

- $\Pr [E_i] \leq p$ for all $i \in [1, n]$,
- The degree of the dependency graph given by $\{E_i\}$ is at most d ,
- $4dp \leq 1$,

then $\Pr [\bar{E}_{[1, n]}] > 0$.

Proof. The proving strategy is proving by induction the following lemma for all possible $S \subseteq [1, n]$.

Lemma 1.6.42. For all $S \subseteq [1, n]$ with $k \notin S$, we have both following being true

$$\Pr [E_k \mid \bar{E}_S] \leq 2p \text{ and } \Pr [\bar{E}_S] \geq (1 - 2p)^{|S|}.$$

Proof. We prove by induction over all possible $|S|$.

When $|S| = 0$, this is vacuously true.

When $|S| = 1$, let $S = \{i\}$. Then $\Pr [\neg E_i] \geq 1 - p \geq 1 - 2p$. For $\Pr [E_k \mid \neg E_i]$:

- If E_k and E_i mutually independent, namely $(i, k) \notin E$, then

$$\Pr [E_k \mid \neg E_i] = \Pr [E_k] \leq p \leq 2p.$$

- Otherwise $(i, k) \in E$, then

$$\Pr [E_k \mid \neg E_i] = \frac{\Pr [E_k \cap \neg E_i]}{\Pr [\neg E_i]} \leq \frac{\Pr [E_k]}{\Pr [\neg E_i]} \leq \frac{p}{1 - p} \leq 2p,$$

where the last inequality holds as $4dp \leq 1$, meaning $p \leq 1/4$.

WLOG let $S' = [1, s - 1]$, $S = [1, s]$, and $s \geq 2$. Supposing the lemma holds for S' with $|S'| = s - 1$, then for S ,

$$\begin{aligned} \Pr [\bar{E}_S] &= \Pr [\neg E_s \mid \bar{E}_{S'}] \cdot \Pr [\bar{E}_{S'}] \\ &= (1 - \Pr [E_s \mid \bar{E}_{S'}]) \cdot \Pr [\bar{E}_{S'}] \\ &\geq (1 - 2p) \cdot (1 - 2p)^{s-1} = (1 - 2p)^s, \end{aligned}$$

where the lower bound follows immediately from the induction hypothesis.

Remark 1.6.43. Alternatively, we have

$$\begin{aligned}
 \Pr \left[\bar{E}_S \right] &= \Pr \left[\neg E_s \mid \bar{E}_{S'} \right] \cdot \Pr \left[\bar{E}_{S'} \right] \\
 &= \Pr \left[\neg E_s \mid \bar{E}_{S'} \right] \cdot \prod_{i \in [1, s-1]} \Pr \left[\neg E_i \mid \bar{E}_{[1, i-1]} \right] \\
 &= \prod_{i \in [1, s]} \left(1 - \Pr \left[E_i \mid \bar{E}_{[1, i-1]} \right] \right) \\
 &\geq (1 - 2p)^s,
 \end{aligned}$$

where the last inequality holds from induction hypothesis.

Let $A \subseteq S$ such that $\forall i \in A, (i, s+1) \in E$, and $B = S \setminus A$. If $B = S$, then E_{s+1} is mutually independent of $\{E_i\}_{i \in S}$,

$$\Pr \left[E_{s+1} \mid \bar{E}_S \right] = \Pr [E_{s+1}] \leq p \leq 2p.$$

Otherwise, by Bayes' rule,

$$\begin{aligned}
 \Pr \left[E_{s+1} \mid \bar{E}_S \right] &= \frac{\Pr \left[E_{s+1} \cap \bar{E}_S \right]}{\Pr \left[\bar{E}_S \right]} = \frac{\Pr \left[E_{s+1} \cap \bar{E}_A \cap \bar{E}_B \right]}{\Pr \left[\bar{E}_A \cap \bar{E}_B \right]} \\
 &= \frac{\Pr \left[E_{s+1} \cap \bar{E}_A \mid \bar{E}_B \right]}{\Pr \left[\bar{E}_A \mid \bar{E}_B \right]}.
 \end{aligned}$$

For numerator, since E_{s+1} is mutually independent of $\{E_i\}_{i \in B}$ by the dependency graph, we have

$$\Pr \left[E_{s+1} \cap \bar{E}_A \mid \bar{E}_B \right] \leq \Pr \left[E_{s+1} \mid \bar{E}_B \right] = \Pr [E_{s+1}] \leq p.$$

For denominator,

$$\begin{aligned}
 \Pr \left[\bar{E}_A \mid \bar{E}_B \right] &= 1 - \Pr \left[\bigcup_{i \in A} E_i \mid \bar{E}_B \right] \\
 &\geq 1 - \sum_{i \in A} \Pr \left[E_i \mid \bar{E}_B \right] \\
 &\geq 1 - d \cdot 2p \\
 &\geq \frac{1}{2},
 \end{aligned}$$

where the first inequality holds from union bound, the second inequality holds from the induction hypothesis, and $|A| \leq d$ by the degree of the dependency graph, and the last inequality holds from the lemma's requirement. Hence, $\Pr \left[E_{s+1} \mid \bar{E}_S \right] \leq 2p$. $\square$

Since lemma 1.6.42 holds for any S , which finishes the proof immediately. $\square$

Exercise 1.6.44 (Exercise 6.17 [MU17]). Use lemma 1.6.41 to show that, if

$$4 \cdot \binom{k}{2} \binom{n}{k-2} \cdot 2^{1-\binom{k}{2}} \leq 1,$$

then it is possible to 2-color the edges of K_n so that it has no monochromatic K_k subgraph.

Proof. Let E_i be the bad event that the i^{th} K_k being monochromatic.

To upper bound the dependency degree d , we can pick an edge from the i^{th} K_k , then uniformly randomly sample the rest of the vertices from all possible vertices. Hence

$$d \leq \binom{k}{2} \binom{n}{k-2}.$$

The probability of K_k being monochromatic is $2^{1-\binom{k}{2}}$. Therefore by lemma 1.6.41, if

$$4 \cdot \binom{k}{2} \binom{n}{k-2} \cdot 2^{1-\binom{k}{2}} \leq 1,$$

then it is possible to 2-color K_n edges such that there is no monochromatic K_k . $\square$

The symmetric Lovász Local Lemma lemma 1.6.41 works best when all events are upper bounded equally, and the dependency degrees are also upper bounded equally. The asymmetric Lovász Local Lemma has likely or rare events, where likely events have smaller dependency degrees, while rare events could depend on many other events.

Onwards we denote $\Gamma_i = \{j \mid (i, j) \in E\}$ for all the neighbors of E_i in the dependency graph.

Lemma 1.6.45 (Asymmetric Lovász Local Lemma [EL75]). *Let $\{E_i\}_{i \in [1, n]}$ be a set of events. If $\{x_i\}_{i \in [1, n]}$ in $[0, 1]$ has*

$$\Pr[E_i] \leq x_i \cdot \prod_{j \in \Gamma_i} (1 - x_j),$$

then

$$\Pr\left[\overline{E}_{[1, n]}\right] \geq \prod_{i \in [1, n]} (1 - x_i).$$

Proof. The proving strategy is similar to lemma 1.6.41 by introducing following lemma for all possible $S \subseteq [1, n]$.

Lemma 1.6.46. *For all $S \subseteq [1, n]$ with $k \notin S$, we have both following being true*

$$\Pr\left[E_k \mid \overline{E}_S\right] \leq x_k \text{ and } \Pr\left[\overline{E}_S\right] \geq \prod_{i \in S} (1 - x_i).$$

Proof. We prove by induction over all possible $|S|$.

When $|S| = 0$, this is vacuously true.

When $|S| = 1$, let $S = \{i\}$. Then

$$\Pr[\neg E_i] \geq 1 - x_i \cdot \prod_{j \in \Gamma_i} (1 - x_j) \geq 1 - x_i.$$

For $\Pr[E_k \mid \neg E_i]$,

- If E_k and E_i are mutually independent, namely $(i, k) \notin E$, then

$$\Pr[E_k \mid \neg E_i] = \Pr[E_k] \leq x_k \cdot \prod_{j \in \Gamma_k} (1 - x_j) \leq x_k.$$

- Otherwise $(i, k) \in E$, then

$$\Pr[E_k \mid \neg E_i] = \frac{\Pr[E_k \cap \neg E_i]}{\Pr[\neg E_i]} \leq \frac{\Pr[E_k]}{\Pr[\neg E_i]} \leq x_k \cdot \prod_{j \in \Gamma_k \setminus \{i\}} (1 - x_j) \leq x_k,$$

where the second inequality holds from $\Pr[\neg E_i] \geq 1 - x_i$.

WLOG let $S' = [1, s-1]$, $S = [1, s]$, and $s \geq 2$. Supposing the lemma holds for any S' with $|S'| = s-1$, then for S ,

$$\Pr[\bar{E}_S] = \Pr[\neg E_s \mid \bar{E}_{S'}] \cdot \Pr[\bar{E}_{S'}] \geq (1 - x_s) \cdot \prod_{i \in S'} (1 - x_i) = \prod_{i \in S} (1 - x_i),$$

which is immediate from induction hypothesis. Alternatively, the same procedure in remark 1.6.43 applies.

Let $A \subseteq S$ such that $\forall i \in A, (i, s+1) \in E$, and $B = S \setminus A$. If $B = S$, then E_{s+1} is mutually independent of $\{E_i\}_{i \in S}$,

$$\Pr[E_{s+1} \mid \bar{E}_S] = \Pr[E_{s+1}] \leq x_{s+1} \cdot \prod_{j \in \Gamma_{s+1}} (1 - x_j) \leq x_{s+1}.$$

Otherwise, by Bayes' rule,

$$\begin{aligned} \Pr[E_{s+1} \mid \bar{E}_S] &= \frac{\Pr[E_{s+1} \cap \bar{E}_S]}{\Pr[\bar{E}_S]} = \frac{\Pr[E_{s+1} \cap \bar{E}_A \cap \bar{E}_B]}{\Pr[\bar{E}_A \cap \bar{E}_B]} \\ &= \frac{\Pr[E_{s+1} \cap \bar{E}_A \mid \bar{E}_B]}{\Pr[\bar{E}_A \mid \bar{E}_B]}. \end{aligned}$$

For numerator, since E_{s+1} is mutually independent of $\{E_i\}_{i \in B}$ by the dependency graph, we have

$$\Pr[E_{s+1} \cap \bar{E}_A \mid \bar{E}_B] \leq \Pr[E_{s+1} \mid \bar{E}_B] = \Pr[E_{s+1}] \leq x_{s+1} \cdot \prod_{j \in \Gamma_{s+1}} (1 - x_j).$$

For denominator, let $|A| = t$, $A = \{a_i\}_{i \in [1, t]}$, and $A_t = \{a_i\}_{i \in [t, t]}$, then

$$\begin{aligned} \Pr[\bar{E}_A \mid \bar{E}_B] &= \prod_{i \in [1, t]} \Pr[\neg E_{a_i} \mid \bar{E}_{A_{i+1}} \cap \bar{E}_B] \\ &= \prod_{i \in [1, t]} \left(1 - \Pr[E_{a_i} \mid \bar{E}_{A_{i+1}} \cap \bar{E}_B]\right) \\ &\geq \prod_{i \in A} (1 - x_i) \\ &\geq \prod_{i \in \Gamma_{s+1}} (1 - x_i), \end{aligned}$$

where the first inequality holds from induction hypothesis, and the second inequality holds from $A \subseteq \Gamma_{s+1}$. Hence $\Pr[E_{s+1} \mid \bar{E}_S] \leq x_{s+1}$. $\square$

Since lemma 1.6.46 holds for any S , we finish the proof. $\square$

Lemma 1.6.47. For integer $x \geq 0$,

$$g(x) = \left(1 - \frac{1}{1+x}\right)^x$$

is monotone decreasing and lower bounded by e^{-1} .

Proof. Let $\varphi(x) = x \ln(1 - 1/(1+x))$, then the first order derivative

$$\varphi'(x) = \ln\left(1 - \frac{1}{1+x}\right) + x \cdot \frac{1}{1 - 1/(1+x)} \cdot \left(-\frac{-1}{(1+x)^2}\right) = \ln\left(1 - \frac{1}{1+x}\right) + \frac{1}{1+x}.$$

Let $t = 1 - 1/(1+x)$, then $\varphi'(x) = \ln t + 1 - t$.

Since $\ln t \leq t - 1$ for all t , then $\varphi'(x) \leq 0$, $g(x)$ is monotone decreasing. As $x \rightarrow \infty$, $g(x) \rightarrow e^{-1}$. $\square$

Exercise 1.6.48 (Exercise 6.18 [MU17]). Use the asymmetric LLL lemma 1.6.45 to show we can improve the symmetric LLL lemma 1.6.41 by replacing the condition $4pd \leq 1$ to $ep(d+1) \leq 1$.

Proof. Let each $x_i = 1/(d+1)$, then

$$x_i \cdot \prod_{j \in \Gamma_i} (1 - x_j) \geq x_i \cdot (1 - x_j)^d = \frac{1}{d+1} \cdot \left(1 - \frac{1}{d+1}\right)^d \geq \frac{1}{e(d+1)} \geq p,$$

where the second inequality holds from lemma 1.6.47.

p is directly applied in lemma 1.6.41, and the weakened condition is immediate. $\square$

Remark 1.6.49. Spencer proved this strengthened result in [Spe77].

Exercise 1.6.50 (Exercise 6.19 [MU17]). Let $G = (V, E)$ be an undirected graph and suppose each $v \in V$ is associated with a set $S(v)$ of $8r$ colors, where $r \geq 1$. Suppose, in addition, that for each $v \in V$ and $c \in S(v)$, there are at most r neighbors u of v with c lies in $S(u)$. Prove that there exists a coloring of G assigning to each vertex v a color from $S(v)$ such that, for any edge $(u, v) \in E$, the color assigned to u and v are different. Hint: use $A_{u,v,c}$ to be the event that u and v are both colored with c , then consider the family of such events.

Proof. Let X_u be the random number for vertex u 's color, and $c \in S(u)$, then $\Pr[X_u = c] = 1/8r$.

Let v be a neighboring vertex of u , since v is also associated with a set $S(v)$ of $8r$ colors, then $\Pr[X_v = c] \leq 1/8r$, as c might not be in $S(v)$. Therefore,

$$\Pr[X_u = c \wedge X_v = c] = \Pr[A_{u,v,c}] \leq \frac{1}{64r^2}.$$

For the dependency graph formed by events $\{A_{u,v,c}\}$, fixing a color $c \in S(u) \cap S(v)$ and an edge $(u, v) \in E$, the event $A_{u,v,c}$ is dependent on all $A_{u',v',c'}$ where either $u' = u$ or $v' = v$. Since u has at most $8r$ colors, and at most $8r^2$ neighbors may share colors with u , therefore the dependency graph degree $d \leq 16r^2$ by symmetry on v side.

By lemma 1.6.41, by $4dp \leq 4 \cdot 16r^2 \cdot 1/64r^2 \leq 1$, there is probability that all edges are colored differently by

$$\Pr \left[\bigcap_{u,v,c} \neg A_{u,v,c} \right] > 0.$$

$\square$

Exercise 1.6.51 (Exercise 6.20 [MU17]). A k -uniform hypergraph is an ordered pair $G = (V, E)$, but edges consist of sets of k (distinct) vertices, instead of just 2. (So a 2-uniform hypergraph is just what we normally call a graph.) A hypergraph is k -regular if all vertices have degree k ; that is, they are in k hypergraph edges.

Show that for sufficiently large k , the vertices of a k -uniform, k -regular hypergraph can be 2-colored so that no edge is monochromatic. What's the smallest value of k you can achieve?

Proof. We form a dependency graph for events of a hypergraph edge being monochromatic. Since the hypergraph is k -regular, then each vertex is appearing in k hypergraph edges. Excluding the hypergraph edge we are discussing, there are $k-1$ other hypergraph edges for the vertex, and therefore the dependency graph degree $d \leq k(k-1)$.

On the other hand, a k -uniform hypergraph edge is monochromatic with probability 2^{1-k} .

- We can try lemma 1.6.41 and by $4dp \leq 1$, then $k \geq 10$.
- We can also try the Spencer result in exercise 1.6.48 that $ep(d+1) \leq 1$, then $k \geq 9$.

$\square$

1.6.7 Algorithmic Lovász Local Lemma

Theorem 1.6.52 (Beck's k -Satisfiability Algorithm [Bec91]). *Consider a k -SAT formula with m clauses, where $k \geq 12$ is an even constant, and each variable appears in up to $2^{\alpha k}$ clauses for a sufficiently small constant $\alpha > 0$, then there is an algorithm finding a satisfying assignment for the formula in expected polynomial time in m .*

Proof. Let $\{x_i\}_{i \in [1, \ell]}$ be the variables, and $\{C_i\}_{i \in [1, m]}$ be the m clauses.

The algorithm has 2 phases: some variables are fixed in the first phase, and the remaining variables are deferred to the second phase. During the first phase variable fixing, a clause C_i is considered *dangerous* if C_i is not yet satisfied, and $k/2$ variables in the clause have been fixed.

In the first phase, the algorithm iterates through all variables, if the variable is not in any dangerous clause, assign it independently and uniformly a value from $\{0, 1\}$.

After the first phase, a clause is *surviving* if it is not yet satisfied, and such surviving clause has no more than $k/2$ variables fixed. A *deferred* variable is a variable not yet fixed in the first phase. In phase 2, we use exhaustive search to assign values to the deferred variables and so to complete a satisfying assignment for the formula.

Remark 1.6.53. The algorithm assigns a subset of variables in the first phase, and the remaining variables are deferred to the second phase. The subset of variables with values assigned in the first phase are chosen such that

- By the Local Lemma, the random partial solution fixed in phase 1 can be extended to a full satisfying assignment without modifying any phase 1 variable assignments.
- The dependency graph H' defined by the deferred variables in phase 2 is with high probability containing small connected subcomponents.

When the dependency graph consists only of small connected subcomponents, a solution for the variables of one component can be found independently of the other components. Therefore, such 2 phase algorithm first breaks the problem into smaller subproblems, and each smaller subproblem can be solved independently with exhaustive search in the phase 2.

We introduce following 2 lemmas for the 2 points in the prior remark.

Lemma 1.6.54. *There is an assignment to the deferred variables such that all the surviving clauses are satisfied.*

Proof. Let $H = (V, E)$ be the dependency graph on m nodes, where $V = [1, m]$, and $(i, j) \in E$ if and only if $C_i \cap C_j \neq \emptyset$. H is the dependency graph for the original problem. Let $H' = (V', E')$ be the dependency graph with $V' \subseteq V, E' \subseteq E$, with $i \in V'$ if and only if C_i is a surviving clause, and $(i, j) \in E'$ if and only if C_i and C_j share deferred variables.

Since a surviving clause has at least $k/2$ deferred variables, then the satisfying probability is at most $2^{-k/2}$. On the other hand, a variable appears up to $2^{\alpha k}$ clauses, then the degree of the dependency graph H' is $d \leq k2^{\alpha k}$.

By symmetric Local Lemma lemma 1.6.41, given constant $\alpha > 0$ that is sufficiently small, if $4 \cdot k2^{\alpha k} \cdot 2^{-k/2} \leq 1$, there exists an assignment that all surviving clauses are satisfied. $\square$

Lemma 1.6.55. *All connected components in H' are of size $O(\log m)$ with probability $1 - o(1)$.*

Proof. We start by bounding the probability of a given clause surviving phase 1. For a given node, it is dangerous with probability at most $2^{-k/2}$, that exactly $k/2$ variables in phase 1 are fixed but the clause is not satisfying. A clause is surviving, either itself is dangerous, or it is sharing at least a deferred variable with a dangerous clause, namely it is neighboring a dangerous clause in H' . Union bounding over the clause and its neighbors in H' , which is at most the dependency graph degree $d \leq k2^{\alpha k}$, a given node surviving phase 1 with probability at most $(d + 1)2^{-k/2}$.

The survival of individual clauses are not independent of each other, as 2 neighboring surviving clauses share at least a deferred variable. We consider a connected component $R \subseteq H$ of r vertices, then we identify a subset of vertices in R such that the survival of the clauses represented by these vertices are independent events.

Definition 1.6.56. A 4-tree S of a connected component R in H is defined as follows:

- S is a rooted tree.
- Any 2 nodes in S is at least distance 4 in H .
- There can be an edge in S only between 2 vertices of distance exactly 4 between them in H .
- Any node in R is either in S or in distance at most 3 from a node in S .

The motivation of introducing 4-tree is that, for any 2 clauses on S , the survival of the clauses are independent. Since any 2 neighboring clauses $u, v \in S$ are of distance exactly 4 in H , if both u and v are surviving, the clause causing u to survive is at least distance 2 in H from the clause causing v to survive. If 2 clauses are at least distance 2 in H , they are not sharing any variables, and their survival are surely independent. Supposing the survival of a large 4-tree is with low probability, this implies the survival of a large connected subcomponent with low probability.

For a 4-tree to survive, the probability is at most $((d+1)2^{-k/2})^{|S|}$. For a clause in a 4-tree, there are at most

$$d + d(d-1) + d(d-1)^2 \leq d^3 - 1$$

clauses in H that are within distance 3. Immediate by the expectation argument, there exists a 4-tree S in R with $|R| = r$ such that $|S| \geq r/d^3$.

The next step is to union bound over all possible sufficiently large 4-tree, such that we can derive an upper bound for the probability of existence of a sufficiently large 4-tree. A crude way to upper bound the number of 4-trees is, for each root, traverse the 4-tree in Euler cycle by visiting each path in 2 directions. There are at most $m(d^4)^{2r/d^3}$ number of traversals by picking a vertex every 4 vertices until $2r/d^3$ vertices are picked, over m roots.

The union bound is

$$md^{8r/d^3} \cdot ((d+1)2^{-k/2})^{r/d^3},$$

when $r \geq c \log m$ for a sufficiently large constant $c > 0$ and a sufficiently small constant $\alpha > 0$, the bound is $o(1)$. $\square$

Therefore, the proof is immediate by lemma 1.6.54 and lemma 1.6.55, that phase 1 partitions the problem into subproblems with $O(k \log m)$ variables with probability $1 - o(1)$, then we just need to run constant number of times of phase 1 before we obtain a good partition. $\square$

We now discuss the constructive proof by Moser [Mos09]. A simpler version was presented in his STOC'09 talk, which was based on the entropy compression argument. The following analysis follows the “entropy compression” idea, which appears explicitly in his doctoral thesis [Mos12] (see the “Incompressibility” section), rather than being the primary framing in the STOC'09 paper.

At the core of the “entropy compression” argument, it is arguing that the algorithm’s expected runtime must be small, otherwise we are able to compress the random bits used by the algorithm significantly.

We begin with a useful fact on the incompressibility of random strings.

Fact 1.6.57 (Incompressibility Principle). For injective $f : \{0, 1\}^t \mapsto \{0, 1\}^*$, if $\mathbf{s} \xleftarrow{\mathbf{R}} \{0, 1\}^t$, then for any integer c ,

$$\Pr [\text{the length of } f(\mathbf{s}) \leq t - c] \leq 2^{1-c}.$$

Proof. There are in total $2^{t-c+1} - 1$ binary strings of length at most $t - c$, while there are 2^t possible inputs in $\{0, 1\}^t$. Hence, if $\mathbf{s} \xleftarrow{\mathbf{R}} \{0, 1\}^t$, then the probability of f outputting at most $t - c$ bits is at most 2^{1-c} . $\square$

There are $2^t - 1$ binary strings of length at most $t - 1$, and the average length is

$$\sum_{i \in [0, t-1]} i \cdot \frac{2^i}{2^t - 1} \geq \sum_{i \in [0, t-1]} i \cdot \frac{2^i}{2^t} \geq t - 2.$$

Therefore, the expected compression length over $\mathbf{s} \xleftarrow{\mathbf{R}} \{0, 1\}^t$ is lower bounded by $t - 2$.

Theorem 1.6.58 (Moser’s k -Satisfiability Algorithm [Mos12]). *For a k -SAT formula with m clauses, and each clause shares variables with at most $2^{k-3} - 1$ other clauses. Then a satisfying assignment can be found in expected polynomial time in m .*

The “entropy compression” argument is rather clean in the specific setting of k -SAT.

We let $\{x_i\}_{i \in [1, \ell]}$ be the ℓ variables, $\{C_i\}_{i \in [1, m]}$ be the m clauses, Γ_i be the set of indices of clauses sharing variables with the clause C_i , Γ_i^+ be the inclusive neighborhood $\Gamma_i \cup \{i\}$, and V_i be the set of indices of variables in C_i .

Informally, the algorithm picks a uniformly random assignment, then look for an unsatisfied C_i . If such C_i exists, we sample a new assignment for the variables in V_i uniformly at random. Doing so may fix C_i , but it may end up unsatisfying C_j for some $j \in \Gamma_i^+$. The algorithm recursively fix the neighboring clauses, and by the end of the recursion, C_i is satisfied with no neighboring clauses damaged. Therefore, the situation is improved by satisfying at least 1 previously unsatisfied clause. Iterating through all m clauses, the algorithm finishes the fixing.

This leads to the question that when will the recursion stop, and is solved by the “entropy compression” argument. We describe the algorithm formally as follows, and bound the expected runtime.

The Moser’s algorithm has 2 functions: `solve` and `local-fix`.

- `solve`
 - Sample a uniformly random variable assignment for $\{x_i\}_{i \in [1, \ell]}$.
 - While there is C_i unsatisfied
 - * Choose the unsatisfied C_i with the smallest i .
 - * Enter i into the execution log with $\lceil \log_2 m \rceil$ bits.
 - * Call `local-fix` on the clause C_i .
- `local-fix`
 - Sample a uniformly random variable assignment for $\{x_j\}_{j \in \Gamma_i}$.
 - While there is $k \in \Gamma_i^+$ such that C_k is unsatisfied
 - * Choose the unsatisfied C_k with the smallest k .
 - * Enter “0” together with k in $k - 3$ bits into the execution log.
 - * Call `local-fix` on the clause C_k .
 - * Enter “1” into the execution log.

Proof. There are 2 distinct ways of describing how the algorithm works.

We can think of the algorithm as being described by the random string of bits it used. It takes ℓ bits for the initial variable assignment at random, then it takes k bits to resample variables for a clause each time `local-fix` is called. We refer to each time `local-fix` is called as a round. Hence, a way to describe the algorithm’s action for j rounds is with the random strings of $\ell + jk$ bits used by the algorithm.

We can also think of the algorithm as being described as the history, the bits in its execution log. It includes the list of clauses called in the main routine and recursion of `local-fix`. For recursion, the log starts with a bit 0 and ends with a bit 1 to mark the start and the end of recursion calls. For the index of the clause in the recursion call, we use $k - 3$ bits instead of $\lceil \log_2 m \rceil$ bits, as $|\Gamma_i^+| \leq 2^{k-3}$. The ℓ bits of final variable assignments are also included in the history, serving as an updatable state: each time a clause’s variables being resampled, the state is updated. Hence, the output execution log for the algorithm’s action for j rounds takes at most $\ell + m \lceil \log_2 m \rceil + j(k - 1)$, as a round of clause variable resampling outputs $k - 1$ bits: 2 bits for start/end marking, and $k - 3$ bits for clause index.

Given $\ell + kj$ resampling bits used by the algorithm, and $\ell + m \lceil \log_2 m \rceil + j(k - 1)$ bits of output execution log, we can view the algorithm as a compression algorithm over random bits. We now show how to recover the resampling bits from the execution log. The execution log determines a sequence of clauses visited by the algorithm, and each visited clause must be unsatisfied before resampling. Only 1 setting of variables unsatisfy a clause. Hence, we work backwards from the final variable assignments. On each clause, we update with unsatisfying variable assignment, such that we know the resampling variables, and therefore we recover the $\ell + jk$ bits of resampling bits.

After proving the compression is injective, now we use the compression argument to bound the expected runtime. Let J be the random variable for the number of rounds in the algorithm. The algorithm takes $\ell + Jk$ uniformly random bits as inputs, outputting at most $\ell + m \lceil \log_2 m \rceil + J(k-1)$ bits of history. By fact 1.6.57,

$$\ell + m \lceil \log_2 m \rceil + J(k-1) \geq \ell + kJ - 1,$$

and when $J = m \lceil \log_2 m \rceil + 1 + i$,

$$\Pr [\text{output length} \leq \ell + m \lceil \log_2 m \rceil + J(k-1)] \leq 2^{-i}.$$

Therefore, $\mathbb{E}[J]$ can be bounded as follows

$$\mathbb{E}[J] \leq m \lceil \log_2 m \rceil + 1 + \sum_{i \geq 0} \frac{i}{2^i} = m \lceil \log_2 m \rceil + 3.$$

Therefore, the expected runtime of the algorithm is $O(m \log m)$. $\square$

Exercise 1.6.59 (Exercise 6.20 [MU17]). In the k -Satisfiability Algorithm using the algorithmic Lovász Local Lemma, we used $\lceil \log_2 m \rceil$ bits in history to represent each clause called in the main routine. Instead, we could record in the history which clauses are initially unsatisfied with an array of m bits. Explain any other changes to the algorithm to properly record a history that one can “reverse” to obtain the initial assignment, and explain how this allows one to modify the proof of theorem 1.6.58 so that only $O(m)$ rounds are needed in expectation.

Proof. Observing each run of `local-fix`, the subroutine visits an unsatisfied clause, resamples its variables, then fix neighboring unsatisfied clauses recursively. The invariant is, all of the visited clauses and their neighboring clauses are satisfied by the end of the recursion.

If we put m bits for the initially unsatisfied clauses, then we know the first unsatisfied clause visited by `local-fix`. We maintain a list of satisfied clauses by the first call to the `local-fix`, update the m bits for unsatisfied clauses, and move on to the remaining unsatisfied clause with smallest index. Repeating $O(m)$ times, we have a full list of visited clauses from history.

Therefore, we use m bits in history for visited unsatisfied clauses from the main routine rather than $m \lceil \log_2 m \rceil$, and therefore replacing all appearances of $m \lceil \log_2 m \rceil$ in the argument, we have the expected runtime being $O(m)$. $\square$

The Algorithmic Lovász Local Lemma by Moser and Tardos [MT10] states that: in the variable-based LLL setting, an LLL existence proof is automatically (Las Vegas) constructive via resampling, with expected polynomial runtime. We describe the asymmetric algorithmic LLL, and the symmetric algorithmic LLL follows from exercise 1.6.48.

Theorem 1.6.60 (Asymmetric Algorithmic Lovász Local Lemma [MT10]). *Let $\{E_i\}_{i \in [1, n]}$ be a set of events in arbitrary probability space that are determined by mutually independent random variables $\{y_i\}_{i \in [1, \ell]}$, and let $G = (V, E)$ be the dependency graph for the events. Suppose there are $\{x_i\}_{i \in [1, n]}$ in $[0, 1]$ such that for $i \in [1, n]$,*

$$\Pr[E_i] \leq x_i \cdot \prod_{j \in \Gamma_i} (1 - x_j),$$

then there exists an assignment $\{y_i\}$ such that $\bar{E}_{[1, n]}$ holds, and a resampling algorithm whose expected number of resampling E_i in finding an assignment is at most $x_i / (1 - x_i)$. Hence, the expected runtime is at most $\sum_i x_i / (1 - x_i)$.

Corollary 1.6.61 (Symmetric Algorithmic Lovász Local Lemma [MT10]). *Let $\{E_i\}_{i \in [1, n]}$ be a set of events in arbitrary probability space that are determined by mutually independent random variables $\{y_i\}_{i \in [1, \ell]}$, and let $G = (V, E)$ be the dependency graph for the events. Supposing the degree of G is at most d , $\Pr[E_i] \leq p$ for $i \in [1, n]$, and $ep(d+1) \leq 1$, then there exists an assignment $\{y_i\}$ such that $\bar{E}_{[1, n]}$ holds, and a resampling algorithm whose expected number of resampling E_i in finding an assignment is at most $1/d$. Hence, the expected runtime is at most n/d .*

The Moser-Tardos algorithmic LLL is even simpler than Moser's k -SAT algorithm, and the algorithm is described as follows. Onwards we say each event E_i as bad events, and the algorithms avoids all the bad events, eventually.

- Moser-Tardos
 - Sample a uniformly random variable assignments for $\{y_i\}$.
 - While there is E_i happening
 - * Randomly pick an happening E_i .
 - * Resample all y_i determines this E_i .

We define $\mathbf{c} \in [1, n]^k$ as the execution log, where the algorithm has k rounds of resampling over events $\{E_i\}$, and each c_i is an index for n events $\{E_i\}$. At the core of the argument is the construction and analysis of a witness tree from resampling each event E_{c_i} in the execution log.

Definition 1.6.62 (Witness Tree). A witness tree is constructed from execution log $\mathbf{c}$ as follows.

- Let c_k be the root.
- Let $T_{j,k}$ be a witness tree from $c_j, \dots, c_k$, and let j goes backwards from $k-1$ to 1.
 - If $c_j \in \Gamma_o^+$ for some $o \in T_{j+1,k}$, then add c_j as the children of the deepest occurrence of such event.
 - If E_{c_j} do not share variables with any event in $T_{j+1,k}$, then drop c_j and move on by $T_{j,k} \leftarrow T_{j+1,k}$.
- Eventually, $T_k \leftarrow T_{1,k}$.

The following facts are immediate by the construction of the witness tree.

Fact 1.6.63. For any events E_i and E_j sharing variables, they are appearing in the same depth of a witness tree.

Fact 1.6.64. A variable x_i is not appearing twice in the same depth of a witness tree.

Intuitively, a witness tree is a tree justifying how we end up selecting E_{c_i} to resample all variables in V_{c_i} .

We introduce an additional fact, that the witness trees from different prefixes of an execution log are different.

Fact 1.6.65. For execution logs $\mathbf{c}_0$ of length k_0 and $\mathbf{c}_1 = \mathbf{c}_0 \parallel \mathbf{c}'$ of length k_1 , the execution trees T_{k_0} and T_{k_1} are different.

Proof. Let the roots from $\mathbf{c}_0$ and $\mathbf{c}_1$ be r_0 and r_1 .

If $r_0 \neq r_1$, then $T_{k_0} \neq T_{k_1}$ by root. Otherwise, $\mathbf{c}_1$ has at least one more occurrence of r_0 , resulting in $T_{k_0} \neq T_{k_1}$. $\square$

Hence, the next lemma for counting the occurrence of resampling E_i in $\mathbf{c}$ is immediate by fact 1.6.65.

Lemma 1.6.66. Let X_i be the random variable of occurrence of resampling E_i in $\mathbf{c}$, and $X_{i,j}$ be the 0-1 random variable indicating the witness tree T_j from $\mathbf{c}_{[1,j]}$, the prefix of $\mathbf{c}$, has root i . Then

$$X_i = \sum_{j \in [1,k]} X_{i,j}.$$

Moreover, let W_i be the set of witness trees rooted by i , then by linearity of expectation,

$$\mathbb{E}[X_i] = \sum_{j \in [1,k]} \mathbb{E}[X_{i,j}] = \sum_{j \in [1,k]} \sum_{T \in W_i} \Pr[T = T_j] = \sum_{T \in W_i} \sum_{j \in [1,k]} \Pr[T = T_j].$$

We want to bound $\Pr[T_j = T]$. We say T appears in $\mathbf{c}$, if $T = T_j$ for T_j constructed from some $\mathbf{c}_{[1,j]}$ prefix of $\mathbf{c}$. Hence, fixing a T , since $\{T_j\}$ are all distinct by fact 1.6.65, $T = T_j$ are all independent, and

$$\Pr[T \text{ appears in } \mathbf{c}] = \sum_{j \in [1,k]} \Pr[T = T_j].$$

Lemma 1.6.67.

$$\Pr [T \text{ appears in } \mathbf{c}] \leq \prod_{j \in T} \Pr [E_j],$$

where the T in the right hand side production is a multiset, such that we capture multiple occurrence of an event being resampled.

Proof. If we traverse the witness tree T in reverse BFS method, namely layering by depth bottom up, then we observe earlier events resampling induces later events resampling, following the Moser-Tardos algorithm's execution order.

For variables $\{y_i\}$, each resampling creates a sequence of assignments, say $y_i^1, \dots, y_i^s$. By fact 1.6.64, we know that y_i^a and y_i^b with $a \neq b$ should be on different depth.

Hence, fixing the randomness from Moser-Tardos algorithm, which is fixing the assignment sequences $\{\{y_i^s\}_s\}_i$, we form a checking algorithm for T as follows. The check algorithm traverses in described way, and checks if events in the current layer are happening: aborts if any event is not happening. Before traversing the next depth, the check algorithm resamples all the variables of the events in the current layer, by moving each y_i^s to y_i^{s+1} .

It becomes obvious that, on input the witness tree T from a successful run of Moser-Tardos algorithm, by fixing the assignment sequences for $\{\{y_i^s\}_s\}_i$, the tree checking algorithm always passes. Therefore

$$\Pr [T \text{ appears in } \mathbf{c}] \leq \Pr [\text{check}(T) \text{ succeeds}].$$

To bound the probability of a tree being valid in the checking algorithm, we use fact 1.6.63 and fact 1.6.64, that the events are not sharing variables in a same depth, such that events in a layer are independent, and each layer is independent of each other. Hence, a tree is valid in the checking algorithm with probability $\prod_{j \in T} \Pr [E_j]$.

Thus, the inequality holds by the coupling arguments above. $\square$

Now we have derived

$$\mathbb{E}[X_i] = \sum_{T \in W_i} \Pr [T \text{ appears in } \mathbf{c}] \leq \sum_{T \in W_i} \prod_{j \in T} \Pr [E_j]$$

by lemma 1.6.66 and lemma 1.6.67, the final step would be bounding the right hand side sum of product. We derive the upper bound by Galton-Watson process, that each children of E_i in Γ_i^+ spawns with probability x_i independently.

Lemma 1.6.68. *Deriving T rooted with i with Galton-Watson process has probability*

$$p_T = \frac{1 - x_i}{x_i} \cdot \prod_{j \in T} \Pr [E_j].$$

Proof. Let each node E_j has a set $S_j \subseteq \Gamma_j^+$ for the indices of children in T , and the probability derivation follows

$$\begin{aligned} p_T &= \prod_{j \in T} \left(\prod_{k \in S_j} x_k \cdot \prod_{k \in \Gamma_j^+ \setminus S_j} (1 - x_k) \right) = \prod_{j \in T} \left(\prod_{k \in S_j} \frac{x_k}{1 - x_k} \cdot \prod_{k \in \Gamma_j^+} (1 - x_k) \right) \\ &= \frac{1 - x_i}{x_i} \cdot \prod_{j \in T} \left(x_j \cdot \prod_{k \in \Gamma_j^+} (1 - x_k) \right) = \frac{1 - x_i}{x_i} \cdot \prod_{j \in T} \Pr [E_j], \end{aligned}$$

where the third equality holds from moving all the $x_k/(1 - x_k)$ terms for $k \in S_j$ from E_j to all the children in S_j . $\square$

The proof for theorem 1.6.60 follows immediately:

Proof. By lemma 1.6.68, we have

$$\mathbb{E}[X_i] = \sum_{T \in W_i} \Pr [T \text{ appears in } \mathbf{c}] \leq \sum_{T \in W_i} \prod_{j \in T} \Pr [E_j] = \frac{x_i}{1 - x_i} \cdot \sum_{T \in W_i} p_T \leq \frac{x_i}{1 - x_i},$$

and $\sum_T p_T \leq 1$ as we are summing over disjoint events. The rest is immediate by the linearity of expectation. $\square$

1.7 Markov Chains and Random Walks

1.7.1 Stirling's Formula

After previous bounds for factorial by corollary 1.5.5 and corollary 1.6.3, we reach Stirling's Formula.

Lemma 1.7.1 (Stirling's Formula). *For $n > 0$,*

$$n! = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n (1 \pm o(1)).$$

In particular, for $n > 0$,

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n \leq n! \leq (\sqrt{2\pi n} + 1) \left(\frac{n}{e}\right)^n.$$

Proof. We start from the fact that factorial is the discrete form of the Gamma function, by

$$n! = \Gamma(n+1) = \int_0^\infty t^n e^{-t} dt.$$

Let $t = n(1+x)$ where $x > -1$, then

$$\begin{aligned} \Gamma(n+1) &= \int_0^\infty t^n e^{-t} dt \\ &= \int_{-1}^\infty (n(1+x))^n \cdot e^{-n(1+x)} \cdot n dx \\ &= n \left(\frac{n}{e}\right)^n \int_{-1}^\infty (1+x)^n e^{-nx} dx \\ &= n \left(\frac{n}{e}\right)^n \int_{-1}^\infty \exp(n(\ln(1+x) - x)) dx \\ &= n \left(\frac{n}{e}\right)^n \int_{-1}^\infty \exp(n \cdot f(x)) dx, \end{aligned}$$

where $dt = n dx$, and we let $f(x) = \log(x+1) - x$. We continue by series expansion:

- $f'(x) = (1+x)^{-1} - 1 = -x/(1+x)$, and $f'(0) = 0$;
- $f''(x) = -(1+x)^{-2}$, and $f''(0) = -1$.

Hence, $f(x) = -x^2/2 + O(x^3)$. We let $-z^2/2 = f(x)$, where

$$z = \operatorname{sgn}(x) \cdot \sqrt{2(x - \ln(x+1))},$$

then by $df(x) = f'(x) dx = -z dz$, $dx = (z + z/x) dz$, and by replacing dx with dz ,

$$\Gamma(n+1) = n \left(\frac{n}{e}\right)^n \int_{-\infty}^\infty \exp\left(-\frac{nz^2}{2}\right) \cdot \left(z + \frac{z}{x}\right) dz.$$

We continue by bounding the term z/x with results from [Top07].

Lemma 1.7.2 ($\log(x+1)$ bounds [Top07]). *For $x > -1$,*

$$\frac{x(6+x)}{6+4x} \geq \ln(x+1) \geq \begin{cases} x - \frac{x^2}{2} & x \geq 0 \\ \frac{x(2+x)}{2+2x} & x \leq 0 \end{cases}.$$

With lemma 1.7.2, we need to prove the following bound to complete the proof.

Corollary 1.7.3. *If $x > -1$, then $1 + 2z/3 \leq z + z/x \leq \max(1, 1 + z)$.*

Proof. First, $z(1 + 1/x) > 0$, since $-1 < x \leq 0$, both z and $1 + 1/x$ are less than 0. Hence, we can bound by $z^2(1 + 1/x)^2$.

When $-1 < x \leq 0$, by lemma 1.7.2

$$z^2 = 2(x - \ln(x + 1)) \leq \frac{x^2}{1 + x} \leq \left(\frac{x}{1 + x}\right)^2,$$

and therefore $z^2(1 + 1/x)^2 \leq 1$. Hence, $z + z/x \leq 1$.

When $x \geq 0$,

$$z^2 = 2(x - \ln(x + 1)) \leq x^2,$$

and therefore $z^2/x^2 \leq 1$. Hence, $z + z/x \leq 1 + z$.

Hence, the upper bound is $\max(1, 1 + z)$.

For lower bound, we have

$$z^2 = 2(x - \ln(x + 1)) \geq \frac{3x^2}{3 + 2x},$$

and

$$\frac{3x^2}{3 + 2x} \cdot \left(\frac{1}{3} + \frac{1}{x}\right)^2 = \frac{x^2}{9 + 6x} + 1 \geq 1.$$

Hence, the lower bound is $1 + 2z/3$, which completes the proof. $\square$

Since

$$\int_{-\infty}^{\infty} \exp(-z^2) dz = \sqrt{\pi} \text{ and } \int_0^{\infty} z \cdot \exp(-z^2) dz = \frac{1}{2},$$

then

$$\int_{-\infty}^{\infty} \exp\left(-\frac{nz^2}{2}\right) dz = \sqrt{\frac{2\pi}{n}} \text{ and } \int_0^{\infty} z \cdot \exp\left(-\frac{nz^2}{2}\right) dz = \frac{1}{n},$$

and therefore the upper bound is derived by,

$$\begin{aligned} \Gamma(n + 1) &= n \left(\frac{n}{e}\right)^n \int_{-\infty}^{\infty} \exp\left(-\frac{nz^2}{2}\right) \cdot \left(z + \frac{z}{x}\right) dz \\ &\leq n \left(\frac{n}{e}\right)^n \cdot \left(\sqrt{\frac{2\pi}{n}} + \frac{1}{n}\right) = (\sqrt{2\pi n} + 1) \left(\frac{n}{e}\right)^n, \end{aligned}$$

and the lower bound is derived by

$$\begin{aligned} \Gamma(n + 1) &= n \left(\frac{n}{e}\right)^n \int_{-\infty}^{\infty} \exp\left(-\frac{nz^2}{2}\right) \cdot \left(z + \frac{z}{x}\right) dz \\ &\geq n \left(\frac{n}{e}\right)^n \int_{-\infty}^{\infty} \exp\left(-\frac{nz^2}{2}\right) \cdot \left(1 + \frac{2z}{3}\right) dz \\ &= n \left(\frac{n}{e}\right)^n \int_{-\infty}^{\infty} \exp\left(-\frac{nz^2}{2}\right) dz \\ &= n \left(\frac{n}{e}\right)^n \cdot \sqrt{\frac{2\pi}{n}} = \sqrt{2\pi n} \left(\frac{n}{e}\right)^n, \end{aligned}$$

which completes the proof ². $\square$

²This derivation is based on Thomas Ahle's post in <https://mathoverflow.net/a/484270>.

1.7.2 Markov Chains and Classification of States

Definition 1.7.4. The stochastic process $\mathbf{X} = \{X_t\}_{t \in T}$ is a collection of random variables. The index t represents time, and the process $\mathbf{X}$ models how a random variable changes over time.

X_t is the state of the process at the t^{th} time step.

- If X_t takes values from a countably infinite set, $\mathbf{X}$ is a discrete-state process.
- If X_t takes values from a finite set, then $\mathbf{X}$ is finite-state.
- If T is a countably infinite set, $\mathbf{X}$ is a discrete-time process.

A discrete stochastic process $X_0, X_1, \dots$ is a time-homogeneous Markov chain if

$$\Pr[X_t = a_t \mid X_{t-1} = a_{t-1}, \dots, X_0 = a_0] = \Pr[X_t = a_t \mid X_{t-1} = a_{t-1}] = P_{a_{t-1}, a_t}.$$

This is called the Markov property, or memoryless property, or Markovian.

Exercise 1.7.5 (Exercise 7.2 [MU17]). Consider the 2-state Markov chain with the following transition matrix

$$\mathbf{P} = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}.$$

Find the minimum expression for $P_{0,0}^t$ in $\mathbf{P}^t$.

Proof. We can derive the following recurrence relation

$$P_{0,0}^t = \langle [P_{0,0}^{t-1}, 1 - P_{0,0}^{t-1}]^T, [p, 1-p]^T \rangle = (2p-1) \cdot P_{0,0}^{t-1} - (p-1).$$

as the $\mathbf{P}^{t-1}$ is diagonally symmetric with same diagonal values. Therefore, $P_{0,0}^t = (1 + (2p-1)^t) / 2$. $\square$

We present a 2-SAT randomized algorithm from Markov chain by [Pap91] that is similar to theorem 1.6.60:

- Start with a random truth assignment.
- Repeat up to $m \cdot 2n^2$ times, terminating on all clauses satisfied:
 - Choose an arbitrary unsatisfied clause.
 - Choose uniformly at random a variable in the clause, and flip it.
- If the loop ends with a satisfying assignment, return it; otherwise, abort.

Lemma 1.7.6 (2-SAT randomized algorithm expected runtime [Pap91]). Assume the 2-SAT formula is satisfiable, and the 2-SAT randomized algorithm is allowed to run until it finds a satisfying assignment. Then the expected number of steps until the algorithm finds a satisfying assignment is at most n^2 .

Proof. One can view the state transition graph as a line of $n+1$ vertices, from 0 to n .

Suppose the algorithm outputs a satisfying assignment S , and we let the initial random assignment be A_0 , A_t be the assignment on the t^{th} step. We define the process $\{X_t\}_{t \in T}$ over the number of matches of A_t with S .

- When $X_{t-1} = 0$, any flip of variable can make A_{t-1} have more match, and thus $\Pr[X_t = 1 \mid X_{t-1} = 0] = 1$.
- When $X_{t-1} = j$, then
 - If the clause picked has 1 unmatched variable, then the probability of making situation better is $1/2$.
 - If the clause picked has 2 unmatched variables, then any flip can make situation better.

Hence, $\Pr[X_t = j+1 \mid X_{t-1} = j] \geq 1/2$.

But this goes against the Markovian property, namely memoryless, as the state transition probability should only be tied to the time step and the states. But in the current model, the probability varies on the number of disagreeing variables in a chosen clause. Hence, we define the following process $\{Y_t\}_{t \in T}$ that $Y_0 = X_0$, with

$$\begin{aligned}\Pr[Y_t = 1 \mid Y_{t-1} = 0] &= 1, \\ \Pr[Y_t = j + 1 \mid Y_{t-1} = j] &= 1/2, \\ \Pr[Y_t = j - 1 \mid Y_{t-1} = j] &= 1/2.\end{aligned}$$

Let $\{Z_i\}_{i \in [0, n]}$ be the random variables for the number of steps to reach n from i matches in process $\{Y_t\}_{t \in T}$. The following claims are immediate:

$$\begin{aligned}\mathbb{E}[Z_n] &= 0, \\ \mathbb{E}[Z_j] &= \mathbb{E}[Z_j \mid Y_{t+1} = j - 1] \cdot P_{j, j-1} + \mathbb{E}[Z_j \mid Y_{t+1} = j + 1] \cdot P_{j, j+1} \\ &= (1 + \mathbb{E}[Z_{j-1}]) / 2 + (1 + \mathbb{E}[Z_{j+1}]) / 2 \\ &= 1 + (\mathbb{E}[Z_{j-1}] + \mathbb{E}[Z_{j+1}]) / 2, \\ \mathbb{E}[Z_0] &= \mathbb{E}[Z_1] + 1.\end{aligned}$$

It is immediate that $\mathbb{E}[Z_1] = \mathbb{E}[Z_2] + 3$. By induction, we derive the following recurrence relation

$$\mathbb{E}[Z_j] = \mathbb{E}[Z_{j+1}] + (2j + 1).$$

By summation, $\mathbb{E}[Z_0] = n^2$, and $\mathbb{E}[Z_j] = n^2 - j^2$. □

The following theorem is immediate from lemma 1.7.6.

Theorem 1.7.7. *The 2-SAT algorithm always returns correct answer on an unsatisfiable formula. If the formula is satisfiable, with probability at least $1 - 2^{-m}$ the algorithm gives a satisfiable assignment.*

Proof. One can segment the $m \cdot 2n^2$ steps in the algorithm into m chunks of $2n^2$ steps, running independently.

Let Z be the random variable for the steps the algorithm needs to take from i matches. By lemma 1.7.6, $\mathbb{E}[Z] \leq n^2$. Immediate from Markov's inequality,

$$\Pr[Z \geq 2n^2] \leq \frac{\mathbb{E}[Z]}{2n^2} \leq \frac{1}{2}.$$

Hence, the probability that the algorithm fails to find a satisfying assignment after m segments is at most 2^{-m} . □

Exercise 1.7.8 (Exercise 7.6 [MU17]). Consider a partially reflecting boundary at 0, rather than a complete reflecting boundary at 0 in lemma 1.7.6. At position 0, with probability 1/2 the walk moves to position 1 and with probability 1/2 the walk stays at 0. Everywhere else is the same. Find the expected number of moves to reach n from position i .

Proof. We continue with reusing $\{Z_i\}_{i \in [0, n]}$ and process $\{Y_t\}_{t \in T}$. The following claims are immediate:

$$\begin{aligned}\mathbb{E}[Z_n] &= 0, \\ \mathbb{E}[Z_j] &= \mathbb{E}[Z_j \mid Y_{t+1} = j + 1] \cdot P_{j, j+1} + \mathbb{E}[Z_j \mid Y_{t+1} = j - 1] \cdot P_{j, j-1} \\ &= 1 + (\mathbb{E}[Z_{j+1}] + \mathbb{E}[Z_{j-1}]) / 2, \\ \mathbb{E}[Z_0] &= \mathbb{E}[Z_0 \mid Y_{t+1} = 0] \cdot P_{0,0} + \mathbb{E}[Z_0 \mid Y_{t+1} = 1] \cdot P_{0,1} \\ &= 1 + (\mathbb{E}[Z_0] + \mathbb{E}[Z_1]) / 2.\end{aligned}$$

It is immediate that $\mathbb{E}[Z_1] = \mathbb{E}[Z_2] + 4$. By induction, we derive the following recurrence relation

$$\mathbb{E}[Z_j] = \mathbb{E}[Z_{j+1}] + (2j + 2).$$

By summation, $\mathbb{E}[Z_n] = n(n + 1)$, and $\mathbb{E}[Z_j] = n(n + 1) - j(j + 1)$. □

Exercise 1.7.9 (Exercise 7.7 [MU17]). Find the expected runtime of the 2-SAT randomized algorithm for lemma 1.7.6, where the input is sampled uniformly at random.

Proof. We follow the $\{Z_j\}_{j \in [0,n]}$ notions of lemma 1.7.6, and let Y be the random variable for the number of mismatches in the initial assignment. Then by conditional expectation,

$$\begin{aligned}
 \mathbb{E}[Z] &= \sum_{j \in [0,n]} \mathbb{E}[Z \mid Y = j] \cdot \Pr[Y = j] = \sum_{j \in [0,n]} \mathbb{E}[Z_j] \cdot \binom{n}{j} \cdot \frac{1}{2^n} = \sum_{j \in [0,n]} (n^2 - j^2) \cdot \binom{n}{j} \cdot \frac{1}{2^n} \\
 &= n^2 - \frac{1}{2^n} \sum_{j \in [0,n]} j^2 \cdot \binom{n}{j} = n^2 - \frac{n}{2^n} \sum_{j \in [1,n]} j \cdot \binom{n-1}{j-1} \\
 &= n^2 - \frac{n}{2} - \frac{n}{2^n} \sum_{k \in [0,n-1]} k \cdot \binom{n-1}{k} = n^2 - \frac{n}{2} - \frac{n(n-1)}{2^n} \sum_{\ell \in [0,n-2]} \binom{n-2}{\ell} \\
 &= n^2 - \frac{n}{2} - \frac{n(n-1)}{4} = \frac{3}{4}n^2 - \frac{1}{4}n.
 \end{aligned}$$

□

After lemma 1.7.6, we present a 3-SAT randomized algorithm from Markov chain. A prototype as follows:

- Pick a random truth assignment.
- Repeat up to m times, terminating on all clauses satisfied:
 - Choose an arbitrary unsatisfied clause.
 - Choose uniformly at random a variable in the clause, and flip it.
- If the loop ends with a satisfying assignment, return it; otherwise, abort.

Again after lemma 1.7.6, we define the Markov chain by the following process $\{Y_t\}_{t \in T}$ by

$$\begin{aligned}
 \Pr[Y_t = 1 \mid Y_{t-1} = 0] &= 1, \\
 \Pr[Y_t = j+1 \mid Y_{t-1} = j] &= 1/3, \\
 \Pr[Y_t = j-1 \mid Y_{t-1} = j] &= 2/3.
 \end{aligned}$$

We proceed by defining the random variables $\{Z_i\}_{i \in [0,n]}$ for the number of steps to reach n matches for the current assignment with a satisfying assignment S , where the current assignment has i matches. By conditional expectation,

$$\begin{aligned}
 \mathbb{E}[Z_n] &= 0, \\
 \mathbb{E}[Z_j] &= \mathbb{E}[Z_j \mid Y_{t+1} = j-1] \cdot P_{j,j-1} + \mathbb{E}[Z_j \mid Y_{t+1} = j+1] \cdot P_{j,j+1} \\
 &= 1 + 2/3 \cdot \mathbb{E}[Z_{j-1}] + 1/3 \cdot \mathbb{E}[Z_{j+1}], \\
 \mathbb{E}[Z_0] &= 1 + \mathbb{E}[Z_1].
 \end{aligned}$$

By induction, $\mathbb{E}[Z_j] = \mathbb{E}[Z_{j+1}] + 2^{j+2} - 3$. But this results in $\mathbb{E}[Z_0] = O(2^n)$, where trying all n variables is also 2^n , which is not impressive. The following 2 observations are helpful:

- It is more likely to make more violation when choosing a random unsatisfied clause and flip a variable at random. Hence, on a random initial assignment, it seems better to try a smaller number of repetition, as it is more likely to drift to a larger number of matches, than a rather large number of repetition.
- Initial assignment of variables with uniform randomness can be viewed as independent Bernoulli trials. Hence, we have some chance of sampling an initial assignment with matches higher than $n/2$.

Based on the observation, it seems better to try for a smaller amount of repetitions on a random initial assignment, and try over a large variety of different random initial assignments.

Given an initial assignment with j mismatches, we let q_j be the probability that the random walk finds a satisfying assignment from this initial assignment in $3n$ steps. The random walk always needs $2k + j$ steps to reach n matches. To lower bound the probability that a satisfying assignment can be found within $2k + j$ steps, we consider $k + j$ steps drifting to more matches, while k steps drifting to less matches, by

$$p_k = \binom{2k+j}{k} \left(\frac{1}{3}\right)^{k+j} \left(\frac{2}{3}\right)^k = \binom{2k+j}{k} \cdot \left(\frac{1}{3}\right)^{2k+j} \cdot 2^k.$$

It turned out that for any k , $p_k \leq q_j$, as

- All exact $2k + j$ steps cases are captured.
- All cases reaching n matches earlier, should drift around the n matches, with a probability at most 1.
- All other cases reaching n matches later, are omitted, while their probability is non-negative.

Since $q_j \geq \max_{2k+j \leq 3n} p_k$, then surely the following holds:

$$q_j \geq \binom{3j}{j} \cdot \frac{2^j}{3^{3j}}.$$

By lemma 1.7.1, we have

$$\begin{aligned} q_j &\geq \binom{3j}{j} \cdot \frac{2^j}{3^{3j}} = \frac{(3j)!}{(2j)! \cdot j!} \cdot \frac{2^j}{3^{3j}} \\ &\geq \frac{\sqrt{2\pi \cdot 3j} \cdot (3j)^{3j}}{(2\sqrt{2\pi \cdot 2j} \cdot (2j)^{2j}) \cdot (2\sqrt{2\pi \cdot j} \cdot j^j)} \cdot \frac{2^j}{3^{3j}} \\ &= \frac{\sqrt{3}}{8\sqrt{\pi}} \cdot \frac{1}{2^j \sqrt{j}} = \frac{c}{2^j \sqrt{j}}. \end{aligned}$$

Therefore, the lower bound for q , the probability for finding a satisfying assignment within $3n$ steps, is

$$\begin{aligned} q &= \sum_{j \in [0, n]} q_j \cdot \Pr[\text{initial assignment has } j \text{ mismatches}] \\ &\geq \sum_{j \in [0, n]} \binom{3j}{j} \cdot \frac{2^j}{3^{3j}} \cdot \binom{n}{j} \cdot \frac{1}{2^n} \geq \frac{c}{2^n} \sum_{j \in [0, n]} \frac{1}{2^j \sqrt{j}} \cdot \binom{n}{j} \\ &\geq \frac{c}{2^n \sqrt{n}} \sum_{j \in [0, n]} \binom{n}{j} \cdot \frac{1}{2^j} = \frac{c}{\sqrt{n}} \cdot \left(\frac{3}{4}\right)^n. \end{aligned}$$

We now formulate the 3-SAT randomized algorithm from random walk as follows:

- Repeat up to m times:
 - Pick a random truth assignment.
 - Repeat up to $3n$ times, terminating on all clauses satisfied:
 - * Choose an arbitrary unsatisfied clause.
 - * Choose uniformly at random a variable in the clause, and flip it.
- If the loop ends with a satisfying assignment, return it; otherwise, abort.

The expected number of rounds is at most $1/q = O(\sqrt{n} (4/3)^n)$, and the expected runtime is $O(n^{3/2} (4/3)^n)$.

Similar to the 2-SAT randomized algorithm, we can make

$$m = b \cdot \frac{2\sqrt{n}}{c} \cdot \left(\frac{4}{3}\right)^n,$$

such that we have b segments, while one segment failing to find a satisfying assignment can be viewed as

$$\Pr \left[Z > \frac{2\sqrt{n}}{c} \cdot \left(\frac{4}{3}\right)^n \right] < \mathbb{E}[Z] \cdot \frac{c}{2\sqrt{n}} \cdot \left(\frac{3}{4}\right)^n \leq \frac{1}{2},$$

by Markov inequality, where Z is a geometrically distributed random variable with success probability q .

Hence, the failing probability of finding a satisfying assignment for 3-SAT is at most 2^{-b} .

Exercise 1.7.10 (Exercise 7.8 [MU17]). Generalize the 3-SAT randomized algorithm to k -SAT algorithm, and write the expected runtime of the algorithm as a function of k .

Proof. We reuse most of the notions in 3-SAT randomized algorithm, where S is a satisfying assignment, $\{A_t\}_{t \in T}$ is the assignment at t^{th} time step, $\{Y_t\}_{t \in T}$ is the Markov chain for the number of matches that

$$\begin{aligned} \Pr[Y_t = 1 \mid Y_{t-1} = 0] &= 1, \\ \Pr[Y_t = j + 1 \mid Y_{t-1} = j] &= 1/k, \\ \Pr[Y_t = j - 1 \mid Y_{t-1} = j] &= 1 - 1/k, \end{aligned}$$

and $\{Z_i\}_{i \in [0, n]}$ are the random variables for the steps to satisfy all clauses, from an initial assignment of i mismatches.

For $k \geq 3$, the random walk is always more likely to drift to less satisfied variables than more satisfied variables. Hence, the prior strategy of trying over a large variety of initial assignments, then repeat a small number of times each for random flipping, remains to apply. Therefore, the 3-SAT algorithm should still apply for k -SAT case.

Let q_i be the probability for finding satisfying assignment from i mismatches in $3n$ steps of random walk. Then

$$q_i \geq \max_{2\ell + i \leq 3n} \binom{2\ell + i}{\ell} \left(\frac{1}{k}\right)^{\ell+i} \left(\frac{k-1}{k}\right)^{\ell},$$

and therefore by lemma 1.7.1,

$$\begin{aligned} q_i &\geq \binom{3i}{i} \left(\frac{1}{k}\right)^{2i} \left(\frac{k-1}{k}\right)^i \\ &\geq \frac{\sqrt{2\pi} \cdot 3i \cdot (3i)^{3i}}{\left(2\sqrt{2\pi} \cdot 2i \cdot (2i)^{2i}\right) \cdot \left(2\sqrt{2\pi} i \cdot i^i\right)} \cdot \frac{(k-1)^i}{k^{3i}} \\ &= \frac{\sqrt{3}}{8\sqrt{\pi}} \cdot \frac{1}{\sqrt{i}} \cdot \left(\frac{27(k-1)}{4k^3}\right)^i = \frac{c}{\sqrt{i}} \cdot \left(\frac{27(k-1)}{4k^3}\right)^i \end{aligned}$$

Let q be the probability for sampling a random initial assignment, then

$$q = \sum_{i \in [0, n]} q_i \cdot \binom{n}{i} \cdot \frac{1}{2^n} \geq \frac{c}{2^n \sqrt{n}} \sum_{i \in [0, n]} \binom{n}{i} \cdot \left(\frac{27(k-1)}{4k^3}\right)^i = \frac{c}{\sqrt{n}} \left(\frac{1}{2} + \frac{27(k-1)}{8k^3}\right)^n.$$

Hence, the expected runtime for the k -SAT algorithm is

$$\frac{3n}{q} \leq \frac{\sqrt{n}}{c} \left(\frac{8k^3}{4k^3 + 27(k-1)}\right)^n \cdot 3n = a \cdot 3n,$$

where a is the expected number of trials to the first satisfying assignment.

We set $m = 2ab$ for at least $1 - 2^{-b}$ probability of outputting satisfying assignment on a satisfiable formula. $\square$

Exercise 1.7.11 (Exercise 7.9 [MU17]). The prior 3-SAT randomized algorithm analysis made pessimistic assumption that the current assignment A_t and the satisfying assignment S differ in one variable in the chosen unsatisfied clause. Instead, suppose that the 2 assignments disagree on 1 variable with probability p , and disagree on at least 2 variables with probability $1 - p$. What is the largest p , such that one can prove that the expected number of steps before the randomized 3-SAT stops is polynomial in n ?

Proof. We reuse most of the notions in 3-SAT randomized algorithm, where S is a satisfying assignment, $\{A_t\}_{t \in T}$ is the assignment at t^{th} time step, $\{Y_t\}_{t \in T}$ is the modified Markov chain for the number of matches that

$$\begin{aligned}\Pr[Y_t = 1 \mid Y_{t-1} = 0] &= 1, \\ \Pr[Y_t = j + 1 \mid Y_{t-1} = j] &= \frac{1}{3} \cdot p + \frac{2}{3} \cdot (1 - p) = \frac{2}{3} - \frac{1}{3} \cdot p, \\ \Pr[Y_t = j - 1 \mid Y_{t-1} = j] &= \frac{2}{3} \cdot p + \frac{1}{3} \cdot (1 - p) = \frac{1}{3} + \frac{1}{3} \cdot p,\end{aligned}$$

and $\{Z_i\}_{i \in [0, n]}$ are the random variables for the steps to satisfy all clauses, from an initial assignment of i mismatches. Then q_j , the probability of finding a satisfying assignment within $3n$ steps for initial j mismatches, has

$$\begin{aligned}q_j &\geq \binom{3j}{j} \left(\frac{2}{3} - \frac{1}{3} \cdot p\right)^{2j} \left(\frac{1}{3} + \frac{1}{3} \cdot p\right)^j \\ &\geq \frac{\sqrt{2\pi \cdot 3j} \cdot (3j)^{3j}}{(2\sqrt{2\pi \cdot 2j} \cdot (2j)^{2j}) \cdot (2\sqrt{2\pi j} \cdot j^j)} \cdot \left(\frac{(2-p)^2(1+p)}{27}\right)^j \\ &= \frac{\sqrt{3}}{8\sqrt{\pi}} \cdot \frac{1}{\sqrt{j}} \cdot \left(\frac{(2-p)^2(1+p)}{4}\right)^j = \frac{c}{\sqrt{j}} \cdot \left(\frac{(2-p)^2(1+p)}{4}\right)^j.\end{aligned}$$

Now the probability q for finding a satisfying assignment within $3n$ steps is

$$q = \sum_{j \in [0, n]} q_j \cdot \binom{n}{j} \cdot \frac{1}{2^n} \geq \frac{c}{2^n \sqrt{n}} \sum_{j \in [0, n]} \binom{n}{j} \cdot \left(\frac{(2-p)^2(1+p)}{4}\right)^j = \frac{c}{\sqrt{n}} \left(\frac{1}{2} + \frac{(2-p)^2(1+p)}{8}\right)^n.$$

In this case, only when $p = 0$ can make $1/q = \text{poly}(n)$. □

Exercise 1.7.12 (Exercise 7.10 [MU17]). A coloring of a graph is an assignment of a color to each of its vertices. A graph is k -colorable if there exists coloring such that no 2 adjacent vertices have the same color. Let G be 3-colorable.

- Show that there exists a 2-coloring for G such that no triangle is monochromatic.
- Consider the following 2-coloring randomized algorithm for G such that there is no monochromatic triangle. Start from any 2-coloring. If there are any monochromatic triangles, choose a random one and flip the color of a random vertex in the triangle. Find an upper bound for the expected number of steps in the algorithm.

Proof. Let the 3-colorable graph G be given a proper 3-coloring such that no pair of adjacent vertices has the same color. Then any triangle in G must be colored in 3 distinct colors; otherwise at least one pair of adjacent vertices would have the same color, contradicting that G is 3-colorable. We randomly merge 2 colors from the 3-coloring into 1 color, then all the triangles are 2-colored, such that none of them are monochromatic.

The algorithm runtime analysis is similar to the one for the 3-SAT randomized algorithm. We define the satisfying coloring S , $\{A_t\}_{t \in T}$ as the coloring at time step t , $\{Y_t\}_{t \in T}$ as the Markov chain for the number of matches that

$$\begin{aligned}\Pr[Y_t = 1 \mid Y_{t-1} = 0] &= 1, \\ \Pr[Y_t = j + 1 \mid Y_{t-1} = j] &= 1/3, \\ \Pr[Y_t = j - 1 \mid Y_{t-1} = j] &= 2/3,\end{aligned}$$

and $\{Z_i\}_{i \in [0, n]}$ as the random variables for the number of steps to match S fully from an initial coloring of i mismatches. We adopt the analysis for the unsuccessful 3-SAT randomized algorithm, and by conditional expectation,

$$\begin{aligned}\mathbb{E}[Z_n] &= 0, \\ \mathbb{E}[Z_j] &= \mathbb{E}[Z_j \mid Y_{t+1} = j-1] \cdot P_{j,j-1} + \mathbb{E}[Z_j \mid Y_{t+1} = j+1] \cdot P_{j,j+1} \\ &= 1 + 2/3 \cdot \mathbb{E}[Z_{j-1}] + 1/3 \cdot \mathbb{E}[Z_{j+1}], \\ \mathbb{E}[Z_0] &= 1 + \mathbb{E}[Z_1].\end{aligned}$$

By induction, $\mathbb{E}[Z_j] = \mathbb{E}[Z_{j+1}] + 2^{j+2} - 3$, and eventually

$$\mathbb{E}[Z_j] = \sum_{k \in [j, n-1]} (2^{k+2} - 3) = 4(2^n - 2^j) - 3(n - j) = (2^{n+2} - 3n) - (2^{j+2} - 3j).$$

Let Z be the random variable for the number of steps in the algorithm, and Y be the random variable for the number of mismatches in the initial assignment, then by conditional expectation,

$$\begin{aligned}\mathbb{E}[Z] &= \sum_{j \in [0, n]} \mathbb{E}[Z \mid Y = j] \cdot \Pr[Y = j] = \sum_{j \in [0, n]} \mathbb{E}[Z_j] \cdot \binom{n}{j} \cdot \frac{1}{2^n} \\ &= (2^{n+2} - 3n) - \frac{1}{2^n} \sum_{j \in [0, n]} (2^{j+2} - 3j) \cdot \binom{n}{j} \\ &= (2^{n+2} - 3n) - 4 \left(\frac{3}{2}\right)^n + \frac{3}{2^n} \sum_{j \in [0, n]} j \cdot \binom{n}{j} \\ &= (2^{n+2} - 3n) - 4 \left(\frac{3}{2}\right)^n + \frac{3n}{2} = 2^{n+2} - 4 \left(\frac{3}{2}\right)^n - \frac{3n}{2}.\end{aligned}$$

□

Theorem 1.7.13 (Chapman-Kolmogorov Equation). *For a time-homogeneous Markov chain,*

$$P_{i,j}^{m+n} = \sum_k P_{i,k}^m \cdot P_{k,j}^n.$$

Proof.

$$\begin{aligned}P_{i,j}^{m+n} &= \Pr[X_{m+n} = j \mid X_0 = i] \\ &= \sum_k \Pr[X_{m+n} = j \mid X_m = k] \cdot \Pr[X_m = k \mid X_0 = i] \\ &= \sum_k \Pr[X_n = j \mid X_0 = k] \cdot \Pr[X_m = k \mid X_0 = i] = \sum_k P_{k,j}^n \cdot P_{i,k}^m.\end{aligned}$$

□

Definition 1.7.14. State j is accessible from state i if there exists some integer $n \geq 0$, $P_{i,j}^n > 0$. If state i, j are accessible from each other, we say they communicate by $i \leftrightarrow j$. The communicating relation defines an equivalence relation:

- Reflexive: For any state i , $i \leftrightarrow i$.
- Symmetric: If $i \leftrightarrow j$, then $j \leftrightarrow i$.
- Transitive: If $i \leftrightarrow j$ and $j \leftrightarrow k$, then $i \leftrightarrow k$.

The communication relation partitions the states into disjoint equivalence classes, which are referred to as communicating classes. It might be possible to move from one class to another, but not in the other way around.

Definition 1.7.15. A Markov chain is irreducible if all states belong to one communicating class.

Hence, a Markov chain is irreducible if for all ordered pair of states (i, j) , the probability that state i can reach state j is nonzero. The next lemma is immediate.

Lemma 1.7.16. A finite Markov chain is irreducible iff its graph representation is a strongly connected graph.

Definition 1.7.17. Let $r_{i,j}^t$ be the probability that, the first transition from state i to state j happens on the t^{th} step:

$$r_{i,j}^t = \Pr \left[X_t = j \wedge \bigwedge_{k \in [1, t-1]} X_k \neq j \mid X_0 = i \right].$$

A state is transient if $\sum_{t \geq 1} r_{i,i}^t < 1$, and it is recurrent if $\sum_{t \geq 1} r_{i,i}^t = 1$.

Definition 1.7.18. Let $T_{i,j}$ be the random variable for the steps to transition from i to j , and $T_{i,i}$ be the random variable for the steps to return to i , such that $T_{i,i} \geq 1$.

Exercise 1.7.19 (Exercise 7.15 [MU17]). Let $P_{i,i}^t$ be the probability that a Markov chain returns to state i when started in state i after t steps. Prove that $\sum_{t \geq 1} P_{i,i}^t$ is unbounded iff state i is recurrent.

Proof. We can express $P_{i,i}^t$ in the following recursive way:

$$P_{i,i}^t = \sum_{k \in [1, t]} r_{i,i}^k \cdot P_{i,i}^{t-k},$$

where $P_{i,i}^0 = 1$. Then the summation can be rearranged to

$$\begin{aligned} \sum_{t \geq 1} P_{i,i}^t &= \sum_{t \geq 1} \sum_{k \in [1, t]} r_{i,i}^k \cdot P_{i,i}^{t-k} \\ &= \left(\sum_{t \geq 1} r_{i,i}^t \right) \cdot \left(\sum_{k \geq 0} P_{i,i}^k \right) = \left(\sum_{t \geq 1} r_{i,i}^t \right) \cdot \left(\sum_{k \geq 1} P_{i,i}^k \right) + \left(\sum_{t \geq 1} r_{i,i}^t \right) \\ &= \left(\sum_{t \geq 1} r_{i,i}^t \right) \cdot \left(1 + \sum_{k \geq 1} P_{i,i}^k \right). \end{aligned}$$

In this way, $\sum_{t \geq 1} P_{i,i}^t$ being unbounded is equivalent to the state i being recurrent. $\square$

Corollary 1.7.20. If a state in a communicating class is transient (recurrent), then all states in the class are transient (recurrent).

Proof. Suppose state i, j are in the same communicating class, and with integers $c_0, c_1 \geq 1$ such that $P_{i,j}^{c_0} > 0, P_{j,i}^{c_1} > 0$. By theorem 1.7.13 and omitting the first $c_0 + c_1$ terms,

$$\sum_{\tau \geq 1} P_{i,i}^\tau \geq \sum_{\tau \geq 1+c_0+c_1} P_{i,i}^\tau = \sum_{t \geq 1} \sum_{u,v} P_{i,u}^{c_0} \cdot P_{u,v}^t \cdot P_{v,i}^{c_1} \geq \sum_{t \geq 1} \sum_k P_{i,k}^{c_0} \cdot P_{k,k}^t \cdot P_{k,i}^{c_1} = \sum_k P_{i,k}^{c_0} \cdot P_{k,i}^{c_1} \cdot \sum_{t \geq 1} P_{k,k}^t \geq P_{i,j}^{c_0} \cdot P_{j,i}^{c_1} \cdot \sum_{t \geq 1} P_{j,j}^t,$$

where the fourth equality for changing summation sequence holds by probability being non-negative.

Then by exercise 1.7.19, if state i is transient, then $\sum_{t \geq 1} P_{j,j}^t$ has to be bounded, that state j is transient. Conversely, if state j is recurrent, then $\sum_{t \geq 1} P_{j,j}^t$ has to be unbounded, that state i is recurrent. The intuition follows.

Suppose state i is transient, then all the other states in the same communicating class have nonzero probability to reach i , and get out of the communicating class from i with nonzero probability, hence all of the states are transient.

On the other hand, supposing state i is recurrent but another state j is transient in the same communicating class, this is a contradiction as state i can reach state j and never get back to itself. Contradiction remains until no state in the communicating class is transient. $\square$

Definition 1.7.21. Let $h_{i,i} = \sum_{t \geq 1} t \cdot r_{i,i}^t$ be the expected time to return to i when starting at state i , and $h_{i,j} = \sum_{t \geq 1} t \cdot r_{i,j}^t$ be the expected time to reach j from i . A recurrent state i is positive recurrent if $h_{i,i} < \infty$, otherwise it is null recurrent.

Corollary 1.7.22. If i is transient, then $h_{i,i} = \infty$.

Proof. Since i may never return by $\sum_{t \geq 1} r_{i,i}^t < 1$, hence $\Pr[T_{i,i} = \infty] > 0$. $\square$

Corollary 1.7.23. If i is recurrent, then $\Pr[T_{i,i} < \infty] = 1$.

Proof. Suppose $\Pr[T_{i,i} = \infty] > 0$ for recurrent i , then it contradicts with recurrence: $\sum_{t \geq 1} r_{i,i}^t = 1$. $\square$

Remark 1.7.24. The convergence of $h_{i,i}$ differs positive recurrence from null recurrence, and $\Pr[T_{i,i} = \infty] = 0$ differs recurrence from transience.

Example 1.7.25. Consider a finite Markov chain of states $[1, 2]$ with $P_{i,j} = 1/2$ for all $i, j \in [1, 2]$. The number of steps for “reaching state j ” can be viewed as a geometrically distributed random variable, and $h_{i,j} = 2$.

However, consider a Markov chain over states $\mathbb{Z}^+$, that $P_{i,1} + P_{i,i+1} = 1$ and $P_{i,1} = (1 + i)^{-1}$. Then

$$r_{1,1}^t = P_{1,1} \cdot \prod_{i \in [1, t-1]} P_{i,i+1} = \frac{1}{t+1} \cdot \prod_{i \in [1, t-1]} \frac{i}{i+1} = \frac{1}{t+1} \cdot \frac{1}{t} = \frac{1}{t} - \frac{1}{t+1}.$$

Hence, $\sum_{t \geq 1} r_{1,1}^t = 1$, yet $h_{1,1} = \sum_{t \geq 1} t \cdot r_{1,1}^t = \infty$.

Theorem 1.7.26. Let i be recurrent, and suppose i can access j . Then j is recurrent, and both $T_{i,j}$ and $T_{j,i}$ are bounded with probability 1.

Proof. If i can access j , then with non-zero probability i can transition to j . Supposing j is transient, by corollary 1.7.20 i, j are in different communicating classes, thus $i \not\leftrightarrow j$, hence $P_{j,i}^\tau = 0$ for any $\tau \geq 1$. Let $P_{i,j}^t > 0$ for some $t \geq 1$, then i returns to itself with probability at most $1 - P_{i,j}^t$, contradicting with the recurrence of i . Thus, j is recurrent.

Since i and j communicate, for one path from i to j of length ℓ , let the probability of taking this path be q , then returning to i before reaching j has probability at most $1 - q$. By corollary 1.7.23, $\Pr[T_{i,i} < \infty] = 1$. Hence, we model this with a geometrically distributed random variable: on a coin flip, either success with probability at least q to reach j in n steps, or fail with probability at most $1 - q$ to return to i in $T_{i,i}$ steps. Hence, $\Pr[T_{i,j} < \infty] = 1$.

Suppose $\Pr[T_{j,i} = \infty] = p > 0$, then $\Pr[T_{i,i} = \infty] \geq p \cdot q > 0$, which is contradicting to i being recurrent. $\square$

Lemma 1.7.27. In a finite Markov chain, there is at least one recurrent state, and all recurrent states are positive recurrent.

Proof. A finite directed graph G can be decomposed into a finite number of strongly connected components (SCCs). Contracting SCCs gives a finite DAG, which has at least one sink SCC. Hence, there must exist a closed SCC.

Such closed SCC with 0 out degree corresponds to a closed communicating class S , such that the subchain is a finite irreducible Markov chain, and the probability of transitioning out is 0.

For $i \in S$, $\sum_{t \geq 1} r_{i,i}^t$ is the probability of i returning to itself. Let m be the integer such that $\Pr[T_{x,i} \leq m] > 0$ for all $x \in S$ by the irreducibility and finiteness of S , then

$$\Pr[T_{i,i} > m] = 1 - \Pr[T_{i,i} \leq m] \leq 1 - \min_{x \in S} \Pr[T_{x,i} \leq m] = 1 - \delta.$$

Moreover, by the Markovian property,

$$\Pr[T_{i,i} > (r+1)m \mid T_{i,i} > rm] \leq 1 - \delta.$$

By the closure of S , we have $\Pr[T_{i,i} > rm] \leq (1 - \delta)^r$ after repeating r blocks of m steps, converging to 0 when $r \rightarrow \infty$, yielding $\sum_{t \geq 1} r_{i,i}^t = 1$. By corollary 1.7.20, all $i \in S$ are recurrent.

For any $i \in S$,

$$\mathbb{E}[T_{i,i}] = \sum_{t \geq 1} t \cdot r_{i,i}^t \leq \sum_{r \geq 0} m \Pr[T_{i,i} > rm] \leq m \sum_{r \geq 0} (1 - \delta)^r \leq \frac{m}{\delta}.$$

Such a bound applies to all sink SCCs, so all recurrent states are positive recurrent. $\square$

1.7.3 Strong Law of Large Numbers (SLLN)

We now introduce the Kolmogorov Strong Law of Large Numbers (SLLN), in contrast to theorem 1.3.11.

Theorem 1.7.28 (Kolmogorov SLLN). *Let $\{X_i\}_{i \geq 1}$ be i.i.d. random variables with $\mathbb{E}[|X_1|] < \infty$, and let $\mu = \mathbb{E}[X_1]$.*

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} X_i = \mu \right] = 1.$$

We adopt the proof structure from [Ete81] and first prove a useful lemma.

Definition 1.7.29. Let $\{A_n\}_{n \geq 1}$ be a sequence of subsets of Ω . The supremum of $\{A_n\}_{n \geq k}$ is

$$\sup_{n \geq k} A_n = \bigcup_{n \geq k} A_n.$$

The limit superior of $\{A_n\}_{n \geq 1}$ is

$$\limsup_{n \rightarrow \infty} A_n = \bigcap_{n \geq 1} \bigcup_{k \geq n} A_k.$$

Equivalently, $\limsup_{n \rightarrow \infty} A_n$ is the event that “for every cutoff n , there exists $N \geq n$ such that A_N holds”. In other words, there are infinitely many indices n such that A_n holds, or equivalently, “ A_n occurs infinitely often”.

Lemma 1.7.30 (Borel-Cantelli Lemma [Fel68a]). *Let $\{A_i\}_{i \geq 1}$ be an infinite sequence of events. If $\sum_{i \geq 1} \Pr[A_i] < \infty$,*

$$\Pr \left[\limsup_{n \rightarrow \infty} A_n \right] = 0. \quad (1.1)$$

Proof. We notice a non-increasing sequence of sets by

$$\limsup_{n \rightarrow \infty} A_n \subseteq \dots \subseteq \bigcup_{k \geq N+1} A_k \subseteq \bigcup_{k \geq N} A_k \subseteq \dots \subseteq \bigcup_{k \geq 2} A_k \subseteq \bigcup_{k \geq 1} A_k.$$

Hence, on an arbitrary N , we upper bound the probability by

$$\Pr \left[\limsup_{n \rightarrow \infty} A_n \right] \leq \Pr \left[\bigcup_{k \geq N} A_k \right] \leq \sum_{k \geq N} \Pr[A_k],$$

the second inequality holds from union bounding, and the right hand side summation approaches to 0 as $N \rightarrow \infty$. $\square$

Remark 1.7.31. The event $\limsup_{n \rightarrow \infty} A_n$ means that no matter how far into the sequence we go, there is always a later A_k that occurs. In other words, A_n keep occurring arbitrarily far out, or equivalently, occur infinitely often.

Hence, eq. (1.1) means that with probability 1, there exists a finite cutoff after which none of the events A_n occur. Equivalently, only finitely many of them occur.

Now we begin the proof for theorem 1.7.28.

Proof. WLOG we let $\mu = \mathbb{E}[X_1] = 0$, $Y_n = X_n \cdot 1_{\leq n}$ by $Y_n = 0$ if $|X_n| > n$, otherwise $Y_n = X_n$, and $Z_n = X_n \cdot 1_{> n} = X_n - Y_n$.

Since $\mathbb{E}[|X_1|] < \infty$, and by $\{X_i\}$ are i.i.d., we have

$$\infty > \mathbb{E}[|X_1|] \geq \sum_{t \geq 1} \Pr[|X_1| \geq t] = \sum_{t \geq 1} \Pr[|X_t| \geq t] \geq \sum_{t \geq 1} \Pr[|X_t| > t]. \quad (1.2)$$

By lemma 1.7.30 and remark 1.7.31, the eq. (1.2) means with probability 1, only finitely many indices i have $|X_i| > i$. Equivalently, with probability 1, there exists a finite cutoff N such that $|X_n| > n$ does not occur for all $n > N$.

Hence, with probability 1, we have $Z_n = 0$ for all $n > N$, so only finitely many Z_n are nonzero. Therefore,

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} Z_i = 0 \right] = 1,$$

and what remains to show is

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} Y_i = 0 \right] = 1.$$

Let $W_i = Y_i - \mathbb{E}[Y_i]$, $S_n = \sum_{i \in [1, n]} W_i$, and $m \in \mathbb{N}^+$, by Chebyshev's inequality and the independence of $\{X_i\}$,

$$\begin{aligned} \Pr \left[\left| \frac{S_{2^m}}{2^m} \right| \geq \varepsilon \right] &\leq \frac{\text{Var}[S_{2^m}]}{\varepsilon^2 \cdot 2^{2m}} \\ &= \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \text{Var}[W_i] = \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \text{Var}[Y_i] \\ &\leq \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_i^2]. \end{aligned}$$

We want to take advantage of Borel-Cantelli lemma 1.7.30 again by showing if

$$\sum_{m \geq 1} \Pr \left[\left| \frac{S_{2^m}}{2^m} \right| \geq \varepsilon \right] \leq \sum_{m \geq 1} \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_i^2] < \infty,$$

then

$$\Pr \left[\limsup_{m \rightarrow \infty} \left| \frac{S_{2^m}}{2^m} \right| \geq \varepsilon \right] = 0,$$

that only a finite number of $m \geq 1$ has $|S_{2^m}| \geq \varepsilon \cdot 2^m$. There are 3 things remains to be proved:

Lemma 1.7.32.

$$\sum_{m \geq 1} \frac{1}{2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_i^2] < \infty.$$

Lemma 1.7.33.

$$\Pr \left[\limsup_{m \rightarrow \infty} \max_{2^m < n \leq 2^{m+1}} \left| \frac{S_n - S_{2^m}}{2^m} \right| \geq \varepsilon \right] = 0.$$

Lemma 1.7.34.

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} \mathbb{E}[Y_i] = 0.$$

We begin with proof for lemma 1.7.34.

Proof. We derive $\mathbb{E}[Y_n] = \mathbb{E}[X_1 \cdot 1_{\leq n}] = -\mathbb{E}[X_1 \cdot 1_{> n}]$, as we simplify by $\mathbb{E}[X_1] = 0$. Hence, by Jensen's inequality,

$$|\mathbb{E}[Y_n]| \leq |\mathbb{E}[X_1 \cdot 1_{> n}]| \leq \mathbb{E}[|X_1| \cdot 1_{> n}].$$

By monotone convergence, we have $\mathbb{E}[|X_1| \cdot 1_{\leq n}] \rightarrow \mathbb{E}[|X_1|]$ as $n \rightarrow \infty$. Since

$$\mathbb{E}[|X_1|] = \mathbb{E}[|X_1| \cdot 1_{> n}] + \mathbb{E}[|X_1| \cdot 1_{\leq n}] < \infty$$

and $|X_1| \cdot 1_{\leq n} \geq 0$, we have $\mathbb{E}[|X_1| \cdot 1_{> n}] \rightarrow 0$ as $n \rightarrow \infty$.

By Cesàro summation, that if $a_n \rightarrow 0$ as $n \rightarrow \infty$, then $1/n \cdot \sum a_n \rightarrow 0$, thus

$$0 \leq \left| \frac{1}{n} \sum_{i \in [1, n]} \mathbb{E}[Y_i] \right| \leq \frac{1}{n} \sum_{i \in [1, n]} |\mathbb{E}[Y_i]| \leq \frac{1}{n} \sum_{i \in [1, n]} \mathbb{E}[|X_1| \cdot 1_{> i}] \rightarrow 0,$$

which completes the proof. $\square$

Now we move on by proving lemma 1.7.32.

Proof. Noticing that for real valued random variable X_k ,

$$Y_k^2 = X_1^2 \cdot 1_{\leq k} \leq \sum_{m \in [1, k]} (2m - 1) \cdot 1_{> m-1},$$

since when $|X_1| \leq k$, the right hand side is $\lceil |X_1| \rceil^2 \geq X_1^2$, and when $|X_1| > k$, the left hand side is 0. Hence,

$$\begin{aligned} \mathbb{E}[Y_k^2] &= \mathbb{E}[X_1^2 \cdot 1_{\leq k}] \\ &\leq \sum_{m \in [1, k]} (2m - 1) \cdot \Pr[|X_1| > m - 1]. \end{aligned} \quad (1.3)$$

Relaxing the summation, then let p_t be the first p such that $2^p \geq t$,

$$\begin{aligned} \sum_{m \geq 1} \frac{1}{2^{2m}} \sum_{t \in [1, 2^m]} \mathbb{E}[Y_t^2] &= \sum_{t \geq 1} \mathbb{E}[Y_t^2] \sum_{m: 2^m \geq t} \frac{1}{2^{2m}} \\ &= \frac{4}{3} \sum_{t \geq 1} \mathbb{E}[Y_t^2] \cdot \frac{1}{2^{2p_t}} \\ &\leq \frac{4}{3} \sum_{t \geq 1} \mathbb{E}[Y_t^2] \cdot \frac{1}{t^2}. \end{aligned} \quad (1.4)$$

Applying the eq. (1.3) to eq. (1.4),

$$\begin{aligned} \sum_{t \geq 1} \frac{1}{t^2} \cdot \mathbb{E}[Y_t^2] &\leq \sum_{t \geq 1} \frac{1}{t^2} \sum_{m \in [1, t]} (2m - 1) \cdot \Pr[|X_1| > m - 1] \\ &= \sum_{m \geq 1} (2m - 1) \cdot \Pr[|X_1| > m - 1] \cdot \sum_{t \geq m} \frac{1}{t^2}. \end{aligned}$$

By integration,

$$\sum_{t \geq m} \frac{1}{t^2} \leq \int_{m-\frac{1}{2}}^{\infty} \frac{dt}{t^2} = \frac{2}{2m-1},$$

then

$$\begin{aligned} \sum_{m \geq 1} (2m - 1) \cdot \Pr[|X_1| > m - 1] \cdot \sum_{t \geq m} \frac{1}{t^2} &\leq 2 \sum_{m \geq 1} \Pr[|X_1| > m - 1] \\ &\leq 2(1 + \mathbb{E}[|X_1|]) < \infty, \end{aligned}$$

which completes the proof. $\square$

We introduce Etemadi Inequality [Ete81] for proof of lemma 1.7.33.

Lemma 1.7.35 (Etemadi Inequality [Ete81]). *Let $\{X_i\}_{i \in [1, n]}$ be i.i.d. real-valued random variables defined over some common probability space, and let $\alpha \geq 0$, then*

$$\Pr \left[\max_{k \in [1, n]} \left| \sum_{i \in [1, k]} X_i \right| \geq 3\alpha \right] \leq 3 \max_{k \in [1, n]} \Pr \left[\left| \sum_{i \in [1, k]} X_i \right| \geq \alpha \right].$$

For lemma 1.7.33, the proof follows.

Proof. We show only finite many m have $\max_{2^m < n \leq 2^{m+1}} |S_n - S_{2^m}| \geq 3\varepsilon \cdot 2^m$ by lemma 1.7.30, then by lemma 1.7.35,

$$\sum_{m \geq 1} \Pr \left[\max_{2^m < n \leq 2^{m+1}} \left| \frac{S_n - S_{2^m}}{2^m} \right| \geq 3\varepsilon \right] \leq 3 \sum_{m \geq 1} \max_{2^m < n \leq 2^{m+1}} \Pr \left[\left| \frac{S_n - S_{2^m}}{2^m} \right| \geq \varepsilon \right]. \quad (1.5)$$

Again by Chebyshev's inequality,

$$\begin{aligned} \max_{2^m < n \leq 2^{m+1}} \Pr \left[\left| \frac{S_n - S_{2^m}}{2^m} \right| \geq \varepsilon \right] &\leq \max_{2^m < n \leq 2^{m+1}} \frac{\text{Var}[S_n - S_{2^m}]}{\varepsilon^2 \cdot 2^{2m}} \\ &= \max_{2^m < n \leq 2^{m+1}} \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [2^m+1, n]} \text{Var}[W_i] \\ &= \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \text{Var}[W_{2^m+i}] \\ &\leq \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_{2^m+i}^2]. \end{aligned} \quad (1.6)$$

The second equality holds by linearity of variance of independent random variables, the third equality holds by the non-negativity of variance, and the last inequality holds by

$$\text{Var}[W_i] = \text{Var}[Y_i] \leq \mathbb{E}[Y_i^2].$$

By lemma 1.7.32, the right hand side upper bound of eq. (1.6) is summable by

$$\sum_{m \geq 1} \frac{1}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_{2^m+i}^2] \leq \sum_{m \geq 1} \frac{4}{\varepsilon^2 \cdot 2^{2(m+1)}} \sum_{i \in [1, 2^{m+1}]} \mathbb{E}[Y_i^2] < \infty,$$

finishes upper bounding the left hand side of eq. (1.5),

$$\sum_{m \geq 1} \Pr \left[\max_{2^m < n \leq 2^{m+1}} \left| \frac{S_n - S_{2^m}}{2^m} \right| \geq 3\varepsilon \right] < \infty,$$

then we conclude lemma 1.7.33 by lemma 1.7.30, that

$$\Pr \left[\limsup_{m \rightarrow \infty} \max_{2^m < n \leq 2^{m+1}} \left| \frac{S_n - S_{2^m}}{2^m} \right| \geq \varepsilon \right] = 0.$$

□

Corollary 1.7.36. For every $\varepsilon > 0$,

$$\Pr \left[\limsup_{n \rightarrow \infty} \left| \frac{S_n}{n} \right| \geq \varepsilon \right] = 0.$$

Proof. By Chebyshev's inequality,

$$\Pr \left[\left| \frac{S_{2^m}}{2^m} \right| \geq \frac{\varepsilon}{2} \right] \leq \frac{4}{\varepsilon^2 \cdot 2^{2m}} \sum_{i \in [1, 2^m]} \mathbb{E}[Y_i^2].$$

Hence,

$$\sum_{m \geq 1} \Pr \left[\left| \frac{S_{2^m}}{2^m} \right| \geq \frac{\varepsilon}{2} \right] < \infty$$

by lemma 1.7.32. Therefore, by lemma 1.7.30,

$$\Pr \left[\limsup_{m \rightarrow \infty} \left| \frac{S_{2^m}}{2^m} \right| \geq \frac{\varepsilon}{2} \right] = 0.$$

Moreover, for $2^m < n \leq 2^{m+1}$, by triangle inequality

$$\left| \frac{S_n}{2^m} \right| \leq \left| \frac{S_n - S_{2^m}}{2^m} \right| + \left| \frac{S_{2^m}}{2^m} \right|,$$

so

$$\left\{ \limsup_{\substack{m, n \rightarrow \infty \\ 2^m < n \leq 2^{m+1}}} \left| \frac{S_n}{2^m} \right| \geq \varepsilon \right\} \subseteq \left\{ \limsup_{m \rightarrow \infty} \max_{2^m < n \leq 2^{m+1}} \left| \frac{S_n - S_{2^m}}{2^m} \right| \geq \frac{\varepsilon}{2} \right\} \cup \left\{ \limsup_{m \rightarrow \infty} \left| \frac{S_{2^m}}{2^m} \right| \geq \frac{\varepsilon}{2} \right\}.$$

Since $2^m < n$ implies $|S_n/n| \leq |S_n/2^m|$, we have

$$\left\{ \limsup_{n \rightarrow \infty} \left| \frac{S_n}{n} \right| \geq \varepsilon \right\} \subseteq \left\{ \limsup_{\substack{m, n \rightarrow \infty \\ 2^m < n \leq 2^{m+1}}} \left| \frac{S_n}{2^m} \right| \geq \varepsilon \right\}.$$

Therefore, by lemma 1.7.33,

$$\Pr \left[\limsup_{n \rightarrow \infty} \left| \frac{S_n}{n} \right| \geq \varepsilon \right] = 0.$$

□

To complete the proof for Kolmogorov SLLN theorem 1.7.28, by corollary 1.7.36,

$$\Pr \left[\limsup_{n \rightarrow \infty} \left| \frac{S_n}{n} \right| \geq \varepsilon \right] = 0.$$

Since $\varepsilon > 0$ was arbitrary, we conclude that

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{S_n}{n} = 0 \right] = 1.$$

On the other hand, by lemma 1.7.34,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} \mathbb{E}[Y_i] = 0.$$

Therefore, on the event $\{\lim_{n \rightarrow \infty} S_n/n = 0\}$, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} Y_i = \lim_{n \rightarrow \infty} \left(\frac{S_n}{n} + \frac{1}{n} \sum_{i \in [1, n]} \mathbb{E}[Y_i] \right) = 0,$$

so

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} Y_i = 0 \right] = 1.$$

Together with

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} Z_i = 0 \right] = 1,$$

and $X_i = Y_i + Z_i$, we conclude that

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1, n]} X_i = 0 \right] = 1.$$

□

Definition 1.7.37 (Stopping Time). Let $\{X_t\}_{t \geq 0}$ be a discrete-time stochastic process. A non-negative integer-valued random variable T is a stopping time if for every $n \geq 0$, whether event $T = n$ occurs is determined only by $X_0, \dots, X_n$.

Definition 1.7.38 (Strong Markov Property for Discrete-Time Markov Chains). Let $\{X_t\}_{t \geq 0}$ be a discrete-time Markov chain. We say that $\{X_t\}_{t \geq 0}$ satisfies the strong Markov property if for every stopping time T with $\Pr[T < \infty] = 1$, conditional on $T = n$ and the history

$$X_0 = x_0, X_1 = x_1, \dots, X_{n-1} = x_{n-1}, X_n = i,$$

the shifted process $\{X_{T+m}\}_{m \geq 0}$ has the same distribution as the original chain started from the state i .

Theorem 1.7.39. Any discrete-time, time-homogeneous Markov chain satisfies the strong Markov property.

Proof. Take any states $x_0, \dots, x_{n-1}, i, j_1, \dots, j_r$, and any time steps $n \geq 0$ and $0 \leq m_1 < \dots < m_r$. We show that conditional on $T = n$ and the history $X_0 = x_0, \dots, X_n = i$, the shifted process has the same finite-dimensional distributions as the original chain started from state i . Since event $T = n$ is in the condition,

$$\begin{aligned} & \Pr[X_{T+m_1} = j_1, \dots, X_{T+m_r} = j_r \mid T = n, X_0 = x_0, \dots, X_{n-1} = x_{n-1}, X_n = i] \\ &= \Pr[X_{n+m_1} = j_1, \dots, X_{n+m_r} = j_r \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}, X_n = i]. \end{aligned}$$

By the Markov property applied successively,

$$\begin{aligned} & \Pr[X_{n+m_1} = j_1, \dots, X_{n+m_r} = j_r \mid X_0 = x_0, \dots, X_{n-1} = x_{n-1}, X_n = i] \\ &= \Pr[X_{n+m_1} = j_1 \mid X_n = i] \cdot \prod_{\ell \in [2, r]} \Pr[X_{n+m_\ell} = j_\ell \mid X_{n+m_{\ell-1}} = j_{\ell-1}]. \end{aligned}$$

By time-homogeneity in definition 1.7.4, this equals

$$\Pr[X_{m_1} = j_1 \mid X_0 = i] \cdot \prod_{\ell \in [2, r]} \Pr[X_{m_\ell - m_{\ell-1}} = j_\ell \mid X_0 = j_{\ell-1}],$$

which is exactly

$$\Pr[X_{m_1} = j_1, \dots, X_{m_r} = j_r \mid X_0 = i].$$

Therefore, conditional on $T = n$ and the history up to time n , the shifted process $\{X_{T+m}\}_{m \geq 0}$ has the same finite-dimensional distributions as the original chain started from state i . Hence it has the same distribution, and the chain satisfies the strong Markov property. $\square$

Theorem 1.7.40 (SLLN for Markov chains). Let $\{X_t\}_{t \in T}$ be a Markov chain, with $X_0 = i$, and states i, j communicating. Let $1_{i,j}^t$ be the indicator random variable for the event $X_t = j$ given $X_0 = i$. Then the limiting fraction of time spent in state j is $h_{j,i}^{-1}$, that is,

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} 1_{i,j}^t = \frac{1}{h_{j,i}} \right] = 1.$$

Proof. If i is transient, and i, j are in same communicating class, then by corollary 1.7.20, j is also transient. Intuitively, there are only finitely many times that i can reach j . We can show by a similar trick in corollary 1.7.20, and let n be a number of steps such that $P_{j,i}^n > 0$, then

$$\infty > \sum_{\tau \geq 1} P_{i,i}^\tau \geq \sum_{\tau \geq n+1} P_{i,i}^\tau = \sum_{t \geq 1} \sum_k P_{i,k}^t \cdot P_{k,i}^n \geq \sum_{t \geq 1} P_{i,j}^t \cdot P_{j,i}^n = P_{j,i}^n \cdot \sum_{t \geq 1} P_{i,j}^t = P_{j,i}^n \cdot \sum_{t \geq 1} \mathbb{E}[1_{i,j}^t].$$

Since $\sum_t \mathbb{E}[1_{i,j}^t] < \infty$, then $\sum_t 1_{i,j}^t = \infty$ with probability 0, hence the left hand side is 0. By corollary 1.7.22, $h_{j,i} = \infty$, thus the right hand side is 0. Hence, the transient case is proved.

If i is recurrent, then i, j are both recurrent by corollary 1.7.20.

We first handle the positive recurrent case, and start with the special case when $i = j$. Suppose that j is positive recurrent, so $h_{j,j} < \infty$. Let $R_0 = 0$, and for each $c \geq 1$, let R_c be the time of the c^{th} visit to state j , where

$$R_c = \min\{t > R_{c-1} \mid X_t = j\},$$

Then the cycle lengths are

$$C_c = R_c - R_{c-1},$$

and consequently, R_c is the total length of the first c cycles:

$$R_c = \sum_{k \in [1, c]} C_k.$$

One checks that each R_c is a stopping time by definition 1.7.37. Then by the strong Markov property theorem 1.7.39 at the successive return times to j and by time-homogeneity, the random variables $\{C_c\}_{c \geq 1}$ are i.i.d., with $\mathbb{E}[C_1] = h_{j,j}$.

Let k_n be the largest number of cycles such that $R_{k_n} \leq n$ by

$$k_n = \max\{k \mid R_k \leq n\}.$$

Hence, $R_{k_n} \leq n < R_{k_n+1}$ by definition, and consequently,

$$\frac{R_{k_n}}{k_n} \leq \frac{n}{k_n} \leq \frac{R_{k_n+1}}{k_n}.$$

Since j is recurrent, we have $k_n \rightarrow \infty$ as $n \rightarrow \infty$. Then, by SLLN theorem 1.7.28,

$$\Pr \left[\lim_{k \rightarrow \infty} \frac{R_k}{k} = \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{c \in [1, k]} C_c = h_{j,j} \right] = 1,$$

and

$$\lim_{k \rightarrow \infty} \frac{R_{k+1}}{k} = \lim_{k \rightarrow \infty} \frac{R_{k+1}}{k+1} \cdot \frac{k+1}{k} = \lim_{k \rightarrow \infty} \frac{R_{k+1}}{k+1}. \quad (1.7)$$

Hence, when $i = j$,

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{k_n}{n} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} 1_{j,j}^t = \frac{1}{h_{j,j}} \right] = 1.$$

We next handle the null recurrent case when $i = j$. Suppose that j is null recurrent, so $h_{j,j} = \infty$. Reuse the same notations $\{R_c\}_{c \geq 0}$, $\{C_c\}_{c \geq 1}$, and $\{k_n\}_{n \geq 1}$. For each positive integer m , define the truncated cycle lengths

$$C_c^{(m)} = \min(C_c, m).$$

Since $\{C_c\}_{c \geq 1}$ are i.i.d., the bounded random variables $\{C_c^{(m)}\}_{c \geq 1}$ are also i.i.d. Hence, by SLLN theorem 1.7.28,

$$\Pr \left[\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{c \in [1, k]} C_c^{(m)} = \mathbb{E}[C_1^{(m)}] \right] = 1.$$

Since $\mathbb{E}[C_1^{(m)}]$ increases monotonically as $m \rightarrow \infty$, by monotone convergence, $\mathbb{E}[C_1^{(m)}] \rightarrow \mathbb{E}[C_1] = h_{j,j} = \infty$ by null recurrence. Then, for any positive integer α , there exists a sufficiently large m such that $\mathbb{E}[C_1^{(m)}] > \alpha$.

$$\Pr \left[\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{c \in [1, k]} C_c^{(m)} > \alpha \right] = 1.$$

Since for each fixed m , $C_c^{(m)} \leq C_c$ pointwise, then

$$R_k = \sum_{c \in [1, k]} C_c \geq \sum_{c \in [1, k]} C_c^{(m)}.$$

Hence, for any fixed positive integer α ,

$$\Pr \left[\liminf_{k \rightarrow \infty} \frac{R_k}{k} > \alpha \right] = 1.$$

Since this holds for every positive integer α , we conclude that

$$\Pr \left[\lim_{k \rightarrow \infty} \frac{R_k}{k} = \infty \right] = 1.$$

Since j is recurrent, we still have $k_n \rightarrow \infty$ as $n \rightarrow \infty$, and $n \geq R_{k_n}$. By prior squeezing argument eq. (1.7), when $i = j$,

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{k_n}{n} = 0 \right] = 1.$$

For the general recurrent case when $i \neq j$, let $T_{i,j}$ be the first hitting time from i to j . By the strong Markov property at time $T_{i,j}$, the path after time $T_{i,j}$ behaves like a fresh copy of the chain started from j . Thus the path consists of an initial segment of length $T_{i,j}$, after which the chain proceeds through successive j -cycles. Let $k_n = \sum_{t \in [1, n - T_{i,j}]} 1_{j,j}^t$ for the cycles before $n - T_{i,j}$ steps, then

$$\frac{R_{k_n} + T_{i,j}}{k_n} \leq \frac{n}{k_n} < \frac{R_{k_n+1} + T_{i,j}}{k_n}.$$

By theorem 1.7.26, $\Pr [T_{i,j} < \infty] = 1$, and thus

$$\Pr \left[\lim_{k \rightarrow \infty} \frac{T_{i,j}}{k} = 0 \right] = 1,$$

hence the same squeezing argument eq. (1.7) as in the $i = j$ case yields

$$\Pr \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} 1_{i,j}^t = \frac{1}{h_{j,j}} \right] = 1.$$

□

Theorem 1.7.41 (Bounded Convergence Theorem). *For a sequence of random variables $\{X_i\}_{i \geq 1}$, if*

$$\Pr \left[\lim_{n \rightarrow \infty} X_n = X \right] = 1$$

and there exists a finite b such that $|X_n| \leq b$ for all n , then

$$\lim_{n \rightarrow \infty} \mathbb{E}[X_n] = \mathbb{E}[X].$$

Corollary 1.7.42 (Cesàro Limit of Return Probability). *Let $\{X_t\}_{t \in T}$ be an irreducible Markov chain, then for any state i, j ,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} P_{i,j}^t = \frac{1}{h_{j,j}}.$$

Proof.

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} P_{i,j}^t = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} \mathbb{E} [1_{i,j}^t] = \lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{1}{n} \sum_{t \in [1, n]} 1_{i,j}^t \right] = \mathbb{E} \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} 1_{i,j}^t \right] = \frac{1}{h_{j,j}},$$

where the first 2 equalities are from linearity of expectation, and the third equality is from theorem 1.7.41. □

Lemma 1.7.43. *Positive recurrence and null recurrence are class properties.*

Proof. By corollary 1.7.20, states in a communicating class S are either all recurrent or all transient. Hence, we focus on recurrent class. Let state i, j be in S , and on $t_0, t_1 > 0$ we have $P_{i,j}^{t_0}, P_{j,i}^{t_1} > 0$. By corollary 1.7.42 and theorem 1.7.13,

$$\begin{aligned} \frac{1}{h_{i,i}} &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} P_{i,i}^t \geq \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [t_0+t_1+1, n]} P_{i,i}^t \\ &\geq P_{i,j}^{t_0} \cdot P_{j,i}^{t_1} \cdot \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n-t_0-t_1]} P_{j,j}^t \\ &= P_{i,j}^{t_0} \cdot P_{j,i}^{t_1} \cdot \lim_{n \rightarrow \infty} \frac{n-t_0-t_1}{n} \cdot \frac{1}{n-t_0-t_1} \sum_{t \in [1, n-t_0-t_1]} P_{j,j}^t \\ &= P_{i,j}^{t_0} \cdot P_{j,i}^{t_1} \cdot \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t \in [1, n]} P_{j,j}^t = P_{i,j}^{t_0} \cdot P_{j,i}^{t_1} \cdot \frac{1}{h_{j,j}}. \end{aligned}$$

The proof idea goes similar to corollary 1.7.20. If i is null recurrent by $h_{i,i}^{-1} = 0$, then j has to be null recurrent; conversely, if j is positive recurrent with $h_{j,j}^{-1} > 0$, then $h_{i,i}^{-1} > 0$ means i is positive recurrent. $\square$

1.7.4 Stationary Distributions

Definition 1.7.44. A state j in a discrete time Markov chain $\mathbf{X} = \{X_t\}$ is periodic if there exists an integer $\Delta > 1$

$$\Pr[X_{t+s} = j \mid X_t = j] = 0$$

unless $\Delta \mid s$. $\mathbf{X}$ is periodic if $\exists \Delta > 1$ such that all states have period Δ . Otherwise, a state or chain is aperiodic.

Example 1.7.45. For a discrete time Markov chain defined over $\mathbb{Z}$, that i transitions to $i \pm 1$ with probability $1/2$ each, the discrete time Markov chain is periodic as $\Delta = 2$.

Definition 1.7.46. An aperiodic, positive recurrent state is ergodic. A Markov chain is ergodic if all states are ergodic.

Corollary 1.7.47. Any finite, irreducible, aperiodic Markov chain is an ergodic Markov chain.

Lemma 1.7.48 (Ergodic Limit of Return Probability). *Let $\{X_t\}_{t \in T}$ be an irreducible ergodic Markov chain, the $\lim_{t \rightarrow \infty} P_{i,i}^t$ exists for any state i , and*

$$\lim_{t \rightarrow \infty} P_{i,i}^t = h_{i,i}^{-1}.$$

Definition 1.7.49 (Discrete Renewal Equation). Discrete convolution for sequences $\mathbf{a} = \{a_i\}_{i \geq 0}$ and $\mathbf{b} = \{b_i\}_{i \geq 0}$ is

$$(\mathbf{a} * \mathbf{b})_n = \sum_{k \in [0, n]} a_k \cdot b_{n-k}.$$

A renewal kernel is a probability mass function $\mathbf{f} = \{f_i\}_{i \geq 1}$ on positive integer such that each $f_i \geq 0$ and $\sum_{i \geq 1} f_i = 1$. Given a forcing sequence $\mathbf{g} = \{g_i\}_{i \geq 0}$, a sequence $\mathbf{u} = \{u_i\}_{i \geq 0}$ solves the renewal equation if

$$u_n = g_n + \sum_{i \in [1, n]} f_i \cdot u_{n-i} = g_n + (\mathbf{f} * \mathbf{u})_{n-1}.$$

Theorem 1.7.50 (Discrete Key Renewal Theorem [Fel68b]). *Let $\mu = \sum_{t \geq 1} t \cdot f_t$, and the span for $\mathbf{f}$ be $d = \gcd\{f_k > 0\}_{k \geq 1}$. If $d = 1$ then $\mathbf{f}$ is aperiodic. Supposing $\mathbf{f}$ is aperiodic and $\mu < \infty$, and $\mathbf{g} = \{g_i\}_{i \geq 0}$ is absolutely summable by $\sum_{i \geq 0} |g_i| < \infty$. Let $\mathbf{u} = \{u_i\}_{i \geq 0}$ be the unique bounded solution to the renewal equation by $u_n = g_n + (\mathbf{f} * \mathbf{u})_{n-1}$, then the limit exists by*

$$\lim_{n \rightarrow \infty} u_n = \frac{1}{\mu} \sum_{i \geq 0} g_i.$$

The proof for lemma 1.7.48 is immediate by theorem 1.7.50.

Proof. Given an irreducible ergodic Markov chain, suppose $X_0 = i$, then $P_{i,i}^n$ can be derived as a renewal equation by

$$P_{i,i}^n = \sum_{t \in [1, n]} r_{i,i}^t \cdot P_{i,i}^{n-t} = \left(\{r_{i,i}^t\}_{t \geq 1} * \{P_{i,i}^t\}_{t \geq 0} \right)_{n-1},$$

and the “forcing sequence” is $r_{i,i}^0 = 1$ follows 0s, as $P_{i,i}^0 = r_{i,i}^0 = 1$, thus the absolute sum is 1. Since the Markov chain is ergodic, then $\{r_{i,i}^t\}_{t \geq 1}$ is aperiodic and $\mu = h_{i,i} < \infty$. Hence,

$$\lim_{n \rightarrow \infty} P_{i,i}^n = \frac{1}{h_{i,i}}.$$

□

Example 1.7.51. Consider a sequence of fair, independent gambling games between 2 gamblers. One wins or loses a dollar with probability 1/2. The state of the chain at time t is the number of dollars won. Start from initial state 0.

We assume that a player loses by $-\ell_1$ dollars and wins by ℓ_2 dollars. Hence, these are the only 2 recurrent states, while the others are transitioning out to either $-\ell_1$ or ℓ_2 . We write P_i^t for the probability of i dollars at the t^{th} steps. Let s be one of the transient states, then $\lim_{t \rightarrow \infty} P_i^t = 0$. Hence, let $q = \lim_{t \rightarrow \infty} P_{-\ell_1}^t$, then $\lim_{t \rightarrow \infty} P_{\ell_2}^t = 1 - q$.

Let W^t be the random variable for the gain of the player at the t^{th} step, where the game has not stopped yet. On one hand, the games are fair, then by linearity of expectation, $\mathbb{E}[W^t] = 0$. On the other hand,

$$\mathbb{E}[W^t] = \sum_{i \in [-\ell_1, \ell_2]} i \cdot P_i^t = -\ell_1 \cdot q + \ell_2 \cdot (1 - q) = 0.$$

Therefore, $q = \ell_2 / (\ell_1 + \ell_2)$.

Definition 1.7.52. A stationary (equilibrium) distribution of a Markov chain is a distribution $\mathcal{D}_\pi$ such that $\pi = \mathbf{P}^T \pi$.

Theorem 1.7.53. Any finite, irreducible, ergodic Markov chain has the following properties:

- The chain has a stationary distribution $\mathcal{D}_\pi$ such that $\pi = [\pi_1, \dots, \pi_n]^T$.
- For all $i, j \in [1, n]$, the $\lim_{t \rightarrow \infty} P_{j,i}^t$ exists, and is independent of j .
- $\pi_i = \lim_{t \rightarrow \infty} P_{j,i}^t = h_{i,i}^{-1}$ for all $j \in [1, n]$.

Remark 1.7.54. The theorem gives 2 interpretations for the stationary distribution $\mathcal{D}_\pi$:

- If we run the finite, irreducible, ergodic Markov chain sufficiently long, then the initial state is forgotten, and the probability of being in state i is given by π_i .
- Given the prior observation of forgetting the initial state, then the limiting probability of $P_{i,i}^t$ being $h_{i,i}^{-1}$ becomes reasonable, that any initial state ends up in state i at the t^{th} step has same probability of $\pi_i = P_{i,i}^t$. Since $h_{i,i}$ is the average returning time from state i , then we should expect to be in state i with probability $h_{i,i}^{-1}$.

The proof for theorem 1.7.53 is immediate by lemma 1.7.48, together with showing $\lim_{t \rightarrow \infty} P_{j,i}^t = \lim_{t \rightarrow \infty} P_{i,i}^t$.

Proof. Given a state j that $j \neq i$, then similar to exercise 1.7.19, we write

$$P_{j,i}^t = \sum_{k \in [1, t]} r_{j,i}^k \cdot P_{i,i}^{t-k}.$$

Then for a $t_1 \leq t$, we have

$$\lim_{t \rightarrow \infty} P_{j,i}^t = \lim_{t \rightarrow \infty} \sum_{k \in [1, t]} r_{j,i}^k \cdot P_{i,i}^{t-k} \geq \lim_{t \rightarrow \infty} \sum_{k \in [1, t_1]} r_{j,i}^k \cdot P_{i,i}^{t-k} = \lim_{t \rightarrow \infty} P_{i,i}^t \sum_{k \in [1, t_1]} r_{j,i}^k \geq (1 - \varepsilon) \cdot \lim_{t \rightarrow \infty} P_{i,i}^t.$$

On the other hand, since $\sum_{k \in [1, t_1]} r_{j,i}^k \geq 1 - \varepsilon$, which leads to

$$\begin{aligned} \lim_{t \rightarrow \infty} P_{j,i}^t &= \lim_{t \rightarrow \infty} \sum_{k \in [1, t_1]} r_{j,i}^k \cdot P_{i,i}^{t-k} + \lim_{t \rightarrow \infty} \sum_{k \in [t_1+1, t]} r_{j,i}^k \cdot P_{i,i}^{t-k} \\ &\leq \lim_{t \rightarrow \infty} \sum_{k \in [1, t_1]} r_{j,i}^k \cdot P_{i,i}^{t-k} + \lim_{t \rightarrow \infty} \sum_{k \in [t_1+1, t]} r_{j,i}^k \\ &\leq \lim_{t \rightarrow \infty} P_{i,i}^t \sum_{k \in [1, t_1]} r_{j,i}^k + \varepsilon \\ &\leq \lim_{t \rightarrow \infty} P_{i,i}^t + \varepsilon. \end{aligned}$$

Since $\varepsilon \geq \sum_{k > t_1} r_{j,i}^k$, as $t_1 < t$ approaches ∞ , $\varepsilon \rightarrow 0$, which proves the limiting probability of $P_{j,i}^t$ is independent of j .

We now show $\sum \pi_i = 1$ and $\langle \pi, [P_{1,i}, \dots, P_{n,i}]^T \rangle = \pi_i$ for all $i \in [1, n]$ to complete the proof. Fix a state $j \in [1, n]$,

$$\sum_{i \in [1, n]} \pi_i = \sum_{i \in [1, n]} \lim_{t \rightarrow \infty} P_{i,i}^t = \sum_{i \in [1, n]} \lim_{t \rightarrow \infty} P_{j,i}^t = 1,$$

and by theorem 1.7.13,

$$\sum_{k \in [1, n]} \pi_k \cdot P_{k,i} = \sum_{k \in [1, n]} \lim_{t \rightarrow \infty} P_{k,k}^t \cdot P_{k,i} = \sum_{k \in [1, n]} \lim_{t \rightarrow \infty} P_{j,k}^t \cdot P_{k,i} = \lim_{t \rightarrow \infty} P_{j,i}^{t+1} = \lim_{t \rightarrow \infty} P_{i,i}^{t+1} = \pi_i.$$

Together with lemma 1.7.48, we complete the proof. $\square$

Remark 1.7.55. Being aperiodic in being ergodic is not necessary for the existence of the stationary distribution, but in this case, it is not the limiting probability, but rather the long term frequency of visiting states.

Theorem 1.7.56. *Let S be a set of states in a finite, irreducible, ergodic Markov chain. In the stationary distribution, the probability of leaving the state is equal to the probability of entering the state.*

Proof. Let there be a state $i \in [1, n]$ for the finite, irreducible, ergodic Markov chain, then by stationary distribution

$$\langle \pi, [P_{1,i}, \dots, P_{n,i}]^T \rangle = \pi_i = \pi_i \cdot \sum_{j \in [1, n]} P_{i,j},$$

that the invariant reaches by

$$\sum_{j \neq i} \pi_j \cdot P_{j,i} = \sum_{j \neq i} \pi_i \cdot P_{i,j}.$$

By generalizing to the subset S , the proof is completed. $\square$

Remark 1.7.57. The theorem 1.7.56 gives another way of deriving stationary distribution of $\mathbf{P}$ on a graph perspective. The cut-set around state i holds an invariant of $\sum_{j \neq i} \pi_j \cdot P_{j,i} = \sum_{j \neq i} \pi_i \cdot P_{i,j}$. If C is a cut-set around $S \subseteq [1, n]$, then in stationary distribution the probability crossing the cut in one direction is equivalent to the probability crossing the cut in the other direction.

Theorem 1.7.58. *Consider a finite, irreducible, ergodic Markov chain with a transition matrix $\mathbf{P}$. If there are nonnegative π such that $\sum_i \pi_i = 1$, and if any pairs $i, j \in [1, n]$ such that $\pi_j \cdot P_{j,i} = \pi_i \cdot P_{i,j}$, then π is a stationary distribution $\mathcal{D}_\pi$ for $\mathbf{P}$. Chains satisfying the condition are time reversible.*

Theorem 1.7.59. *An irreducible, aperiodic Markov chain belongs to one of the following two categories:*

- The chain is ergodic: For any pair of states i, j , $\lim_{t \rightarrow \infty} P_{j,i}^t$ exists and is independent of j , and $\mathcal{D}_\pi$ has $\pi_i = \lim_{t \rightarrow \infty} P_{j,i}^t > 0$.
- The chain has no positive recurrent state: For any state i , $\lim_{t \rightarrow \infty} P_{i,i}^t = 0$, and the chain has no $\mathcal{D}_\pi$.

Proof. Immediate by lemma 1.7.43, states in an irreducible ergodic Markov chain are either all positive recurrent or all null recurrent. Hence, by lemma 1.7.48, if the irreducible aperiodic Markov chain is positive recurrent,

$$\lim_{t \rightarrow \infty} P_{i,i}^t = \frac{1}{h_{i,i}}.$$

Showing the limiting probability $P_{j,i}^t$ is independent of j can be done by the same sandwiching trick in theorem 1.7.53.

Otherwise, the limiting recurrent probability of a null recurrent Markov chain has to be 0. By corollary 1.7.42,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in [1,n]} P_{j,i}^t = 0,$$

since $h_{i,i} = \infty$ for null recurrent states. $\lim_{t \rightarrow \infty} P_{j,i}^t > 0$ contradicts the Cesàro summation, hence $\lim_{t \rightarrow \infty} P_{j,i}^t = 0$. $\square$

Remark 1.7.60. Irreducibility and ergodicity leads to convergence, existence and uniqueness of stationary distribution. For both finite and countably infinite Markov chains, both needs to be irreducible and aperiodic, but in finite case, positive recurrence derive from the finiteness, hence the ergodicity derive from aperiodicity and positive recurrence, while the countably infinite case need to decide if the states are positive recurrent.

Exercise 1.7.61 (Exercise 7.11 [MU17]). An $n \times n$ matrix $\mathbf{P}$ is called stochastic if all entries are nonnegative and the sum of each row is 1. It is doubly stochastic if the sum of each column is 1. Show that the uniform distribution is a stationary distribution for any Markov chain represented by a doubly stochastic matrix.

Proof. Supposing

$$\sum_{j \in [1,n]} P_{i,j} = \sum_{j \in [1,n]} P_{j,i} = 1$$

for a doubly stochastic matrix, then the stationary distribution $\{\pi_i\}_{i \in [1,n]}$ has

$$\sum_{j \in [1,n]} \pi_j \cdot P_{j,i} = \frac{1}{n} = \pi_i$$

for an arbitrary $i \in [1, n]$, which completes the proof. $\square$

Exercise 1.7.62 (Exercise 7.12 [MU17]). Let X_n be the sum of n independent fair dice rolls. Show that for any $k \geq 2$,

$$\lim_{n \rightarrow \infty} \Pr[X_n \equiv 0 \pmod{k}] = \frac{1}{k}.$$

Proof. Let the current accumulation be X_n , $Y_n = X_n \pmod{k}$ be the accumulation modulo k , and $P_{i,j}$ be the transition probability over k states of Y_n , then $P_{i,j}$ can be represented as follows

$$P_{i,j} = \sum_{\substack{\alpha \in [1,6] \\ i+\alpha \equiv j \pmod{k}}} \frac{1}{6}.$$

Hence, $P_{a,b} = P_{b,c}$ if $a + \alpha \equiv b \pmod{k}$ and $b + \alpha \equiv c \pmod{k}$. Moreover, we discover the invariant of

$$\sum_{j \in [0,k-1]} P_{j,i} = \sum_{j \in [0,k-1]} P_{0,j} = 1,$$

making the uniform distribution the stationary distribution by exercise 1.7.61.

The chain is irreducible, as any 2 states can be reached from each other. The chain is also aperiodic, by existence of path of length $k-1$: rolling 1 for exactly $k-2$ times and 2 for a time, and path of length k : rolling 1 for k times. Since $\gcd(k, k-1) = 1$, the chain is aperiodic, and hence it is ergodic. By theorem 1.7.53,

$$\lim_{n \rightarrow \infty} \Pr[X_n \equiv 0 \pmod{k}] = \lim_{n \rightarrow \infty} P_{0,0}^n = \frac{1}{k}.$$

$\square$

Exercise 1.7.63 (Exercise 7.13 [MU17]). Consider a finite n state Markov chain with stationary distribution π and transition probability $P_{i,j}$. Starting at time 0 and running for m states, obtaining a sequence of states $X_0, X_1, \dots, X_m$. Consider the states in reverse order, $X_m, \dots, X_1, X_0$.

- Argue that given X_{t+1} , X_t is independent of $X_{t+2}, \dots, X_m$, that the reverse sequence is Markovian.
- Argue that for the reverse sequence, the transition probability is given by

$$Q_{i,j} = \frac{\pi_j \cdot P_{j,i}}{\pi_i}.$$

- Show that if the original Markov chain is time reversible, that $\pi_i \cdot P_{i,j} = \pi_j \cdot P_{j,i}$, then $Q_{i,j} = P_{i,j}$: the state follows the same transition probability, whether viewed in forward order or reverse order.

Proof. We already know the Markovian by

$$\Pr[X_{t+1} = s_{t+1} \mid X_t = s_t, \dots, X_1 = s_1, X_0 = s_0] = \Pr[X_{t+1} = s_{t+1} \mid X_t = s_t].$$

By Bayes' rule,

$$\Pr[X_t = s_t \mid X_{t+1} = s_{t+1}, \dots, X_m = s_m] = \frac{\Pr[X_t = s_t, \dots, X_m = s_m]}{\Pr[X_{t+1} = s_{t+1}, \dots, X_m = s_m]},$$

the numerator can be decomposed by

$$\Pr[X_t = s_t, \dots, X_m = s_m] = \Pr[X_t = s_t] \cdot \Pr[X_{t+1} = s_{t+1} \mid X_t = s_t] \cdots \Pr[X_m = s_m \mid X_{m-1} = s_{m-1}],$$

and the denominator can be decomposed by

$$\Pr[X_{t+1} = s_{t+1}, \dots, X_m = s_m] = \Pr[X_{t+1} = s_{t+1}] \cdot \Pr[X_{t+2} = s_{t+2} \mid X_{t+1} = s_{t+1}] \cdots \Pr[X_m = s_m \mid X_{m-1} = s_{m-1}].$$

Hence, we prove the independence by

$$\Pr[X_t = s_t \mid X_{t+1} = s_{t+1}, \dots, X_m = s_m] = \frac{\Pr[X_t = s_t, X_{t+1} = s_{t+1}]}{\Pr[X_{t+1} = s_{t+1}]} = \Pr[X_t = s_t \mid X_{t+1} = s_{t+1}].$$

Given 2 states i, j , when $t \rightarrow \infty$,

$$\lim_{t \rightarrow \infty} \Pr[X_t = j \mid X_{t+1} = i] = \lim_{t \rightarrow \infty} \frac{\Pr[X_t = j, X_{t+1} = i]}{\Pr[X_{t+1} = i]} = \lim_{t \rightarrow \infty} \frac{\Pr[X_{t+1} = i \mid X_t = j] \cdot \Pr[X_t = j]}{\Pr[X_{t+1} = i]},$$

hence

$$Q_{i,j} = \frac{\pi_j \cdot P_{j,i}}{\pi_i}.$$

If the chain is time reversible, then it is straightforward that $Q_{i,j} = P_{i,j}$ by swapping in $\pi_j \cdot P_{j,i} = \pi_i \cdot P_{i,j}$. $\square$

Exercise 1.7.64 (Exercise 7.18 [MU17]). Consider a random walk over 1/2/3-dimensional integer lattice, is each state transient, null recurrent, or positive recurrent? ³

Proof. Consider the 1-dimensional random walk, and the particle returns to 0 only on even length steps. Hence,

$$P_{0,0}^{2m} = \frac{1}{2^{2m}} \cdot \binom{2m}{m} \geq \frac{1}{2^{2m}} \cdot \frac{\sqrt{2\pi \cdot 2m} \cdot (2m)^{2m}}{(\sqrt{2\pi \cdot m} + 1)^2 \cdot m^{2m}} \geq \frac{1}{\sqrt{2\pi \cdot 2m}},$$

³A drunk man will find his way home, but a drunk bird may get lost forever. — Shizuo Kakutani

where the second inequality holds from lemma 1.7.1, and the third inequality holds from $\sqrt{2\pi \cdot 2m} > \sqrt{2\pi \cdot m} + 1$. By exercise 1.7.19, that

$$\sum_{t \geq 1} P_{0,0}^t = \sum_{m \geq 1} P_{0,0}^{2m} \geq \frac{1}{2\sqrt{\pi}} \cdot \sum_{m \geq 1} \frac{1}{\sqrt{m}}$$

is unbounded, and hence the 1-dimensional random walk is at least recurrent by corollary 1.7.20.

When we expand to 2-dimensional random walk, we want to show the probability of $P_{0^2,0^2}^{2m}$ from $P_{0,0}^{2t}$ by

$$\begin{aligned} P_{0^2,0^2}^{2m} &= \sum_{a+b=m} \frac{1}{2^{2m}} \cdot \binom{2m}{2a} \cdot P_{0,0}^{2a} \cdot P_{0,0}^{2b} \\ &= \sum_{a+b=m} \frac{1}{2^{2m}} \cdot \binom{2m}{2a} \cdot \frac{1}{2^{2m}} \cdot \binom{2a}{a} \binom{2b}{b} \\ &= \frac{1}{4^{2m}} \cdot \sum_{a+b=m} \frac{(2m)!}{(2b)! \cdot (2a)!} \cdot \frac{(2a)!}{(a!)^2} \cdot \frac{(2b)!}{(b!)^2} \\ &= \frac{1}{4^{2m}} \cdot \sum_{a+b=m} \binom{2m}{m} \binom{m}{a} \binom{m}{b} = \frac{1}{4^{2m}} \cdot \binom{2m}{m}^2 = \left(P_{0,0}^{2m}\right)^2, \end{aligned}$$

hence the 2-dimensional random walk is also at least recurrent by corollary 1.7.20 and

$$\sum_{t \geq 1} P_{0^2,0^2}^t = \sum_{m \geq 1} P_{0^2,0^2}^{2m} = \sum_{m \geq 1} \left(P_{0,0}^{2m}\right)^2 \geq \frac{1}{4\pi} \cdot \sum_{m \geq 1} \frac{1}{m} = \infty.$$

For a random walk of dimension $d \geq 3$, the exact probability of recurrence after $2m$ moves is

$$\begin{aligned} P_{0^d,0^d}^{2m} &= \sum_{a_1+\dots+a_d=m} \frac{1}{d^{2m}} \cdot \binom{2m}{2a_1} \dots \binom{2a_{d-1}+2a_d}{2a_{d-1}} \cdot P_{0,0}^{2a_1} \dots P_{0,0}^{2a_d} \\ &= \frac{1}{(2d)^{2m}} \cdot \sum_{a_1+\dots+a_d=m} \binom{2m}{2a_1; \dots; 2a_{d-1}} \cdot \binom{2a_1}{a_1} \dots \binom{2a_d}{a_d} \\ &= \frac{1}{(2d)^{2m}} \cdot \sum_{a_1+\dots+a_d=m} \frac{(2m)!}{(a_1!)^2 \dots (a_d!)^2} \end{aligned}$$

that is choosing each operation out of $2m$ moves over d dimensions, then each dimension contains only even moves, with the probability of returning to 0 after the even moves; or equivalently, for $2m$ moves, each choose 1 out of $2 \times d$ directions, and each opposite directions should have same amount of moves. We simplify further by noticing

$$\sum_{a_1+\dots+a_d=m} \binom{m}{a_1; \dots; a_d} = d^m,$$

and hence

$$P_{0^d,0^d}^{2m} = \frac{1}{2^{2m}} \cdot \binom{2m}{m} \sum_{a_1+\dots+a_d=m} \binom{m}{a_1; \dots; a_d}^2 \cdot \frac{1}{d^{2m}} \leq \frac{1}{2^{2m}} \cdot \binom{2m}{m} \max_{a_1+\dots+a_d=m} \binom{m}{a_1; \dots; a_d} \cdot \frac{1}{d^m}.$$

For fixed dimension d , the multinomial coefficient is maximized when the a_i are as balanced as possible, and

$$\max_{a_1+\dots+a_d=m} \binom{m}{a_1; \dots; a_d} = \Theta\left(d^m m^{-(d-1)/2}\right),$$

which follows from lemma 1.7.1⁴. Hence,

$$P_{0^d,0^d}^{2m} = \Theta\left(m^{-d/2}\right).$$

⁴<https://www.statslab.cam.ac.uk/~jrn10/Markov/s16.pdf>

Hence 3-dimensional (and above) random walk is transient by corollary 1.7.20.

To decide the 1-dimensional and 2-dimensional random walks are null recurrent or positive recurrent, recall the computation of $P_{0,0}^{2m}$ from exercise 1.7.19, that

$$P_{0,0}^{2m} = \sum_{k \in [1, m]} r_{0,0}^{2k} \cdot P_{0,0}^{2m-2k},$$

we can (no I can't) construct as follows

$$p(x) - 1 = \sum_{m \geq 1} P_{0,0}^{2m} \cdot x^m = \sum_{m \geq 1} \sum_{k \in [1, m]} r_{0,0}^{2k} \cdot x^k \cdot P_{0,0}^{2m-2k} \cdot x^{m-k} = \left(\sum_{k \geq 1} r_{0,0}^{2k} \cdot x^k \right) \cdot \left(\sum_{t \geq 0} P_{0,0}^{2t} \cdot x^t \right) = r(x) \cdot p(x),$$

then $r(x) = 1 - 1/p(x)$. Since

$$\frac{1}{\sqrt{1-4x}} = \sum_{t \geq 0} \binom{2t}{t} x^t, \quad (1.8)$$

then $p(x) = (1-x)^{-1/2}$, and hence $r(x) = 1 - \sqrt{1-x}$, yielding

$$\begin{aligned} r_{0,0}^{2m} &= -\binom{1/2}{m} \cdot (-1)^m = (-1)^{m-1} \cdot \frac{(1/2) \cdots (1/2 - (m-1))}{m!} \\ &= \frac{1}{2^m} \cdot \frac{1 \cdots (2m-3)}{m!} = \frac{1}{2^m} \cdot \frac{(2m-2)!}{2^{m-1} \cdot (m-1)! \cdot m!} \\ &= \frac{1}{2^{2m}} \cdot \frac{2}{2m-1} \cdot \binom{2m-1}{m} = \frac{1}{2^{2m}} \cdot \frac{1}{2m-1} \cdot \binom{2m}{m} \end{aligned}$$

for $m \geq 1$, and

$$h_{0,0} = \sum_{m \geq 1} r_{0,0}^{2m} \cdot 2m \geq \sum_{m \geq 1} \frac{1}{2^{2m}} \cdot \binom{2m}{m} \geq \frac{1}{2\sqrt{\pi}} \cdot \sum_{m \geq 1} \frac{1}{\sqrt{m}},$$

which is unbounded, meaning the whole 1-dimensional random walk is null recurrent by lemma 1.7.43.

The 2-dimensional random walk is also null recurrent. If it is positive recurrent, then by lemma 1.7.48, the limit of $P_{0^2,0^2}^{2m}$ should be positive, yet $\left(P_{0,0}^{2m}\right)^2$ approaches 0, which is a contradiction, hence it is null recurrent. $\square$

After exercise 1.7.64 and by lemma 1.7.1, we derive a finer bound than corollary 1.6.3.

Lemma 1.7.65.

$$\frac{2}{\sqrt{\pi n}} \cdot 2^{2n} \geq \binom{2n}{n} \geq \frac{1}{2\sqrt{\pi n}} \cdot 2^{2n}.$$

Exercise 1.7.66 (Exercise 7.17 [MU17]). Consider the Markov chain over $\mathbb{N}$: For $i \geq 1$, the particle move to $i+1$ with probability p , and to $i-1$ with probability $1-p$ for $i \geq 1$, and from 0 to 1 with probability 1.

Show that the chain is transient when $p > 1/2$, positive recurrent when $p < 1/2$, and null recurrent when $p = 1/2$.

Proof. The particle returns to 0 only on even number of steps. Hence, we want to count all possible $2m$ steps random walk over $\mathbb{Z}$, such that the particle never go to negative side, to derive $r_{0,0}^{2m}$ and $P_{0,0}^{2m}$.

There are 2 relevant combinatorics tools: Catalan number and Dyck path. A Dyck path is a series of ups and downs. The path will begin and end on the same level; and as the path moves from left to right it will rise and fall, never dipping below the height it began on. Hence, the Dyck path models the desired random walk well.

Catalan number C_n counts unique Dyck paths of n ups and n downs. Noticing that for n ups and n downs, C_{n-1} is the number of unique Dyck paths that first returns after $2n$ steps. Hence, we can count C_n by

$$C_n = \sum_{k \in [1, n]} C_{k-1} C_{n-k},$$

that leads by first returns from 2 to $2n$ steps, and $C_0 = 1$. We derive C_i with generating function method in the same way as exercise 1.7.64, by constructing

$$\frac{1}{x} (f(x) - 1) = \sum_{n \geq 1} C_n \cdot x^{n-1} = \sum_{n \geq 1} \sum_{k \in [1, n]} C_{k-1} \cdot x^{k-1} \cdot C_{n-k} \cdot x^{n-k} = \left(\sum_{t \geq 0} C_t \cdot x^t \right)^2 = f(x)^2,$$

Then $f(x) = (1 - \sqrt{1 - 4x}) \cdot (2x)^{-1}$, as $f(x)$ approaches 1 when $x = 0$. By previous eq. (1.8) in exercise 1.7.64,

$$1 - \sqrt{1 - 4x} = - \sum_{n \geq 1} \binom{1/2}{n} \cdot (-4x)^n = \sum_{n \geq 1} \frac{1}{2n-1} \binom{2n}{n} x^n = \sum_{n \geq 0} \frac{2}{n+1} \binom{2n}{n} x^{n+1} = 2x \sum_{n \geq 0} \frac{1}{n+1} \binom{2n}{n} x^n = 2x \sum_{n \geq 0} C_n x^n.$$

Hence, $r_{0,0}^{2n} = C_{n-1} \cdot p^{n-1} \cdot (1-p)^n$, and we march on to derive $h_{0,0}$:

$$\begin{aligned} h_{0,0} &= \sum_{n \geq 1} C_{n-1} \cdot (1-p)^n \cdot p^{n-1} \cdot 2n \\ &= \sum_{n \geq 0} C_n \cdot (1-p)^{n+1} \cdot p^n \cdot (2n+2) \\ &= \sum_{n \geq 0} \frac{1}{n+1} \binom{2n}{n} \cdot (1-p)^{n+1} \cdot p^n \cdot (2n+2) \\ &= 2(1-p) \sum_{n \geq 0} \binom{2n}{n} (1-p)^n p^n. \end{aligned}$$

Again by eq. (1.8),

$$h_{0,0} = \frac{2-2p}{\sqrt{1-4p(1-p)}} = \frac{2-2p}{|1-2p|}.$$

- When $p < 1/2$, $h_{0,0}$ is bounded, and the chain is positive recurrent.
- When $p = 1/2$, $h_{0,0}$ is infinity, and the chain is null recurrent.
- When $p > 1/2$, $h_{0,0}$ is not valid, and the chain is transient.

Another way of showing this is to use the cut-set way by theorem 1.7.56. ⁵ Assuming π exists, then we can derive

$$\begin{aligned} \pi_0 &= (1-p)\pi_1, \\ \pi_i &= \pi_{i-1} + (1-p)\pi_{i+1}, \end{aligned}$$

and we have $(1-p)\pi_{i+1} = \pi_i$.

- When $p < 1/2$, $\pi_0 = (1-2p)/(2-2p)$, hence the chain is positive recurrent.
- When $p = 1/2$, $\sum_i \pi_i = 1$ but π_i are 0, by lemma 1.7.48 the chain is null recurrent.
- When $p > 1/2$, $\sum_i \pi_i$ diverges for any positive π_0 , hence no stationary distribution exists.

□

Lemma 1.7.67. A Dyck path is a series of ups and downs. The path will begin and end on the same level; and the path never dips below the height it began on. Catalan number C_n counts Dyck paths of n ups and n downs by

$$C_n = \frac{1}{n+1} \binom{2n}{n}$$

with the property $C_0 = 1$ and

$$C_n = \sum_{k \in [1, n]} C_{k-1} C_{n-k}.$$

⁵For further read on other methods, <https://chihaozhang.com/teaching/SP2025/lec4.pdf>

Exercise 1.7.68 (Exercise 7.19 [MU17]). Consider example 1.7.51. Show that the expected number of games is $\ell_1 \ell_2$.

Proof. Let X_j be the random variable for the number of steps to be absorbed for a particle to start at state j . Then

$$\begin{aligned}\mathbb{E}[X_j] &= \mathbb{E}[X_j \mid \text{move to } j+1] \cdot r_{j,j+1}^1 + \mathbb{E}[X_j \mid \text{move to } j-1] \cdot r_{j,j-1}^1 \\ &= \frac{1}{2} \mathbb{E}[X_{j+1}] + \frac{1}{2} \mathbb{E}[X_{j-1}] + 1\end{aligned}$$

for $j \in [-\ell_1 + 1, \ell_2 - 1]$, and $\mathbb{E}[X_{-\ell_1}] = \mathbb{E}[X_{\ell_2}] = 0$. Now we have the linear system for $k \in [-\ell_1 + 1, \ell_2 - 1]$

$$2\mathbb{E}[X_k] = \mathbb{E}[X_{k-1}] + \mathbb{E}[X_{k+1}] + 2$$

and solve by

$$\begin{aligned}\mathbb{E}[X_{-\ell_1+1}] &= \frac{1}{2} \mathbb{E}[X_{-\ell_1+2}] + 1 \\ &\dots \\ \mathbb{E}[X_0] &= \frac{\ell_1}{\ell_1 + 1} \mathbb{E}[X_1] + \ell_1 \\ &\dots \\ \mathbb{E}[X_{\ell_2-1}] &= \frac{\ell_1 + \ell_2 - 1}{\ell_1 + \ell_2} \mathbb{E}[X_{\ell_2}] + (\ell_1 + \ell_2 - 1)\end{aligned}$$

Hence, for $i \in [-\ell_2, \ell_1]$, we have $\mathbb{E}[X_i] = (\ell_1 + i) \cdot (\ell_2 - i)$ to complete the proof. $\square$

Exercise 1.7.69 (Exercise 7.20 [MU17]). Consider the gambler's ruin where games are unfair with losing probability $2/3$. Suppose one starts with i dollars by either reaching n or 0. Let W_t be the amount one gained after t games.

- Show that $\mathbb{E}[2^{W_{t+1}}] = \mathbb{E}[2^{W_t}]$.
- Determine the probability of finishing with 0 and the probability of finishing with n when starting at position i .
- Generalize when losing probability is $p > 1/2$.

Proof. We write p_t^j to denote the probability of being at position j at t^{th} step. Moreover, $r_{k,k-1}^1 = 2/3$ and $r_{k,k+1}^1 = 1/3$ for each $k \in [1, n-1]$, we have the following invariant

$$2^k = \frac{2}{3} \cdot 2^{k-1} + \frac{1}{3} \cdot 2^{k+1} = r_{k,k-1}^1 \cdot 2^{k-1} + r_{k,k+1}^1 \cdot 2^{k+1},$$

and hence $\mathbb{E}[2^{W_t}] = \mathbb{E}[2^{W_{t+1}}]$ by

$$\mathbb{E}[2^{W_{t+1}}] = 2^0 \cdot p_t^0 + \sum_{k \in [1, n-1]} \left(r_{k,k-1}^1 \cdot 2^{k-1} + r_{k,k+1}^1 \cdot 2^{k+1} \right) \cdot p_t^k + 2^n \cdot p_t^n = \sum_{k \in [0, n]} 2^k \cdot p_t^k = \mathbb{E}[2^{W_t}].$$

Since $\mathbb{E}[2^{W_0}] = \mathbb{E}[2^{W_t}]$ for arbitrary $t > 0$, then $\mathbb{E}[2^{W_0}] = 2^i = \mathbb{E}[2^{W_t}]$.

Let q_i be the probability of absorbing to n from state i , then

$$\lim_{t \rightarrow \infty} \mathbb{E}[2^{W_t}] = q_i \cdot 2^n + (1 - q_i) \cdot 2^0 = 2^i,$$

hence $q_i = (1 - 2^i) / (1 - 2^n)$.

In general, for losing probability p , we want to find similar invariant $\mathbb{E}[q^{W_t}] = \mathbb{E}[q^{W_{t+1}}]$ such that

$$p \cdot q^{t-1} + (1 - p) \cdot q^{t+1} = q^t,$$

and solve $(1 - p)x^2 - x + p = 0$ for $x = p/(1 - p)$, which yields

$$q_k = \frac{1 - x^k}{1 - x^n}.$$


When $p = 1/2$, $q_k = k/n$ as in example 1.7.51. $\square$

Exercise 1.7.70 (Exercise 7.27 [MU17]). Suppose that we are given n records, $R_1, R_2, \dots, R_n$ kept in some order. The cost of accessing the j^{th} record in the order is j . Thus, if we had four records ordered as R_2, R_4, R_3, R_1 , then the cost of accessing R_4 would be 2 and the cost of accessing R_1 would be 4.

Suppose further that, at each step, record R_j is accessed with probability p_j , with each step being independent of other steps. If we knew the values of the p_j in advance, we would keep the R_j in decreasing order with respect to p_j . But if we don't know the p_j in advance, we might use the “move to front” heuristic: at each step, put the record that was accessed at the front of the list. We assume that moving the record can be done with no cost and that all other records remain in the same order. For example, if the order was R_2, R_4, R_3, R_1 before R_3 was accessed, then the order at the next step would be R_3, R_2, R_4, R_1 .

The order of the records can be thought of as the state of a Markov chain. Give the stationary distribution of this chain. Also, let X_k be the cost for accessing the k^{th} requested record. Determine an expression for $\lim_{k \rightarrow \infty} \mathbb{E}[X_k]$.

Proof. By the building process, we observe a pattern: To build towards a target $\overline{a_1 \dots a_n}$, we build from $\overline{a_n}$ to $\overline{a_{n-1}a_n}$, then to $\overline{a_k \dots a_n}$, finally $\overline{a_1 \dots a_n}$. We illustrate here by enforcing each state to contain substring $\overline{a_k \dots a_n}$ at most, and




notice the absorbing relation

$$\overline{a_k \dots a_n} = \overline{a_1 \dots a_n} \cup \bigcup_{i \in [2, k]} (\overline{a_i \dots a_n} \setminus \overline{a_{i-1} \dots a_n}),$$

then we can use cutset to inductively derive the stationary distribution for $\overline{a_1 \dots a_n}$. For simplicity we say $\pi_k = \pi_{\overline{a_k \dots a_n}}$.


The simplest case is as follows,



and hence the stationary distribution of $\overline{a_{n-1}a_n}$ is

$$\pi_{n-1} = \frac{p_{a_{n-1}}}{p_{a_{n-1}} + p_{a_n}} \cdot \pi_n = \frac{p_{a_{n-1}}}{p_{a_{n-1}} + p_{a_n}}.$$

Relaxing $\overline{a_{n-1}a_n}$ to $\overline{a_{n-1}a_n \setminus a_{n-2} \dots a_n}$ and $\overline{a_{n-2} \dots a_n}$, we have



and hence the stationary distribution for $\overline{a_{n-2} \dots a_n}$ is

$$\pi_{n-2} = \frac{p_{a_{n-2}}}{p_{a_{n-2}} + \dots + p_{a_n}} \cdot \pi_{n-1}.$$

By cutset, for $\overline{a_k \dots a_n}$, the inflowing probability is its stationary distribution, then

$$p_{a_k} \cdot (\pi_{k+1} - \pi_k) + \left(1 - \sum_{i \in [k+1, n]} p_{a_i}\right) \cdot \pi_k = \pi_k,$$

and therefore

$$\pi_k = \frac{p_{a_k}}{p_{a_k} + \dots + p_{a_n}} \cdot \pi_{k+1}.$$

Hence, by induction, the resulting stationary distribution is

$$\pi_1 = p_{a_1} \cdot \frac{p_{a_2}}{1 - p_{a_1}} \cdots \frac{p_{a_{n-1}}}{1 - \sum_{i \in [1, n-2]} p_i} \cdot \frac{p_{a_n}}{1 - \sum_{i \in [1, n-1]} p_i} = \prod_{j \in [1, n]} p_j \cdot \prod_{k \in [1, n]} \left(\sum_{i \in [k, n]} p_{a_i} \right)^{-1}.$$

To derive the $\mathbb{E}[X_k]$ as $k \rightarrow \infty$, we eventually reach the stationary distribution after infinity steps, and hence we want to derive the expected distance to the first place in the list, for each record. By prior cutset method and state building, the probability of having i^{th} record, and j^{th} record in any position after i^{th} record, is

$$\pi_{i\bar{j}} = \frac{p_i}{p_i + p_j}.$$

Let Y_i be random variable for the order of i^{th} record in the queue, then by linearity of expectation,

$$\mathbb{E}[Y_i] = 1 + \sum_{j \neq i} \pi_{j\bar{i}} = 1 + \sum_{j \neq i} \frac{p_j}{p_i + p_j}.$$

Hence, by conditional expectation, we have

$$\lim_{k \rightarrow \infty} \mathbb{E}[X_k] = \lim_{k \rightarrow \infty} \sum_{i \in [1, n]} p_i \cdot \mathbb{E}[X_k \mid k^{\text{th}} \text{ query is on } i^{\text{th}} \text{ record}] = \sum_{i \in [1, n]} p_i \cdot \mathbb{E}[Y_i] = 1 + \sum_{i \in [1, n]} p_i \sum_{j \neq i} \frac{p_j}{p_i + p_j}.$$


□

Remark 1.7.71. Buddy, why torture me with Tsetlin Library exercise 1.7.70 [Tse63].

Exercise 1.7.72 (Exercise 7.28 [MU17]). Consider the following variation of the discrete time queue. Time is divided into fixed-length steps. At the beginning of each time step, a customer arrives with probability λ . At the end of each time step, if the queue is nonempty then the customer at the front of the line completes service with probability μ .

- Explain under what conditions a stationary distribution exists, and find the stationary distribution when it exists.
- Now consider the variation where we change the order of incoming arrivals and service. That is: at the beginning of each time step, if the queue is nonempty then a customer is served with probability μ ; and at the end of a time step a customer arrives with probability λ . How does this change the stationary distribution?

Proof. We show the vanilla Markov chain of the infinite discrete queue that each step only 1 of the 2 things happens.




By cutset $\lambda\pi_i = \mu\pi_{i+1}$, π exists only on $\lambda < \mu$ when the chain is positively recurrent,

$$\pi_i = \left(1 - \frac{\lambda}{\mu}\right) \cdot \left(\frac{\lambda}{\mu}\right)^i,$$

otherwise when $\lambda = \mu$, the chain is null recurrent, $\lambda > \mu$ the chain is transient.

When the chain introduces a new customer at the beginning of a time step, and processes a customer at the end of the time step, the Markov chain follows, where remaining unchanged can be either no customer coming and no customer processed, or 1 customer coming and 1 customer processed. Otherwise, transitioning to the next state means 1 customer coming, and transitioning to the previous state means 1 customer being processed.




By cutset $\lambda(1-\mu)\pi_i = (1-\lambda)\mu\pi_{i+1}$, π exists only on $\lambda < \mu$ when the chain is positively recurrent,

$$\pi_i = \left(1 - \frac{\lambda(1-\mu)}{\mu(1-\lambda)}\right) \cdot \left(\frac{\lambda(1-\mu)}{\mu(1-\lambda)}\right)^i,$$

and the chain being null recurrent on $\lambda = \mu$, being transient on $\lambda > \mu$.

When we reverse the order, say on new time step with probability μ the nonempty queue a customer is processed, and a customer will arrive with probability λ at the end of the time step, the new Markov chain follows.



The only difference lies in the empty queue, that follows the vanilla discrete time queue behavior, as the empty queue does not have customers to consume, and either with probability λ to accept new customer at the end of the time step, or remain empty with probability $1 - \lambda$.

By cutset, we have

$$\begin{aligned}\lambda\pi_0 &= (1-\lambda)\mu\pi_1 \\ \lambda(1-\mu)\pi_i &= (1-\lambda)\mu\pi_{i+1},\end{aligned}$$

hence for $i \geq 1$,

$$\pi_i = \frac{\pi_0}{1-\mu} \cdot \left(\frac{\lambda(1-\mu)}{\mu(1-\lambda)}\right)^i,$$

and by $\sum_{i \geq 0} \pi_i = 1$, we have

$$\pi_0 = 1 - \frac{\lambda}{\mu}.$$

Therefore, for $i \geq 1$,

$$\pi_i = \frac{\pi_0}{1-\mu} \cdot \left(\frac{\lambda(1-\mu)}{\mu(1-\lambda)}\right)^i,$$

with the existence condition for π being $\lambda < \mu$, in which case the chain is positively recurrent. $\square$

1.7.5 Random Walks on Undirected Graphs

Definition 1.7.73. A random walk on G is a Markov chain defined by the sequence of moves of a particle between vertices of G . The position of the particle at a particular time step is the state of the chain. If the particle is at vertex i with $d(i)$ outgoing edges, then the probability that the particle follows edge (i, j) to neighboring vertex j is $d(i)^{-1}$.

Lemma 1.7.74. A random walk on an undirected graph G is aperiodic iff G is not bipartite, which means there is no 2-coloring such that the endpoints of each edge are colored differently.

Proof. A graph G being not bipartite means it has cycles of odd length. If a graph G is bipartite and undirected, then the length of the path going back to itself is always even, such that the Markov chain is of period 2. Otherwise, undirected G is not bipartite, each vertex has an odd-length path back to itself, then the Markov chain is aperiodic. $\square$

Similar to theorem 1.7.53, A random walk on a finite, undirected, connected, and non-bipartite graph G defines a Markov chain that converges to a stationary distribution that is determined by the degree sequences of the graph.

Theorem 1.7.75. A random walk on a finite, undirected, connected, and non-bipartite G has a stationary distribution π where

$$\pi_v = \frac{d(v)}{2|E|}.$$

Proof. Since $\sum_{v \in V} d(v) = 2|E|$, then π defines a proper distribution of particle appearing over V .

On the other hand, by definition of stationary distribution in theorem 1.7.56, for $u \in V$ and its neighbors $\Gamma(u)$,

$$\sum_{v \in \Gamma(u)} \frac{d(v)}{2|E|} \cdot \frac{1}{d(v)} = \frac{d(u)}{2|E|} = \frac{d(u)}{2|E|} \cdot \sum_{v \in \Gamma(u)} \frac{1}{d(u)},$$

where the second equality holds from $|\Gamma(u)| = d(u)$. $\square$

Exercise 1.7.76 (Exercise 7.14 [MU17]). Show that the Markov chain defined by a random walk on a finite, undirected, non-bipartite, connected graph G is time reversible.

Proof. When G is finite, then theorem 1.7.75 kicks in, by definition of time reversibility in exercise 1.7.63,

$$\frac{d(v)}{2|E|} \cdot \frac{1}{d(v)} = \frac{1}{2|E|} = \frac{d(u)}{2|E|} \cdot \frac{1}{d(u)}.$$

$\square$

Lemma 1.7.77 (Commute Time). For $G = (V, E)$, if $(u, v) \in E$, the commute time $h_{u,v} + h_{v,u} \leq 2|E|$.

Proof. Let D be the set of directed edges: for every $(u, v) \in E$, we have 2 directed edges $u \rightarrow v$ and $v \rightarrow u$ in D .

We can view the random walk over G as a Markov chain of state space D , where the state of the Markov chain at the t^{th} time step is the directed edge taken by the random walk in its t^{th} transition. The Markov chain has $2|E|$ states, and the stationary distribution $\pi_{u \rightarrow v} = (2|E|)^{-1}$ for all $(u, v) \in E$, by cutset theorem 1.7.56

$$\pi_{u \rightarrow v} = \frac{1}{d(u)} \sum_{w \in \Gamma(u)} \pi_{w \rightarrow u}$$

for all $(u, v) \in E$ and $\Gamma(u)$ is the set of neighboring vertices of u , $\pi_{u \rightarrow v} = (2|E|)^{-1}$ is a solution. By theorem 1.7.53,

$$\pi_{u \rightarrow v} = h_{u \rightarrow v, u \rightarrow v}^{-1} = (2|E|)^{-1},$$

hence $h_{u \rightarrow v, u \rightarrow v} = 2|E|$: after taking the directed path $u \rightarrow v$, the expected number of steps to retake $u \rightarrow v$ is $2|E|$. By Markovian (memoryless) of the random walk, once it reaches v , it forgets about it reaches by $u \rightarrow v$, hence the expected time to start from v and reach u and take (u, v) to v is $2|E|$.

Since $h_{u,v}$ will not rule out the walk reaching v without taking (u, v) , then it will not force the walk to go back to u and take (u, v) , even if it reached v before, hence $h_{v,u} + h_{u,v} \leq h_{u \rightarrow v, u \rightarrow v} = 2|E|$. $\square$

Exercise 1.7.78 (Exercise 7.22 [MU17]). A cat and a mouse each independently take a random walk on a connected, undirected, non-bipartite graph G . They start at the same time on different nodes, and each makes 1 transition at each time step. The cat eats the mouse if they are ever at the same node at a time step. Let $|V| = n$ and $|E| = m$. Show an upper bound of $O(m^2n)$ on the expected time before the cat eats the mouse.

Proof. Construct a graph G^2 from G by making vertices $V^2 = (V, V)$, and edges E^2 for all $((u_0, u_1), (v_0, v_1))$ such that $(u_0, v_0) \in E$ and $(u_1, v_1) \in E$. Counting the degrees of each vertices $(a, b) \in V^2$ double counts $|E^2|$, then

$$2|E^2| = \sum_{(a,b) \in V^2} d(a,b) = \sum_{a \in V} d(a) \sum_{b \in V} d(b) = 4|E|^2,$$

such that $|E^2| = 2|E|^2$.

We show that, for every $(a, b) \in V^2$, there exists paths of $O(|V|)$ length to a $(v, v) \in V^2$ in the connected, undirected, non-bipartite graph. First, for every $u, v \in V$, the path in between over G is $O(|V|)$ length. Then, for paths in between a, v and b, v , we use ℓ_1, ℓ_2 to denote the paths lengths,

- If both ℓ_1, ℓ_2 are even or odd, then there exists path in between $(a, b), (v, v)$ of length $\max(\ell_1, \ell_2) = O(|V|)$, as the shorter one can keep looping in between v and $u \in \Gamma(v)$ for $2k$ steps, such that they eventually meet at (v, v) .
- Otherwise, since G is non-bipartite, there exists odd length cycle in G , such that the odd length path can be transformed into an even length path, and the 2 lengths can be calibrated to be the same as previously described. The cycle is of $O(|V|)$ length, and original paths are $O(|V|)$ lengths, hence the resulting length is $O(|V|)$.

After obtaining such path s from (a, b) to (v, v) of length $\ell = O(|V|)$, we bound $h_{(a,b),(v,v)}$ by

$$h_{(a,b),(v,v)} \leq h_{(a,b),(v,v)} + h_{(v,v),(a,b)} \leq \sum_{i \in [1, \ell-1]} (h_{(s_i, s_{i+1})} + h_{(s_{i+1}, s_i)}) \leq \ell \cdot 2|E^2| = 4\ell|E|^2.$$

where the second inequality holds by enforcing the path s from (a, b) to (v, v) , similar to the idea of lemma 1.7.77, and the third inequality is immediate by lemma 1.7.77. Hence, the expected time $h_{(a,b),(v,v)} = O(|E|^2|V|)$. $\square$

Definition 1.7.79 (Covering Time). The covering time of $G = (V, E)$ is the max of all $v \in V$ of the expected time to visit all vertices by a random walk starting from v .

Lemma 1.7.80 (Covering Time Upper Bound). The covering time C_G of $G = (V, E)$ is upper bounded by $2|E|(|V| - 1)$.

Proof. One can consider traversing all vertices with a MST in G , such that there are $|V| - 1$ directed paths in the MST. Then we bound similarly in lemma 1.7.77 and exercise 1.7.78 by enforcing the paths

$$C_G \leq \sum_{i \in [1, |V|-1]} (h_{v_i, v_{i+1}} + h_{v_{i+1}, v_i}) \leq 2|E|(|V| - 1).$$

$\square$

We introduce Matthew's theorem, connecting the hitting time $h_{u,v}$ to the covering time. We write $H(n) = \sum_{i \leq n} 1/i$.

Lemma 1.7.81 (Matthews' Theorem [Mat88]). The covering time C_G of $G = (V, E)$ with $|V| = n$ is bounded by

$$H(n-1) \min_{u,v \in V, u \neq v} h_{u,v} \leq C_G \leq H(n-1) \max_{u,v \in V, u \neq v} h_{u,v}.$$

Proof. The proof utilizes balls-and-bins model, and the path enforcing technique of lemma 1.7.77 and exercise 1.7.78.

We begin by enforcing a traversal order $\{Z_i\}_{i \in [1, n]}$ over V sampled from uniform permutation over $[1, n]$, write B as the maximal hitting time $h_{u,v}$ for convenience, T_i be the first time step when all the first i vertices $Z_1, \dots, Z_i$ have been visited in the order⁶, and X_i be the state of the Markov chain at the i^{th} time step.

⁶If Z_j are visited for $j \in [1, k] \setminus \{i\}$, then when Z_i is visited, $T_i = T_{i+1} = \dots = T_k$.

Given all the prior setup, we are interested in the time intervals $\{T_i - T_{i-1}\}_{i \in [2, n]}$. Given the random walk history,

$$Y_i = \mathbb{E} [T_i - T_{i-1} \mid Z_1, \dots, Z_i; X_1, \dots, X_{T_{i-1}}]$$

for $i \in [2, n]$ ⁷, and the expected time to cover the graph starting from u is

$$C_G = \sum_{i \in [2, n]} Y_i + \mathbb{E} [T_1].$$

For T_1 , one has $1/n$ probability having $Z_1 = u$, otherwise u is reached by random walk. By conditional expectation,

$$\begin{aligned} \mathbb{E} [T_1] &= \mathbb{E} [T_1 \mid Z_1 = u] \Pr [Z_1 = u] \\ &\quad + \mathbb{E} [T_1 \mid Z_1 \neq u] \Pr [Z_1 \neq u] \\ &\leq \left(1 - \frac{1}{n}\right) B, \end{aligned}$$

where the second inequality holds by B upper bounding the expectation of hitting time, and $T_1 = 0$ when $Z_1 = u$.

For Y_i , if i has already been visited previously, then $Y_i = 0$ as $T_i = T_{i-1}$, otherwise $Y_i \leq B$ regardless.

Since the permutation $\{Z_i\}_{i \in [1, n]}$ is sampled uniformly random and independent of the random walk, then any one of $\{Z_k\}_{i \in [1, k]}$ can be the last visited for first k entries with probability $1/k$ ⁸, and hence

$$\begin{aligned} \sum_{i \in [2, n]} Y_i + \mathbb{E} [T_i] &\leq \sum_{i \in [2, n]} \frac{1}{i} B + \left(1 - \frac{1}{n}\right) B \\ &= \sum_{i \in [1, n-1]} \frac{1}{i} B \\ &= H(n-1)B. \end{aligned}$$

The lower bound is proven in the same way, by setting B to be the minimal of the expected hitting time. $\square$

Exercise 1.7.82 (Exercise 7.24 [MU17]). A lollipop graph on n vertices is a clique $K_{n/2}$ connected to a $n/2$ long path. u is the node being both clique and path, and v be the other end of the path. Show the C_G starting from v is $\Theta(n^2)$, and starting from u is $\Theta(n^3)$.

Proof. We bound the covering time from v by enforcing the path from v to u , then to other $x \in K_{n/2}$. By lemma 1.7.80,

$$C_G \leq h_{v,u} + \sum_{x \in K_{n/2}, x \neq u} (h_{x,u} + h_{u,x}).$$

For $h_{v,u}$, let $i \in [1, n/2 - 1]$ be vertices on the (u, v) path, we have

$$\begin{aligned} h_{1,u} &= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (1 + h_{2,u}) \\ &\quad \dots \\ h_{n/2-1,u} &= \frac{1}{2} \cdot (1 + h_{n/2-2,u}) + \frac{1}{2} \cdot (1 + h_{v,u}) \\ h_{v,u} &= 1 + h_{n/2-1,u}, \end{aligned}$$

hence $h_{i,u} = i \cdot (n - i)$, and $h_{v,u} = n^2/4$.

⁷Coupon-collector on the first k entries, that does not strictly enforces traverse by the permutation, but enforces by visiting first k entries.

⁸ Z_k being last visited w.r.t. the fixed random walk is only related to $\{Z_i\}_{i \in [1, k]}$, by symmetry of random permutation, the probability is $1/k$.

For $h_{u,x}$ for some $x \in K_{n/2}$, let $k = n/2$ and $x' \neq x$ that $x' \in K_{n/2}$, then

$$\begin{aligned} h_{u,x} &= \frac{1}{k} \cdot 1 + \frac{k-2}{k} \cdot (1 + h_{x',x}) + \frac{1}{k} \cdot (1 + h_{1,x}) \\ h_{x',x} &= \frac{1}{k-1} \cdot 1 + \frac{k-3}{k-1} \cdot (1 + h_{x',x}) + \frac{1}{k-1} \cdot (1 + h_{u,x}) \\ h_{1,x} &= \frac{1}{2} \cdot (1 + h_{u,x}) + \frac{1}{2} \cdot (1 + h_{2,x}) \\ &\dots \\ h_{v,x} &= 1 + h_{n/2-1,x}, \end{aligned}$$

then $h_{u,x} = (2k^2 + 3)/2k = n + 3/n$.

For the rest of hitting time from $x \neq u$, $x \in K_{n/2}$, there is a vanilla way of computing by

$$\begin{aligned} h_{x,u} &= \frac{1}{k-1} \cdot 1 + \frac{k-2}{k-1} \cdot (1 + h_{x',u}) \\ h_{x',u} &= \frac{1}{k-1} \cdot 1 + \frac{k-2}{k-1} \cdot (1 + h_{x',u}), \end{aligned}$$

then $h_{x,u} = h_{x',u} = k - 1$.

Another way of deriving $h_{x,u}$ is by taking advantage of theorem 1.7.75, that the lollipop graph is non-bipartite as $K_{n/2}$ has odd length cycle, then

$$\pi_{u,u}^{-1} = \frac{2|E|}{d(u)} = h_{u,u} = \frac{1}{d(u)} \sum_{w \in \Gamma(u)} (h_{w,u} + 1) = 1 + \frac{1}{k} \sum_{w \in \Gamma(u)} h_{w,u},$$

where $d(u) = |\Gamma(u)| = k$, hence $\sum_{w \in \Gamma(u)} h_{w,u} = 2|E| - k$. Since $|E| = k(k+1)/2$, and $h_{1,u} = n - 1 = 2k - 1$, hence

$$\sum_{w \in \Gamma(u), w \neq 1} h_{w,u} = 2|E| - k - h_{1,u} = (k-1)^2.$$

Therefore, the covering time from v is $\Theta(n^2)$.

To derive the C_G from u , $h_{u,v}$ is upper bounding the covering time for the path part, then

$$C_G \leq h_{u,v} + \sum_{x \in K_{n/2}, x \neq u} (h_{x,u} + h_{u,x}),$$

and we have the formulas

$$\begin{aligned} h_{x,v} &= \frac{1}{k-1} \cdot (1 + h_{u,x}) + \frac{k-2}{k-1} \cdot (1 + h_{x,v}) \\ h_{u,v} &= \frac{1}{k} \cdot (1 + h_{1,v}) + \frac{k-1}{k} \cdot (1 + h_{x,v}) \\ h_{1,v} &= \frac{1}{2} \cdot (1 + h_{u,v}) + \frac{1}{2} \cdot (1 + h_{2,v}) \\ &\dots \\ h_{n/2-1,v} &= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (1 + h_{n/2-2,v}), \end{aligned}$$

then $h_{k-i,v} = i \cdot (k^2 + k - i)$, and hence $h_{u,v} = k^3$, which explains C_G from u is $\Theta(n^3)$. $\square$

Exercise 1.7.83 (Exercise 7.26 [MU17]). Let n equidistant points be marked on a circle. WLOG, think of the points as being labeled clockwise from 0 to $n - 1$. Initially, a wolf begins at 0 and there is one sheep at each of the remaining $n - 1$ points. The wolf takes a random walk on the circle. For each step, it moves with probability $1/2$ to one neighboring point and with probability $1/2$ to the other neighboring point. At the first visit to a point, the wolf eats a sheep if there is still one there. Which sheep is most likely to be the last eaten?

Proof. Brain teaser. Say for any sheep $j \in [1, n-1]$ is eaten last, then view this as a Gambler's Ruin example 1.7.51 by states $j, j+1 \cdots n-1, 0, 1, \cdots, j-1, j$ and absorb at both j 's. Now the requirement is, visit all states before absorbing.

Visiting $j+1$ first is mutually exclusive with visiting $j-1$ first, and hence the rest is visiting $j+1$ or $j-1$, then absorb to the j on the other end, which is of probability $1/(n-1)$ in total.

Such $1/(n-1)$ probability holds for any $j \in [1, n-1]$, and hence all sheeps eaten last is equal probability. $\square$

Exercise 1.7.84 (Exercise 7.30 [MU17]). Let (u, v) be an edge in a hypercube of $N = 2^n$ nodes.

- Prove the expected time between the traversal of (u, v) edge is nN .
- We consider the time between transitions from u to v in a different way. After moving from u to v , the walk must first return to u . When it returns to u , the walk might next move to v , or it might move to another neighbor of u , in which case it must return to u again before moving to v for there to be a transition from u to v .

Use symmetry and the above description to prove the following recurrence:

$$Nn = \sum_{i \geq 1} \frac{1}{n} \left(\frac{n-1}{n} \right)^{i-1} (i(h_{u,v} + 1)) = n(h_{u,v} + 1).$$

- Conclude that $h_{u,v} = N - 1$.
- Using the result on the hitting time of adjacent vertices and Matthews' theorem, show the cover time is $O(Nn^2)$.
- As a much more challenging problem, try to prove that the maximum hitting time between any two vertices for the random walk on the hypercube is $O(N)$, and that the cover time is correspondingly $O(Nn)$.

Proof. Each vertex in an n -dimensional hypercube has n edges, then $2|E| = nN$. By lemma 1.7.77,

$$h_{u \rightarrow v, u \rightarrow v} = 2|E| = nN.$$

To derive $h_{u \rightarrow v, u \rightarrow v}$ again from geometric distribution, consider u going to v , or $n-1$ other neighbors in $\Gamma(u) \setminus \{v\}$. It takes one step to go from u to any $w \in \Gamma(u)$. To get back to u from $w \in \Gamma(u)$, $h_{w,u} = h_{v,u} = h_{u,v}$ by symmetry.

Hence, a trial contains going back from neighbor $w \in \Gamma(u)$, then depart from u to one of the neighbors in $\Gamma(u)$. By Markovian of random walk, the probability of going from u to v is n^{-1} , hence the expected steps to take (u, v) is

$$\sum_{i \geq 1} \frac{1}{n} \left(\frac{n-1}{n} \right)^{i-1} \cdot i(h_{u,v} + 1) = n(h_{u,v} + 1).$$


By prior result on $h_{u \rightarrow v, u \rightarrow v}$ by lemma 1.7.77, we conclude that $h_{u,v} = N - 1$.

Consider any $p, q \in V$ for the hypercube, that $(p, q) \notin E$. Similar to exercise 1.7.78, we upper bound by

$$h_{p,q} \leq \sum_{(a,b) \in P} h_{a,b} = O(Nn),$$

where there is a path of length $O(n)$ connecting p and q , as there are at most $O(n)$ bits different in between p and q . By lemma 1.7.81, $H(N-1) = O(n)$, hence the covering time is upper bounded by $O(n^2N)$.

For departing from u , we can model the hitting time to v by the Hamming distance as follows.



Then the linear system for the hitting time is

$$\begin{aligned} h_{0,n} &= 1 + h_{1,n} \\ h_{i,n} &= 1 + \frac{i}{n} \cdot h_{i-1,n} + \frac{n-i}{n} \cdot h_{i+1,n} \\ h_{n-1,n} &= 1 + \frac{n-1}{n} \cdot h_{n-2,n}, \end{aligned}$$

and let $d_i = h_{i,n} - h_{i+1,n}$ with $d_{n-1} = h_{n-1,n}$, then

$$\begin{aligned} d_0 &= 1 \\ (n-i)d_i &= n + id_{i-1}. \end{aligned}$$

Since

$$\begin{aligned} (n-i)\binom{n-1}{i-1} &= i\binom{n-1}{i} \\ n\binom{n-1}{i} &= (n-i)\binom{n}{i}, \end{aligned}$$

then

$$\begin{aligned} (n-i)\binom{n-1}{i}d_i &= n\binom{n-1}{i} + i\binom{n-1}{i}d_{i-1} \\ &= (n-i)\binom{n}{i} + (n-i)\binom{n-1}{i-1}d_{i-1}, \end{aligned}$$

and eventually

$$s_i = \binom{n-1}{i}d_i = \binom{n}{i} + \binom{n-1}{i-1}d_{i-1} = \binom{n}{i} + s_{i-1}.$$

Since $s_0 = 1$, hence

$$s_k = \sum_{i \in [0,k]} \binom{n}{i} = 2^n - \sum_{i \in [k+1,n]} \binom{n}{i},$$

and

$$d_k = \binom{n-1}{k}^{-1} \sum_{i \in [0,k]} \binom{n}{i} = \binom{n-1}{k}^{-1} \left(2^n - \sum_{i \in [k+1,n]} \binom{n}{i} \right),$$

and finally

$$\begin{aligned} h_{0,n} &= \sum_{k \in [0,n-1]} d_k \\ &= \sum_{k \in [0,n-1]} \binom{n-1}{k}^{-1} \sum_{i \in [0,k]} \binom{n}{i}. \end{aligned}$$

Since $d_{n-1} = 2^n - 1$, $d_{n-2} \leq 2^n/(n-1)$, and

$$\begin{aligned} \sum_{k \in [2,n-3]} \binom{n-1}{k}^{-1} &\leq \sum_{k \in [2,n-3]} \binom{n-1}{2}^{-1} \\ &= \frac{2(n-4)}{(n-1)(n-2)} = O(n^{-1}). \end{aligned}$$

then $h_{0,n} = O(N)$. By lemma 1.7.81 again, the covering time is $O(nN)$. $\square$

1.7.6 Parrondo's Paradox

Parrondo's paradox shows that, sometimes 2 losing games combined together makes a winning game.

- Consider a game A with a biased coin A , winning with probability p_A and losing with probability $1 - p_A$.
- Consider a game B , that the coin flipped depends on the state of the game.

Let w be the number of games won so far, and ℓ be the number of games lost so far, then $w - \ell$ is the winning.

- If $w - \ell \equiv 0 \pmod{3}$, then use coin B winning with probability p_B and losing with probability $1 - p_B$.
- Otherwise, use coin C winning with probability p_C and losing with probability $1 - p_C$.

To show whether game B is winning or losing, we use the following lemma.

Lemma 1.7.85. *For any sequence s of moves that starts at 0 and ends at 3 before reaching -3 , let $f(s)$ be a 1-to-1 mapping to s such that $f(s)$ negates every number starting from the last 0 in the sequence s , then*

$$\frac{\Pr[s \text{ occurs}]}{\Pr[f(s) \text{ occurs}]} = \frac{p_B p_C^2}{(1 - p_B)(1 - p_C)^2}.$$

Proof. We can view the game B by a Markov chain over $[-3, 3]$ that starts from 0, with ± 3 being 2 absorbing states. Similar to the Gambler's ruin argument in example 1.7.51, if it is more likely to reach -3 before reaching 3, then the game is a losing game.

Now consider the 1-to-1 mapping. Let

- a_1 be the number of moves $0 \rightarrow 1$,
- a_2 be the number of moves from $0 \rightarrow -1$,
- a_3 be the number of moves $1 \rightarrow 2, 2 \rightarrow 3, -2 \rightarrow -1, -1 \rightarrow 0$,
- a_4 be the number of moves $2 \rightarrow 1, 1 \rightarrow 0, -1 \rightarrow -2, -2 \rightarrow -3$

in s . Then

$$\Pr[s \text{ occurs}] = p_B^{a_1} (1 - p_B)^{a_2} p_C^{a_3} (1 - p_C)^{a_4}.$$

Consider $f(s)$ reaching -3 before reaching 3 that is 1-to-1 mapped to s .

- Since the first move after last 0 in s is $0 \rightarrow 1$, then the first move after last 0 in $f(s)$ is $0 \rightarrow -1$, hence $a'_1 = a_1 - 1$. By symmetry of $a_1 + a_2 = a'_1 + a'_2$, we have $a'_2 = a_2 + 1$.
- After the last $0 \rightarrow 1$ in s , there is no $1 \rightarrow 0$ moves. The remaining moves in s are $1 \rightarrow 2, 2 \rightarrow 1$, and $2 \rightarrow 3$. For $f(s)$, the remaining moves are $-1 \rightarrow -2, -2 \rightarrow -1$, and $-2 \rightarrow -3$.

All the moves $-1 \rightarrow -2$ then $-2 \rightarrow -1$ in $f(s)$ are self-cancelling, symmetric to the $1 \rightarrow 2$ then $2 \rightarrow 1$ on s side.

What remains is $1 \rightarrow 2 \rightarrow 3$ in s , and $-1 \rightarrow -2 \rightarrow -3$ in $f(s)$. Hence, $a'_3 = a_3 - 2$, and $a'_4 = a_4 + 2$.

Therefore,

$$\Pr[f(s) \text{ occurs}] = p_B^{a_1-1} (1 - p_B)^{a_2+1} p_C^{a_3-2} (1 - p_C)^{a_4+2},$$

and the ratio $\eta = p_B p_C^2 (1 - p_B)^{-1} (1 - p_C)^{-2}$. □

Suppose $p_A = 0.49, p_B = 0.09, p_C = 0.74$, then both game A and game B are both losing, as $\eta \approx 0.8$.

However, with a fair coin D with $p_D = 0.5$, we combine the game by winning on coin D plays game A , and losing on coin D plays game B . In this case, $p'_B = (p_A + p_B)/2$ and $p'_C = (p_A + p_C)/2$, and eventually $\eta > 1$.

The difference lies in the fair coin D breaks the structure of game B , that is rather likely to lose on any game with $w - \ell \equiv 0 \pmod{3}$. If one managed to get over the barrier in game B with game A , which is close to a fair game, then it is likely to win the next 2 games.

1.8 Continuous Distributions and the Poisson Process

TODO: ...

1.9 The Normal Distribution

TODO: ...

1.10 Entropy, Randomness, and Information

1.10.1 The Entropy Function

Definition 1.10.1 (Discrete Entropy). The entropy in bits of a discrete random variable X is

$$H(X) = - \sum_x \Pr[X = x] \cdot \log_2 \Pr[X = x],$$

where the summation is over the support of X , namely over all values x in the range of X . Equivalently,

$$H(X) = \mathbb{E} \left[\log_2 \frac{1}{\Pr[X]} \right].$$

Definition 1.10.2 (Binary Entropy). The binary entropy function $H(p)$ for a random variable that assumes only two possible outcomes, one of which occurs with probability p , is

$$H(p) = -p \log_2 p - (1-p) \log_2 (1-p).$$

Exercise 1.10.3 (Exercise 10.5 [MU17]). If p is chosen uniformly at random from real interval $[0, 1]$, derive $\mathbb{E}[H(p)]$.

Proof. Since $H(p) = -p \log_2 p - (1-p) \log_2 (1-p)$ is symmetric, hence we derive the part on the left hand side by

$$- \int_0^1 p \log_2 p \, dp = - \frac{1}{\ln 2} \int_0^1 p \ln p \, dp = - \frac{1}{\ln 2} \left(\frac{1}{2} p^2 \ln p \Big|_0^1 - \int_0^1 \frac{1}{2} p^2 \cdot \frac{1}{p} \, dp \right) = \frac{1}{4 \ln 2},$$

where the second equality holds by integration by parts. Therefore, $\mathbb{E}[H(p)] = 1/(2 \ln 2)$. $\square$

Exercise 1.10.4 (Exercise 10.2 [MU17]). Consider an n -sided die, where the i^{th} face comes up with probability p_i . Show that the entropy of a die roll is maximized when each face comes up with equal probability.

Proof. Let X be the random variable for the face showing up for the n -sided die. The probability of the face showing up can be decided by a set of probabilities $\{p_i\}_{i \in I}$ where $|I| = n-1$, such that $k \notin I$ and $I \cup \{k\} = [1, n]$.

On an index $j \in I$, take partial derivative over p_j for $H(X)$,

$$\frac{\partial}{\partial p_j} H(X) = - \frac{\partial}{\partial p_j} (p_j \log_2 p_j) - \frac{\partial}{\partial p_j} (p_k \log_2 p_k) = - \frac{1}{\ln 2} ((\ln p_j + 1) - (\ln p_k + 1)) = \log_2 \frac{p_k}{p_j},$$

where $p_k = 1 - \sum_{i \in I} p_i$. By symmetry, k can be arbitrary over $[1, n]$, hence if the partial derivative is 0 for any pair of $j \neq k$, the only solution would be $p_i = 1/n$ for any $i \in [1, n]$. $\square$

Remark 1.10.5. For a fair n -faced die, the entropy is $\log_2 n$ bits of randomness.

Exercise 1.10.6 (Exercise 10.4 [MU17]). Let $p_k = \left(k \ln^2 k\right)^{-1}$, and show that $S = \sum_{k \geq 2} p_k$ is finite. Let integer-valued discrete random variable X has distribution $\Pr[X = k] = p_k/S$ for $k \geq 2$, show that $H(X)$ is unbounded.

Proof. Let $g(k) = -(\ln k)^{-1}$, then $g'(k) = \left(k \ln^2 k\right)^{-1}$. Hence,

$$S \leq \int_{\frac{3}{2}}^{\infty} \frac{1}{k \ln^2 k} \, dk = \frac{-1}{\ln k} \Big|_{\frac{3}{2}}^{\infty} = \frac{1}{\ln 3 - \ln 2}, \quad (1.9)$$

which is upper bounded by a bounded value. For the entropy, we have

$$H(X) = \sum_{k \geq 2} \frac{\log_2 (Sk \ln^2 k)}{Sk \ln^2 k} = \frac{1}{\ln 2} \sum_{k \geq 2} \frac{\ln (Sk \ln^2 k)}{Sk \ln^2 k} = \frac{1}{\ln 2} \sum_{k \geq 2} \left(\left\lfloor \frac{\ln S}{Sk \ln^2 k} \right\rfloor_1 + \left\lfloor \frac{1}{Sk \ln k} \right\rfloor_2 + \left\lfloor \frac{2 \ln \ln k}{Sk \ln^2 k} \right\rfloor_3 \right).$$

The sum over the first term is bounded by eq. (1.9) or by definition of X . For the sum over the third term, the sum over the first few terms with $2 \leq k < 6$ is bounded. For $k \geq 6$, the sum over the third term is bounded by

$$\sum_{k \geq 6} \frac{2 \ln \ln k}{S k \ln^2 k} \leq \frac{2}{S} \int_5^\infty \frac{\ln \ln k}{k \ln^2 k} dk = \frac{2}{S} \int_{\ln 5}^\infty \frac{\ln u}{u^2} du = \frac{2}{S} \left(\left. \frac{-\ln u}{u} \right|_{\ln 5}^\infty + \int_{\ln 5}^\infty \frac{1}{u^2} du \right) = -\frac{2}{S} \cdot \frac{\ln u + 1}{u} \Big|_{\ln 5}^\infty,$$

where $u = \ln k$, and the summation is finitely bounded. The sum over the second term is

$$\sum_{k \geq 2} \frac{1}{S k \ln k} \geq \frac{1}{S} \int_2^\infty \frac{1}{k \ln k} dk = \frac{1}{S} \int_{\ln 2}^\infty \frac{1}{u} du = \frac{1}{S} \ln u \Big|_{\ln 2}^\infty,$$

which is unbounded. We conclude that $H(X)$ is unbounded. $\square$

Lemma 1.10.7. Let X_1, X_2 be independent random variables, and let $Y = (X_1, X_2)$. Then $H(Y) = H(X_1) + H(X_2)$.

Proof. Let X_1 has support Ω_1 , and X_2 has support Ω_2 . We show by

$$\begin{aligned} H(Y) &= - \sum_{x_1 \in \Omega_1, x_2 \in \Omega_2} \Pr[X_1 = x_1, X_2 = x_2] \cdot \log_2 \Pr[X_1 = x_1, X_2 = x_2] \\ &= - \sum_{x_1 \in \Omega_1, x_2 \in \Omega_2} \Pr[X_1 = x_1] \cdot \Pr[X_2 = x_2] \cdot (\log_2 \Pr[X_1 = x_1] + \log_2 \Pr[X_2 = x_2]) \\ &= - \sum_{x_1 \in \Omega_1, x_2 \in \Omega_2} (\Pr[X_1 = x_1] \cdot \log_2 \Pr[X_1 = x_1] \cdot \Pr[X_2 = x_2] + \Pr[X_2 = x_2] \cdot \log_2 \Pr[X_2 = x_2] \cdot \Pr[X_1 = x_1]) \\ &= - \sum_{x_1 \in \Omega_1} \Pr[X_1 = x_1] \cdot \log_2 \Pr[X_1 = x_1] - \sum_{x_2 \in \Omega_2} \Pr[X_2 = x_2] \cdot \log_2 \Pr[X_2 = x_2] = H(X_1) + H(X_2), \end{aligned}$$

where the second and the last equality holds by the independence between X_1 and X_2 . $\square$

Exercise 1.10.8 (Exercise 10.6 [MU17]). The conditional entropy $H(Y | X)$ is defined by

$$H(Y | X) = - \sum_{x,y} \Pr[X = x, Y = y] \cdot \log_2 \Pr[Y = y | X = x].$$

If $Z = (X, Y)$, show that $H(Z) = H(X) + H(Y | X)$.

Proof. This can be seen as an extended version of lemma 1.10.7. We show by

$$\begin{aligned} H(Z) &= - \sum_{x,y} \Pr[X = x, Y = y] \cdot \log_2 \Pr[X = x, Y = y] \\ &= - \sum_{x,y} \Pr[X = x, Y = y] \cdot (\log_2 \Pr[X = x] + \log_2 \Pr[Y = y | X = x]) \\ &= - \sum_{x,y} (\Pr[X = x] \cdot \log_2 \Pr[X = x] + \Pr[X = x, Y = y] \cdot \log_2 \Pr[Y = y | X = x]) \\ &= H(X) + H(Y | X), \end{aligned}$$

where the third equality holds by summing $\Pr[X = x, Y = y]$ over y preserves the marginal probability of X . $\square$

Corollary 1.10.9 (Chain Rule of Conditional Entropy).

$$H(X_1, \dots, X_n) = \sum_{i \in [1, n]} H(X_i | X_1, \dots, X_{i-1}).$$

Corollary 1.10.10 (Bayes' Rule of Conditional Entropy).

$$H(Y | X) = H(X | Y) + H(Y) - H(X).$$

Definition 1.10.11 (Mutual Information). The mutual information of a pair of random variables (X, Y) is defined by

$$I(X; Y) = H(X) - H(X | Y) = H(Y) - H(Y | X).$$

Remark 1.10.12. For mutual information in definition 1.10.11, we can expand the expression by

$$\begin{aligned} H(Y) - H(Y | X) &= - \sum_y \Pr[Y = y] \cdot \log_2 \Pr[Y = y] + \sum_{x,y} \Pr[Y = y, X = x] \cdot \log_2 \Pr[Y = y | X = x] \\ &= - \sum_{x,y} \Pr[Y = y, X = x] \cdot \log_2 \Pr[Y = y] + \sum_{x,y} \Pr[Y = y, X = x] \cdot \log_2 \Pr[Y = y | X = x] \\ &= - \sum_{x,y} \Pr[Y = y, X = x] \cdot (\log_2 \Pr[Y = y] - \log_2 \Pr[Y = y | X = x]) \\ &= - \sum_{x,y} \Pr[Y = y, X = x] \cdot (\log_2 \Pr[Y = y] + \log_2 \Pr[X = x] - \log_2 \Pr[Y = y, X = x]) \\ &= \sum_{x,y} \Pr[Y = y, X = x] \cdot \log_2 \frac{\Pr[Y = y, X = x]}{\Pr[Y = y] \cdot \Pr[X = x]}. \end{aligned}$$

Definition 1.10.13 (KL Divergence [KL51]). The Kullback-Leibler (KL) divergence is a statistical distance defined as

$$D_{\text{KL}}(P \parallel Q) = \sum_x P(x) \cdot \log_2 \frac{P(x)}{Q(x)}.$$

Remark 1.10.14 (Mutual Information as KL Divergence). In KL divergence by definition 1.10.13 and mutual information with remark 1.10.12, if we consider $\mathcal{D}_{(X,Y)}$, that is the joint distribution of (X, Y) , and $\mathcal{D}_X \times \mathcal{D}_Y$, which is the marginal distribution of (X, Y) , mutual information fits into KL divergence by

$$D_{\text{KL}}(\mathcal{D}_{(X,Y)} \parallel \mathcal{D}_X \times \mathcal{D}_Y) = I(X; Y).$$

Moreover, we can again expand the expression in remark 1.10.12 by

$$\begin{aligned} H(Y) - H(Y | X) &= \sum_{x,y} \Pr[Y = y, X = x] \cdot \log_2 \frac{\Pr[Y = y, X = x]}{\Pr[Y = y] \cdot \Pr[X = x]} \\ &= \sum_{x,y} \Pr[Y = y | X = x] \cdot \Pr[X = x] \cdot \log_2 \frac{\Pr[Y = y | X = x]}{\Pr[Y = y]} \\ &= \sum_x \Pr[X = x] \cdot \left(\sum_y \Pr[Y = y | X = x] \cdot \log_2 \frac{\Pr[Y = y | X = x]}{\Pr[Y = y]} \right) \\ &= \mathbb{E}[D_{\text{KL}}(\mathcal{D}_{Y|X} \parallel \mathcal{D}_Y)], \end{aligned}$$

where $\mathcal{D}_{Y|X}$ is conditional distribution of Y given X with distribution $\mathcal{D}_X$.

Lemma 1.10.15 (Gibbs' Inequality). Given 2 distributions P and Q over the same support,

$$D_{\text{KL}}(P \parallel Q) \geq 0.$$

Proof. By Jensen's inequality theorem 1.2.5,

$$\sum_{x \in \Omega} P(x) \cdot \log_2 \frac{Q(x)}{P(x)} \leq \log_2 \left(\sum_{x \in \Omega} P(x) \cdot \frac{Q(x)}{P(x)} \right) = 0.$$

□

Corollary 1.10.16 (Monotonicity of Conditional Entropy).

$$H(Y | X) \leq H(Y).$$

1.10.2 Entropy and Binomial Coefficients

For lemma 1.7.1, we give a tighter version as follows, to be used soon.

Lemma 1.10.17 (Tighter Stirling's Formula).

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n \leq n! \leq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n \cdot \exp\left(\frac{1}{12n}\right).$$

Lemma 1.10.18. For an integer $nq \in [0, n]$, we have

$$\frac{2^{nH(q)}}{2\sqrt{n}} \leq \binom{n}{nq} \leq 2^{nH(q)}.$$

Proof. We begin with the upper bound. Since

$$1 = (q + (1 - q))^n \geq \binom{n}{nq} q^{nq} (1 - q)^{n(1-q)},$$

then

$$\log_2 \binom{n}{nq} \leq -nq \log_2 q - n(1 - q) \log_2 (1 - q) = nH(q).$$

The lower bound is trivial when $n = 1$, and it is also trivial when $q = 0$ or 1 for $n \geq 2$. We derive the lower bound for $n \geq 2$ and $q \in (0, 1)$ using the following upper and lower bounds by lemma 1.10.17:

$$\begin{aligned} n! &\geq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n \\ (nq)! &\leq \sqrt{2\pi nq} \cdot \left(\frac{nq}{e}\right)^{nq} \cdot \exp\left(\frac{1}{12nq}\right) \\ (n(1 - q))! &\leq \sqrt{2\pi n(1 - q)} \cdot \left(\frac{n(1 - q)}{e}\right)^{n(1-q)} \cdot \exp\left(\frac{1}{12n(1 - q)}\right), \end{aligned}$$

and therefore

$$\begin{aligned} \binom{n}{nq} &\geq \frac{1}{\sqrt{2\pi nq(1 - q)}} \cdot \exp\left(-\frac{1}{12nq(1 - q)}\right) \cdot \left(\frac{1}{q}\right)^{nq} \cdot \left(\frac{1}{1 - q}\right)^{n(1-q)} \\ &= \frac{2^{nH(q)}}{\sqrt{2\pi nq(1 - q)}} \cdot \exp\left(-\frac{1}{12nq(1 - q)}\right). \end{aligned} \tag{1.10}$$

Let $k = q(1 - q)$ such that $0 < k \leq 1/4$, then

$$\frac{d}{dk} \left(\frac{1}{\sqrt{k}} \cdot \exp\left(-\frac{1}{12nk}\right) \right) = -\frac{1}{2k^{3/2}} \cdot \exp\left(-\frac{1}{12nk}\right) + \frac{1}{\sqrt{k}} \cdot \exp\left(-\frac{1}{12nk}\right) \cdot \frac{1}{12nk^2} = \frac{1}{k^{3/2}} \cdot \exp\left(-\frac{1}{12nk}\right) \cdot \left(\frac{1}{12nk} - \frac{1}{2}\right),$$

which is negative for $k \geq (6n)^{-1}$, and $k \geq (n - 1)/n^2 \geq (6n)^{-1}$ when $q \in (0, 1)$, so eq. (1.10) is minimized when $k = 1/4$, giving

$$\binom{n}{nq} \geq \frac{2^{nH(q)}}{\sqrt{\pi n/2}} \cdot \exp\left(-\frac{1}{3n}\right).$$

Moreover,

$$\left(\frac{2\sqrt{n}}{\sqrt{\pi n/2}}\right) \cdot \exp\left(-\frac{1}{3n}\right) = \sqrt{\frac{8}{\pi}} \cdot \exp\left(-\frac{1}{3n}\right) > 1$$

for all $n \geq 1$, then

$$\binom{n}{nq} \geq \frac{2^{nH(q)}}{2\sqrt{n}}.$$

□

Corollary 1.10.19. When $0 \leq q \leq 1/2$,

$$\frac{2^{nH(q)}}{2\sqrt{n}} \leq \binom{n}{\lceil nq \rceil}, \quad \binom{n}{\lfloor nq \rfloor} \leq 2^{nH(q)}.$$

When $1/2 \leq q \leq 1$,

$$\frac{2^{nH(q)}}{2\sqrt{n}} \leq \binom{n}{\lfloor nq \rfloor}, \quad \binom{n}{\lceil nq \rceil} \leq 2^{nH(q)}.$$

1.10.3 Entropy: A Measure of Randomness

Definition 1.10.20 (Variable-Length Fair-Bit Extraction Function). Let $|y|$ be the number of bits in a sequence of bits y . An extraction function Ext takes as input the value of a random variable X , and outputs a sequence of bits. For every sequence of bits y with $|y| = k$,

$$\Pr[\text{Ext}(X) = y \mid |\text{Ext}(X)| = k] = 2^{-k},$$

whenever $\Pr[|\text{Ext}(X)| = k] > 0$.

Remark 1.10.21. We can view the input X as the outcomes of n biased coin flips, and the output y as k fair coin flips. In this sense, if the biased coin comes up heads with probability p , the entropy of the input is $nH(p)$, which is the maximum possible expected number of fair bits that any extraction function can output.

There is no guarantee that Ext is efficient (runs in polynomial time).

Example 1.10.22. Consider a construction of Ext for X uniformly distributed over $[0, 11]$. The elements in $[0, 7]$ can be represented in 3 bits, while we can associate the 4 probabilities in $[8, 11]$ with 2 bits. Associate each input $[0, 11]$ to each 2 or 3 bit output, such that they output with the same probability. Hence, the probability of outputting 3 bits is $2/3$, while the probability of outputting 2 bits is $1/3$, and

$$\Pr[\text{Ext}(X) = y \mid |\text{Ext}(X)| = k] = 2^{-k}$$

for k being 2 or 3, such that definition 1.10.20 is satisfied.

Theorem 1.10.23. Suppose the value of a random variable X is uniform over $[0, m-1]$, such that $H(X) = \log_2 m$, then there is an extraction function for X that outputs on average at least $\lfloor \log_2 m \rfloor - 1 = \lfloor H(X) \rfloor - 1$ independent and unbiased bits.

Proof. Suppose m is a power of 2, the extraction function trivially outputs the value's $\lfloor \log_2 m \rfloor$ bits representation.

In the general case where m is not a power of 2, we first show in an inductive way. Let $\alpha = \lfloor \log_2 m \rfloor$, Y be the random variable for the length of output bits, and for the elements in $[2^\alpha, m-1]$, the theorem 1.10.23 still holds, that

$$\mathbb{E}[Y] \geq \frac{2^\alpha}{m} \alpha + \frac{m-2^\alpha}{m} (\lfloor \log_2 (m-2^\alpha) \rfloor - 1) = \alpha - \frac{m-2^\alpha}{m} (\alpha - \lfloor \log_2 (m-2^\alpha) \rfloor + 1).$$

Let $\beta = \lfloor \log_2 (m-2^\alpha) \rfloor$, then since

$$\frac{2^{\beta+1}}{2^{\beta+1} + 2^\alpha} \geq \frac{m-2^\alpha}{m} \geq \frac{2^\beta}{2^\beta + 2^\alpha},$$

we have

$$\mathbb{E}[Y] \geq \alpha - \frac{m-2^\alpha}{m} (\alpha - \beta + 1) \geq \alpha - \frac{2^{\beta+1}}{2^{\beta+1} + 2^\alpha} (\alpha - \beta + 1) = \alpha - \frac{1}{1 + 2^{\alpha-\beta-1}} (\alpha - \beta + 1) \geq \alpha - 1,$$

as $2^{\alpha-\beta-1} \geq \alpha - \beta$ holds for any $\alpha > \beta$.

We show in an explicit construction. Consider k distinct indices $\{\beta_i\}_{i \in [1, k]} \subseteq [0, \alpha]$ in descending order, satisfying

$$m = \sum_{i \in [1, k]} 2^{\beta_i}, \quad s_i = \sum_{j \in [1, i]} 2^{\beta_j},$$

such that $[0, m-1]$ is partitioned into k subsets, and the i^{th} subset is $[s_{i-1}, s_i - 1]$, of size 2^{β_i} , then

$$\mathbb{E}[Y] = \sum_{i \in [1, k]} \frac{2^{\beta_i}}{m} \beta_i,$$

and since $\beta_1 = \alpha$,

$$\mathbb{E}[Y] - (\alpha - 1) = \frac{2^{\beta_1}}{m} + \sum_{i \in [2, k]} \frac{2^{\beta_i}}{m} (\beta_i - (\alpha - 1)) = \frac{1}{m} \left(2^{\beta_1} - \sum_{i \in [2, k]} 2^{\beta_i} ((\alpha - 1) - \beta_i) \right).$$

Considering

$$2^{\beta_1} - \sum_{i \in [2, k]} 2^{\beta_i} ((\alpha - 1) - \beta_i) \geq 2^\alpha - \sum_{i \in [0, \alpha-2]} 2^i \cdot (\alpha - 1 - i) = 2^\alpha - ((2^\alpha - 1) - \alpha) = \alpha + 1 > 0,$$

then $\mathbb{E}[Y] > \alpha - 1$. □

Exercise 1.10.24 (Exercise 10.9 [MU17]). We have shown in theorem 1.10.23 that an extraction function for a random variable X uniform over $[0, m-1]$ extracts on average at least $\lfloor H(X) \rfloor - 1$ unbiased bits. Show that there is an extraction function for $\{X_i\}_{i \in [1, k]}$ independent random variables uniform over $[0, m-1]$ extracting on average at least $k \lfloor H(X) \rfloor - 1$ unbiased bits.

Proof. Apparently, one can invoke k times the extraction function in theorem 1.10.23, and by linearity of expectation, outputting on average at least $k \lfloor H(X) \rfloor - k$ unbiased bits. The idea is to construct W over $[0, m^k - 1]$ by

$$W = X_1 + mX_2 + m^2X_3 + \dots + m^{k-1}X_k,$$

and $H(W) = H(X_1, \dots, X_k) = k \log_2 m$. The average number of output bits by theorem 1.10.23 is

$$\lfloor k \log_2 m \rfloor - 1 \geq k \lfloor \log_2 m \rfloor - 1.$$

□

Exercise 1.10.25 (Exercise 10.10 [MU17]). Suppose we have an oracle of generating independent fair coin flips, and m is not a power of 2.

- Show that one cannot uniformly sample over $[0, m-1]$ with exactly k fair coin flips, for any fixed k .
- Show that one cannot uniformly sample over $[0, m-1]$ with at most k fair coin flips, for any fixed k .
- Give an algorithm using fair coins to sample uniformly over $[0, m-1]$, with expected flips at most $2 \lceil \log_2 m \rceil$.

Proof. We first rule out the case that $2^k < m$, as it is impossible to associate 2^k outcomes with m numbers.

If one tries to uniformly sample over $[0, m-1]$ with exactly k fair coins, $m \nmid 2^k$ as m is not a power of 2, and one cannot associate m numbers with 2^k outcomes in an equal-sized partition manner.

In the at most k flips case, pad each execution to exactly k flips: if the algorithm stops after ℓ flips with $\ell < k$, then all $2^{k-\ell}$ possible continuations are counted as producing the same outcome. Hence each output a corresponds to some integer number N_a of the 2^k length- k bit strings, so

$$\Pr[\text{outputting } a] = N_a \cdot 2^{-k}.$$

Uniformity requires $N_a/2^k = 1/m$ for every a , so $m \mid 2^k$, which is impossible because m is not a power of 2, failing in the same manner as the exact k tosses.

An algorithm sampling uniformly over $[0, m-1]$ is rejection sampling: keep tossing $\lceil \log_2 m \rceil$ fair coins, until the outputting integer is within range $[0, m-1]$. A round being rejected has probability at most $1/2$, and the expected number of rounds is therefore at most 2, thus the expected runtime is at most $2 \lceil \log_2 m \rceil$. □

Exercise 1.10.26 (Exercise 10.11 [MU17]). Suppose we have an oracle of generating independent fair coin flips.

- Give an algorithm using the fair coin that simulates flipping a biased coin that comes up heads with probability p . The expected number of flips should be at most 2.
- Give an algorithm using the coin to sample uniformly over $[0, m - 1]$ with expected flips at most $\lceil \log_2 m \rceil + 2$.

Proof. Write p into its binary form, such that there exists bits $\{b_i\}_{i \geq 1}$, to binary decompose p into the form $\sum_{i \geq 1} b_i / 2^i$. Keep tossing independent fair coins $\{X_i\}_{i \geq 1}$ until $X_i \neq b_i$. If $X_i > b_i$, the biased coin comes up tails, otherwise heads. Since $\Pr[X_i \neq b_i] = 1/2$, the expected number of sampling is at most 2, as there might be finite number of $\{b_i\}_{i \geq 1}$.

A first attempt at sampling uniformly over $[0, m - 1]$ is an extension of the biased coin sampling, with probability $m/2^\alpha$ the biased coin comes up heads, where $\alpha = \lceil \log_2 m \rceil$. In each attempt, fair coins are tossed sequentially, and they are used both to simulate the biased coin, and maintained as the prefix of an α -bit integer.

- If the biased coin comes up tails, then any completion of the current prefix ties to an integer at least m , so restart.
- If the biased coin comes up heads, the generated number is within $[0, m - 1]$; keep tossing until α fair coins have been sampled, and output the resulting α -bit integer.

Let $A_i = (Y_i, Z_i)$ be the random variable for the i^{th} biased coin simulation, where Y_i is the number of fair coins tossed, and Z_i is the outcome of the biased coin simulation. Let T be a stopping time for $A_1, A_2, \dots$ on the first heads of the biased coin. Then the total number of fair coin flips is

$$W = \sum_{i \in [1, T]} Y_i + (\alpha - Y_T) = \sum_{i \in [1, T-1]} Y_i + \alpha.$$

Since Z_i is heads with probability $m/2^\alpha$, then $\mathbb{E}[T] = 2^\alpha/m \leq 2$. Moreover, for each Y_i ,

$$\mathbb{E}[Y_i] = \sum_{k \in [1, \alpha]} \frac{k}{2^k} + \frac{\alpha}{2^\alpha} = 2 - \frac{2}{2^\alpha}.$$

Since $\mathbb{E}[T]$ and $\mathbb{E}[Y_i]$ are both bounded, by theorem 1.13.18,

$$\mathbb{E}\left[\sum_{i \in [1, T]} Y_i\right] = \mathbb{E}[Y_i] \cdot \mathbb{E}[T] = \frac{2^\alpha}{m} \cdot \left(2 - \frac{2}{2^\alpha}\right).$$

By linearity of expectation,

$$\mathbb{E}[W] = \frac{2^\alpha}{m} \cdot \left(2 - \frac{2}{2^\alpha}\right) + \alpha - \mathbb{E}[Y_T] \leq \alpha + 2 \cdot \frac{2^\alpha - 1}{m} - 1 \leq \alpha + 3$$

where the second inequality holds by $\mathbb{E}[Y_T] \geq 1$, and the last inequality holds by $2^{\alpha-1} < m \leq 2^\alpha$.

This falls short of the desired $\alpha + 2$ bound, because rejected attempts discard all previously sampled bits. We now give a recycling version that keeps the remaining randomness after rejection. Consider the following algorithm:

- Let $c \leftarrow 0$ and $v \leftarrow 1$. Keep looping:
 - Sample $b \stackrel{\text{R}}{\leftarrow} \{0, 1\}$, and update $c \leftarrow c \times 2 + b$, $v \leftarrow v \times 2$.
 - Continue looping if $v < m$, otherwise if $v \geq m$:
 - * If $c \geq m$, update $c \leftarrow c - m$, $v \leftarrow v - m$.
 - * Otherwise, output c .

The invariant here is c is always uniform over $[0, v - 1]$, and the algorithm is rejection sampling over $[0, m - 1]$.

Indeed, if c is uniform over $[0, v - 1]$, then after sampling $b \xleftarrow{R} \{0, 1\}$, the value $2c + b$ is uniform over $[0, 2v - 1]$. On $v \geq m$, conditioning on $c < m$ gives a uniform output over $[0, m - 1]$; conditioning on $c \geq m$ makes $c - m$ uniform over $[0, v - m - 1]$, preserving the invariant.

The new algorithm always runs α steps up to the first comparison against m , then let T be the step for the output. We have the following observation: At the t^{th} step, we have 2^t possible bit strings, separated into 2 classes:

- Resolved prefixes: On reading the prefix of a t -bit string, the algorithm stops at some bit and outputs a sample. Any continuation for a resolved prefix is considered stopped.
- Unresolved prefixes: The algorithm does not output after consuming all of the t bits as inputs.

Lemma 1.10.27. *After the first α bits, the unresolved prefixes at the t^{th} step form a possibly empty interval $[\lfloor 2^t / m \rfloor m, 2^t - 1]$.*

Proof. We prove by induction. Let $q_t = \lfloor 2^t / m \rfloor$. When $t = \alpha$, the unresolved prefixes are $[m, 2^\alpha - 1]$, since $q_\alpha = 1$.

Suppose the claim holds for $t = t'$. After tossing one more fair coin, each unresolved prefix has two continuations, so before the next rejection step the unresolved interval doubles to

$$[2q_{t'}m, 2^{t'+1} - 1].$$

Once the algorithm sees an entire block of size m in the interval at the bottom of the interval, the block is removed, as that goes to the new resolved prefixes. Hence the unresolved interval that remains is

$$[q_{t'+1}m, 2^{t'+1} - 1].$$

This proves the induction step. □

With lemma 1.10.27, we upper bound the number of unresolved prefixes at the t^{th} step v_t by m , and thus

$$\begin{aligned} \mathbb{E}[T] &= \sum_{t \geq 0} \Pr[T > t] = \alpha + \sum_{t \geq \alpha} \Pr[T > t] \\ &= \alpha + \Pr[T > \alpha] + \sum_{t > \alpha} \Pr[T > t] \\ &\leq \alpha + \frac{2^\alpha - m}{2^\alpha} + \sum_{t \geq \alpha+1} \frac{m}{2^t} = \alpha + 1. \end{aligned}$$

The last equality can be roughly relaxed to an upper bound of $\alpha + 2$, as $m/2^t \leq 1/2$ for $t > \alpha$, and the sum of the last 2 terms is upper bounded by 2. □

Theorem 1.10.28. *Suppose a coin comes up heads with probability $p > 1/2$. For any constant $\delta > 0$ and sufficiently large n ,*

- *There exists an extraction function that on input sequences of n independent flips, outputs an average of at least $(1 - \delta)nH(p)$ independent random bits.*
- *The average number of bits output by any extraction function on an input of n independent flips is at most $nH(p)$.*

Proof. We begin by describing an extraction function that extracts on average at least $(1 - \delta)nH(p)$ random bits from n flips of the biased coin. After n flips, the number of heads coming up concentrates around np , which means the outcome is most likely to be one of the $\binom{n}{np}$ sequences with np heads up.

Conditioned that k heads coming up, the sequence with k heads is uniformly distributed over all $\binom{n}{k}$ sequences, which can be bijectively mapped to $[0, \binom{n}{k} - 1]$. We apply the randomness extraction function for uniform random variable in theorem 1.10.23 to extract on average $\lfloor \log_2 \binom{n}{k} \rfloor - 1$ fair bits. Moreover, by lemma 1.10.18, we bound the number of fair bits within the range of $[(1 - \delta)nH(q), nH(q)]$, where $q = k/n$.

We now formalize the derivation of the bounds for the number of fair bits extracted from n flips of biased coins with probability p coming up heads. Let Z be a random variable representing the number of heads flipped, and let B be the random variable representing the bits extracted by the extraction function, then by conditional expectation,

$$\mathbb{E}[B] = \sum_{x \in [0, n]} \Pr[Z = x] \cdot \mathbb{E}[B \mid Z = x],$$

and by theorem 1.10.23,

$$\mathbb{E}[B \mid Z = x] \geq \left\lfloor \log_2 \binom{n}{x} \right\rfloor - 1.$$

Now consider the fact that the number of heads concentrates around np , to lower bound the average number of fair bits extracted, we consider only $x \in [(p - \varepsilon)n, \lfloor (p + \varepsilon)n \rfloor]$, where $\varepsilon \leq \min(p - 1/2, 1 - p)$. By monotonicity,

$$\mathbb{E}[B \mid Z = x] \geq \left\lfloor \log_2 \binom{n}{x} \right\rfloor - 1 \geq \left\lfloor \log_2 \binom{n}{\lfloor n(p + \varepsilon) \rfloor} \right\rfloor - 1 = \mathbb{E}[B \mid Z = \lfloor n(p + \varepsilon) \rfloor]. \quad (1.11)$$

By standard Chernoff's bound, let $\gamma = \varepsilon/p$ and $\mu = np$,

$$\Pr[|Z - np| \geq \varepsilon n] \leq \left(\frac{e^{-\gamma}}{(1 - \gamma)^{1-\gamma}} \right)^\mu + \left(\frac{e^\gamma}{(1 + \gamma)^{1+\gamma}} \right)^\mu \leq \exp\left(-\frac{n\varepsilon^2}{3p}\right) + \exp\left(-\frac{n\varepsilon^2}{2p}\right).$$

Alternatively, by KL Divergence driven Chernoff's bound in exercise 1.4.4,

$$\Pr[|Z - np| \geq \varepsilon n] \leq 2 \exp(-2n\varepsilon^2).$$

Moreover, we can further lower bound eq. (1.11) by theorem 1.10.23, that

$$\begin{aligned} \mathbb{E}[B \mid Z = \lfloor n(p + \varepsilon) \rfloor] &\geq \left\lfloor \log_2 \binom{n}{\lfloor (p + \varepsilon)n \rfloor} \right\rfloor - 1 \geq \log_2 \binom{n}{\lfloor (p + \varepsilon)n \rfloor} - 2 \\ &\geq \log_2 \frac{2^{nH(p+\varepsilon)}}{2\sqrt{n}} - 2 = nH(p + \varepsilon) - \frac{1}{2} \log_2 n - 3. \end{aligned} \quad (1.12)$$

Therefore,

$$\begin{aligned} \mathbb{E}[B] &\geq \sum_{x \in [\lceil (p-\varepsilon)n \rceil, \lfloor (p+\varepsilon)n \rfloor]} \Pr[Z = x] \cdot \mathbb{E}[B \mid Z = x] \\ &\geq \Pr[|Z - np| \leq \varepsilon n] \cdot \mathbb{E}[B \mid Z = \lfloor (p + \varepsilon)n \rfloor] \\ &\geq (1 - 2 \exp(-2n\varepsilon^2)) \cdot \left(nH(p + \varepsilon) - \frac{1}{2} \log_2 n - 3 \right), \end{aligned}$$

where the last inequality holds by eq. (1.12). We conclude that, for any $\delta > 0$, we have

$$\mathbb{E}[B] \geq (1 - \delta) nH(p)$$

for sufficiently small ε and sufficiently large n .

We now show no extraction function on average obtains more than $nH(p)$ fair bits. Suppose input bits $\mathbf{x}$ occurs with probability q , and the extraction function outputs $\mathbf{y}$ corresponding to input $\mathbf{x}$, then by definition 1.10.20,

$$\Pr[\text{Ext}(X) = \mathbf{y} \mid |\text{Ext}(X)| = k] = 2^{-k},$$

that conditioned the output number of bits from the extraction function is k , the probability of outputting $\mathbf{y}$ is 2^{-k} . Hence, for a fixed $\mathbf{y}$ with k bits, the probability of outputting $\mathbf{y}$ is immediate by the probability of outputting k bits,

$$\Pr[\text{Ext}(X) = \mathbf{y}] = \Pr[\text{Ext}(X) = \mathbf{y}, |\text{Ext}(X)| = k] = 2^{-k} \cdot \Pr[|\text{Ext}(X)| = k] \leq 2^{-k}.$$

Moreover, since

$$\Pr[\text{Ext}(X) = \mathbf{y}] \geq \Pr[X = \mathbf{x}] = q,$$

then $q \leq 2^{-k}$, giving an upper bound on $k \leq -\log_2 q$. Therefore,

$$\mathbb{E}[B] = \sum_{\mathbf{x}} \Pr[X = \mathbf{x}] \cdot \mathbb{E}[B \mid X = \mathbf{x}] \leq - \sum_{\mathbf{x}} \Pr[X = \mathbf{x}] \cdot \log_2 \Pr[X = \mathbf{x}] = H(X) = nH(p),$$

where the second inequality is by upper bounding the output length, and the last equality holds by lemma 1.10.7. $\square$

Exercise 1.10.29 (Exercise 10.12 [MU17]). Consider the extraction function $\mathcal{A}$ whose input is a sequence of n independent flips of a coin that comes up heads with probability $p > 1/2$. Break the sequence $X_1, \dots, X_n$ into $\lfloor n/2 \rfloor$ pairs, such that $A_i = (X_{2i-1}, X_{2i})$, and consider the pairs in order. If A_i is heads and tails, it outputs a 0; if A_i is tails and heads, it outputs a 1; otherwise, move up to the next pair.

- Show that the bits extracted by $\mathcal{A}$ are independent and unbiased, and the expected number of bits extracted is

$$\lfloor n/2 \rfloor 2p(1-p) \approx np(1-p).$$

- We derive another set of flips $Y_1, \dots$ from $X_1, \dots, X_n$: Let j, k be 1, and repeat until $j > \lfloor n/2 \rfloor$: If A_j is both heads, Y_k is heads and increment j and k ; if A_j is both tails, Y_k is tails and increment j and k ; otherwise increment j .

The intuition here is to take some randomness $\mathcal{A}$ cannot use effectively and reuse it. Show that the bits produced by running $\mathcal{A}$ over $Y_1, \dots$ are independent and unbiased, and further argue that they are independent of those produced by running $\mathcal{A}$ over $X_1, \dots, X_n$.

- We derive another set of flips $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ from $X_1, \dots, X_n$: Z_i is heads if A_i is both heads or tails, otherwise tails. Show that the bits produced by running $\mathcal{A}$ over $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ are independent and unbiased, and further argue that they are independent of those produced by running $\mathcal{A}$ over $X_1, \dots, X_n$, and $Y_1, \dots$.
- After we run $\mathcal{A}$ over $Y_1, \dots$ and $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$, we can recursively derive two further sequences from each of the sequences, namely $Y_1, \dots$ and $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$, in the same way, run $\mathcal{A}$ on those, and so on.

Let $A(p)$ be the average number of bits extracted for each flip in the sequence $X_1, \dots, X_n$, in the limit as the length of the sequence $X_1, \dots, X_n$ goes to infinity. Let $q = 1 - p$, and argue that $A(p)$ satisfies the recurrence

$$A(p) = pq + \frac{p^2 + q^2}{2} A\left(\frac{p^2}{p^2 + q^2}\right) + \frac{1}{2} A(p^2 + q^2).$$

- Show that $H(p)$ satisfies the same recurrence as $A(p)$.

Proof. First, for the bits directly extracted by $\mathcal{A}$, since the A_i are mutually independent, and for $A_i = (X_{2i-1}, X_{2i})$, the biased coin tosses are mutually independent, each output bit by $\mathcal{A}$ is independent of the others. The probability of a bit sampled by $\mathcal{A}$ being 1 is the same as the probability of a bit sampled by $\mathcal{A}$ being 0; both are $p(1-p)$, requiring a head and a tail in A_i , so the bits are unbiased. Since a bit is output with probability $2p(1-p)$ by $\mathcal{A}$, by linearity of expectation, the expected number of bits extracted is $\lfloor n/2 \rfloor 2p(1-p)$.

To show the bits produced by $\mathcal{A}$ over $Y_1, \dots$ are independent and unbiased, we first show $Y_1, \dots$ are independent. Since $X_1, \dots, X_n$ are independent, then $A_1, \dots, A_{\lfloor n/2 \rfloor}$ are independent. Since the retained A_i are both heads or both tails, then $Y_1, \dots$ are independent, so are the output by $\mathcal{A}$ over $Y_1, \dots$. To show the outputs by $\mathcal{A}$ from $Y_1, \dots$ are unbiased, since each Y_i can be modeled by a biased coin with probability $p^2/(p^2 + q^2)$ coming up heads, applying the same argument for $\mathcal{A}$ producing unbiased bits over $X_1, \dots, X_n$, the outputs by $\mathcal{A}$ over $Y_1, \dots$ are unbiased.

To show the bits produced by $\mathcal{A}$ over $Y_1, \dots$ are independent of the bits produced by $\mathcal{A}$ over $X_1, \dots, X_n$, observe that the former uses only the same-faced A_i , while the latter uses only the distinct-faced A_i . In the latter case, HT and TH are equally likely, and the corresponding A_i are independent across index i . In the former case, the same-faced

A_i are independent across index i , and the values of these A_i are independent of the orders of the distinct-faced A_i . Hence, the bits extracted by $\mathcal{A}$ over $Y_1, \dots$ are independent of the bits extracted by $\mathcal{A}$ over $X_1, \dots, X_n$.

To show the bits produced by $\mathcal{A}$ over $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ are independent and unbiased, we show $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ are independent of each other, as $A_1, \dots, A_{\lfloor n/2 \rfloor}$ are independent. Each Z_i can be modeled by a biased coin coming up heads with probability $p^2 + q^2$, then the bits extracted by $\mathcal{A}$ are independent and unbiased.

To show the bits produced by $\mathcal{A}$ over $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ are independent of the bits from $\mathcal{A}$ over $X_1, \dots, X_n$, for a bit produced by $\mathcal{A}$ over $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$, supposing it is extracted through (Z_i, Z_{i+1}) , which are determined by (A_i, A_{i+1}) , then it is possible that the bit overlaps with the bits extracted over A_i and/or A_{i+1} . Since the A_i are independent of each other, then all the other bits extracted over $X_1, \dots, X_n$ are independent of the bit extracted over (Z_i, Z_{i+1}) . Conditioned that the bit extracted over (Z_i, Z_{i+1}) overlaps with the bit extracted from A_i , then A_i can be HT or TH, the bit extracted from A_i can be 0 or 1, and Z_i is tails regardless, and thus the bit extracted over (Z_i, Z_{i+1}) is 1, which is independent of the bit extracted from A_i . Symmetrically, if the bit extracted over (Z_i, Z_{i+1}) overlaps with the bit extracted over A_{i+1} , the same argument applies, that the bit extracted over (Z_i, Z_{i+1}) is 0.

To show the bits produced by $\mathcal{A}$ over $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$ are independent of the bits from $\mathcal{A}$ over $Y_1, \dots$, the strategy is similar to the last independence argument. Supposing a bit is extracted through (Z_i, Z_{i+1}) , determined by (A_i, A_{i+1}) , then by A_i being independent of each other, it suffices to only argue that the bit is independent of the bit extracted from $Y_1, \dots$ using a Y_j that comes from A_i or A_{i+1} . Supposing A_j is overlapped, regardless of A_j being HH or TT, Z_j is heads, and the bit extracted from (Z_j, Z_{j+1}) is 0, independent of the bit extracted using Y_j . Symmetrically, if the overlap is with the second coordinate, then that coordinate is heads, and the bit extracted from the Z pair is 1.

For $Y_1, \dots$, the biased coin comes up heads with probability $p^2/(p^2 + q^2)$, and the expected number of retained Y_i is $(p^2 + q^2)n/2$. For $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$, the biased coin has probability $p^2 + q^2$ of coming up heads. For $\mathcal{A}$ directly extracting from $X_1, \dots, X_n$, the expected number of bits extracted is $\lfloor n/2 \rfloor 2pq \approx npq$. Thus, when $n \rightarrow \infty$, the average number of bits extracted per flip approaches

$$pq + \frac{p^2 + q^2}{2} A\left(\frac{p^2}{p^2 + q^2}\right) + \frac{1}{2} A(p^2 + q^2).$$

To show that $A(p) = H(p)$, we have

$$\begin{aligned} \frac{p^2 + q^2}{2} H\left(\frac{p^2}{p^2 + q^2}\right) &= -p^2 \log_2 p - q^2 \log_2 q + \frac{p^2 + q^2}{2} \log_2 (p^2 + q^2), \\ \frac{1}{2} H(p^2 + q^2) &= -\frac{p^2 + q^2}{2} \log_2 (p^2 + q^2) - pq \log_2 (2pq), \end{aligned}$$

and finally

$$\begin{aligned} A(p) &= pq - p^2 \log_2 p - q^2 \log_2 q - pq \log_2 (2pq) \\ &= -p^2 \log_2 p - q^2 \log_2 q - pq \log_2 pq \\ &= -(p^2 + pq) \log_2 p - (q^2 + pq) \log_2 q \\ &= -p \log_2 p - q \log_2 q = H(p). \end{aligned}$$

□

Remark 1.10.30. Exercise 1.10.29 is called the Peres extractor [von51, Eli72, Per92]. The von Neumann method [von51] extracts the bits pairwise from $X_1, \dots, X_n$, but throws away the randomness from the structure of the bit string, and there are two leftover sources of randomness: One is the structure of both heads or both tails, which gives $Y_1, \dots$, and the other is the structure of being equal or not, which gives $Z_1, \dots, Z_{\lfloor n/2 \rfloor}$. The von Neumann method took care of the symmetry case of a head and a tail, and the two leftovers are taken care of in the Peres extractor.

Exercise 1.10.31 (Exercise 10.13 [MU17]). Suppose we only have a biased 6-sided die with entropy $h > 0$ instead of a fair coin. Modify the extraction function in theorem 1.10.28 so that it extracts, on average, almost h random bits per roll from a sequence of die rolls.

Proof. We consider a variant construction based on theorem 1.10.28, where we roll the die n times. Conditioned on seeing n_k occurrences of the k^{th} face, the sequence is uniformly distributed over all

$$N = \binom{n}{n_1; n_2; n_3; n_4; n_5; n_6}$$

sequences. We apply the theorem 1.10.23 extractor to extract the random variable uniform over $[0, N - 1]$.

Let Z_k be the random variable for the number of times the k^{th} face comes up, p_k be the probability that the k^{th} face comes up, and B be the random variable for the bits extracted. Define

$$\mathcal{T}_\varepsilon = \left\{ (n_1, \dots, n_6) : \sum_{k \in [1,6]} n_k = n, \text{ and } \forall k \in [1,6], |n_k - np_k| \leq \varepsilon n \right\}.$$

By conditional expectation,

$$\begin{aligned} \mathbb{E}[B] &= \sum_{n_1, \dots, n_6} \mathbb{E}[B \mid Z_1 = n_1, \dots, Z_6 = n_6] \cdot \Pr[Z_1 = n_1, \dots, Z_6 = n_6] \\ &\geq \min_{(n_1, \dots, n_6) \in \mathcal{T}_\varepsilon} \mathbb{E}[B \mid Z_1 = n_1, \dots, Z_6 = n_6] \cdot \left(1 - \sum_{k \in [1,6]} \Pr[|Z_k - np_k| \geq \varepsilon n] \right) \\ &\geq \min_{(n_1, \dots, n_6) \in \mathcal{T}_\varepsilon} \left(\left\lfloor \log_2 \binom{n}{n_1; \dots; n_6} \right\rfloor - 1 \right) \cdot \left(1 - \sum_{k \in [1,6]} \Pr[|Z_k - np_k| \geq \varepsilon n] \right), \end{aligned}$$

where the second inequality holds by union bounding, and the third inequality holds by theorem 1.10.23.

We now consider a variant of lemma 1.10.18 for the biased 6-sided die.

Lemma 1.10.32. For integers $\{nq_i\}_{i \in [1,6]}$ in $[0, n]$ and $\sum_{i \in [1,6]} q_i = 1$, we have

$$\frac{1}{2^{n^{5/2}}} \cdot 2^{nH(q_1, \dots, q_6)} \leq \binom{n}{nq_1; \dots; nq_6} \leq 2^{nH(q_1, \dots, q_6)},$$

where $H(q_1, \dots, q_6) = -\sum_{i \in [1,6]} q_i \log_2 q_i$.

Proof. We begin with the upper bound. Since

$$1 = (q_1 + \dots + q_6)^n \geq \binom{n}{nq_1; \dots; nq_6} \prod_{i \in [1,6]} q_i^{nq_i},$$

then we complete the upper bound proof by

$$\log_2 \binom{n}{nq_1; \dots; nq_6} \leq -n \sum_{i \in [1,6]} q_i \log_2 q_i = nH(q_1, \dots, q_6).$$

The lower bound is trivial for $n \in [1, 6]$. Let $I \subseteq [1, 6]$ be the subset of indices such that $nq_i > 0$. By lemma 1.10.17,

$$n! \geq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n, \quad (nq_i)! \leq \sqrt{2\pi nq_i} \left(\frac{nq_i}{e}\right)^{nq_i} \exp\left(\frac{1}{12nq_i}\right),$$

and therefore

$$\binom{n}{nq_1; \dots; nq_6} \geq \left(\frac{1}{\sqrt{2\pi n}}\right)^{|I|-1} \cdot \prod_{i \in I} \left(\frac{1}{\sqrt{q_i}} \cdot \exp\left(-\frac{1}{12nq_i}\right)\right) \cdot \prod_{i \in I} q_i^{-nq_i} = \frac{2^{nH(q_1, \dots, q_6)}}{(2\pi n)^{(|I|-1)/2}} \cdot \prod_{i \in I} \left(\frac{1}{\sqrt{q_i}} \cdot \exp\left(-\frac{1}{12nq_i}\right)\right).$$

By AM-GM inequality,

$$\prod_{i \in I} q_i \leq \left(\frac{1}{|I|} \sum_{i \in I} q_i \right)^{|I|} \leq |I|^{-|I|},$$

then

$$\binom{n}{nq_1; \dots; nq_6} \geq \frac{2^{nH(q_1, \dots, q_6)}}{(2\pi n)^{(|I|-1)/2}} \cdot \prod_{i \in I} \left(\frac{1}{\sqrt{q_i}} \cdot \exp \left(-\frac{1}{12nq_i} \right) \right) \geq \frac{2^{nH(q_1, \dots, q_6)} \cdot |I|^{|I|/2}}{(2\pi n)^{(|I|-1)/2}} \cdot \prod_{i \in I} \exp \left(-\frac{1}{12nq_i} \right).$$

On the other hand, since each $nq_i \geq 1$ for $i \in I$, then

$$\begin{aligned} \binom{n}{nq_1; \dots; nq_6} &\geq \frac{2^{nH(q_1, \dots, q_6)} \cdot |I|^{|I|/2}}{(2\pi n)^{(|I|-1)/2}} \cdot \prod_{i \in I} \exp \left(-\frac{1}{12nq_i} \right) \\ &\geq \frac{2^{nH(q_1, \dots, q_6)} \cdot |I|^{|I|/2}}{(2\pi n)^{(|I|-1)/2}} \cdot \exp \left(-\frac{|I|}{12} \right) \\ &= \frac{2^{nH(q_1, \dots, q_6)}}{n^{(|I|-1)/2}} \cdot \sqrt{2\pi} \cdot \left(\frac{|I|}{2\pi \exp(1/6)} \right)^{|I|/2}, \end{aligned}$$

and for $|I| \in [1, 6]$ the coefficient

$$\sqrt{2\pi} \cdot \left(\frac{|I|}{2\pi \exp(1/6)} \right)^{|I|/2} \geq \frac{1}{2},$$

which completes the proof. $\square$

By the KL-divergence-driven Chernoff bound in exercise 1.4.4,

$$\Pr [|Z_k - np_k| \geq \varepsilon n] \leq 2 \exp(-2n\varepsilon^2).$$

Together with lemma 1.10.32,

$$\begin{aligned} \mathbb{E}[B] &\geq \min_{(n_1, \dots, n_6) \in \mathcal{T}_\varepsilon} \left(\left\lfloor \log_2 \binom{n}{n_1; \dots; n_6} \right\rfloor - 1 \right) \cdot \left(1 - \sum_{k \in [1, 6]} \Pr [|Z_k - np_k| \geq \varepsilon n] \right) \\ &\geq \min_{(n_1, \dots, n_6) \in \mathcal{T}_\varepsilon} \left(nH(q_1, \dots, q_6) - \frac{5}{2} \log_2 n - 3 \right) \cdot (1 - 12 \exp(-2n\varepsilon^2)), \end{aligned}$$

where $q_k = n_k/n$. Since H is continuous over $[0, 1]^6$, there exists a constant $\gamma_\varepsilon \rightarrow 0$ as $\varepsilon \rightarrow 0$, such that

$$\begin{aligned} \mathbb{E}[B] &\geq \min_{(n_1, \dots, n_6) \in \mathcal{T}_\varepsilon} \left(nH(q_1, \dots, q_6) - \frac{5}{2} \log_2 n - 3 \right) \cdot (1 - 12 \exp(-2n\varepsilon^2)) \\ &\geq \left(n(h - \gamma_\varepsilon) - \frac{5}{2} \log_2 n - 3 \right) \cdot (1 - 12 \exp(-2n\varepsilon^2)), \end{aligned}$$

proving the lower bound.

To show an extraction function extracts on average at most nh fair bits, the idea is similar to the theorem 1.10.28. Let X be the random variable of the outcome of n die rolls, then for an input-output pair $(\mathbf{x}, \mathbf{y})$ with $|\mathbf{y}| = k$,

$$\Pr[X = \mathbf{x}] \leq \Pr[\text{Ext}(X) = \mathbf{y}] \leq 2^{-k},$$

the second inequality was discussed in theorem 1.10.28 from the conditional probability and definition 1.10.20, then

$$k = \mathbb{E}[B | X = \mathbf{x}] \leq -\log_2 \Pr[X = \mathbf{x}].$$

By conditional expectation,

$$\mathbb{E}[B] = \sum_{\mathbf{x}} \Pr[X = \mathbf{x}] \cdot \mathbb{E}[B | X = \mathbf{x}] \leq - \sum_{\mathbf{x}} \Pr[X = \mathbf{x}] \cdot \log_2 \Pr[X = \mathbf{x}] = H(X) = nh,$$

where the last equality is from lemma 1.10.7, and we complete the proof. $\square$

1.10.4 Entropy: A Limit of Data Compression

Example 1.10.33. Consider a biased coin that comes up heads with probability $3/4$. Let HH map to 0, HT map to 10, TH map to 110, and TT map to 111. Then on average, the number of bits we use for each pair of flips is

$$1 \cdot \frac{9}{16} + 2 \cdot \frac{3}{16} + 3 \cdot \frac{3}{16} + 3 \cdot \frac{1}{16} = \frac{27}{16} < 2.$$

This is an example of compression, as on average a biased coin flip can be represented with less than 1 fair coin flip.

Moreover, it is worth noting that this construction allows for breaking a sequence of biased coin flips into pairs, and concatenating the compression results. The result can be uniquely decoded by simply parsing from left to right. For example, 011110 stands for HHTTHT. The construction preserves the property that the bit representation of any pair of coin flips is not a prefix of the bit representation of any other pair of coin flips.

Representations with this property are called prefix codes.

Exercise 1.10.34 (Exercise 10.15 [MU17]). We wish to compress a sequence of i.i.d. random variables $X_1, \dots$, where each X_j takes on one of n possible values, into a prefix code like example 1.10.33. We map the i^{th} value among all n values to a codeword, which is a sequence of ℓ_i bits. Prove that the ℓ_i must satisfy

$$\sum_{i \in [1, n]} 2^{-\ell_i} \leq 1.$$

Proof. Consider a binary tree of depth $\ell_{\max} = \max\{\ell_i\}_{i \in [1, n]}$, then the length of a codeword is the distance to the root, and a codeword is a prefix of the path to leaves: 0 stands for the left subtree, and 1 stands for the right subtree. For a codeword of length ℓ_i , there are $2^{\ell_{\max} - \ell_i}$ leaves whose root-to-leaf paths begin with that codeword. These sets of leaves are disjoint, since otherwise one codeword would be a prefix of another codeword, contradicting the prefix code property.

Now that every leaf in the binary tree is owned by at most 1 codeword, the number of owned leaves are

$$\sum_{i \in [1, n]} 2^{\ell_{\max} - \ell_i} \leq 2^{\ell_{\max}},$$

which completes the proof. □

Remark 1.10.35. The exercise 1.10.34 is also called Kraft-McMillan inequality [Kra49, McM56].

Definition 1.10.36. A compression function Com takes as input a sequence of n coin flips, given as an element of $\{H, T\}^n$, and outputs a sequence of bits, such that each input sequence of n flips yields distinct output sequences.

We consider the case of compressing the outcome of a sequence of biased coin flips, similar to theorem 1.10.28.

Theorem 1.10.37. Suppose a coin comes up heads with probability $p > 1/2$. For any constant $\delta > 0$ and n sufficiently large,

- There exists a compression function Com such that the expected number of output bits on an input sequence of n independent coin flips is at most $(1 + \delta)nH(p)$.
- The expected number of output bits by any compression function on an input sequence of n independent coin flips is at least $(1 - \delta)nH(p)$.

Proof. We begin with the upper bound, with an explicit construction of a compression function.

Let $\varepsilon > 0$ be a sufficiently small constant with $p - \varepsilon > 1/2$. On input a sequence of independent coin flips, output the first bit as a flag bit, 0 when the input has at least $\lceil n(p - \varepsilon) \rceil$ heads, 1 otherwise. When the first flag bit is 1, the compression function translates a head to a 1, and a tail to a 0, which requires $n + 1$ bits to output. When the first flag bit is 0, let the number of coin flip sequences be m , and map a sequence to a value in $[0, m - 1]$.

If there are less than $\lceil n(p - \varepsilon) \rceil$ heads, we upper bound the probability using Chernoff's bound by

$$\Pr[X - np < -n\varepsilon] \leq \left(\frac{e^{-\gamma}}{(1 - \gamma)^{1-\gamma}} \right)^\mu \leq \exp\left(-\frac{\mu\gamma^2}{2}\right) = \exp\left(-\frac{n\varepsilon^2}{2p}\right),$$

where $\gamma = \varepsilon/p$, and $\mu = np$.

If there are at least $\lceil n(p - \varepsilon) \rceil$ heads, we upper bound the number of coin flip sequences by

$$\sum_{k \in [\lceil n(p-\varepsilon) \rceil, n]} \binom{n}{k} \leq \sum_{k \in [\lceil n(p-\varepsilon) \rceil, n]} \binom{n}{\lceil n(p-\varepsilon) \rceil} \leq \frac{n}{2} \binom{n}{\lceil n(p-\varepsilon) \rceil} \leq \frac{n}{2} \cdot 2^{nH(p-\varepsilon)} = 2^{nH(p-\varepsilon) + \log_2 n - 1} \leq 2^{\lfloor nH(p-\varepsilon) + \log_2 n \rfloor},$$

where the second and the third inequality holds by $p - \varepsilon > 1/2$, and the fourth inequality holds by corollary 1.10.19. Including the flag bit, it takes at most $nH(p - \varepsilon) + \log_2 n + 1$ bits for the sequences with at least $\lceil n(p - \varepsilon) \rceil$ heads.

Let B be the number of output bits of the compression function. The upper bound of its expectation is

$$\begin{aligned} \mathbb{E}[B] &\leq \exp\left(-\frac{n\varepsilon^2}{2p}\right) \cdot (n + 1) + \left(1 - \exp\left(-\frac{n\varepsilon^2}{2p}\right)\right) \cdot (nH(p - \varepsilon) + \log_2 n + 1) \\ &= 1 + \exp\left(-\frac{n\varepsilon^2}{2p}\right) \cdot n + \left(1 - \exp\left(-\frac{n\varepsilon^2}{2p}\right)\right) \cdot (nH(p) + n\gamma_\varepsilon + \log_2 n) \\ &= 1 + nH(p) + \boxed{\exp\left(-\frac{n\varepsilon^2}{2p}\right) \cdot n(1 - H(p))} + \boxed{\left(1 - \exp\left(-\frac{n\varepsilon^2}{2p}\right)\right) \cdot (n\gamma_\varepsilon + \log_2 n)}, \end{aligned}$$

where the γ_ε is a constant similar to the one in exercise 1.10.31, since H is continuous and monotone over $[1/2, 1]$, $\gamma_\varepsilon \rightarrow 0$ as $\varepsilon \rightarrow 0$. For ε sufficiently small, when n is sufficiently large, the first boxed term approaches to 0, and the second boxed term is upper bounded by $\delta nH(p)$ for some constant δ .

We now show the lower bound. It suffices to consider a compression function minimizing the expected output number of bits. We observe a fact: if an input string $\mathbf{s}_1$ is more likely than $\mathbf{s}_2$, then the output number of bits on input $\mathbf{s}_2$ should be at least as long as the one of $\mathbf{s}_1$. Otherwise, swapping the two outputs decreases the expectation. Since the probability of deriving a specific sequence with k heads is $p^k(1-p)^{n-k}$, it is more likely for a string to have more heads. Hence, by Chernoff's bound,

$$\Pr[X > \lceil n(p + \varepsilon) \rceil] \leq \Pr[X - np > n\varepsilon - 1] \leq \left(\frac{e^\gamma}{(1 + \gamma)^{1+\gamma}} \right)^\mu \leq \exp\left(-\frac{\mu\gamma^2}{3}\right) = \exp\left(-\frac{n(\varepsilon - 1/n)^2}{3p}\right),$$

where $\mu = np$, $\gamma = \varepsilon/p - 1/n$, and ε sufficiently small so that $p + \varepsilon < 1$.

Another fact: for a random variable uniform over $[1, m]$, on average any compression function outputs at least $\log_2 m - 3$ bits. By mapping to $[0, m - 1]$, and let ℓ be the first integer that $2 + 2 + 4 + \dots + 2^\ell \geq m$, hence $\ell \geq \log_2 m - 1$. Hence, the average output bit of the compression function is at least $\ell - 2$. Averaged over sequences with $\lceil n(p + \varepsilon) \rceil$ heads, the compression function outputs at least

$$\log_2 \left(\binom{n}{\lceil n(p + \varepsilon) \rceil} \right) - 3 \geq \log_2 \frac{2^{nH(p + \varepsilon)}}{2\sqrt{n}} - 3 = nH(p + \varepsilon) - \frac{1}{2} \log_2 n - 4$$

bits. By the ordering of an optimal compression function, sequences with fewer heads have average output length at least this large. The expected number of output bits of the compression function is lower bounded by ignoring the strings with more than $\lceil n(p + \varepsilon) \rceil$ heads, then

$$\left(1 - \exp\left(-\frac{n(\varepsilon - 1/n)^2}{3p}\right)\right) \cdot \left(nH(p + \varepsilon) - \frac{1}{2} \log_2 n - 4\right).$$

By choosing a sufficiently small ε , and then a sufficiently large n , we can lower bound the expected number of output bits to $(1 - \delta)nH(p)$. $\square$

Exercise 1.10.38 (Exercise 10.14 [MU17]). Suppose we only have a biased 6-sided die with entropy $h > 0$ instead of a fair coin. Modify the compression function in theorem 1.10.37 so that it compresses a sequence of die rolls to almost nh bits on average.

Proof. We begin with the lower bound by reusing lemma 1.10.32 and the ideas in exercise 1.10.31 and theorem 1.10.37. Let Z_i be the random variable for the number of times the i^{th} face comes up, p_i be the probability that the i^{th} face comes up, and B be the random variable for the bits of the compression outcome. Recall the $\mathcal{T}_\varepsilon$ defined in exercise 1.10.31, by union bound over Chernoff's bound in exercise 1.4.4,

$$\Pr[(Z_1, \dots, Z_6) \in \mathcal{T}_\varepsilon] = 1 - \Pr[(Z_1, \dots, Z_6) \notin \mathcal{T}_\varepsilon] \geq 1 - \sum_{i \in [1,6]} \Pr[|Z_i - np_i| \geq n\varepsilon] \geq 1 - 12 \exp(-2n\varepsilon^2).$$

Moreover, by previous discussion in theorem 1.10.37, a compression function minimizing the expected output bits outputs fewer bits on inputs that are more likely than the others, then let $(n_1, \dots, n_6)$ be the face counts that

$$\prod_{i \in [1,6]} p_i^{n_i} = \max_{(k_1, \dots, k_6) \in \mathcal{T}_\varepsilon} \prod_{i \in [1,6]} p_i^{k_i},$$

and conditioned on $(k_1, \dots, k_6)$, the inputs are uniform over $\binom{n}{n_1, \dots, n_6}$ sequences, hence within the $\mathcal{T}_\varepsilon$, the compression function outputs at least

$$\mathbb{E}[B \mid Z_1 = k_1, \dots, Z_6 = k_6] \geq \log_2 \binom{n}{n_1; \dots; n_6} - 3 \geq \log_2 \left(\frac{1}{2n^{5/2}} \cdot 2^{nH(q_1, \dots, q_6)} \right) - 3 = nH(q_1, \dots, q_6) - \frac{5}{2} \log_2 n - 4,$$

where the second inequality holds by lemma 1.10.32, and $q_i = n_i/n$. Since this choice maximizes the probability of each individual sequence with these counts, every other sequence in $\mathcal{T}_\varepsilon$ is less likely, so its output length is at least the average output length over this type. Then

$$\begin{aligned} \mathbb{E}[B] &\geq \sum_{(k_1, \dots, k_6) \in \mathcal{T}_\varepsilon} \mathbb{E}[B \mid Z_1 = k_1, \dots, Z_6 = k_6] \cdot \Pr[Z_1 = k_1, \dots, Z_6 = k_6] \\ &\geq \left(nH(q_1, \dots, q_6) - \frac{5}{2} \log_2 n - 4 \right) \cdot (1 - 12 \exp(-2n\varepsilon^2)) \\ &= \left(n(h - \gamma_\varepsilon) - \frac{5}{2} \log_2 n - 4 \right) \cdot (1 - 12 \exp(-2n\varepsilon^2)), \end{aligned}$$

where γ_ε is a constant such that $\gamma_\varepsilon \rightarrow 0$ when $\varepsilon \rightarrow 0$ as H is continuous over $[0, 1]^6$. Hence, for a sufficiently small ε , and a sufficiently large n , there exists δ such that the lower bound $(1 - \delta)nh$ is proved.

Now for the upper bound, we follow the main architecture of theorem 1.10.28, by introducing a flag bit separating the case where the dice rolling sequence is not in $\mathcal{T}_\varepsilon$ and the case where the dice rolling sequence is in $\mathcal{T}_\varepsilon$. If some count deviates too far, then output in its vanilla form where each face is represented in 3 bits. Otherwise, we count the number of possible sequences in $\mathcal{T}_\varepsilon$ to be m , and map a sequence to a value in $[0, m - 1]$, and we have

$$m = \sum_{(k_1, \dots, k_6) \in \mathcal{T}_\varepsilon} \binom{n}{k_1; \dots; k_6} \leq (2\varepsilon n + 1)^5 \binom{n}{n_1; \dots; n_6} \leq (2\varepsilon n + 1)^5 \cdot 2^{nH(q_1, \dots, q_6)},$$

where $n_1, \dots, n_6$ maximizes $\binom{n}{k_1; \dots; k_6}$ over $\mathcal{T}_\varepsilon$, the last inequality holds by lemma 1.10.32, and $q_i = n_i/n$. Hence,

$$\begin{aligned} \mathbb{E}[B] &\leq 1 + 12 \exp(-2\varepsilon^2 n) \cdot 3n + \left(\log_2 \left((2\varepsilon n + 1)^5 \cdot 2^{nH(q_1, \dots, q_6)} \right) + 1 \right) \\ &= 2 + 12 \exp(-2\varepsilon^2 n) \cdot 3n + 5 \log_2 (2\varepsilon n + 1) + nH(q_1, \dots, q_6) \\ &= 2 + 12 \exp(-2\varepsilon^2 n) \cdot 3n + 5 \log_2 (2\varepsilon n + 1) + n(h + \gamma_\varepsilon), \end{aligned}$$

where γ_ε is a constant such that $\gamma_\varepsilon \rightarrow 0$ when $\varepsilon \rightarrow 0$ as H is continuous over $[0, 1]^6$. Hence, for a sufficiently small ε , and a sufficiently large n , there exists δ such that the upper bound $(1 + \delta)nh$ is proved. $\square$

Exercise 1.10.39 (Exercise 10.16 [MU17]). We wish to compress a sequence of i.i.d. random variables $X_1, \dots$, where each X_j takes on one of n values. The i^{th} value occurs with probability p_i , where $p_1 \geq p_2 \geq \dots \geq p_n$. The compressed result follows. Let $T_i = \sum_{j \in [1, i-1]} p_j$, and let the i^{th} codeword to be the first $\ell_i = \lceil \log_2(1/p_i) \rceil$ bits of T_i . Show that it is a prefix code example 1.10.33. Let z be the average number of bits used for each X_j . Show that $H(X) \leq z \leq H(X) + 1$.

Proof. We observe that, for the binary representation of p_i , the first positive bit is the ℓ_i^{th} bit. Since the first positive bit is the i^{th} bit for any value in $[2^{-i}, 2^{-i+1})$, proving the claim. Moreover, as $p_1 \geq \dots \geq p_n$, then $\ell_1 \leq \dots \leq \ell_n$, and it suffices to show that any pair of neighboring adjacent codewords does not violate the prefix code property.

On a pair of T_i and $T_{i+1} = T_i + p_i$, since $T_i \leq T_i + 2^{-\ell_i} \leq T_i + p_i$, the first ℓ_i bits are distinct for the i^{th} and $(i+1)^{\text{th}}$ codeword, proving the prefix code property.

Now for the average number of bits output by the compression function,

$$z = \sum_{i \in [1, n]} p_i \lceil \log_2(1/p_i) \rceil \geq - \sum_{i \in [1, n]} p_i \log_2 p_i = H(X),$$

as $\lceil \log_2(1/p_i) \rceil \geq -\log_2 p_i$. Moreover, by $\lceil \log_2(1/p_i) \rceil \leq -\log_2 p_i + 1$, then

$$z = \sum_{i \in [1, n]} p_i \lceil \log_2(1/p_i) \rceil \leq \sum_{i \in [1, n]} p_i (-\log_2 p_i + 1) = H(X) + 1,$$

and we proved $H(X) \leq z \leq H(X) + 1$. $\square$

Exercise 1.10.40 (Exercise 10.3 [MU17]). Flip a fair coin repeatedly X times until the first heads occurs. Find $H(X)$.

Now a friend flips the fair coin repeatedly until the first heads occurs. One wants to determine how many flips required, and is allowed to ask a series of yes-no questions of the following form: give the friend a set of integers, and the friend answers “yes” iff the number of flips is in the set, and “no” otherwise. Find a strategy such that the expected number of questions asked before determining the number of flips is $H(X)$. Give an intuitive explanation of why one cannot come up with a strategy that would ask fewer than $H(X)$ questions on average.

Proof. X follows geometrically distributed random variable with success probability $1/2$, then $\mathbb{E}[X] = 2$, and

$$H(X) = - \sum_{i \geq 1} \Pr[X = i] \cdot \log_2 \Pr[X = i] = \sum_{i \geq 1} \frac{1}{2^i} \cdot i = 2.$$

Intuitively, on the k^{th} question, ask if $X = k$, then the number of questions is same distributed as X . For a strategy, it is impossible to have a same terminating transcript of the yes-no answers from distinct values of X . Moreover, a terminating transcript cannot be continued, as it converges to a single value, ruling out other values, namely other suffixes. Hence, they form a prefix code. By exercise 1.10.39, at least $H(X) = 2$ questions are needed on average. $\square$

Exercise 1.10.41 (Exercise 10.17 [MU17]). Arithmetic coding is a standard compression method. In the case where the string to be compressed is a sequence of biased coin flips, it can be described as follows. For a sequence of i.i.d. Bernoulli trials $X_1, \dots, X_n$ with success probability $1 - p$. The sequences can be ordered lexicographically, so that for $\mathbf{x}, \mathbf{y} \in \{0, 1\}^n$, we say $\mathbf{x} < \mathbf{y}$ if $x_i = 0$ and $y_i = 1$ in the first coordinate i that $x_i \neq y_i$. If $z_{\mathbf{x}}$ is the number of zeroes in the string $\mathbf{x}$, define $p(\mathbf{x}) = p^{z_{\mathbf{x}}}(1-p)^{n-z_{\mathbf{x}}}$, and $q(\mathbf{x}) = \sum_{\mathbf{y} < \mathbf{x}} p(\mathbf{y})$.

- Suppose we are given $\mathbf{x}$ sequentially, explain how to compute $q(\mathbf{x})$ in $O(n)$ time.
- Argue that $[q(\mathbf{x}), q(\mathbf{x}) + p(\mathbf{x})]$ are disjoint subintervals over $[0, 1)$, show that $\bar{X} = (X_1, \dots, X_n)$ can be represented by any point in $[q(\bar{X}), q(\bar{X}) + p(\bar{X})]$, and a codeword can be chosen in the interval by the first $\lceil \log_2(1/p(\bar{X})) \rceil + 1$ bits, such that the codeword is a prefix code.
- On a codeword, show how to decompress to determine the corresponding $\bar{X}$.
- Using a Chernoff's bound, argue that $\log_2(1/p(\bar{X}))$ is close to $nH(p)$ with high probability.

Proof. To derive $q(\mathbf{x})$ in $O(n)$, count $i \in [1, n]$: when $x_i = 0$, move forward; otherwise, add $p \prod_{j \in [1, i-1]} q_j$ to $q(\mathbf{x})$, where q_j is the probability corresponding to x_j . Eventually, one sub-hypercube at a time, and

$$q(\mathbf{x}) = p \sum_{i \in [1, n]} x_i \prod_{j \in [1, i-1]} q_j.$$

Let $\mathbf{x} + 1$ be defined as adding by 1 to the number with binary representation $\mathbf{x}$, then $q(\mathbf{x})$ is monotonically increasing for each $\mathbf{x} \in [0, 2^n - 1]$, such that $[q(\mathbf{x}), q(\mathbf{x}) + p(\mathbf{x})]$ are disjoint, as $q(\mathbf{x} + 1) = q(\mathbf{x}) + p(\mathbf{x})$, and $q(\mathbf{1}) + p(\mathbf{1}) = 1$.

By exercise 1.10.39, the first positive bit of $p(\mathbf{x})$ is $\ell_{\mathbf{x}} = \lceil \log_2(1/p(\mathbf{x})) \rceil$, then there exists a dyadic interval of length $2^{-(\ell_{\mathbf{x}}+1)}$ in $[q(\mathbf{x}), q(\mathbf{x}) + p(\mathbf{x})]$, as $2^{-\ell_{\mathbf{x}}} \leq p(\mathbf{x})$. The codeword can be chosen by the lower bound of the dyadic interval, requiring $\ell_{\mathbf{x}} + 1$ bits. If a codeword $\mathbf{c}_0$ is a prefix of a codeword $\mathbf{c}_1$, then $\mathbf{c}_1$ is in a dyadic interval of $\mathbf{c}_0$, while the exterior intervals are disjoint, which is a contradiction.

Decoding the codeword c is in the reverse order of deriving $q(\mathbf{x})$: Count $i \in [1, n]$, and let the left endpoint $a = 0$. If $c \geq a + p \prod_{j \in [1, i-1]} q_j$, the i^{th} bit is 1, $q_i \leftarrow 1 - p$, and add $p \prod_{j \in [1, i-1]} q_j$ to a ; otherwise, the i^{th} bit is 0 and $q_i \leftarrow p$.

Let Z be the number of tails in $[0, n]$, then $\mathbb{E}[Z] = np$, and

$$Y = \log_2(1/p(\bar{X})) = -\log_2 p \cdot Z - \log_2(1 - p) \cdot (n - Z) = -n \log_2(1 - p) + Z \log_2 \frac{1 - p}{p},$$

such that $\mathbb{E}[Y] = nH(p)$. If $p = 1/2$, $Y = nH(p) = n$ deterministically. Otherwise, as $Z = \sum_{i \in [1, n]} Z_i$ for independent Z_i being 1 on the i^{th} flip, we can apply Chernoff's bound, such that

$$\Pr \left[|Y - nH(p)| \geq \delta np \left| \log_2 \frac{1 - p}{p} \right| \right] = \Pr [|Z - np| \geq \delta np] \leq \left(\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right)^{np} + \left(\frac{e^{-\delta}}{(1 - \delta)^{1 - \delta}} \right)^{np} \leq 2 \exp \left(-\frac{\delta^2}{3} np \right).$$

□

1.10.5 Shannon's Theorem

Consider the following type of channel.

Definition 1.10.42. The input to a binary symmetric channel with parameter p is a sequence of bits $x_1, \dots$ and the output is a sequence of bits $y_1, \dots$ such that $\Pr[x_i = y_i] = 1 - p$ independently for each i .

Consider encoding functions that bring redundancy to help protect against the introduction of errors.

Definition 1.10.43. A (k, n) encoding function $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ takes as input a sequence of k bits and outputs a sequence of n bits. Conversely, a (k, n) decoding function $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^k$ takes as input a sequence of n bits and outputs a sequence of k bits.

Exercise 1.10.44 (Exercise 10.18 [MU17]). Alice wants to send Bob the result of a fair coin flip over a binary symmetric channel that flips each bit with probability $p < 1/2$. To avoid errors in transmission, she encodes heads as a sequence of $2k + 1$ zeroes and tails as a sequence of $2k + 1$ ones.

- Consider the case where $k = 1$. For each possible received sequence of 3 bits, determine the probability that Alice flipped a heads conditioned on Bob receiving that sequence.
- Bob decodes by examining the 3 bits. If two or three of the bits are 0, Bob decides the corresponding coin flip was heads. Prove that this rule minimizes the probability of error for each flip.
- Argue that, for general k , Bob minimizes the probability of error by deciding the flip was heads if at least $k + 1$ of the bits are 0.
- Give a formula for the probability that Bob makes an error that holds for general k .

Proof. Let $n = 2k + 1$, $\mathbf{y} \in \{0, 1\}^n$, $\text{nzero}(\mathbf{y})$ be the number of zeroes of the $\mathbf{y}$, X be a fair coin outcome, tails on $X = 0$, otherwise heads, and $\bar{Y} = (Y_1, \dots, Y_n)$ be the received bits. Hence,

$$\begin{aligned}\Pr[\bar{Y} = \mathbf{y} \mid X = 1] &= (1 - p)^{\text{nzero}(\mathbf{y})} \cdot p^{n - \text{nzero}(\mathbf{y})}, \\ \Pr[\bar{Y} = \mathbf{y} \mid X = 0] &= p^{\text{nzero}(\mathbf{y})} \cdot (1 - p)^{n - \text{nzero}(\mathbf{y})},\end{aligned}$$

and by Bayes' Rule,

$$\Pr[X = 1 \mid \bar{Y} = \mathbf{y}] = \frac{(1 - p)^{\text{nzero}(\mathbf{y})} \cdot p^{n - \text{nzero}(\mathbf{y})}}{(1 - p)^{\text{nzero}(\mathbf{y})} \cdot p^{n - \text{nzero}(\mathbf{y})} + p^{\text{nzero}(\mathbf{y})} \cdot (1 - p)^{n - \text{nzero}(\mathbf{y})}}.$$

Consider decoding the flip: if at least m bits are 0, then we have heads; otherwise we have tails. Then

$$\begin{aligned}\Pr[\text{error}] &= \Pr[X = 0 \wedge \text{nzero}(\bar{Y}) \geq m] + \Pr[X = 1 \wedge \text{nzero}(\bar{Y}) < m] \\ &= \frac{1}{2} \sum_{i \in [m, n]} \binom{n}{i} \cdot p^i \cdot (1 - p)^{n-i} + \frac{1}{2} \sum_{i \in [0, m-1]} \binom{n}{i} \cdot (1 - p)^i \cdot p^{n-i} \\ &= \frac{1}{2} + \frac{1}{2} \sum_{i \in [m, n]} \binom{n}{i} \cdot (p^i \cdot (1 - p)^{n-i} - (1 - p)^i \cdot p^{n-i}).\end{aligned}$$

Since $p < 1/2$, then $p^i \cdot (1 - p)^{n-i} < (1 - p)^i \cdot p^{n-i}$ for $i \geq k + 1$, which proves the error minimizing when $m = k + 1$. $\square$

Lemma 1.10.45. *For a binary symmetric channel with parameter $p < 1/2$ and for any constant $\delta, \gamma > 0$, when n is sufficiently large, for any $k \leq n(1 - H(p) - \delta)$ and a uniformly distributed k -bit input message, there exist (k, n) encoding and decoding functions such that the probability the receiver fails to obtain the correct message is at most γ .*

Proof. We begin by describing a probabilistic algorithm constructing 2^k codewords in $\{0, 1\}^n$, one for each message, and the inefficient encoding and decoding functions.

Let $\{X_i\}_{i \in [0, 2^k - 1]}$ be the distinct codewords sampled uniformly at random from $\{0, 1\}^n$. The encoding function $\text{Enc} : \{0, 1\}^k \rightarrow \{0, 1\}^n$ is a lookup table mapping the k -bit inputs to their corresponding codewords. The decoding function $\text{Dec} : \{0, 1\}^n \rightarrow \{0, 1\}^k \cup \{\perp\}$ maintains the look up table, iterates through the codewords, and checks if a codeword differing from the received n bits in between $\lceil n(p - \varepsilon) \rceil, \lfloor n(p + \varepsilon) \rfloor$ bits. If there is only one codeword in this interval, output the corresponding k -bit message; otherwise, decoding function fails and outputs $\perp$. Since the decoding runtime is exponential in k , the decoding is not efficient, and we are not requiring it to be efficient.

When it comes to decoding failure or failure to obtain the correct message, it falls into 2 cases:

- The noisy channel made less than $n(p - \varepsilon)$ or more than $n(p + \varepsilon)$ errors.
- There are other codeword(s) that are differing from the received n bits in between $\lceil n(p - \varepsilon) \rceil, \lfloor n(p + \varepsilon) \rfloor$ bits.

The goal is to prove the existence of a set of codewords, whose probability of failures above is at most γ . The strategy here is probabilistic method, via averaging argument lemma 1.6.12, if probability of failures above on average among all $\{X_i\}_{i \in [0, 2^k - 1]}$ is at most γ , then there must exist a set of codewords whose failure probability is at most γ . The first failure event can be upper bounded via concentration inequalities, while the second failure event can be upper bounded via counting argument.

We continue by introducing notation and random variable formulations. Let $\Delta(\mathbf{s}_0, \mathbf{s}_1)$ be the Hamming weight of $\mathbf{s}_0 \oplus \mathbf{s}_1$ for the number of positions $\mathbf{s}_0$ and $\mathbf{s}_1$ differing from each other. We say $(\mathbf{s}_0, \mathbf{s}_1)$ pair has weight

$$w(\mathbf{s}_0, \mathbf{s}_1) = p^{\Delta(\mathbf{s}_0, \mathbf{s}_1)} \cdot (1 - p)^{n - \Delta(\mathbf{s}_0, \mathbf{s}_1)},$$

corresponding to the probability of $\mathbf{s}_0$ being altered into $\mathbf{s}_1$ by the binary symmetric channel. Let $\{S_i\}_{i \in [0, 2^k - 1]}$ be the sets of n bits that decodes to X_i , which are n bits that are close to X_i , and are far from other X_j . Let $\{W_i\}_{i \in [0, 2^k - 1]}$ be the set of probabilities that X_i fails in decoding correctly, where

$$W_i = \sum_{\mathbf{r} \notin S_i} w(X_i, \mathbf{r}).$$

It can also be expressed in a way that, let $1_{i,\mathbf{r}}$ be an indicator 0/1 random variable that is 1 iff $\mathbf{r} \notin S_i$, then

$$W_i = \sum_{\mathbf{r} \in \{0,1\}^n} 1_{i,\mathbf{r}} \cdot w(X_i, \mathbf{r}).$$

Note that, both $\{S_i\}_{i \in [0, 2^k - 1]}$ and $\{W_i\}_{i \in [0, 2^k - 1]}$ are *random variables* dependent on $\{X_i\}_{i \in [0, 2^k - 1]}$.

WLOG, we analyze with respect to X_0 , as by symmetry the following results applies to other X_i . Write

$$W_0 = \sum_{\mathbf{r} \in T_1} 1_{0,\mathbf{r}} \cdot w(X_0, \mathbf{r}) + \sum_{\mathbf{r} \in T_2} 1_{0,\mathbf{r}} \cdot w(X_0, \mathbf{r}), \quad (1.13)$$

where $T_1 = \{\mathbf{s} : |\Delta(X_0, \mathbf{s}) - np| > n\varepsilon\}$, and $T_2 = \{\mathbf{s} : |\Delta(X_0, \mathbf{s}) - np| \leq n\varepsilon\}$.

For the first term, let $\{Z_i\}_{i \in [1, n]}$ be independent Bernoulli trials with success probability p , then $Z = \sum_{i \in [1, n]} Z_i$ is the random variable for the number of bits flipped by the channel, and by Chernoff's bound, for $0 < \varepsilon < p$,

$$\sum_{\mathbf{r} \in T_1} 1_{0,\mathbf{r}} \cdot w(X_0, \mathbf{r}) = \sum_{\mathbf{r} \in T_1} w(X_0, \mathbf{r}) = \Pr[|Z - np| > n\varepsilon] \leq \left(\frac{e^\theta}{(1 + \theta)^{1+\theta}} \right)^{np} + \left(\frac{e^{-\theta}}{(1 - \theta)^{1-\theta}} \right)^{np} \leq 2 \exp\left(-\frac{\mu\theta^2}{3}\right), \quad (1.14)$$

where $\theta = \varepsilon/p$ and $\mu = np$. For any threshold $0 < \varepsilon < p$, there is n sufficiently large such that $2 \exp(-\mu\theta^2/3) \leq \gamma/2$.

For the second term, consider the following symmetry: On a set of codewords $\{X_i\}_{i \in [0, 2^k - 1]}$, we have $\{X'_i\}_{i \in [0, 2^k - 1]}$ such that $X'_i \leftarrow X_i \oplus X_0$, and $X'_0 = \mathbf{0}$. Then $\{S'_i\}_{i \in [0, 2^k - 1]}$ can be obtained by XORing X_0 with elements in $\{S_i\}_{i \in [0, 2^k - 1]}$, and the W_0 can be expressed by

$$W_0 = \sum_{\mathbf{r} \in T'_1} 1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r}) + \sum_{\mathbf{r} \in T'_2} 1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r}).$$

where $T'_1 = \{\mathbf{s} : |\Delta(\mathbf{0}, \mathbf{s}) - np| > n\varepsilon\}$, and $T'_2 = \{\mathbf{s} : |\Delta(\mathbf{0}, \mathbf{s}) - np| \leq n\varepsilon\}$. The symmetry holds for the claim in the first term, while the symmetry allows us to argue the second term with respect to $X'_0 = \mathbf{0}$, then

$$\sum_{\mathbf{r} \in T'_2} 1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})$$

is a random variable depending on $\{X'_i\}_{i \in [1, 2^k - 1]}$. Moreover, for a $\mathbf{r} \in T'_2$, averaging over $\{X'_i\}_{i \in [1, 2^k - 1]}$,

$$\mathbb{E}[1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})] = w(\mathbf{0}, \mathbf{r}) \cdot \Pr[\exists i \in [1, 2^k - 1], |\Delta(X'_i, \mathbf{r}) - np| \leq n\varepsilon].$$

Supposing we have

$$\Pr[\exists i \in [1, 2^k - 1], |\Delta(X'_i, \mathbf{r}) - np| \leq n\varepsilon] \leq \frac{\gamma}{2},$$

then the second term can be upper bounded by

$$\sum_{\mathbf{r} \in T'_2} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})] \leq \sum_{\mathbf{r} \in T'_2} w(\mathbf{0}, \mathbf{r}) \cdot \frac{\gamma}{2} \leq \frac{\gamma}{2}.$$

We continue by upper bounding the number of possible n bits around $\mathbf{r}$, such that we upper bound the second term:

$$\sum_{k \in [\lceil n(p-\varepsilon) \rceil, \lfloor n(p+\varepsilon) \rfloor]} \binom{n}{k} \leq (2\varepsilon n + 1) \cdot \binom{n}{\lfloor n(p+\varepsilon) \rfloor} \leq (2\varepsilon n + 1) \cdot 2^{nH(p+\varepsilon)},$$

threshold ε is chosen such that $\lceil n(p + \varepsilon) \rceil \leq n/2$, and the second inequality holds by corollary 1.10.19. The probability of a particular codeword X'_i with $i > 0$ having a Hamming distance $[[n(p - \varepsilon)], [n(p + \varepsilon)]]$ to $\mathbf{r}$ causing a decoding failure, when $\mathbf{0}$ is sent is at most

$$(2\varepsilon n + 1) \cdot 2^{n(H(p+\varepsilon)-1)} = (2\varepsilon n + 1) \cdot 2^{n(H(p)+\gamma_\varepsilon-1)},$$

where γ_ε is a constant similar to the one appearing in theorem 1.10.28 and theorem 1.10.37. Union bounding over all other codewords, the probability of other codeword failing decoding $\mathbf{0}$ is upper bounded by

$$\sum_{\mathbf{r} \in T'_2} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})] \leq (2^k - 1) \cdot (2\varepsilon n + 1) \cdot 2^{n(H(p)+\gamma_\varepsilon-1)} \leq (2\varepsilon n + 1) \cdot 2^{n(\gamma_\varepsilon-\delta)}, \quad (1.15)$$

where the last inequality holds by $k \leq n(1 - H(p) - \delta)$. Since eq. (1.15) approaches 0 as $n \rightarrow \infty$ with $\delta > \gamma_\varepsilon$, then it is at most $\gamma/2$ on a sufficiently large n , and by symmetry, we conclude

$$\sum_{\mathbf{r} \in T_2} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(X_0, \mathbf{r})] \leq \frac{\gamma}{2}. \quad (1.16)$$

Combining eqs. (1.13), (1.14) and (1.16), on a sufficiently large n , the probability of decoding incorrectly X_0 is

$$\mathbb{E}[W_0] = \sum_{\mathbf{r} \in \{0,1\}^n} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(X_0, \mathbf{r})] \leq \gamma.$$

Since the k -bit message is uniform over $\{0, 1\}^k$, the probability of decoding incorrectly is averaging over all $\mathbb{E}[W_i]$ by

$$2^{-k} \sum_{i \in [0, 2^k - 1]} \mathbb{E}[W_i] \leq \gamma.$$

By an averaging argument lemma 1.6.12, there must exist a set of codewords $\{\mathbf{x}_i\}_{i \in [0, 2^k - 1]}$, such that

$$\sum_{i \in [0, 2^k - 1]} (W_i \mid X_0 = \mathbf{x}_0, \dots, X_{2^k-1} = \mathbf{x}_{2^k-1}) \leq 2^k \gamma.$$

□

The decoding failure probability in lemma 1.10.45 is obtained by averaging over all k -bit messages and all possible codeword samplings. We will obtain a stronger result than lemma 1.10.45 in theorem 1.10.46: for all $i \in [0, 2^k - 1]$,

$$(W_i \mid X_0 = \mathbf{x}_0, \dots, X_{2^k-1} = \mathbf{x}_{2^k-1}) \leq \gamma$$

holds simultaneously.

Theorem 1.10.46 (Shannon's Theorem [Sha48a, Sha48b]). *For a binary symmetric channel with parameter $p < 1/2$ and for any constants $\delta, \gamma > 0$, when n is sufficiently large:*

- *for any $k \leq n(1 - H(p) - \delta)$, there exist (k, n) encoding and decoding functions such that the probability the receiver fails to obtain the correct message is at most γ for every k -bit input message; and*
- *there are no (k, n) encoding and decoding functions with $k \geq n(1 - H(p) + \delta)$ such that the probability decoding correctly is at least γ for a k -bit input message chosen uniformly at random.*

Proof. WLOG we let $\{\mathbf{x}_i\}_{i \in [0, 2^k - 1]}$ be sorted in increasing order of W_i . Hence, for $i \in [0, 2^{k-1} - 1]$, each $\mathbf{x}_i$ has $W_i \leq 2\gamma$, or otherwise we contradict by

$$\sum_{i \in [2^{k-1}, 2^k - 1]} (W_i \mid X_0 = \mathbf{x}_0, \dots, X_{2^k-1} = \mathbf{x}_{2^k-1}) \geq 2\gamma \cdot 2^{k-1} = 2^k \gamma.$$

In this way, we prove there exist encoding and decoding functions for $k - 1$ bits messages over 2^{k-1} codewords, and the probability of each codeword being decoded incorrectly is simultaneously at most 2γ , when $k \leq n(1 - H(p) - \delta)$. Since δ and γ can be any constant, the first part of the proof is completed.

Having the first part of the proof finished, we move on to the second part of the proof, which is the converse of lemma 1.10.45. We begin by an observation in lemma 1.10.45 to build an intuition: The decoding function maps the received n bits to the only one codeword within $[[n(p - \varepsilon)], [n(p + \varepsilon)]]$ many errors, and the number of such n bits is lower bounded by

$$\sum_{k \in [[n(p - \varepsilon)], [n(p + \varepsilon)]]} \binom{n}{k} \geq \binom{n}{[np]} \geq \frac{2^{nH(p)}}{2\sqrt{n}},$$

where the last inequality holds by corollary 1.10.19. Since there are 2^k possible messages, and $k \geq n(1 - H(p) + \delta)$, on a sufficiently large n ,

$$2^k \cdot \frac{2^{nH(p)}}{2\sqrt{n}} \geq \frac{2^{n\delta}}{2\sqrt{n}} \cdot 2^n \geq 2^n.$$

Hence, given $k \geq n(1 - H(p) + \delta)$ and n sufficiently large, conditioned that the number of bits flipped by the channel is within $[[n(p - \varepsilon)], [n(p + \varepsilon)]]$, there is no such set of codewords always decode successfully, namely it is getting too dense, such that in the “unique decoding” range defined in lemma 1.10.45, there are more than 1 valid codeword.

More specifically, we can show how to rule out lemma 1.10.45. Let C_i be the set of received n -bit strings that are correctly decoded to the codeword $\mathbf{x}_i$ in lemma 1.10.45, then the probability of successful decoding is

$$\frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \sum_{\mathbf{r} \in C_i} w(\mathbf{x}_i, \mathbf{r})$$

by averaging over all 2^k messages. Since for any $\mathbf{r} \in C_i$, it is within $[[n(p - \varepsilon)], [n(p + \varepsilon)]]$ many bit flips from $\mathbf{x}_i$,

$$w(\mathbf{x}_i, \mathbf{r}) = p^{\Delta(\mathbf{x}_i, \mathbf{r})} \cdot (1 - p)^{n - \Delta(\mathbf{x}_i, \mathbf{r})} \leq p^{n(p - \varepsilon)} \cdot (1 - p)^{n(1 - p + \varepsilon)} = \left(\frac{1 - p}{p}\right)^{n\varepsilon} \cdot 2^{-nH(p)},$$

then the averaging probability is upper bounded by

$$\begin{aligned} \frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \sum_{\mathbf{r} \in C_i} w(\mathbf{x}_i, \mathbf{r}) &\leq \frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \left(\frac{1 - p}{p}\right)^{n\varepsilon} \cdot 2^{-nH(p)} \cdot |C_i| \\ &\leq \frac{1}{2^k} \cdot \left(\frac{1 - p}{p}\right)^{n\varepsilon} \cdot 2^{-nH(p)} \cdot 2^n = \left(\left(\frac{1 - p}{p}\right)^\varepsilon \cdot 2^{-\delta}\right)^n, \end{aligned} \quad (1.17)$$

where the last inequality holds by $k \geq n(1 - H(p) + \delta)$, so conditioned that

$$\left(\frac{1 - p}{p}\right)^\varepsilon \cdot 2^{-\delta} < 1, \quad (1.18)$$

then when n is sufficiently large, the probability is upper bounded by γ .

We proceed by analyzing the general case. For the set of codewords $\{\mathbf{x}_i\}_{i \in [0, 2^k - 1]}$, let S_i be the set of n bits that are uniquely decoded to $\mathbf{x}_i$, and $S_{i,1} \cap S_{i,2} = \emptyset$, $S_i = S_{i,1} \cup S_{i,2}$ be defined as follows: $S_{i,1} = \{\mathbf{r} \in S_i : |\Delta(\mathbf{x}_i, \mathbf{r}) - pn| > \varepsilon n\}$, and $S_{i,2} = \{\mathbf{r} \in S_i : |\Delta(\mathbf{x}_i, \mathbf{r}) - pn| \leq \varepsilon n\}$. Then the probability of successful decoding is

$$\frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \sum_{\mathbf{r} \in S_i} w(\mathbf{x}_i, \mathbf{r}) = \frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \sum_{\mathbf{r} \in S_{i,1}} w(\mathbf{x}_i, \mathbf{r}) + \frac{1}{2^k} \sum_{i \in [0, 2^k - 1]} \sum_{\mathbf{r} \in S_{i,2}} w(\mathbf{x}_i, \mathbf{r}),$$

where the first term is upper bounded by Chernoff's bound with the same bound in eq. (1.14), and the second term is upper bounded in the same procedure of eq. (1.17), conditioned that eq. (1.18) holds. $\square$

Exercise 1.10.47 (Exercise 10.19 [MU17]). Consider the following channel. The sender can send a symbol from $[0, 4]$, the channel introduces errors: when the symbol k is sent, the recipient receives $k + 1 \pmod 5$ with probability $1/2$, or $k - 1 \pmod 5$ with probability $1/2$. The errors on each symbol are mutually independent of each other.

Define the encoding and decoding functions for this channel by: A (j, n) encoding function maps a number in $[0, j - 1]$ into a sequence in $[0, 4]^n$. A (j, n) decoding function maps a sequence in $[0, 4]^n$ into a number in $[0, j - 1]$.

There are $(1, 1)$ encoding and decoding functions with zero probability of error, as the decoder maps 1 and 4 to 0, and maps 0 and 2 to 1. Hence at least 1 bit can be sent without error per channel use. Show that:



- There are $(5, 2)$ encoding and decoding functions with zero probability of error. Argue that more than one bit of information can be sent per use of the channel.
- If there exists (j, n) encoding and decoding functions with zero probability of error, then $n \geq \log_2 j / (\log_2 5 - 1)$.

Proof. Consider the $(5, 2)$ encoding and decoding function. The $(5, 2)$ encoder encodes a with a pair $(a, 2a \pmod 5)$. The $(5, 2)$ decoder performs by the following lookup table

encode	received	decode to
(0, 0)	(1, 1), (1, 4), (4, 1), (4, 4)	0
(1, 2)	(2, 3), (2, 1), (0, 3), (0, 1)	1
(2, 4)	(3, 0), (3, 3), (1, 0), (1, 3)	2
(3, 1)	(4, 2), (4, 0), (2, 2), (2, 0)	3
(4, 3)	(0, 4), (0, 2), (3, 4), (3, 2)	4

The encoding scheme can also be instantiated by $(a, 3a \pmod 5)$ ⁹.

In fact, the channel operates over a cycle graph C_5 , mapping a to $a \pm 1 \pmod 5$ each with probability $1/2$. Moreover, since we want a set of codewords avoiding overlapping received symbols on the decoding side, we can construct the following graph for the single coordinate $(1, 1)$ encoding and decoding, by pairing up symbols bringing “confusion” after transmitted through the channel. Hence, $\{0, 1\}$ is an independent set of size 2 over the C_5 illustrated as follows, and it happens to be the codewords for $(1, 1)$ encoding and decoding.




On a pair of C_5 confusion graphs, we want to define a graph product such that any neighboring (i, j) brings overlaps on the decoding side. On a pair of (u, v) and (u', v') , they are adjacent iff

- $u = u', v$ and v' are adjacent on C_5 confusion graph;
- $v = v', u$ and u' are adjacent on C_5 confusion graph;
- u and u', v and v' are both adjacent on C_5 confusion graph.

We then define the symbol $\boxtimes$ such that $C_5 \boxtimes C_5$ is illustrated as follows:

⁹The story is, I brute-forced adding constant, additive inverse, multiplicative inverse, and finally scalar multiplication.



The new codewords form an independent set of size 5 on $C_5 \boxtimes C_5$, each C_5 is ordered as 0, 2, 4, 1, 3 so that adjacency corresponds to difference $\pm 2 \pmod 5$.

When it comes to general (j, n) encoding and decoding functions, supposing it exists, then there must be $j \cdot 2^n \leq 5^n$, such that for each of the j values, the decoding boundary is 2^n , and the total area of decoding is no more than 5^n , or confusion must exist, and hence

$$n \geq \frac{\log_2 j}{\log_2 5 - 1}.$$

□

Remark 1.10.48. The exercise 1.10.47 is the classical zero-error capacity viewpoint of Shannon [Sha56]. The pentagon case C_5 was solved by Lovász [Lov79].

Exercise 1.10.49 (Exercise 10.20 [MU17]). A binary erasure channel transfers a sequences of n bits. Each bit either arrives successfully without error or fails to arrive, and is replaced with a “?” symbol. Failure occur independently with probability p . Define (k, n) encoding and decoding functions for BEC similar to ones in the BSC, except that $\text{Dec} : \{0, 1, ?\}^n \rightarrow \{0, 1\}^k \cup \{\perp\}$. Prove that, for any $p > 0$ and any constants $\delta, \gamma > 0$, if n is sufficiently large, then there exists (k, n) encoding and decoding functions with $k \leq n(1 - p - \delta)$ such that the probability that the receiver fails to obtain the correct message is at most γ for every possible k -bit message.

Proof. The proof architecture is similar to lemma 1.10.45 and theorem 1.10.46, with a same setup of the randomized codewords, and the same encoding function. The decoder finds a codeword that has within $\lceil n(p - \varepsilon) \rceil, \lfloor n(p + \varepsilon) \rfloor$ number of "?", and the other symbols must match. If there is one such codeword, output the corresponding message; otherwise, reject by $\perp$.

We reuse a few notations in lemma 1.10.45: Applying Hamming weight Δ over $\mathbf{x}_i \in \{0, 1\}^n$ and $\mathbf{s}_i \leftarrow \mathbf{x}_i \oplus \mathbf{r}_i$ where $\mathbf{r}_i \in \{0, ?\}^n$, and $\oplus$ overrides ? to a symbol in $\{0, 1\}$. Then the weight defined in between $\mathbf{x}_i$ and $\mathbf{s}_i$ is

$$w(\mathbf{x}_i, \mathbf{s}_i) = p^{\Delta(\mathbf{x}_i, \mathbf{s}_i)} \cdot (1 - p)^{n - \Delta(\mathbf{x}_i, \mathbf{s}_i)}.$$

Let $\{S_i\}_{i \in [0, 2^k - 1]}$ be the set of $\{0, 1, ?\}^n$ decodable to X_i , $\{W_i\}_{i \in [0, 2^k - 1]}$ be the set of probabilities that X_i fails in decoding correctly, where

$$W_i = \sum_{\mathbf{r} \in \{0, ?\}^n, X_i \oplus \mathbf{r} \notin S_i} w(X_i, X_i \oplus \mathbf{r}).$$

It can also be expressed in the way that, let $1_{i, \mathbf{r}}$ be an indicator 0/1 random variable that is 1 iff $X_i \oplus \mathbf{r} \notin S_i$, then

$$W_i = \sum_{\mathbf{r} \in \{0, ?\}^n} 1_{i, \mathbf{r}} \cdot w(X_i, X_i \oplus \mathbf{r}) = \sum_{\mathbf{r} \in T_1} 1_{i, \mathbf{r}} \cdot w(X_i, X_i \oplus \mathbf{r}) + \sum_{\mathbf{r} \in T_2} 1_{i, \mathbf{r}} \cdot w(X_i, X_i \oplus \mathbf{r}),$$

by letting $T_1 = \{\mathbf{r} \in \{0, ?\}^n : |\Delta(\mathbf{0}, \mathbf{r}) - np| > n\varepsilon\}$ and $T_2 = \{\mathbf{r} \in \{0, ?\}^n : |\Delta(\mathbf{0}, \mathbf{r}) - np| \leq n\varepsilon\}$.

WLOG we analyze X_0 , as by symmetry the same analysis hold for all other X_i .

The first term is bounded in the same way as eq. (1.14) via Chernoff's bound. Let $\{Z_i\}_{i \in [1, n]}$ be the independent Bernoulli trials with success probability p , then $Z = \sum_{i \in [1, n]} Z_i$ is the number of "?"s, for $0 < \varepsilon < p$,

$$\sum_{\mathbf{r} \in T_1} 1_{0, \mathbf{r}} \cdot w(X_0, X_0 \oplus \mathbf{r}) = \sum_{\mathbf{r} \in T_1} w(X_0, X_0 \oplus \mathbf{r}) = \Pr[|Z - np| > n\varepsilon] \leq \left(\frac{e^\theta}{(1 + \theta)^{1 + \theta}} \right)^{np} + \left(\frac{e^{-\theta}}{(1 - \theta)^{1 - \theta}} \right)^{np} \leq 2 \exp\left(-\frac{\mu\theta^2}{3}\right),$$

For any threshold $0 < \varepsilon < p$, there is n sufficiently large such that $2 \exp(-\mu\theta^2/3) \leq \gamma/2$.

It suffices to argue with respect to $X_0 = \mathbf{0}$ for the second term by the same symmetry discussed in lemma 1.10.45, then the random variable

$$\sum_{\mathbf{r} \in T_2} 1_{0, \mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})$$

is dependent on $\{X_i\}_{i \in [1, 2^k - 1]}$. For $\mathbf{r} \in T_2$, averaging over $\{X_i\}_{i \in [1, 2^k - 1]}$,

$$\mathbb{E}[1_{0, \mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})] = w(\mathbf{0}, \mathbf{r}) \cdot \Pr[\exists i \in [1, 2^k - 1], |\Delta(X_i, \mathbf{r}) - np| \leq n\varepsilon].$$

Supposing we have

$$\Pr[\exists i \in [1, 2^k - 1], |\Delta(X_i, \mathbf{r}) - np| \leq n\varepsilon] \leq \frac{\gamma}{2},$$

then the second term can be upper bounded by

$$\sum_{\mathbf{r} \in T_2} \mathbb{E}[1_{0, \mathbf{r}} \cdot w(\mathbf{0}, \mathbf{r})] \leq \sum_{\mathbf{r} \in T_2} w(\mathbf{0}, \mathbf{r}) \cdot \frac{\gamma}{2} \leq \frac{\gamma}{2}.$$

We continue by upper bounding the number of possible n bits to derive $\mathbf{r}$ by $2^{\Delta(\mathbf{0}, \mathbf{r})} \leq 2^{n(p + \varepsilon)}$, then the probability of a particular codeword X_i with $i > 0$ being able to produce $\mathbf{r}$ is at most $2^{n(p + \varepsilon - 1)}$. Union bounding over all other codewords, the probability of other codeword failing decoding $\mathbf{0}$ is upper bounded by

$$(2^k - 1) \cdot 2^{n(p + \varepsilon - 1)} \leq 2^k \cdot 2^{n(p + \varepsilon - 1)} \leq 2^{n(\varepsilon - \delta)},$$

where the last inequality is by $k \leq n(1 - p - \delta)$. Since $2^{n(\varepsilon - \delta)} \rightarrow 0$ as $n \rightarrow \infty$ if $\delta > \varepsilon$, then it is at most $\gamma/2$ on an n . By symmetry, we conclude

$$\sum_{\mathbf{r} \in T_2} \mathbb{E}[1_{0, \mathbf{r}} \cdot w(X_0, X_0 \oplus \mathbf{r})] \leq \frac{\gamma}{2}.$$

Combining the separated cases, we have

$$\mathbb{E}[W_0] = \sum_{\mathbf{r} \in T_1} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(X_0, X_0 \oplus \mathbf{r})] + \sum_{\mathbf{r} \in T_2} \mathbb{E}[1_{0,\mathbf{r}} \cdot w(X_0, X_0 \oplus \mathbf{r})] \leq \gamma.$$

By symmetry and linearity of expectation,

$$\sum_{i \in [0, 2^k - 1]} \mathbb{E}[W_i] \leq 2^k \gamma.$$

By the same ranking argument in theorem 1.10.46, there exists a set of codewords with at least 2^{k-1} codewords with $W_i \leq 2\gamma$, and we choose the half of the set of codewords with small W_i , and we complete the proof. $\square$

Exercise 1.10.50 (Exercise 10.21 [MU17]). In lemma 1.10.45 and theorem 1.10.46, we let decoder find a codeword with Hamming distance to the bits received within $[n(p - \varepsilon), n(p + \varepsilon)]$. Instead, let decoder find the codeword by looking for the codeword with least number of differences, and break ties arbitrarily. Show how to modify the proof for lemma 1.10.45 and theorem 1.10.46 to obtain a similar result.

Proof. Just use the Hamming ball of radius $\lfloor n(p + \varepsilon) \rfloor$ around the codeword. The “too far” Chernoff case still decays exponentially fast, and the “in radius” case should approach 0 as $n \rightarrow \infty$. $\square$

1.11 The Monte Carlo Method

1.11.1 The Monte Carlo Method

1.11.2 DNF Counting

1.11.3 From Approximate Sampling to Approximate Counting

1.11.4 The Markov Chain Monte Carlo Method (MCMC)

TODO: ...

1.12 Coupling of Markov Chains

The motivation of coupling is to decide the rate of convergence to an ergodic Markov chain's stationary distribution.

1.12.1 Variation Distance and Mixing Time

Definition 1.12.1 (Variation Distance). The variation distance between two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ on a countable state space S is given by

$$d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2) = \frac{1}{2} \sum_{x \in S} |\mathcal{D}_1(x) - \mathcal{D}_2(x)|.$$

The factor $1/2$ in definition 1.12.1 guarantees that the variation distance is in $[0, 1]$.

Lemma 1.12.2. For any $A \subseteq S$, let $\mathcal{D}(A) = \sum_{x \in A} \mathcal{D}(x)$, then

$$d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2) = \max_{A \subseteq S} |\mathcal{D}_1(A) - \mathcal{D}_2(A)|.$$

Proof. Let S^+ be all the $x \in S$ that $\mathcal{D}_1(x) \geq \mathcal{D}_2(x)$, and $S^- = S \setminus S^+$, such that $\forall x \in S^-, \mathcal{D}_1(x) < \mathcal{D}_2(x)$. Then

$$(\mathcal{D}_1(S^+) - \mathcal{D}_2(S^+)) + (\mathcal{D}_2(S^-) - \mathcal{D}_1(S^-)) = \sum_{a \in S} |\mathcal{D}_1(a) - \mathcal{D}_2(a)| = 2d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2).$$

Since $\mathcal{D}_1(S^+) + \mathcal{D}_1(S^-) = \mathcal{D}_2(S^+) + \mathcal{D}_2(S^-) = 1$,

$$\mathcal{D}_1(S^+) - \mathcal{D}_2(S^+) = \mathcal{D}_2(S^-) - \mathcal{D}_1(S^-) = d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2).$$

Moreover, by the definition of S^+ and S^- , such that $\forall x \in S^+, \mathcal{D}_1(x) \geq \mathcal{D}_2(x)$, then $\forall A \subseteq S$,

$$\mathcal{D}_1(A) - \mathcal{D}_2(A) \leq \mathcal{D}_1(S^+) - \mathcal{D}_2(S^+).$$

Since $\mathcal{D}_1(S^+) - \mathcal{D}_2(S^+)$ is maximized, so is $\mathcal{D}_2(S^-) - \mathcal{D}_1(S^-)$ by symmetry, the proof is concluded. $\square$

Remark 1.12.3. The alternative characterization of variation distance by lemma 1.12.2 gives a strong statement, that supposing $d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2) < \varepsilon$, then for any $x \in S$, $|\mathcal{D}_1(x) - \mathcal{D}_2(x)| < \varepsilon$.

Corollary 1.12.4. Let $m(a) = \min(\mathcal{D}_1(a), \mathcal{D}_2(a))$ for all $a \in S$, then

$$d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2) = 1 - m(S).$$

Proof. The result is immediate by $d_{\text{TV}}(\mathcal{D}_1, \mathcal{D}_2) = \mathcal{D}_1(S^+) - \mathcal{D}_2(S^+)$ and $m(S) = \mathcal{D}_1(S^-) + \mathcal{D}_2(S^+)$. $\square$

Definition 1.12.5 (Mixing Time). Let π be the stationary distribution of an ergodic Markov chain with state space S . Let p_x^t represent the distribution of the state of the chain starting at state x after t steps. We define

$$\begin{aligned} \Delta_x(t) &= d_{\text{TV}}(\mathcal{D}_\pi, p_x^t) \\ \Delta(t) &= \max_{x \in S} \Delta_x(t) = \max_{x \in S} d_{\text{TV}}(\mathcal{D}_\pi, p_x^t), \end{aligned}$$

that is, $\Delta_x(t)$ is the variation distance between the stationary distribution and p_x^t , and $\Delta(t)$ is the maximum of $\Delta_x(t)$ over all states x . We also define

$$\begin{aligned} \tau_x(\varepsilon) &= \min\{t \mid \Delta_x(t) \leq \varepsilon\} \\ \tau(\varepsilon) &= \max_{x \in S} \tau_x(\varepsilon) = \max_{x \in S} \min\{t \mid \Delta_x(t) \leq \varepsilon\}, \end{aligned}$$

that is, $\tau_x(\varepsilon)$ is the first step t , at which the variation distance between p_x^t and the stationary distribution is at most ε , and $\tau(\varepsilon)$ is the maximum of $\tau_x(\varepsilon)$ over all states x .

While $\tau(\varepsilon)$ is a function of ε , it is generally called the mixing time of the Markov chain.

Definition 1.12.6 (Rapidly Mixing [JS89]). A Markov chain with state space S is rapidly mixing if its mixing time $\tau(\varepsilon)$ is $\text{poly}(\log(1/\varepsilon), \log|S|)$. We view $\log|S|$ as the problem size, since a state can be encoded using $O(\log|S|)$ bits.

1.12.2 Coupling

Coupling of Markov chains is a general technique for bounding the mixing time of a Markov chain.

Definition 1.12.7 (Markov Chain Coupling). A coupling of a Markov chain M_t with state space S is a Markov chain $Z_t = (X_t, Y_t)$ on state space $S \times S$ such that

$$\begin{aligned}\Pr[X_{t+1} = x' \mid Z_t = (x, y)] &= \Pr[M_{t+1} = x' \mid M_t = x]; \\ \Pr[Y_{t+1} = y' \mid Z_t = (x, y)] &= \Pr[M_{t+1} = y' \mid M_t = y].\end{aligned}$$

Remark 1.12.8. A coupling consists of two copies of the Markov chain M_t running simultaneously, each copy behaves exactly like the original Markov chain in terms of its transition probabilities. The joint distribution preserves the marginal distribution of each copy, but the two copies may be correlated.

Remark 1.12.9. To show a chain is close to stationarity, it suffices to couple it with a copy started from the stationary distribution. When the two chains reach the same state (the *coupling time* T), they can be made to evolve identically thereafter. In particular, if they have coupled by time T , then the first chain is exactly equal to a stationary sample at time T . Thus, the only obstruction to being close to stationarity is the probability that they have not yet met.

Lemma 1.12.10 (Mixing-Time Coupling Lemma). Let $Z_t = (X_t, Y_t)$ be a coupling for a Markov chain M_t on a state space S . Suppose that there exists a T such that, for every $x, y \in S$,

$$\Pr[X_T \neq Y_T \mid X_0 = x, Y_0 = y] \leq \varepsilon, \quad (1.19)$$

then $\tau(\varepsilon) \leq T$, that the variation distance between the distribution over S after T steps and $\mathcal{D}_\pi$ is at most ε .

Proof. Consider the coupling when initial state x is fixed and $y \leftarrow \mathcal{D}_\pi$. We first prove the following core lemma.

Lemma 1.12.11. In the coupling $Z_t = (X_t, Y_t)$, where initial state x is fixed and $y \leftarrow \mathcal{D}_\pi$, then

$$\Delta_x(T) \leq \Pr[X_T \neq Y_T].$$

Proof. We begin by lower bounding $p_x^T(A)$, where $A \subseteq S$, by

$$\begin{aligned}p_x^T(A) &= \Pr[X_T \in A] \\ &= \Pr[X_T \in A \wedge Y_T = X_T] + \Pr[X_T \in A \wedge Y_T \neq X_T] \\ &\geq \Pr[X_T \in A \wedge Y_T = X_T] = \Pr[Y_T \in A \wedge Y_T = X_T],\end{aligned}$$

where the last equality holds by the same event. We continue the lower bounding by

$$\begin{aligned}p_x^T(A) &\geq \Pr[Y_T \in A \wedge Y_T = X_T] \\ &= 1 - \Pr[Y_T \notin A \vee Y_T \neq X_T] \\ &\geq 1 - \Pr[Y_T \notin A] - \Pr[Y_T \neq X_T] = \Pr[Y_T \in A] - \Pr[Y_T \neq X_T] \\ &= \mathcal{D}_\pi(A) - \Pr[Y_T \neq X_T],\end{aligned}$$

where the third inequality holds by union bound. Symmetric to the prior derivation, we have

$$p_x^T(S \setminus A) \geq \mathcal{D}_\pi(S \setminus A) - \Pr[Y_T \neq X_T],$$

or equivalently,

$$\mathcal{D}_\pi(A) \geq p_x^T(A) - \Pr[Y_T \neq X_T].$$

Now that for all $A \subseteq S$, we have

$$|p_x^T(A) - \mathcal{D}_\pi(A)| \leq \Pr[X_T \neq Y_T],$$

then by the alternative characterization of variation distance lemma 1.12.2, we complete the proof by

$$\Delta_x(T) = d_{TV}(p_x^T, \mathcal{D}_\pi) = \max_{A \subseteq S} |p_x^T(A) - \mathcal{D}_\pi(A)| \leq \Pr[X_T \neq Y_T].$$

□

Since $\Pr[X_T \neq Y_T] \leq \varepsilon$ by eq. (1.19), then by lemma 1.12.11, $\Delta_x(T) \leq \varepsilon$ for all $x \in S$. Hence, by definition 1.12.5,

$$T \geq \tau_x(\varepsilon) = \min\{t \mid \Delta_x(t) \leq \varepsilon\}$$

holds for all $x \in S$, and therefore $\tau(\varepsilon) \leq T$. $\square$

Example 1.12.12 (Coupling for Card Shuffling). A card-shuffling Markov chain for n cards is defined over $n!$ states, where a state transitions to another by choosing a position $j \xleftarrow{R} [1, n]$, and moving the j^{th} card to the deck top. It is obvious that the stationary distribution is the uniform distribution.

To analyze the mixing time $\tau(\varepsilon)$ of the chain by lemma 1.12.10, we obtain 2 copies of the chain, and the coupling is constructed by sampling a card C uniformly at random, and moving C to the deck top in both copies. The coupling is valid: recall remark 1.12.8, that the marginal distribution is preserved, while the joint distribution may be correlated.

Moreover, when C is moved to the deck top, then it is always in the same position in both copies, and the copies are coupled once every card has been moved to the top for at least once. Let $Z_t = (X_t, Y_t)$ be the coupling, then $\Pr[X_T \neq Y_T]$ is the probability that after T steps, there exists cards still not chosen.

We continue by analyzing the coupon collector's problem: with probability $(1 - n^{-1})^N$, a specific card is not chosen after N trials. By theorem 1.5.23, let $N = n \ln n + cn$, and by union bounding over all n cards,

$$\Pr[X_N \neq Y_N] \leq n \left(1 - \frac{1}{n}\right)^N \leq n \exp\left(-\frac{N}{n}\right) = e^{-c},$$

where the second inequality holds by $1 + x \leq e^x$. When $N = n \ln(n/\varepsilon)$, we have $\Pr[X_N \neq Y_N] \leq \varepsilon$.

We conclude by lemma 1.12.10, that the variation distance between the uniform distribution and the distribution of the state of the chain after $n \ln(n/\varepsilon)$ steps is at most ε , and $\tau(\varepsilon) \leq n \ln(n/\varepsilon)$.

Exercise 1.12.13 (Exercise 12.2 [MU17]). Considering the Markov chain for shuffling n cards as in example 1.12.12. Suppose that, instead of running the chain for a fixed number of steps, we stop the chain at the first step where every card has been moved to the top at least once. Show that, at the stopping time, the state of the chain is uniformly distributed over $n!$ possible permutations of the cards.

Proof. By the coupling in example 1.12.12, fix an initial state x for X_0 and let Y_0 be sampled uniformly from $n!$ decks. Let N be the first time that every card has been moved to the top at least once, then $X_N = Y_N$. We conclude that X_N is uniformly distributed over $n!$ possible decks. $\square$

Exercise 1.12.14 (Exercise 12.3 [MU17]). Considering the Markov chain for shuffling n cards as in example 1.12.12. Show that if the chain runs for only $(1 - \varepsilon)n \ln n$ steps for some constant $\varepsilon > 0$, then the variation distance is $1 - o(1)$.

Proof. We started by lower bounding $d_{TV}(p_x^N, \mathcal{D}_\pi)$, where $N = (1 - \varepsilon)n \ln n$, x be an initial state. Since

$$d_{TV}(p_x^N, \mathcal{D}_\pi) = \max_{A \subseteq S} |p_x^N(A) - \mathcal{D}_\pi(A)| \geq |p_x^N(B) - \mathcal{D}_\pi(B)|$$

for all $B \subseteq S$, the lower bound can be derived from lower bounding $|p_x^N(B) - \mathcal{D}_\pi(B)|$.

Let the cards be initially unshuffled and ordered, then we define a subset of states by A_ℓ such that the last ℓ cards preserve the relative order. We let X_i be the 0-1 indicator random variable for the i^{th} card not being picked, and $X = \sum_i X_i$ be the random variable for the number of untouched cards.

Since $X \geq \ell$ leads to A_ℓ occurring, then $\{X \geq \ell\} \subseteq A_\ell$, and

$$\Pr[X \geq \ell] \leq p_x^N(A_\ell).$$

By the linearity of expectation,

$$\mathbb{E}[X] = \sum_{i \in [1, n]} \mathbb{E}[X_i] = n \left(1 - \frac{1}{n}\right)^N \leq n \cdot \exp(-(1 - \varepsilon) \ln n) = n^\varepsilon.$$

We let $\ell = n^{\varepsilon/2}$, and hence for any sufficiently large n ,

$$\Pr[X < \ell] \leq \Pr\left[X \leq \frac{1}{2}n^{\varepsilon}\right].$$

Moreover, for $i \neq j \in [1, n]$,

$$\text{Cov}[X_i, X_j] = \mathbb{E}[X_i X_j] - \mathbb{E}[X_i] \mathbb{E}[X_j] = \left(1 - \frac{2}{n}\right)^N - \left(1 - \frac{1}{n}\right)^{2N} < 0,$$

hence

$$\text{Var}[X] = \sum_{i \in [1, n]} \text{Var}[X_i] + 2 \sum_{j < k \in [1, n]} \text{Cov}[X_j, X_k] \leq \sum_{i \in [1, n]} \text{Var}[X_i] \leq \sum_{i \in [1, n]} \mathbb{E}[X_i] = \mathbb{E}[X],$$

then by Chebyshev's inequality,

$$\Pr[X < \ell] \leq \Pr\left[X \leq \frac{1}{2}n^{\varepsilon}\right] \leq \frac{4\text{Var}[X]}{n^{2\varepsilon}} \leq \frac{4}{n^{\varepsilon}}.$$

On the other hand, $\mathcal{D}_{\pi}(A_{\ell}) = (\ell!)^{-1} = (n^{\varepsilon/2})^{-1}$. Hence,

$$d_{\text{TV}}(p_x^N, \mathcal{D}_{\pi}) = \max_{A \subseteq S} |p_x^N(A) - \mathcal{D}_{\pi}(A)| \geq |p_x^N(A_{\ell}) - \mathcal{D}_{\pi}(A_{\ell})| \geq 1 - \frac{4}{n^{\varepsilon}} - \frac{1}{n^{\varepsilon/2!}},$$

which completes the proof for $d_{\text{TV}}(p_x^N, \mathcal{D}_{\pi}) = 1 - o(1)$. $\square$

Example 1.12.15 (Coupling for Random Walks on the Hypercube). Consider the following Markov chain defined on a n -dimensional hypercube. At each step, it picks a coordinate $i \xleftarrow{R} [1, n]$. The new state $\mathbf{x}'$ is obtained from the current state $\mathbf{x}$ by keeping all coordinates the same, except possibly for x_i . x_i is set to 0 with probability 1/2, and to 1 with probability 1/2. The stationary distribution of the Markov chain is uniform distribution.

A coupling for the Markov chain can be obtained by enforcing the same move on the two copies: both changing x_i to 0 or 1. The marginal probability of the coupling still follows the transition probability of the chain.

Moreover, once both copies make the same move on x_i , they are surely agreeing on the i^{th} coordinate, and they will couple after all n coordinates have each been chosen at least once. Immediate from the coupon collector's problem, and by the same analysis in example 1.12.12, $\tau(\varepsilon) \leq n \ln(n/\varepsilon)$.

Exercise 1.12.16 (Exercise 12.11 [MU17]). Show that the Markov chain for sampling all independent sets of fixed size $k \leq n/(3\Delta + 3)$ in a graph of n vertices and maximum degree Δ is ergodic and has a uniform stationary distribution. A move is made from the independent set X_t by choosing a vertex $v \xleftarrow{R} X_t$ and $w \xleftarrow{R} V$. If $w \notin X_t$ and $X_t \setminus \{v\} \cup \{w\}$ is an independent set, then $X_{t+1} = X_t \setminus \{v\} \cup \{w\}$; otherwise, $X_{t+1} = X_t$.

Proof. Consider an independent set $A \subseteq V$ with $|A| = k$, and let $\Gamma^+(A)$ be the inclusive neighborhood of A , such that

$$|\Gamma^+(A)| \leq k(\Delta + 1).$$

Suppose we want to transition from A to B , another independent set of fixed size with k , we know that

$$|\Gamma^+(A \cup B)| \leq |\Gamma^+(A)| + |\Gamma^+(B)| \leq 2k(\Delta + 1).$$

Since $n \geq 3k(\Delta + 1)$, then $|V \setminus \Gamma^+(A \cup B)| \geq k(\Delta + 1)$.

We can greedily construct an independent set $C \subseteq V \setminus \Gamma^+(A \cup B)$: at each step choose a vertex, put it into C , and delete it together with all its neighbors. Since each step deletes at most $\Delta + 1$ vertices, we can greedily select at least k vertices from $V \setminus \Gamma^+(A \cup B)$ into C . Hence there exists an independent set C of size k .

Then for any transition in between states of independent sets of fixed size, one can first transition to an independent set of fixed size $C \subseteq V \setminus \Gamma^+(A \cup B)$, then transition to B . Hence, the irreducibility is proved.

Positive recurrence is immediate by finiteness and irreducibility, and aperiodicity is immediate by the existence of odd length loop, namely the loop back to a state itself. Hence, the ergodicity is proved by corollary 1.7.47. $\square$

Example 1.12.17 (Independent Sets of Fixed Size). For exercise 1.12.16, consider a coupling $Z_t = (X_t, Y_t)$, requiring a perfect matching M between vertices in $X_t \setminus Y_t$ and $Y_t \setminus X_t$ at each step. We choose a transition for X_t by sampling $v \stackrel{R}{\leftarrow} X_t$ and $w \stackrel{R}{\leftarrow} V$ and perform the move. If $v \in Y_t$, then perform the move over v, w , and Y_t ; otherwise, perform the move over $M(v), w$, and Y_t .

An alternative way of coupling is established by sampling $v \stackrel{R}{\leftarrow} X_t$ and $w \stackrel{R}{\leftarrow} V$ and perform the move for X_t , if $v \in Y_t$, perform the move over v, w , and Y_t ; otherwise, sample $v' \stackrel{R}{\leftarrow} Y_t \setminus X_t$ and perform the move over v', w , and Y_t .

The coupling still remains valid, by

$$\begin{aligned} \Pr[v' \text{ chosen} \wedge v' \in X_t \cap Y_t] &= \Pr[v' \text{ chosen} \mid v \stackrel{R}{\leftarrow} X_t \text{ and } v \in Y_t] \Pr[v \stackrel{R}{\leftarrow} X_t \text{ and } v \in Y_t] = \frac{1}{|X_t \cap Y_t|} \cdot \frac{|X_t \cap Y_t|}{k}, \\ \Pr[v' \text{ chosen} \wedge v' \in Y_t \setminus X_t] &= \Pr[v' \text{ chosen} \mid v \stackrel{R}{\leftarrow} X_t \text{ and } v \notin Y_t] \Pr[v \stackrel{R}{\leftarrow} X_t \text{ and } v \notin Y_t] = \frac{1}{|Y_t \setminus X_t|} \cdot \frac{|X_t \setminus Y_t|}{k}, \end{aligned}$$

such that $v \in Y_t$ and $w \in V$ are still chosen with probability $1/kn$. Note that, in the coupling, it is possible to transition in one chain and stay put in the other.

Let $D_t = |X_t \setminus Y_t|$ be the random variable for the difference size between X_t and Y_t , and D_t can change by at most 1 at a step. We want to show D_t is more likely to decrease, and we want to upper bound the probability of $D_t > 0$ for a sufficiently large t .

Suppose $D_t > 0$. In order to make $D_{t+1} = D_t + 1$, the vertex v to be removed has to be sampled from $X_t \cap Y_t$, while the vertex w to be introduced can only cause one of the two states to transition, such that¹⁰

$$w \in \Gamma^+((X_t \setminus Y_t) \cup (Y_t \setminus X_t)).$$

Then it follows that

$$\Pr[D_{t+1} = D_t + 1 \mid D_t > 0] \leq \frac{k - D_t}{k} \cdot \frac{2D_t(\Delta + 1)}{n} \quad \text{11.}$$

In order to make $D_{t+1} = D_t - 1$, the vertex v to be removed has to be sampled from $X_t \setminus Y_t$, while the vertex w to be introduced must cause both states to transition, such that

$$w \notin \Gamma^+(X_t \cup Y_t \setminus \{v, v'\}),$$

hence

$$\Pr[D_{t+1} = D_t - 1 \mid D_t > 0] \geq \frac{D_t}{k} \cdot \frac{n - (\Delta + 1)(k + D_t - 2)}{n}.$$

We thus have, for $D_t > 0$,

$$\begin{aligned} \mathbb{E}[D_{t+1} \mid D_t] &= D_t + \Pr[D_{t+1} = D_t + 1 \mid D_t > 0] - \Pr[D_{t+1} = D_t - 1 \mid D_t > 0] \\ &\leq D_t + \frac{k - D_t}{k} \cdot \frac{2D_t(\Delta + 1)}{n} - \frac{D_t}{k} \cdot \frac{n - (\Delta + 1)(k + D_t - 2)}{n} \\ &= D_t \left(1 - \frac{n - (\Delta + 1)(3k - D_t - 2)}{kn} \right) \\ &\leq D_t \left(1 - \frac{n - (\Delta + 1)(3k - 3)}{kn} \right), \end{aligned}$$

and once $D_t = 0$, $\mathbb{E}[D_{t+1} \mid D_t] = 0$. By conditional expectation,

$$\mathbb{E}[D_{t+1}] = \mathbb{E}[\mathbb{E}[D_{t+1} \mid D_t]] \leq \mathbb{E}[D_t] \left(1 - \frac{n - (\Delta + 1)(3k - 3)}{kn} \right).$$

¹⁰ w can be strengthened from causing one of the two states to transition, by observing if $w \in (X_t \setminus Y_t) \cup (Y_t \setminus X_t)$, the transition state's shared component with the state staying put is still of size $k - D_t$, by removing 1 shared element and reintroducing 1 shared element. If w is sampled over the neighbors of $(X_t \setminus Y_t) \cup (Y_t \setminus X_t)$, then the transition state is surely decreasing the size of the shared component with the state staying put.

¹¹If we strengthen the condition for elements to introduce, then the $\Delta + 1$ should be Δ by considering only the neighboring elements.

We then conclude that

$$\begin{aligned} \Pr[D_t \geq 1] &\leq \mathbb{E}[D_t] \leq \mathbb{E}[D_0] \left(1 - \frac{n - (\Delta + 1)(3k - 3)}{kn}\right)^t \\ &\leq k \left(1 - \frac{n - (\Delta + 1)(3k - 3)}{kn}\right)^t \\ &\leq k \cdot \exp\left(-t \cdot \frac{n - (\Delta + 1)(3k - 3)}{kn}\right), \end{aligned}$$

the first inequality holds from Markov inequality, the third holds by $\mathbb{E}[D_0] \leq k$, and the last holds by $1 + x \leq e^x$. It converges to 0 when $k \leq n/(3\Delta + 3)$, then in this case,

$$\tau(\varepsilon) \leq \frac{kn \ln(k/\varepsilon)}{n - (\Delta + 1)(3k - 3)}.$$

It turns out that $\tau(\varepsilon)$ is polynomial in n (which is logarithmic in the number of independent sets of fixed size k) and $\ln(\varepsilon^{-1})$, hence it is rapidly mixing.

Exercise 1.12.18 (Exercise 12.8 [MU17]). Consider the random walk on a non-bipartite, connected graph with $|V| = n$, where each vertex has the same degree $n > d > n/2$. Show that

$$\tau(\varepsilon) \leq \frac{\ln \varepsilon}{\ln(1 - (2d - n)/d)}.$$

Proof. By theorem 1.7.75, the stationary distribution of such a graph is uniform.

A coupling can be established by sampling one copy's initial state from the stationary distribution, while the other's initial state is any fixed state. Suppose they are on states $u \neq v$, then on transition, construct a perfect matching M between $\Gamma(u)$ and $\Gamma(v)$, such that $\Gamma(u) \cap \Gamma(v)$ are matched to themselves, while the other neighboring vertices are randomly 1-to-1 matched. We randomly choose 1 out of d options in M , and since $|\Gamma(u) \cap \Gamma(v)| \geq 2d - n$, then

$$\Pr[X_{t+1} = Y_{t+1} \mid X_t \neq Y_t, X_0 = x, Y_0 = y] \geq \frac{2d - n}{d}$$

for all initial states $x \neq y$, hence by induction

$$\Pr[X_t \neq Y_t \mid X_0 = x, Y_0 = y] \leq \left(1 - \frac{2d - n}{d}\right)^t = \varepsilon,$$

then

$$t = \frac{\ln \varepsilon}{\ln(1 - (2d - n)/d)}.$$

By lemma 1.12.10, $\Delta_x(t) \leq \Pr[X_t \neq Y_t]$, then it is immediate that $\tau(\varepsilon) \leq t$. $\square$

Exercise 1.12.19 (Exercise 12.9 [MU17]). Consider a Markov chain on n points $[0, n - 1]$ lying in order on a circle. At each step, the chain stays at the current point with probability $1/2$, or moves to the next point in the clockwise direction with probability $1/2$. Show that for any $\varepsilon > 0$, the mixing time $\tau(\varepsilon)$ is $O(n^2 \ln(1/\varepsilon))$.

Proof. It is clear by the cutset argument that the stationary distribution is the uniform distribution.

Consider a coupling where each copy operates oppositely, namely if one stays, the other moves to the next point. Let the distance be the number of steps to take clockwise from a chain state to the other, then it can be modeled by the gambler's ruin example 1.7.51: states are $[0, n]$, absorbing states are 0 and n , moving ± 1 with probability $1/2$.

By exercise 1.7.68, if the current distance is i , the expected number of moves is $i \cdot (n - i) \leq n^2/4$. Let T be the random variable for the number of steps before coupling, then by Markov's inequality, for all initial states $x \neq y$,

$$\Pr\left[T \geq \frac{n^2}{2} \mid X_0 = x, Y_0 = y\right] \leq \frac{1}{2}.$$

By Markov property,

$$\Pr \left[T \geq n^2 \mid T \geq \frac{n^2}{2}, X_0 = x, Y_0 = y \right] \leq \frac{1}{2},$$

and by induction,

$$\Pr \left[T \geq \frac{k}{2} n^2 \mid X_0 = x, Y_0 = y \right] \leq 2^{-k}.$$

Hence,

$$\Pr [X_T \neq Y_T \mid X_0 = x, Y_0 = y] = \Pr \left[T \geq \frac{\log(\varepsilon^{-1})}{2} n^2 \mid X_0 = x, Y_0 = y \right] \leq \varepsilon.$$

Immediate from lemma 1.12.10, we have $\tau(\varepsilon) \leq T = \log(\varepsilon^{-1})n^2/2$, hence $\tau(\varepsilon) = O(\ln(\varepsilon^{-1})n^2)$. $\square$

Exercise 1.12.20 (Exercise 12.12 [MU17]). We improve the coupling in example 1.12.17: if an attempt is made to move $v \in X_t \setminus Y_t$ to a vertex w , then the same attempt is made with the matched vertex in the other chain; if, however, if an attempt is made to move $v \in X_t \cap Y_t$, we no longer attempt to make the same move.

- Assume there are sets $S_1 = \Gamma^+(X_t \setminus Y_t)$, and $S_2 = \Gamma^+(Y_t \setminus X_t)$, each of exactly $D_t(\Delta + 1)$ distinct vertices. Assume further that $S_1 \cap S_2 = \emptyset$. Suppose that we match up elements in S_1 and S_2 in a 1-to-1 fashion.

Argue that the moves can be coupled that, when one chain attempts and fails to move v to a vertex in S_1 , it also attempts and fails to move v to the matching vertex in S_2 on the other chain.

Similarly, argue that the moves can be coupled that, when one chain attempts and succeeds in moving v to a vertex in S_1 , it also attempts and succeeds in moving v to the matching vertex in S_2 on the other chain.

Show that the coupling gives

$$\Pr [D_{t+1} = D_t + 1] \leq \frac{k - D_t}{k} \cdot \frac{D_t(\Delta + 1)}{n}. \quad (1.20)$$

- In the general case, S_1 and S_2 are not necessarily disjoint or of same size. Show that in the general case, by pairing up failing moves as much as possible, the number of choices for w that can increase D_t is $\max(|S_1|, |S_2|) \leq D_t(\Delta + 1)$. Then argue that eq. (1.20) holds in all cases.
- Use this coupling to derive a polynomial bound on $\tau(\varepsilon)$ that holds for any $k \leq n/(2\Delta + 2)$.

Proof. Consider in the optimal case, where S_1 and S_2 are $D_t(\Delta + 1)$ large, $S_1 \cap S_2 = \emptyset$, and M is the 1-to-1 mapping. Let M' be the 1-to-1 mapping between $X_t \setminus Y_t$ and $Y_t \setminus X_t$. The chain is coupled by sampling $v \xleftarrow{R} X_t$ and $w \xleftarrow{R} V$:

- If $v \in X_t \setminus Y_t$, then let $v' = M'(v)$ for Y_t ; otherwise, let $v' = v$.
- If $w \in S_1$, then let $w' = M(w)$ for Y_t ; if $w \in S_2$, then let $w' = M^{-1}(w)$ for Y_t ; otherwise, let $w' = w$.

Perform the move removing v for w on X_t and removing v' for w' on Y_t .

When we want to remove $v \in X_t \cap Y_t$ for chain X_t , the transition fails when $w \in \Gamma^+(X_t \setminus \{v\})$, hence the transitions for X_t fails when $w \in S_1 = \Gamma^+(X_t \setminus Y_t)$. Symmetrically, by the mapping M , if $w' = M(w) \in S_2$, transition for Y_t fails.

Conversely, when $v \in X_t \cap Y_t$, for chain X_t , transition succeeds when $w \in V \setminus \Gamma^+(X_t \setminus \{v\})$, hence the transition for X_t succeeds when $w \in S_2$. Symmetrically, by mapping M , if $w' = M^{-1}(w) \in S_1$, transition for Y_t succeeds.

In order for $D_{t+1} = D_t + 1$, it must be at time t the vertex v chosen from $X_t \cap Y_t$, and w must be chosen from S_2 , or otherwise either $w = w'$, or $w \in S_1$ fails to transition. Hence, by the same upper bound argument in example 1.12.17,

$$\Pr [D_{t+1} = D_t + 1] \leq \frac{|X_t \cap Y_t|}{k} \cdot \frac{|S_2|}{n} = \frac{k - D_t}{k} \cdot \frac{D_t(\Delta + 1)}{n}.$$

Moving on the general case. We reuse previous S_1, S_2, M' , and let M be the mapping between S_1 and S_2 , such that $S_1 \cap S_2$ are mapped to itself, bijectively map elements between S_1 and S_2 as many as possible, and the remaining elements in the larger set map to itself.

In this way, if $v \in X_t \cap Y_t$, failing moves of moving v for $w \in S_1$ in X_t and moving v for $w' \in S_2$ in Y_t are maximally paired. WLOG let S_2 be the larger set, as the symmetric case can be analyzed in the same way.

- If $w \in M(S_1 \setminus S_2)$, then the successful moves are also maximally paired.
- If $w \in (S_2 \setminus S_1) \setminus M(S_1 \setminus S_2)$, then w is mapped to itself because S_2 is larger, only X_t transitions successfully.
- Otherwise, for $w \in V \setminus (S_1 \cup S_2)$, then either both chains transition to same w , or neither chain transitions.

D_t increases only in the first 2 cases for at most $\max(|S_1|, |S_2|) \leq D_t(\Delta + 1)$ elements. Hence eq. (1.20) holds by

$$\Pr[D_{t+1} = D_t + 1] \leq \frac{|X_t \cap Y_t|}{k} \cdot \frac{\max(|S_1|, |S_2|)}{n} = \frac{k - D_t}{k} \cdot \frac{D_t(\Delta + 1)}{n}.$$

To decrease D_t , it has to be $u \in X_t \setminus Y_t$ moving to $w \in V \setminus \Gamma^+(X_t \cup Y_t \setminus \{u, u'\})$ in X_t . Continuing from example 1.12.17,

$$\Pr[D_{t+1} = D_t - 1 \mid D_t > 0] \geq \frac{D_t}{k} \cdot \frac{n - (\Delta + 1)(k + D_t - 2)}{n}.$$

By the same analysis,

$$\begin{aligned} \mathbb{E}[D_{t+1} \mid D_t] &= D_t + \Pr[D_{t+1} = D_t + 1 \mid D_t > 0] - \Pr[D_{t+1} = D_t - 1 \mid D_t > 0] \\ &\leq D_t + \frac{k - D_t}{k} \cdot \frac{D_t(\Delta + 1)}{n} - \frac{D_t}{k} \cdot \frac{n - (\Delta + 1)(k + D_t - 2)}{n} \\ &= D_t \left(1 - \frac{n - (\Delta + 1)(2k - 2)}{kn} \right), \end{aligned}$$

then by the same Markov inequality argument in example 1.12.17,

$$\begin{aligned} \Pr[D_t \geq 1] &\leq \mathbb{E}[D_t] = \mathbb{E}[D_0] \left(1 - \frac{n - (\Delta + 1)(2k - 2)}{kn} \right)^t \\ &\leq k \left(1 - \frac{n - (\Delta + 1)(2k - 2)}{kn} \right)^t \\ &\leq k \cdot \exp \left(-t \cdot \frac{n - (\Delta + 1)(2k - 2)}{kn} \right), \end{aligned}$$

which converges to 0 when $k \leq n/(2\Delta + 2)$, and we have

$$\tau(\varepsilon) \leq \frac{kn \ln(k/\varepsilon)}{n - (\Delta + 1)(2k - 2)}.$$

□

Remark 1.12.21. The exercise 1.12.20 is a finer analysis of example 1.12.17 derived from [BD97].

Exercise 1.12.22 (Exercise 12.14 [MU17]). Consider the following variant of shuffling a deck of n cards: at each step, two cards are chosen independently and uniformly at random from the deck, and their positions are exchanged.

- Argue the following being equivalent: at each step, a specific card is chosen uniformly at random from the deck, and a position $i \xleftarrow{R} [1, n]$ sampled independently. The i^{th} card is exchanged with the specific chosen card.
- Consider the coupling where the two choices of card and position are the same for both copies of the chain. Let X_t be the number of cards whose positions differ in the two copies of the chain. Show X_t is nonincreasing,

$$\Pr[X_{t+1} \leq X_t - 1 \mid X_t > 0] \geq \left(\frac{X_t}{n} \right)^2,$$

and the expected time until $X_t = 0$ is $O(n^2)$, regardless of the initial states of the two chains.

Proof. It suffices to show that, for a specific card, sampling $i \xleftarrow{R} [1, n]$ and the i^{th} card in the deck being this specific card has the same probability as choosing this card out of uniform randomness.

For every coupled move,

- If the chosen card is in the same position in both decks, then $X_{t+1} = X_t$.
- Otherwise, if the chosen card is in different positions in both decks,
 - If the i^{th} position cards are the same in both decks, then $X_{t+1} = X_t$.
 - Otherwise, the coupled move at least introduces 1 correction, hence $X_{t+1} \leq X_t - 1$.

Therefore the coupled move makes X_t nonincreasing over time, and correction only happens in both the card chosen and the position chosen being in disagreeing section, hence

$$\Pr[X_{t+1} \leq X_t - 1 \mid X_t > 0] \geq \left(\frac{X_t}{n}\right)^2.$$

Conditioned on $X_t = k$, the steps to decrease X_t is geometrically distributed, with success probability $(k/n)^2$. Let Y_k be the steps to decrease from k differences, and $Y = \sum_k Y_k$ be the steps before $X_t = 0$,

$$\mathbb{E}[Y] \leq \sum_{k \in [1, n]} \left(\frac{n}{k}\right)^2 = O(n^2),$$

which completes the proof. $\square$

1.12.3 Variation Distance is Nonincreasing

Lemma 1.12.23 (Optimal Coupling Lemma). *Given distribution σ_X and σ_Y on a state space S , let $Z = (X, Y)$ be a random variable on $S \times S$, where X is distributed according to σ_X and Y is distributed according to σ_Y . Then*

$$\Pr[X \neq Y] \geq d_{\text{TV}}(\sigma_X, \sigma_Y).$$

Moreover, there exists a joint distribution $Z = (X, Y)$, where X is distributed according to σ_X and Y is distributed according to σ_Y , for which the equality holds.

Proof. The proof idea of the first part is the same as lemma 1.12.11, where we replace $\mathcal{D}_\pi$ with σ_Y .

Alternatively, we can prove by

$$\begin{aligned} \Pr[X = x] &= \Pr[X = x \wedge X = Y] + \Pr[X = x \wedge X \neq Y] \\ &= \Pr[Y = x \wedge X = Y] + \Pr[X = x \wedge X \neq Y] \\ &\leq \Pr[Y = x] + \Pr[X = x \wedge X \neq Y]. \end{aligned}$$

By symmetry, we have

$$\begin{aligned} \Pr[X = x] - \Pr[Y = x] &\leq \Pr[X = x \wedge X \neq Y], \\ \Pr[Y = x] - \Pr[X = x] &\leq \Pr[Y = x \wedge X \neq Y]. \end{aligned}$$

Let $S^+, S^- \subseteq S$ be the sets, such that $S^- = S \setminus S^+$, $\forall a \in S^+$, $\sigma_X(a) \geq \sigma_Y(a)$, and $\forall a \in S^-$, $\sigma_X(a) < \sigma_Y(a)$, then

$$\begin{aligned} 2d_{\text{TV}}(\sigma_X, \sigma_Y) &= (\sigma_X(S^+) - \sigma_Y(S^+)) + (\sigma_Y(S^-) - \sigma_X(S^-)) \\ &\leq \sum_{a \in S^+} \Pr[X = a \wedge X \neq Y] + \sum_{a \in S^-} \Pr[Y = a \wedge X \neq Y] \\ &\leq \Pr[X \neq Y] + \Pr[X \neq Y]. \end{aligned}$$

Or alternatively, by corollary 1.12.4, for each $a \in S$, both X and Y are a with probability at most $m(a)$

$$\Pr[X = Y = a] \leq m(a),$$

where the equality holds by matching X and Y as much as possible, hence we complete the proof by corollary 1.12.4,

$$\Pr[X = Y] \leq m(S) = 1 - d_{TV}(\sigma_X, \sigma_Y).$$

The equality case follows by the construction matching X and Y as much as possible in the joint distribution. If $m(S) = 1$, the equality holds by X and Y are same distribution. Otherwise, let $Z = (X, Y)$,

$$\Pr[X = x, Y = y] = \begin{cases} m(x) & x = y, \\ \frac{(\sigma_X(x) - m(x))(\sigma_Y(y) - m(y))}{1 - m(S)} & x \neq y. \end{cases}$$

The equality holds by corollary 1.12.4,

$$\Pr[X \neq Y] = 1 - m(S) = d_{TV}(\sigma_X, \sigma_Y).$$

The marginal distribution is preserved as follows: if $m(x) = \sigma_X(x)$, then $\Pr[X = x] = \sigma_X(x)$; otherwise $m(x) = \sigma_Y(x)$,

$$\Pr[X = x] = m(x) + \frac{\sigma_X(x) - m(x)}{1 - m(S)} \sum_{y \neq x} (\sigma_Y(y) - m(y)) = m(x) + \sigma_X(x) - m(x) = \sigma_X(x).$$

The case for Y holds by symmetry. □

Corollary 1.12.24. *Let (X_t, Y_t) be two copies of a Markov chain M_t on a state space S . For any initial states $x, y \in S$ and any time $t \geq 0$, there exists a coupling of the two copies, such that*

$$d_{TV}(p_x^t, p_y^t) = \Pr[X_t \neq Y_t].$$

Proof. There exists a coupled pair of random variables (U, V) of marginal distributions p_x^t and p_y^t

$$\Pr[U \neq V] = d_{TV}(p_x^t, p_y^t)$$

by lemma 1.12.23. We now extend this optimal coupling of the endpoints to a coupling of full trajectories.

First sample the endpoint pair (U, V) . Conditioned on $U = u$, sample a path

$$X_0, \dots, X_{t-1}, X_t$$

by the transition probability of M_t , conditioned on the events $X_t = u$ and $X_0 = x$.

Similarly, conditioned on $V = v$, sample a path

$$Y_0, \dots, Y_{t-1}, Y_t$$

conditioned on the events $Y_t = v$ and $Y_0 = y$.

For any path $X_0 = x, \dots, X_{t-1} = x_{t-1}, X_t = u$, the conditional probability is

$$\Pr[X_1 = x_1, \dots, X_{t-1} = x_{t-1} \mid X_0 = x, X_t = u] = \frac{1}{p_x^t(u)} \prod_{i \in [1, t]} P_{x_{i-1}, x_i}.$$

Since $p_x^t(u) > 0$ by U having marginal distribution p_x^t , the conditional probability is well-defined.

By construction, the marginal distribution of $(X_0, \dots, X_t)$ is that of the chain M_t started from x , and the marginal distribution of $(Y_0, \dots, Y_t)$ is that of the chain M_t started from y .

Moreover, since $X_t = U$ and $Y_t = V$, then

$$\Pr[U \neq V] = \Pr[X_t \neq Y_t] = d_{TV}(p_x^t, p_y^t),$$

and we have constructed a coupling of two copies of the Markov chain reaching optimal coupling at time t . □

Lemma 1.12.25. For any ergodic Markov chain, $\Delta(T+1) \leq \Delta(T)$.

Proof. Keep in mind that $\Delta_x(t)$ is a concrete deterministic value of variation distance between p_x^t and $\mathcal{D}_\pi$. The proof idea is to couple optimally at time T , then show that $\Delta_x(T+1)$ has a lower upper bound than $\Pr[X_T \neq Y_T]$.

Fix an initial state $x \in S$. If we sample $y \leftarrow \mathcal{D}_\pi$ for the initial state of p_y^T , then

$$\Delta_x(T) = d_{\text{TV}}(p_x^T, p_y^T).$$

By corollary 1.12.24, there exists a coupling for the chains such that they are optimally coupled at time T , we have

$$\Pr[X_T \neq Y_T] = \Delta_x(T).$$

Next, we couple the chains X_T and Y_T by: if $X_T = Y_T$, for all $s \geq T$, $X_s = Y_s$; otherwise, they transition independently. By the sticky step-by-step coupling, since $\{X_T = Y_T\} \subseteq \{X_{T+1} = Y_{T+1}\}$,

$$\Pr[X_T \neq Y_T] \geq \Pr[X_{T+1} \neq Y_{T+1}].$$

Hence, $\Delta_x(T)$ is non-increasing by

$$\Delta_x(T) = \Pr[X_T \neq Y_T] \geq \Pr[X_{T+1} \neq Y_{T+1}] \geq \Delta_x(T+1),$$

where the last inequality holds by lemma 1.12.11. Since it holds for all $x \in S$, we have $\Delta(T) \geq \Delta(T+1)$. $\square$

Exercise 1.12.26 (Exercise 12.13 [MU17]). Consider a Markov chain with state space S and a stationary distribution $\mathcal{D}_\pi$. For any nonnegative integer t define

$$\bar{\Delta}(t) = \max_{x, y \in S} d_{\text{TV}}(p_x^t, p_y^t).$$

- For any positive s and t , show $\bar{\Delta}(s+t) \leq \bar{\Delta}(s) \cdot \bar{\Delta}(t)$.
- For any positive s and t , show $\Delta(s+t) \leq \Delta(s) \cdot \bar{\Delta}(t)$.
- For any positive t , show $\Delta(t) \leq \bar{\Delta}(t) \leq 2\Delta(t)$.

Proof. Fix initial states $(X_0, Y_0) = (x, y)$. By corollary 1.12.24, there exists a coupling for the chains such that they are optimally coupling at time t , then

$$\Pr[X_t \neq Y_t] = d_{\text{TV}}(p_x^t, p_y^t).$$

Consider any coupling for the chains in such a way that, whenever $X_t = Y_t$, then $X_{t+i} = Y_{t+i}$ for all $i \geq 1$. Then

$$\begin{aligned} \Pr[X_{t+s} \neq Y_{t+s}] &= \Pr[X_{t+s} \neq Y_{t+s} \mid X_t \neq Y_t] \Pr[X_t \neq Y_t] + \Pr[X_{t+s} \neq Y_{t+s} \mid X_t = Y_t] \Pr[X_t = Y_t] \\ &= \Pr[X_{t+s} \neq Y_{t+s} \mid X_t \neq Y_t] \Pr[X_t \neq Y_t]. \end{aligned}$$

Conditioned on $(X_t, Y_t) = (x', y')$, again by corollary 1.12.24, there exists a coupling for the chains such that they are optimally coupling at time $t+s$, then

$$\Pr[X_{t+s} \neq Y_{t+s} \mid X_t = x', Y_t = y'] = d_{\text{TV}}(p_{x'}^s, p_{y'}^s) \leq \max_{x' \neq y' \in S} d_{\text{TV}}(p_{x'}^s, p_{y'}^s) = \bar{\Delta}(s), \quad (1.21)$$

Hence,

$$\Pr[X_{t+s} \neq Y_{t+s} \mid X_t \neq Y_t] \leq \bar{\Delta}(s). \quad (1.22)$$

Therefore, we prove $\bar{\Delta}(t+s) \leq \bar{\Delta}(t) \cdot \bar{\Delta}(s)$ by showing the following holds for all initial states x and y ,

$$d_{\text{TV}}(p_x^{t+s}, p_y^{t+s}) \leq \Pr[X_{t+s} \neq Y_{t+s}] \leq \bar{\Delta}(s) \cdot d_{\text{TV}}(p_x^t, p_y^t) \leq \bar{\Delta}(s) \cdot \bar{\Delta}(t).$$

Fix initial states $(X_0, Y_0) = (x, y)$, where $y \leftarrow \mathcal{D}_\pi$. Let (X_s, Y_s) be optimally coupled at time s by corollary 1.12.24, then

$$\Pr[X_s \neq Y_s] = \Delta_x(s).$$

Continue with the same sticky coupling after time s . Conditioned on $(X_s, Y_s) = (x', y')$, by corollary 1.12.24, there exists an optimal coupling for the next t steps. Then, again by corollary 1.12.24,

$$\Pr[X_{t+s} \neq Y_{t+s} \mid X_s \neq Y_s] \leq \bar{\Delta}(t)$$

by eq. (1.21) and eq. (1.22). Hence, $\Delta(t+s) \leq \Delta(s) \cdot \bar{\Delta}(t)$ holds by the same argument.

It is immediate from the triangle inequality that $\bar{\Delta}(t) \leq 2\Delta(t)$. To see $\Delta(t) \leq \bar{\Delta}(t)$, by theorem 1.7.13, for any t

$$\mathcal{D}_\pi(a) = \sum_{y \in S} \mathcal{D}_\pi(y) \cdot p_y^t(a),$$

then by lemma 1.12.2, let $S^+ \subseteq S$ be the set of all elements in S such that $p_x^t(a) \geq \mathcal{D}_\pi(a)$, we have

$$\begin{aligned} \Delta(t) &= p_x^t(S^+) - \mathcal{D}_\pi(S^+) \\ &= p_x^t(S^+) - \sum_{y \in S} \mathcal{D}_\pi(y) \cdot p_y^t(S^+) = \sum_{y \in S} \mathcal{D}_\pi(y) \cdot (p_x^t(S^+) - p_y^t(S^+)) \\ &\leq \sum_{y \in S} \mathcal{D}_\pi(y) \cdot d_{\text{TV}}(p_x^t, p_y^t) \leq \bar{\Delta}(t). \end{aligned}$$

Alternatively, by definition 1.12.1,

$$\begin{aligned} \Delta(t) &= \frac{1}{2} \sum_{a \in S} |p_x^t(a) - \mathcal{D}_\pi(a)| \\ &= \frac{1}{2} \sum_{a \in S} \left| p_x^t(a) - \sum_{y \in S} \mathcal{D}_\pi(y) \cdot p_y^t(a) \right| = \frac{1}{2} \sum_{a \in S} \left| \sum_{y \in S} \mathcal{D}_\pi(y) \cdot (p_x^t(a) - p_y^t(a)) \right| \\ &\leq \frac{1}{2} \sum_{a \in S} \sum_{y \in S} |\mathcal{D}_\pi(y) \cdot (p_x^t(a) - p_y^t(a))| = \sum_{y \in S} \mathcal{D}_\pi(y) \cdot \frac{1}{2} \sum_{a \in S} |p_x^t(a) - p_y^t(a)| \\ &= \sum_{y \in S} \mathcal{D}_\pi(y) \cdot d_{\text{TV}}(p_x^t, p_y^t) \leq \bar{\Delta}(t). \end{aligned}$$

□

Remark 1.12.27. The lemma 1.12.25 and exercise 1.12.26¹² are deriving from [Dob40], and well noted in [LP17].

Lemma 1.12.28.

$$d_{\text{TV}}(p_x^{t+s}, p_y^{t+s}) \leq d_{\text{TV}}(p_x^t, p_y^t).$$

Proof. Let $A \subseteq S$ be all the elements $a \in S$ such that $p_x^{t+s}(a) \geq p_y^{t+s}(a)$. By theorem 1.7.13, we have

$$d_{\text{TV}}(p_x^{t+s}, p_y^{t+s}) = p_x^{t+s}(A) - p_y^{t+s}(A) = \sum_{c \in S} (p_x^t(c) - p_y^t(c)) p_c^s(A).$$

Let $B \subseteq S$ be all the elements $a \in S$ such that $p_x^t(a) \geq p_y^t(a)$, then

$$d_{\text{TV}}(p_x^{t+s}, p_y^{t+s}) \leq \sum_{c \in B} (p_x^t(c) - p_y^t(c)) p_c^s(A) \leq p_x^t(B) - p_y^t(B) = d_{\text{TV}}(p_x^t, p_y^t),$$

where the second inequality holds by $p_c^s(A) \leq 1$, and the last equality holds by lemma 1.12.2. □

¹²Quite a bit of the unsticking came from reading Daskalakis's <https://people.csail.mit.edu/costis/6896sp11/>.

1.12.4 Geometric Convergence

Theorem 1.12.29 (Geometric Convergence). *Let $\mathbf{P}$ be the transition matrix for a finite, irreducible, aperiodic Markov chain, m_j be the smallest entry in the j^{th} column of $\mathbf{P}$, and $m = \sum_j m_j$. Then, for all $x \in S$ and t ,*

$$\Delta_x(t) = d_{\text{TV}}(p_x^t, \mathcal{D}_\pi) \leq (1 - m)^t.$$

Proof. Given m_j is the smallest entry in the j^{th} column, then in one step the chain reaches state j with probability at least m_j from every state. Consider a coupling similar to the Optimal Coupling lemma 1.12.23, such that the two copies of the chain both move to state j with probability m_j . Since $\sum_j P_{i,j} = 1$, then $m \leq 1$.

Concretely, a possible construction for $m < 1$ follows

$$\Pr[X_{t+1} = x, Y_{t+1} = y \mid X_t = x', Y_t = y'] = \begin{cases} m_x & x = y, \\ \frac{(P_{x',x} - m_x)(P_{y',y} - m_y)}{1 - m} & x \neq y. \end{cases}$$

The chains couple with probability m at a step, and the probability they haven't coupled after t steps is $(1 - m)^t$. By lemma 1.12.11, $\Delta_x(t) \leq \Pr[X_t \neq Y_t] = (1 - m)^t$. $\square$

A side product for theorem 1.12.29 is the following corollary.

Corollary 1.12.30. *If $m = 1$ in theorem 1.12.29, then all rows in $\mathbf{P}$ are identical, and $p_x^t = \mathcal{D}_\pi$ for all x and $t \geq 1$.*

Proof. Since m_j is the smallest entry in the j^{th} column, then for each column, there is no $P_{i,j}$ that is smaller than m_j .

On the other hand, by $m = \sum_j m_j = 1$ and $\sum_j P_{i,j} = 1$ for any i , there is no m_j that is smaller than $P_{i,j}$.

Therefore, any $m_j = P_{i,j}$ for all i , which proves all rows in $\mathbf{P}$ are identical.

Moreover, when $m_j = p$, then by the cutset, the outflowing probability is $(1 - p)\pi_j$, while inflowing probability is $(1 - \pi_j)p$, and we have $\pi_j = p$. Since any state reaches state j with probability $P_{i,j} = m_j = p$, this applies to any target state $j \in S$, and hence reaching stationarity is only 1 step. $\square$

Theorem 1.12.31 (Frobenius's Coin Problem). *Given $m, n \in \mathbb{Z}^+$ with $\gcd(m, n) = 1$, every integer $k > mn - m - n$ can be represented by $am + bn$ where $a, b \in \mathbb{N}$.*

Proof. We first show that $mn - m - n$ cannot be represented by $am + bn$ where $a, b \in \mathbb{N}$. Suppose there are $a, b \in \mathbb{N}^+$, such that $mn - m - n = am + bn$, then

$$(a + 1)m = (m - 1 - b)n. \quad (1.23)$$

But we notice that $mn - m - n < mn$, then $a < n$ and $b < m$. Since m and n are coprime, then there is no a and b such that $a + 1 \leq n$, $m > m - 1 - b \geq 0$, and eq. (1.23) is satisfied.

We now show for every $k > mn - m - n$, there is a non-negative linear combination representation $am + bn$.¹³ Since m and n are coprime, then $\{cn\}_{c \in [0, m-1]}$ represents different residue classes modulo m , and $(m - 1)n$ is the first element among the linear representation $am + bn$ that

$$(m - 1)n \equiv -n \pmod{m}.$$

Hence, for any $q \not\equiv -n \pmod{m}$, there always exists $b \in [0, m - 2]$ such that $bn \equiv q \pmod{m}$, and $bn < (m - 1)n$. Moreover, the elements $d \in [(m - 1)n - m + 1, (m - 1)n - 1]$ lie in distinct residue classes, so for each d , there is always $bn \equiv d \pmod{m}$. Since $d < (m - 1)n$, then each $d = am + bn$ for some a .

Since every $d \in [(m - 1)n - m + 1, (m - 1)n]$ is representable, and adding m preserves representability, any integer $k > mn - m - n$ is representable. $\square$

¹³The idea derives from <https://math.stackexchange.com/a/66978/794321>.

Exercise 1.12.32 (Exercise 12.6 [MU17]). The theorem 1.12.29 is useful only if there exists at least one column j in the transition matrix with $m_j > 0$. Argue that for any finite, aperiodic, irreducible Markov chain, there exists a time T such that every entry of $\mathbf{P}^T$ is nonzero.

Proof. Since the chain is aperiodic, then for any state j , there are 2 paths back to itself of length ℓ_{j_0} and ℓ_{j_1} , such that

$$\gcd(\ell_{j_0}, \ell_{j_1}) = 1.$$

By theorem 1.12.31, let $t_j = (\ell_{j_0} - 1)(\ell_{j_1} - 1)$, then there always exist loops starting and ending in j of any length $\ell \geq t_j$.

Since the chain is irreducible, then all states are in the same communicating class, such that for any pair of $i \neq j$, there is a time $t_{i,j}$ such that i reaches j , and $P_{i,j}^{t_{i,j}} > 0$.

Given the threshold t_j for self looping path length and $t_{i,j}$ for i to reach j , then there must exist a time

$$T > \max_{i,j} (t_{i,j} + t_j),$$

such that all entries in $\mathbf{P}^T$ are positive. □

A more general result than theorem 1.12.29 is the following: suppose we upper-bound $\tau(c)$ for some constant $c \leq 1/2$, then we can bootstrap a bound for $\tau(\varepsilon)$ for any $\varepsilon > 0$.

Theorem 1.12.33. Let M_t be a finite, irreducible, aperiodic Markov chain with $\tau(c) \leq T$ for some $c \leq 1/2$, then for M_t ,

$$\tau(\varepsilon) \leq \left\lceil \frac{\ln \varepsilon}{\ln(2c)} \right\rceil T.$$

Proof. The proving strategy is similar to exercise 1.12.26. Since for all initial states $x \in S$, we have

$$\Delta_x(T) = d_{TV}(p_x^T, \mathcal{D}_\pi) \leq c,$$

then by the triangle inequality, $d_{TV}(p_x^T, p_y^T) \leq 2c$ for all $x, y \in S$, and hence $\bar{\Delta}(T) \leq 2c$.

Therefore, on a pair of initial states x, y , we couple the copies of M_t such that (X_T, Y_T) are optimally coupled by

$$\Pr[X_T \neq Y_T] = d_{TV}(p_x^T, p_y^T)$$

from lemma 1.12.23.

Then we couple the chains such that: if $X_T = Y_T$, they stay together onwards; otherwise, (X_{2T}, Y_{2T}) are optimally coupled conditioned that $(X_T, Y_T) = (x', y')$ by

$$\Pr[X_{2T} \neq Y_{2T} \mid X_T = x', Y_T = y'] = d_{TV}(p_{x'}^T, p_{y'}^T).$$

By induction, we have

$$\Pr[X_{kT} \neq Y_{kT} \mid X_0 = x, Y_0 = y] \leq (2c)^k,$$

then $\Delta_x(kT) \leq (2c)^k$ for all $x \in S$ by lemma 1.12.10. If $(2c)^k = \varepsilon$, then

$$k \leq \left\lceil \frac{\ln \varepsilon}{\ln(2c)} \right\rceil,$$

and we have $\tau(\varepsilon) \leq kT$ by lemma 1.12.10. □

1.12.5 Application in Approximately Sampling Proper Colorings

We now present a Markov chain Monte Carlo (MCMC) process that generates an almost uniformly at random sample of a proper coloring of a graph, then use a coupling technique to show that it is rapidly mixing.

The Markov chain on proper coloring works as follows: At each step, choose a vertex v and a color ℓ uniformly at random. Recolor vertex v with color ℓ if the new coloring is proper, that v is not neighboring a vertex colored ℓ , or otherwise let the state of the chain be unchanged.

The Markov chain is aperiodic, as there exists loop of length 1 back to itself; and it is irreducible if $c \geq \Delta + 2$, as recoloring in any order can be unblocked by recoloring a later unrecoloring neighbor to an unconflicting color: the degree of the graph is at most Δ , then $|\Gamma^+(v)| \leq \Delta + 1$ for every vertex, such that the recoloring is always unblocked. Therefore, the Markov chain is ergodic and has stationary distribution being uniform distribution.

When $c \geq 4\Delta + 1$, we can use a trivial coupling for (X_t, Y_t) by choosing the same vertex and color at each step, allowing us to almost uniformly sample colorings efficiently.

Theorem 1.12.34. *For an n -vertex graph with maximum degree Δ , the mixing time of the graph-coloring Markov chain satisfies*

$$\tau(\epsilon) \leq \left\lceil \frac{nc}{c - 4\Delta} \ln \left(\frac{n}{\epsilon} \right) \right\rceil,$$

provided that $c \geq 4\Delta + 1$.

Proof. The proof strategy is similar to example 1.12.17. Fix initial states x, y be distinct c -colorings. Let D_t be the set of vertices with different colors in the two chains at time t , and let $d_t = |D_t|$.

Consider any vertex $v \in D_t$, but colored same after a move. There are at least $c - 2\Delta$ colors that have not appeared on the at most 2Δ neighbors among 2 chains, then

$$\Pr[d_{t+1} = d_t - 1 \mid d_t > 0] \geq \frac{d_t}{n} \cdot \frac{c - 2\Delta}{c}.$$

Consider any vertex $v \in V \setminus D_t$, but colored differently after a move. It has to be the case that the move succeeded in one chain, but not the other. Moreover, v must have at least one neighbor in D_t , such that the move's color may work in one chain, but not the other. Hence, for every $w \in D_t$, it can affect at most Δ neighbors with 2 colors, and

$$\Pr[d_{t+1} = d_t + 1 \mid d_t > 0] \leq \frac{\Delta d_t}{n} \cdot \frac{2}{c}.$$

By the linearity of expectation,

$$\mathbb{E}[d_{t+1} \mid d_t] \leq d_t + \frac{\Delta d_t}{n} \cdot \frac{2}{c} - \frac{d_t}{n} \cdot \frac{c - 2\Delta}{c} = d_t \cdot \left(1 - \frac{c - 4\Delta}{nc} \right),$$

then by the Markov inequality argument,

$$\Pr[d_t \geq 1] \leq \mathbb{E}[d_t] \leq \mathbb{E}[d_0] \cdot \left(1 - \frac{c - 4\Delta}{nc} \right)^t \leq n \left(1 - \frac{c - 4\Delta}{nc} \right)^t \leq n \exp \left(-\frac{c - 4\Delta}{nc} \cdot t \right).$$

The proof is immediate by lemma 1.12.10. □

We then show how to improve the coupling to reduce the number of colors necessary to $2\Delta + 1$. On one hand, for $v \in D_t$, if some neighboring vertices are same colored in both chains, then the color options should be much greater than $c - 2\Delta$. On the other hand, for $v \in V \setminus D_t$, we can decrease the number of bad moves of increasing differences by a careful coupling.

Theorem 1.12.35. *For an n -vertex graph with maximum degree Δ , the mixing time of the graph-coloring Markov chain satisfies*

$$\tau(\epsilon) \leq \left\lceil \frac{nc}{c - 2\Delta} \ln \left(\frac{n}{\epsilon} \right) \right\rceil,$$

provided that $c \geq 2\Delta + 1$.

Proof. The proof strategy is similar to exercise 1.12.20.

We reuse all notions in theorem 1.12.34, and let $A_t = V \setminus D_t$ be the set of vertices that are colored same in both chains. For $v \in A_t$, $d'(v) = |\Gamma(v) \cap D_t|$, that is the number of neighboring vertices colored differently in both chains; for $w \in D_t$, $d'(w) = |\Gamma(w) \cap A_t|$, that is the number of neighboring vertices colored same in both chains.

Hence, we explore the invariant that

$$m = \sum_{v \in A_t} d'(v) = \sum_{w \in D_t} d'(w),$$

which can be viewed as the number of edges between vertices in D_t and vertices in A_t .

For $v \in D_t$, a move on v sends v to A_{t+1} if both chains recolor v to the same valid color. Taking into account the neighbors of v that are colored the same in both chains, the forbidden color options for v have union size at most $2\Delta - d'(v)$, since those $d'(v)$ neighbors are double-counted. Therefore,

$$\Pr[d_{t+1} = d_t - 1 \mid d_t > 0] \geq \frac{1}{n} \sum_{v \in D_t} \frac{c - 2\Delta + d'(v)}{c} = \frac{1}{n} \cdot \frac{d_t(c - 2\Delta) + m}{c}.$$

For $w \in A_t$, but colored distinctly after move by $w \in D_{t+1}$, consider the coupling in a similar way in exercise 1.12.20.

Suppose in the particular case where $d'(w) = 1$, then instead of assigning w the same color and causing 2 potential bad moves, on c_1 and c_2 being the disagreeing colors of neighbor v of w , if w is assigned one of the two colors in X_t , assign the other color to w in Y_t . In this case, 2 bad moves are halved down to 1, as either they both fail, or they both move through and cause 1 more difference.

Moving to the general case. Let $S_1(w)$ be the colors of neighboring vertices of w over X_t that are in D_t , and $S_2(w)$ be the colors of neighboring vertices of w over Y_t that are in D_t . Using the same strategy in exercise 1.12.20, we pair the distinct colors between $S_1(w)$ and $S_2(w)$ as much as possible: if $c_1 \in S_1(w)$ is chosen for one chain, $c_2 \in S_2(w)$ is chosen for the other; otherwise they are mapped to the same color. In this way, the bad moves are reduced down to $\max(|S_1(w)|, |S_2(w)|) \leq d'(w)$. As a result, we have

$$\Pr[d_{t+1} = d_t + 1 \mid d_t > 0] \leq \frac{1}{n} \sum_{w \in A_t} \frac{d'(w)}{c} = \frac{m}{cn}.$$

By the linearity of expectation,

$$\mathbb{E}[d_{t+1} \mid d_t] \leq d_t + \frac{m}{cn} - \frac{1}{n} \cdot \frac{d_t(c - 2\Delta) + m}{c} = d_t \cdot \left(1 - \frac{c - 2\Delta}{nc}\right),$$

then by the Markov inequality argument,

$$\Pr[d_t \geq 1] \leq \mathbb{E}[d_t] \leq \mathbb{E}[d_0] \cdot \left(1 - \frac{c - 2\Delta}{nc}\right)^t \leq n \cdot \exp\left(-\frac{c - 2\Delta}{nc} \cdot t\right).$$

The proof is immediate by lemma 1.12.10. □

Remark 1.12.36. By definition 1.12.6, the theorem 1.12.34 and theorem 1.12.35 are rapidly mixing.

Remark 1.12.37. The theorem 1.12.34 and theorem 1.12.35 are from [Jer95].

1.12.6 Path Coupling

The motivation of using “path coupling” is that, it starts with coupling between pairs of states differ in small ways, then generalize to arbitrary pairs of states. It may be easier to construct coupling of adjacent states, and any pairs of states can be connected with a path, then we couple all pairs of states along the path, such that the expected distance contracts at each step; by summing along the path, this implies contraction for any pairs of states.

Similar to example 1.12.17, exercise 1.12.20, theorem 1.12.34, and theorem 1.12.35, we conclude by utilizing the expected-distance bound for bounding disagreement probability via a Markov inequality argument, hence bounding the mixing-time with lemma 1.12.10. Eventually, this makes coupling argument easier to be established.

Example 1.12.38 (Markov Chain for Independent Sets). Consider a Markov chain for sampling independent sets over a graph with $\Delta \leq 4$. Let X_t be the independent set at time t . At a step, the chain samples $(u, v) \xleftarrow{R} E$, and

- With probability $1/3$, $X_{t+1} = X_t \setminus \{u, v\}$.
- With probability $1/3$, let $Y = X_t \setminus \{u\} \cup \{v\}$. If Y is an independent set, $X_{t+1} = Y$; otherwise, $X_{t+1} = X_t$.
- With probability $1/3$, let $Y = X_t \setminus \{v\} \cup \{u\}$. If Y is an independent set, $X_{t+1} = Y$; otherwise, $X_{t+1} = X_t$.

Similar to exercise 1.12.16, the stationary distribution of X_t is uniform over all independent sets: By the cutset theorem 1.7.56, the total probability of adjacent states transitioning to the current state is the same as the total probability that the current state transitioning to adjacent states. Since $P_{a,b} = (3|E|)^{-1}$ for any pair of adjacent states (a, b) , then uniform distribution is a trivial solution.

We start by coupling a pair of states (X_t, Y_t) that differ by just one vertex, and this coupling can be extended to all pairs of states. Let $Z_t = (X_t \setminus Y_t) \cup (Y_t \setminus X_t)$, and we say a vertex v is bad if $v \in Z$; otherwise v is good. Let $d_t = |Z_t|$, then d_t is the random variable counting bad vertices.

We continue by assuming $d_t = 1$, and WLOG let $X_t = I$, $Y_t = I \cup \{x\}$. We apply a simple coupling by performing the same move on both chains, and show $\mathbb{E}[d_{t+1} \mid d_t] \leq d_t$ when $d_t = 1$, or equivalently, $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] \leq 0$.

Since a change in d_t is caused only by a move involving x or the neighbors of x , we restrict our attention to the edges sampled adjacent to x , neighbors of x , or neighbors of neighbors of x ¹⁴. Let δ_z be a random variable over $\{-1, 1\}$ for a vertex z between time t and $t+1$: $\delta_z = 1$ when z goes from good to bad, and $\delta_z = -1$ when z goes from bad to good. By linearity of expectation,

$$\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] = \mathbb{E}\left[\sum_{w \in \Gamma^+(\Gamma^+(x))} \delta_w \mid d_t = 1\right] = \sum_{w \in \Gamma^+(\Gamma^+(x))} \mathbb{E}[\delta_w \mid d_t = 1].$$

We shall show that when $\Delta \leq 4$, $\mathbb{E}[\delta_w \mid d_t = 1] \leq 0$. For each vertex $y \in \Gamma(x)$:

- If $|\Gamma(y) \cap I| \geq 2$, only moves on edge (x, y) changes $d_{t+1} - d_t$, either removing y and adding x , or removing $\{x, y\}$. With probability $2 \cdot (3|E|)^{-1}$, if the (x, y) edge is chosen, then $d_{t+1} - d_t$ decreases by 1.
- If $|\Gamma(y) \cap I| = 0$, for moves on edges with y ,
 - For increasing $d_{t+1} - d_t$, for $w \in \Gamma(y) \setminus \{x\}$, removing w and adding x works on X_t but fails on Y_t . Since $|\Gamma(y) \setminus \{x\}| \leq 3$ by $\Delta \leq 4$, by union bound, $d_{t+1} - d_t$ increases with probability at most $3 \cdot (3|E|)^{-1} = |E|^{-1}$.
 - For decreasing $d_{t+1} - d_t$, any option among 3 possible moves on (x, y) decreases bad vertices. With probability $|E|^{-1}$, if the (x, y) edge is chosen, then $d_{t+1} - d_t$ decreases by 1.

Hence, for moves on edges with y , $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] \leq 0$.

- If $|\Gamma(y) \cap I| = 1$, and let $z \in I \cap \Gamma(y)$, for moves on edges with y ,
 - For increasing $d_{t+1} - d_t$, removing z and adding y works on X_t but not Y_t . With probability $(3|E|)^{-1}$, $d_{t+1} - d_t$ increases by 2 if (y, z) edge is chosen and removes z and adds y .
 - For decreasing $d_{t+1} - d_t$, both adding x and removing y , and removing $\{x, y\}$ works. With probability $2 \cdot (3|E|)^{-1}$, if the (x, y) edge is chosen, then $d_{t+1} - d_t$ decreases by 1.

Hence, for move on edge with y , $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] = 0$.

¹⁴Self noting: need to consider the case where an edge is chosen with one end in I , and the other is a neighbor of x .

Since for every pair of edge, we have

$$\sum_{w \in \Gamma^+(\Gamma^+(x))} \mathbb{E} [\delta_w \mid d_t = 1, \text{move on edge } (a, b)] \leq 0,$$

as was shown in the previous discussion of $y \in \Gamma(x)$, then

$$\begin{aligned} \mathbb{E} [d_{t+1} - d_t \mid d_t = 1] &= \sum_{w \in \Gamma^+(\Gamma^+(x))} \mathbb{E} [\delta_w \mid d_t = 1] \\ &= \frac{1}{|E|} \sum_{w \in \Gamma^+(\Gamma^+(x))} \sum_{(a,b) \in E} \mathbb{E} [\delta_w \mid d_t = 1, \text{move on edge } (a, b)] \\ &= \frac{1}{|E|} \sum_{(a,b) \in E} \sum_{w \in \Gamma^+(\Gamma^+(x))} \mathbb{E} [\delta_w \mid d_t = 1, \text{move on edge } (a, b)] \leq 0. \end{aligned}$$

We now generalize the coupling for $d_t > 1$ case. To couple (X_t, Y_t) where $d_t > 1$, we create a chain of states

$$Z_0 = X_t, \dots, Z_{d_t} = Y_t,$$

where neighboring pair of states (Z_{i-1}, Z_i) has $|Z_{i-1} \setminus Z_i| + |Z_i \setminus Z_{i-1}| = 1$, that each successive Z_i is obtained from Z_{i-1} by either removing an element in $X_t \setminus Y_t$ or adding an element in $Y_t \setminus X_t$.

Our coupling arises as follows. When a move is made in $Z_0 = X_t$, the coupling for $d_t = 1$ case gives a move for the state Z_1 . This move in Z_1 can similarly be coupled with a move in state Z_2 , and so on, until the move in Z_{d_t-1} yields a move for $Z_{d_t} = Y_t$. Let Z'_i be the state after the move is made from state Z_i , and define

$$\Delta(Z'_{i-1}, Z'_i) = |Z'_{i-1} \setminus Z'_i| + |Z'_i \setminus Z'_{i-1}|.$$

Since $\mathbb{E} [d_{i+1} - d_i \mid d_i = 1] \leq 0$, then

$$\mathbb{E} [\Delta(Z'_{i-1}, Z'_i)] \leq 1.$$

Moreover, note that $(Z'_0, Z'_{d_t}) = (X_{t+1}, Y_{t+1})$, then by the triangle inequality for sets

$$(A \setminus C) \cup (C \setminus A) \subseteq (A \setminus B) \cup (B \setminus C) \cup (B \setminus A) \cup (C \setminus B),$$

as for any element $a \in A \setminus C$, either $a \in B$, then $a \in B \setminus C$; or $a \notin B$, then $a \in A \setminus B$. Then

$$\Delta(Z'_0, Z'_{d_t}) = |Z'_0 \setminus Z'_{d_t}| + |Z'_{d_t} \setminus Z'_0| \leq \sum_{i \in [1, d_t]} |Z'_{i-1} \setminus Z'_i| + \sum_{i \in [1, d_t]} |Z'_i \setminus Z'_{i-1}|,$$

and by linearity of expectation,

$$\mathbb{E} [d_{t+1} \mid d_t] = \mathbb{E} [\Delta(Z'_0, Z'_{d_t})] \leq \sum_{i \in [1, d_t]} \mathbb{E} [\Delta(Z'_{i-1}, Z'_i)] \leq d_t,$$

and eventually

$$\mathbb{E} [d_{t+1}] = \mathbb{E} [\mathbb{E} [d_{t+1} \mid d_t]] \leq \mathbb{E} [d_t].$$

Suppose $\mathbb{E} [d_{t+1} \mid d_t] \leq \beta d_t$ for $\beta < 1$, the rest is taken over by lemma 1.12.10, given a sufficient number of steps and a Markov inequality argument.

Remark 1.12.39. The example 1.12.38 derives from [BD97].

The example 1.12.38 illustrates the philosophy of path coupling: global distance is controlled by understanding local disagreements. Rather than constructing a coupling for arbitrary pairs of states, we first choose a pre-metric on the state space, whose edges represent elementary discrepancies between configurations.

In example 1.12.38, the natural choice is to connect two states when they differ at a vertex, so that the induced path metric is Hamming distance. Once such a pre-metric is fixed, any pair of states can be joined by a shortest path of discrepancies, and contraction on adjacent pairs propagates to contraction for all pairs. Thus the pre-metric is the structure that makes the path-coupling argument possible: it identifies the right local geometry of the state space.

Definition 1.12.40 (Pre-Metric). A pre-metric on Ω is a positively-weighted connected undirected graph such that all edges are shortest paths. More specifically, if (x, y) are adjacent in the pre-metric, the weight of the edge should be the least weight of any path from x to y in the pre-metric.

Theorem 1.12.41 (Path Coupling Theorem [BD97]). Suppose there exists a coupling defined for all adjacent pair of states in the pre-metric such that for all adjacent (X, Y) ,

$$\mathbb{E}[d(X', Y') \mid X, Y] \leq (1 - \alpha) d(X, Y),$$

where d is the shortest path metric, then this coupling can be extended to a coupling between all pair of states that also satisfy the contraction inequality.

Exercise 1.12.42 (Exercise 12.7 [MU17]). Continuing from example 1.12.38, we have shown there is $\beta < 1$ such that

$$\mathbb{E}[d_{t+1} \mid d_t] \leq \beta d_t.$$

Let d^* be the maximum distance over all possible pairs of initial states for the coupling.

- Give an upper bound for $\tau(\varepsilon)$ in terms of β and d^* .
- Suppose we have $\mathbb{E}[d_{t+1} \mid d_t] \leq d_t$, $d_{t+1} \in [d_t - 1, d_t + 1]$, and $\Pr[d_t \neq d_{t+1}] \geq \gamma$.
Give an upper bound for $\tau(\varepsilon)$ in terms of β , d^* , and γ .
- Show the mixing time of the graph coloring chain in theorem 1.12.34 and theorem 1.12.35, is $\text{poly}(|V|, \ln(1/\varepsilon))$, even when the number of colors is only 2Δ .
- Show the mixing time of the Markov chain for independent sets in example 1.12.38 is $\text{poly}(|V|, \ln(1/\varepsilon))$.

Proof. Suppose we have the $\beta < 1$, that $\mathbb{E}[d_{t+1} \mid d_t] \leq \beta d_t$, then

$$\mathbb{E}[d_{t+1}] = \mathbb{E}[\mathbb{E}[d_{t+1} \mid d_t]] \leq \beta \cdot \mathbb{E}[d_t].$$

Since d^* is the maximum distance over all possible pairs of states for the coupling, then

$$\mathbb{E}[d_t] \leq \beta^t \cdot d^*.$$

By the same Markov inequality argument, we have

$$\Pr[d_t \geq 1] \leq \mathbb{E}[d_t] \leq \beta^t \cdot d^*.$$

Supposing $\varepsilon = \beta^t \cdot d^*$, then

$$\tau(\varepsilon) \leq \left\lceil \frac{\ln \varepsilon - \ln d^*}{\ln \beta} \right\rceil.$$

If $\Pr[d_{t+1} \neq d_t] \geq \gamma$, then let $p^+ + p^- = \Pr[d_{t+1} \neq d_t]$, where $p^+ = \Pr[d_{t+1} = d_t + 1]$, and $p^- = \Pr[d_{t+1} = d_t - 1]$. Consider d_t as a Markov chain for random walking over $[0, d^*]$, where $d_0 = d^*$, state 0 is the absorbing state, and the transition probability from state d^* to state $d^* - 1$ is $p^+ + p^- = \Pr[d_{t+1} \neq d_t]$. Let $h_{i,0}$ be the expected number of steps to be absorbed in 0, then we have the following linear system

$$\begin{aligned} (p^+ + p^-)h_{1,0} &= 1 + p^+h_{2,0} + p^-h_{0,0} \\ (p^+ + p^-)h_{2,0} &= 1 + p^+h_{3,0} + p^-h_{1,0} \\ &\dots \\ (p^+ + p^-)h_{d^*-1,0} &= 1 + p^+h_{d^*,0} + p^-h_{d^*-2,0} \\ (p^+ + p^-)h_{d^*,0} &= 1 + p^+h_{d^*-1,0} + p^-h_{d^*-1,0}, \end{aligned}$$

where $h_{0,0} = 0$.

The chain is lazy, as it stays put with probability at most $1 - \gamma$. We continue by deriving the expected number of active moves. Let $q = \Pr[d_{i+1} \neq d_i]$, then conditioned that a move is made,

$$\begin{aligned} q^+ &= \Pr[d_{i+1} = d_i + 1 \mid d_{i+1} \neq d_i] = p^+/q \\ q^- &= \Pr[d_{i+1} = d_i - 1 \mid d_{i+1} \neq d_i] = p^-/q, \end{aligned}$$

and let $g_{i,0}$ be the expected number of active moves, then

$$\begin{aligned} g_{1,0} &= 1 + q^+ g_{2,0} + q^- g_{0,0} \\ g_{2,0} &= 1 + q^+ g_{3,0} + q^- g_{1,0} \\ &\dots \\ g_{d^*-1,0} &= 1 + q^+ g_{d^*,0} + q^- g_{d^*-2,0} \\ g_{d^*,0} &= 1 + g_{d^*-1,0}, \end{aligned}$$

where $g_{0,0} = h_{0,0} = 0$. Since $q^+ + q^- = 1$, and $q^- \geq 1/2$, then consider a Markov chain with $q^+ = q^- = 1/2$, and

$$g_{i,0} = (2d^* - i)i,$$

and therefore $g_{d^*,0} = (d^*)^2$. Since $q^- \geq 1/2$, the real $g_{d^*,0} \leq (d^*)^2$.

Since $\Pr[d_{t+1} \neq d_t] \geq \gamma$, then the expected wait time before an active move is at most γ^{-1} , and therefore

$$h_{d^*,0} \leq \gamma^{-1} (d^*)^2.$$

Hence, when $L = \lceil 2\gamma^{-1} (d^*)^2 \rceil$, let T be the random variable for the steps before being absorbed, by Markov inequality

$$\Pr[T \geq L] \leq \frac{\mathbb{E}[T]}{L} \leq \frac{1}{2}.$$

Similar to exercise 1.12.19, by the Markov property, $\Pr[T \geq kL] \leq 2^{-k}$. Then when $\varepsilon = 2^{-k}$,

$$\tau(\varepsilon) \leq \left\lceil \log \frac{1}{\varepsilon} \right\rceil L. \quad (1.24)$$

For path coupling in the Markov chain sampling graph coloring, we first show a pre-metric over all valid coloring, such that any valid pair of adjacent valid coloring differ at the color of a single vertex. Supposing that X_0 and Y_0 are disagreeing only at v , let c_1 and c_2 be the colors of v at X_0 and Y_0 , and apply the prior coupling in theorem 1.12.35: if we change $w \in \Gamma(v)$ in X_t to one of the colors in $\{c_1, c_2\}$, we change w in Y_t to the other color; otherwise, apply the same move on both chains.

One thing to beware is, in theorem 1.12.34 and theorem 1.12.35, the distance function d_t is measuring the number of disagreeing vertices, while the current d_t is measuring the minimal number of steps to transition from one valid coloring to the other, as the pre-metric is defined on Ω of valid colorings. Therefore, after a bad move in the coupling, where v is c_1 , w is c_2 in X_t , while v is c_2 , w is c_1 in Y_t , the d_t is 3 rather than 2¹⁵.

We continue bounding $\mathbb{E}[d_{t+1} \mid d_t = 1]$. Since v can be corrected with at least $c - \Delta$ colors, then

$$\Pr[d_{t+1} = d_t - 1 \mid d_t = 1] \geq \frac{1}{n} \cdot \frac{c - \Delta}{c}.$$

On the other hand, a bad move happens only on $\Gamma(v)$ with 1 color configuration on both chains, hence

$$\Pr[d_{t+1} = d_t + 2 \mid d_t = 1] \leq \frac{\Delta}{n} \cdot \frac{1}{c}.$$

¹⁵Self noting: check Daskalakis's slides <https://people.csail.mit.edu/costis/6896sp11/lec7s.pdf>.

Therefore,

$$\mathbb{E}[d_{t+1} \mid d_t = 1] \leq 1 + 2 \cdot \frac{\Delta}{n} \cdot \frac{1}{c} - \frac{1}{n} \cdot \frac{c - \Delta}{c} = 1 - \frac{c - 3\Delta}{nc} = \beta.$$

Now let the maximal distance be d^* , and construct a path joined by adjacent valid states over Ω with the pre-metric by $Z_0 = X_t$ and $Z_{d_t} = Y_t$. When a move is made on Z_i , the move from coupling is made on Z_{i+1} , and let Z'_i be the state after the move is made on state Z_i . By linearity of expectation,

$$\mathbb{E}[d_{t+1} \mid d_t] = \mathbb{E}\left[\Delta(Z'_0, Z'_{d_t})\right] \leq \sum_{i \in [1, d_t]} \mathbb{E}[\Delta(Z'_{i-1}, Z'_i)] \leq d_t \cdot \beta,$$

where $\Delta(Z'_{i-1}, Z'_i)$ is defined as the distance between Z'_{i-1} and Z'_i in the pre-metric, and finally,

$$\mathbb{E}[d_t] = \mathbb{E}[\mathbb{E}[d_t \mid d_{t-1}]] \leq \mathbb{E}[d_{t-1}] \cdot \beta \leq d^* \cdot \beta^t,$$

which gives a similar result to theorem 1.12.34 and theorem 1.12.35 but $c \geq 3\Delta + 1$.

Now we use Hamming distance as pre-metric: the states along the path are not all valid states, yet it still serves as a path coupling. Let $\tilde{\Omega}$ be the set for all possible coloring states, and $\Omega \subseteq \tilde{\Omega}$, then any adjacent states (X_0, Y_0) are still differing at the color of a vertex, but they are not necessarily valid states.

A move is made by choosing a vertex v and a color c , if c is not appearing in the neighboring $\Gamma(v)$. One observation is, if $X_0 \in \Omega$, then the chain stays in Ω ; otherwise, the chain will turn into Ω at some time step, and it will stay in Ω . Hence, coupling any pair of $(X_0, Y_0) \in \Omega^2$ over the enlarged state space $\tilde{\Omega}$ still preserves the marginal probabilities, while the path may pass through $\tilde{\Omega} \setminus \Omega$ does not hurt the argument, because the path is only used to define the metric for the path coupling theorem.

Applying the same coupling in theorem 1.12.35, conditioned that (X_0, Y_0) are adjacent in the Hamming metric, then on the bad move, d_t increases by 1. Therefore, with the same probability bounds on distance changing,

$$\mathbb{E}[d_{t+1} \mid d_t = 1] \leq 1 - \frac{c - 2\Delta}{nc} = \beta,$$

and by the same path coupling argument, $\mathbb{E}[d_t] \leq \mathbb{E}[d_{t-1}] \cdot \beta$ gives the same result as theorem 1.12.35 for $c \geq 2\Delta + 1$.

For $c = 2\Delta$, $\mathbb{E}[d_{t+1} \mid d_t] \leq d_t$. We can loosely lower bound by

$$\Pr[d_{t+1} \neq d_t] \geq \gamma = (nc)^{-1},$$

as there is at least 1 move fixing a disagreeing vertex, then the eq. (1.24) is poly $(|V| = n, \log(1/\epsilon))$.

For independent set sampling in example 1.12.38, we can also loosely lower bound by

$$\Pr[d_{t+1} \neq d_t] \geq \gamma = (3n)^{-1},$$

as a disagreeing vertex v can be removed by choosing an edge (u, v) , and removing $\{u, v\}$. By eq. (1.24), the mixing time bound is poly $(|V| = n, \log(1/\epsilon))$. $\square$

Exercise 1.12.43 (Exercise 12.16 [MU17]). Consider the Markov chain sampling independent set in a graph $G = (V, E)$. On a step, sample $v \xleftarrow{R} V$ and $b \xleftarrow{R} \{0, 1\}$. If $v \in X_t$, $X_{t+1} = X_t \setminus \{v\}$ on $b = 1$, if $v \notin X_t$ and $X_t \cup \{v\}$ is independent, $X_{t+1} = X_t \cup \{v\}$ on $b = 0$; otherwise, $X_{t+1} = X_t$.

- Show that the stationary distribution of the chain is uniform over all independent sets.
- Consider the chain over cycles and line graphs, where a line graph over n vertices is constructed by connecting vertex $1, \dots, n$ with an edge, and a cycle graph is the same with an addition of an additional edge $(1, n)$.

Devise a coupling (X_t, Y_t) for the Markov chain on line graphs and cycle graphs: if $d_t = |X_t \setminus Y_t| + |Y_t \setminus X_t|$, then at each step the coupling is at least as likely to reduce d_t as to increase d_t . Show that it is rapidly mixing.

- For the cases of line graphs and cycle graphs, derive the exact formulas for the number of independent sets.

Proof. The stationary argument is similar to the one in example 1.12.38: for any pair of adjacent state of independent sets, which are differing by a vertex, the transition probability $P_{a,b} = (2n)^{-1}$, then uniform distribution is a trivial solution for stationary distribution.

We use Hamming metric as the pre-metric: though we use the enlarged state space $\tilde{\Omega}$ of all subsets of V , for the initial states $(X_0, Y_0) \in \Omega^2$, the coupling Markov chains still preserve their marginal probabilities, hence the coupling is still valid. This follows the same argument in exercise 1.12.42. Now consider the same coupling by taking on same sampled v and b on both (X_t, Y_t) chains, where (X_t, Y_t) are adjacent states. WLOG let w be the disagreeing vertex, $X_t = I$, and $Y_t = I \cup \{w\}$. We are interested in $v \in \Gamma^+(w)$, as $v \in V \setminus \Gamma^+(w)$ does not change d_t .

- If $v = w$, any b decreases the difference, and $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1, v = w] = -1$.
- If $v \in \Gamma(w)$, then by $\Delta \leq 2$, we know $|\Gamma(v) \cap I| \leq 1$.
 - If $|\Gamma(v) \cap I| = 0$, then the distance is unchanged when $b = 1$, and the distance increases by 1 when $b = 0$. Hence, conditioned that $v \in \Gamma(w)$ and $|\Gamma(v) \cap I| = 0$, $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] = 1/2$.
 - If $|\Gamma(v) \cap I| = 1$, then no operation changes d_t . Conditioned that $v \in \Gamma(w)$ and $|\Gamma(v) \cap I| = 1$, $\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] = 0$.

Moreover, since $\Delta \leq 2$, then $|\Gamma(w)| \leq 2$. Let X be the random variable for the vertex to choose, then $\Pr[X = v] = n^{-1}$. By linearity of expectation,

$$\mathbb{E}[d_{t+1} - d_t \mid d_t = 1] = \mathbb{E}[\mathbb{E}[d_{t+1} - d_t \mid d_t = 1, X]] \leq 0.$$

With the same trick from exercise 1.12.42, we loosely lower bound by $\Pr[d_{t+1} \neq d_t] \geq (2n)^{-1} = \gamma$, as there is a choice out of $2n$ choices to fix the disagreement at w , then eq. (1.24) is poly $(|V| = n, \log(1/\epsilon))$.

For the number of independent sets in line graphs and cycle graphs, this is basically dynamic programming. Let L_i be the number of independent sets of line graph with i vertices. By including the first vertex or not, we have

$$L_i = L_{i-1} + L_{i-2}.$$

Let C_i be the number of independent sets of cycle graph with i vertices. If the first vertex is not chosen, then there are L_{i-1} independent sets, while if the first vertex is chosen, then there are L_{i-3} independent sets, hence

$$C_i = L_{i-1} + L_{i-3}.$$

□

Exercise 1.12.44 (Exercise 12.18 [MU17]). Use path coupling to simplify example 1.12.17 and exercise 1.12.20.

Proof. We enhance the result in exercise 1.12.16, by showing that the Markov chain over fixed size independent sets of size $k \leq n/(2\Delta + 2)$ is ergodic. Let $I = X_t \cap Y_t$, $A_t = X_t \setminus I$, and $B_t = Y_t \setminus I$. Consider following 2 easier corner case:

- If $|B_t| = |B_t \setminus \Gamma(A_t)|$, then $B_t \cap \Gamma(X_t) = \emptyset$, and the transition takes $|A_t| = |B_t|$ steps by moving from A_t to B_t .
- If $|B_t| = |B_t \cap \Gamma(A_t)|$, since $|A_t| \leq |X_t| \leq k$, then $|V \setminus \Gamma^+(X_t)| \geq n/2$, and B_t can move to $V \setminus \Gamma^+(X_t)$.

The transition can take $2|A_t|$ steps by moving A_t to an independent set in $V \setminus \Gamma^+(X_t)$, then moving back into B_t .

In general case, X_t can be broken into 3 parts: I , $A_t \cap \Gamma(B_t)$, and $A_t \setminus \Gamma(B_t)$. Since $|A_t| \leq k$ and $|V \setminus \Gamma^+(X_t)| \geq n/2$, then there exists a way to move A_t to $V \setminus \Gamma^+(X_t)$ in $|A_t|$ steps, such that $B_t \setminus \Gamma(A_t)$ are all occupied. The remaining steps are transitioning occupied vertices back into $B_t \setminus \Gamma(A_t)$.

The uniform distribution being stationary distribution argument is immediate by the cutset, that the neighboring pair of independent sets transition to each other with transition probability $1/kn$ everywhere, and the chain has odd length loop of going back to itself, hence it is aperiodic.

Let D_t be the pre-metric over all valid k -sized independent sets, which is the least number of steps to transition from X_t to Y_t . Conditioned that $D_t = 1$, let $X_t = I \cup \{u\}$ and $Y_t = I \cup \{v\}$. At a step, sample $w \xleftarrow{r} X_t$ and $x \xleftarrow{r} V$, and a transition happens when $w \in X_t$ and $\Gamma(x) \cap X_t = \emptyset$.

If $w \neq u$, $w_x = w_y = w$; otherwise $(w_x, w_y) = (u, v)$. Let $S_1 = \Gamma^+(u)$, $S_2 = \Gamma^+(v)$, and M be the mapping between S_1 and S_2 , such that elements in $S_1 \cap S_2$ are mapped to themselves, elements are mapped bijectively between $S_1 \setminus S_2$ and $S_2 \setminus S_1$ as many as possible, bijectively map u and v , with the remaining mapped to itself. Hence, if $x \in S_1 \cup S_2$, $z_x = x$ and $z_y = M(x)$; otherwise, $z_x = z_y = x$.

The transition step is coupled by: Try removing w_x and adding z_x to X_t , and removing w_y and adding z_y to Y_t . To decrease the distance, $(w_x, w_y) = (u, v)$ and $x \notin \Gamma^+(X_t \cup Y_t \setminus \{u, v\})$, then the probability is lower bounded by

$$\Pr[D_{t+1} = D_t - 1 \mid D_t = 1] \geq \frac{1}{k} \cdot \frac{n - (\Delta + 1)(k - 1)}{n}.$$

To increase the distance, $w \in I$ and $z_x \in S_2$, such that the number of moves where at least a chain transitions is upper bounded by $\max(|S_1|, |S_2|) \leq \Delta + 1$ by exercise 1.12.20, hence the probability is upper bounded by

$$\Pr[D_{t+1} = D_t + 1 \mid D_t = 1] \leq \frac{k - 1}{k} \cdot \frac{\Delta + 1}{n}.$$

By linearity of expectation,

$$\mathbb{E}[D_{t+1} \mid D_t = 1] \leq 1 - \frac{1}{k} \cdot \frac{n - (\Delta + 1)(k - 1)}{n} + \frac{k - 1}{k} \cdot \frac{\Delta + 1}{n} = 1 - \frac{n - 2(\Delta + 1)(k - 1)}{kn}.$$

Again by path coupling argument, let

$$Z_0 = X_t, \dots, Z_{d_t} = Y_t$$

be a chain of states where neighboring states have pre-metric distance 1. The move in Z_0 gives a coupled move for Z_1 , and eventually leads to a coupled move for Z_{d_t} . Let Z'_i be the state after the move is made from Z_i , and let

$$\Delta(Z'_{i-1}, Z'_i) = |Z'_{i-1} \setminus Z'_i|.$$

By linearity of expectation,

$$\mathbb{E}[D_{t+1} \mid D_t] = \sum_{i \in [1, D_t]} \mathbb{E}[\Delta(Z'_{i-1}, Z'_i)] \leq D_t \cdot \left(1 - \frac{n - 2(\Delta + 1)(k - 1)}{kn}\right),$$

and we conclude the proof. $\square$

1.13 Martingales

1.13.1 Martingales

Definition 1.13.1 (Martingale). A sequence of random variables $Z_0, Z_1, \dots$ is a martingale with respect to another sequence $X_0, X_1, \dots$ if for all $n \geq 0$, the following conditions hold:

- Z_n is a function of $X_0, \dots, X_n$.
- $\mathbb{E}[|Z_n|] < \infty$.
- $\mathbb{E}[Z_{n+1} \mid X_0, \dots, X_n] = Z_n$.

A sequence of random variables $Z_0, Z_1, \dots$ is called a martingale when it is a martingale with respect to itself, that is

- $\mathbb{E}[|Z_n|] < \infty$.
- $\mathbb{E}[Z_{n+1} \mid Z_0, \dots, Z_n] = Z_n$.

Example 1.13.2 (Martingale from Fair Games). Consider a gambler who plays a sequence of fair games. Let X_i be the winning on the i^{th} game, that is either +1 or -1, and Z_i be the gambler's total winning of the first i games. Since $Z_0 = 0$ and

$$Z_i = Z_0 + \sum_{j \in [1, i]} X_j,$$

then by linearity of expectation,

$$\mathbb{E}[Z_{i+1} \mid X_1, \dots, X_i] = Z_i + \mathbb{E}[X_{i+1} \mid X_1, \dots, X_i] = Z_i,$$

where the second equality holds as $\mathbb{E}[X_{i+1} \mid X_1, \dots, X_i] = 0$ on any input of $X_1, \dots, X_i$ by fair game.

Thus, $Z_1, \dots, Z_n$ is a martingale with respect to the sequence $X_1, \dots, X_n$.

Exercise 1.13.3 (Exercise 13.2 [MU17]). Let $X_0, X_1, \dots$ be a sequence of random variables, and let $S_i = \sum_{j \in [1, i]} X_j$. Show that $S_0, S_1, \dots$ is a martingale with respect to $X_0, X_1, \dots$, then for all $i \neq j$, $\mathbb{E}[X_i X_j] = 0$.

Proof. If $S_0, S_1, \dots$ is a martingale with respect to $X_0, X_1, \dots$, then by definition 1.13.1,

$$S_i = \mathbb{E}[S_{i+1} \mid X_0, \dots, X_i] = S_i + \mathbb{E}[X_{i+1} \mid X_0, \dots, X_i],$$

where the second equality holds by linearity of expectation, then for all $i \geq 0$,

$$\mathbb{E}[X_{i+1} \mid X_0, \dots, X_i] = 0.$$

Hence, on an $i \geq 1$ and any $j < i$, conditioned that $X_j = x$,

$$\mathbb{E}[X_i \mid X_j = x] = \mathbb{E}[\mathbb{E}[X_i \mid X_0, \dots, X_j = x, \dots, X_{i-1}]] = 0.$$

Therefore, we conclude the proof by

$$\begin{aligned} \mathbb{E}[X_i X_j] &= \mathbb{E}[\mathbb{E}[X_i X_j \mid X_j]] \\ &= \sum_{x \in \Omega} \mathbb{E}[X_i X_j \mid X_j = x] \Pr[X_j = x] \\ &= \sum_{x \in \Omega} x \mathbb{E}[X_i \mid X_j = x] \Pr[X_j = x] = 0. \end{aligned}$$

□

Exercise 1.13.4 (Exercise 13.1 [MU17]). Show that if $Z_0, Z_1, \dots$ is a martingale with respect to $X_0, X_1, \dots$, then it is a martingale with respect to itself.

Proof. By definition 1.2.6 and lemma 1.2.7, we use the tower property

$$\mathbb{E}[\mathbb{E}[Y | X] | Z] = \mathbb{E}[Y | Z]$$

for random variables X, Y, Z such that Z is determined by X , or X carries at least as much information as Z , or Z can be obtained from X from some function.

Since $Z_0, \dots, Z_n$ is a martingale with respect to $X_0, \dots, X_n$, then by definition 1.13.1, for all $i \in [0, n]$,

$$Z_i = \mathbb{E}[Z_{i+1} | X_0, \dots, X_i].$$

By the prior fact,

$$\mathbb{E}[\mathbb{E}[Z_{i+1} | X_0, \dots, X_i] | Z_0, \dots, Z_i] = \mathbb{E}[Z_{i+1} | Z_0, \dots, Z_i].$$

Yet by definition 1.13.1,

$$\mathbb{E}[\mathbb{E}[Z_{i+1} | X_0, \dots, X_i] | Z_0, \dots, Z_i] = \mathbb{E}[Z_i | Z_0, \dots, Z_i] = Z_i.$$

Therefore we conclude by $\mathbb{E}[Z_{i+1} | Z_0, \dots, Z_i] = Z_i$. $\square$

Exercise 1.13.5 (Exercise 13.4 [MU17]). Let $X_1, X_2, \dots$ be independent identically distributed random variables with expectation 0 and variance $\sigma^2 < \infty$. Let

$$Z_n = \left(\sum_{i \in [1, n]} X_i \right)^2 - n\sigma^2.$$

Show that $Z_1, Z_2, \dots$ is a martingale.

Proof. Let $Y_n = \sum_{i \in [1, n]} X_i$, then $Z_n = Y_n^2 - n\sigma^2$. Since $\mathbb{E}[X_i] = 0$, then $\sigma^2 = \mathbb{E}[X_i^2]$. We show as follows: For $n \geq 1$,

$$\begin{aligned} \mathbb{E}[Z_{n+1} | X_1, \dots, X_n] &= \mathbb{E}[Z_n + X_{n+1}^2 + 2X_{n+1}Y_n - \sigma^2 | X_1, \dots, X_n] \\ &= Z_n + \mathbb{E}[X_{n+1}^2] + 2\mathbb{E}[X_{n+1}Y_n] - \sigma^2 \\ &= Z_n + \mathbb{E}[X_{n+1}^2] + 2\mathbb{E}[X_{n+1}]\mathbb{E}[Y_n] - \sigma^2 \\ &= Z_n + \mathbb{E}[X_{n+1}^2] - \sigma^2 = Z_n, \end{aligned}$$

where the second and the third equality are from the independence, and the last two equalities are from expectation being 0. By exercise 1.13.4, we conclude the proof. $\square$

Definition 1.13.6 (Doob Martingale [Doo40]). Let $X_0, \dots, X_n$ be a sequence of random variables, and Y be a random variable depending on $X_0, \dots, X_n$ with $\mathbb{E}[|Y|] < \infty$, then for $i \in [0, n]$,

$$Z_i = \mathbb{E}[Y | X_0, \dots, X_i]$$

gives a martingale with respect to $X_0, \dots, X_n$, as

$$\begin{aligned} \mathbb{E}[Z_{i+1} | X_0, \dots, X_i] &= \mathbb{E}[\mathbb{E}[Y | X_0, \dots, X_{i+1}] | X_0, \dots, X_i] \\ &= \mathbb{E}[Y | X_0, \dots, X_i] \\ &= Z_i, \end{aligned}$$

where the second equality holds by lemma 1.2.7 and definition 1.2.6.

Remark 1.13.7. One can view the Doob Martingale as a sequence of refined predictions to Y , where each element Z_i is the expectation of Y when the values of $X_0, \dots, X_i$ are known. Hence, the refined predictions gradually using more information on the values of the random variables $X_0, \dots, X_n$.

Exercise 1.13.8 (Exercise 13.3 [MU17]). Let $X_0 = 0$ and for $j \geq 0$, X_{j+1} is chosen uniformly over real interval $[X_j, 1]$. Show that, for $k \geq 0$, the sequence $Y_k = 2^k(1 - X_k)$ is a martingale.

Proof. Since Y_{k+1} depends only on Y_k , then

$$\mathbb{E}[Y_{k+1} \mid Y_0, \dots, Y_k] = \mathbb{E}[Y_{k+1} \mid Y_k].$$

Moreover, since $\mathbb{E}[X_{k+1} \mid X_k] = (1 + X_k)/2$, then

$$\mathbb{E}[Y_{k+1} \mid X_k] = \mathbb{E}[2^{k+1}(1 - X_{k+1}) \mid X_k] = 2^{k+1} - 2^{k+1} \mathbb{E}[X_{k+1} \mid X_k] = 2^k(1 - X_k) = Y_k.$$

By $\mathbb{E}[Y_k \mid Y_k] = Y_k$, we show that

$$\mathbb{E}[\mathbb{E}[Y_{k+1} \mid X_k] \mid Y_k] = \mathbb{E}[Y_{k+1} \mid Y_k] = Y_k,$$

which concludes the proof. $\square$

Lemma 1.13.9. If $Z_0, Z_1, \dots$ is a martingale with respect to $X_0, X_1, \dots$, then $\mathbb{E}[Z_n] = \mathbb{E}[Z_0]$ for any $n \geq 0$.

Proof. By definition 1.13.1, for any $n \geq 1$

$$\mathbb{E}[Z_n \mid X_0, \dots, X_{n-1}] = Z_{n-1}.$$

Given that

$$\mathbb{E}[Z_{n-1} \mid X_0, \dots, X_{n-2}] = Z_{n-2},$$

we have

$$\begin{aligned} Z_{n-2} &= \mathbb{E}[Z_{n-1} \mid X_0, \dots, X_{n-2}] \\ &= \mathbb{E}[\mathbb{E}[Z_n \mid X_0, \dots, X_{n-1}] \mid X_0, \dots, X_{n-2}] \\ &= \mathbb{E}[Z_n \mid X_0, \dots, X_{n-2}]. \end{aligned}$$

By induction, we have $Z_0 = \mathbb{E}[Z_n \mid X_0]$, then

$$\mathbb{E}[Z_0] = \mathbb{E}[\mathbb{E}[Z_n \mid X_0]] = \mathbb{E}[Z_n],$$

which concludes the proof. $\square$

Example 1.13.10 (Galton-Watson Process). The Galton-Watson process was used in theorem 1.6.60, but was originally used to statistically investigate the extinction of family names. Let G_t be the random variable for the individuals in the t^{th} generation, and $X_{t,k}$ be the number of offspring of the k^{th} individual in the t^{th} generation. Each individual gives birth to offspring independently, and their numbers of offspring are identically distributed. Hence,

$$G_t = \sum_{k \in [1, G_{t-1}]} X_{t-1,k},$$

and let $\mu = \mathbb{E}[X_{t,k}]$. By linearity of expectation,

$$\mathbb{E}[G_{t+1} \mid G_1, \dots, G_t] = \mu G_t.$$

Let $M_t = \mu^{-t} \cdot G_t$, then

$$\mathbb{E}[M_{t+1} \mid G_1, \dots, G_t] = \mathbb{E}[\mu^{-(t+1)} \cdot G_{t+1} \mid G_1, \dots, G_t] = \mu^{-t} \cdot G_t = M_t,$$

and we have

$$\mathbb{E}[\mathbb{E}[M_{t+1} \mid G_1, \dots, G_t] \mid M_1, \dots, M_t] = \mathbb{E}[M_{t+1} \mid M_1, \dots, M_t].$$

On the other hand,

$$\mathbb{E}[\mathbb{E}[M_{t+1} \mid G_1, \dots, G_t] \mid M_1, \dots, M_t] = \mathbb{E}[M_t \mid M_1, \dots, M_t] = M_t,$$

which makes $M_1, \dots$ a martingale.

Example 1.13.11 (Pólya's Urn). Suppose there are black balls and white balls in an urn, that are identical except for the colors. Each time pick a ball, and put it back with another identical ball with same color. For simplicity, we start at round 2, and ends at round n where there are n balls in the urn. Let X_n be the number of black balls in the urn, and $Z_n = X_n/n$ be the ratio of black balls in the urn. Then for all $n \geq 2$,

$$\mathbb{E}[Z_{n+1} \mid Z_2, \dots, Z_n] = Z_n \cdot \frac{nZ_n + 1}{n+1} + (1 - Z_n) \cdot \frac{nZ_n}{n+1} = Z_n,$$

which makes $Z_2, Z_3, \dots$ a martingale.

1.13.2 Stopping Times

By lemma 1.13.9, if the number of games is fixed initially, then the expected gain from the fair games is 0, if $Z_0 = 0$. But suppose the number of games is not fixed initially, for example the gambler choose to play a random number of games, by deciding when to quit based on the outcome of the games already played.

Definition 1.13.12 (Stopping Time). A nonnegative integer-valued random variable T is a stopping time for the sequence $Z_0, Z_1, \dots$ if the probability of the event $T = n$ is independent of the variables $\{Z_{n+j} \mid Z_0, \dots, Z_n\}_{j \geq 1}$, which are the random variables $Z_{n+1}, Z_{n+2}, \dots$ conditioned on the values of $Z_0, \dots, Z_n$.

A stopping time is corresponding to a strategy of when to stop based only on the outcome seen so far. But given that one can stop by strategy such as “winning reaches 10 for the first time”, it would be helpful to characterize the conditions on the stopping time T that maintain the property $\mathbb{E}[Z_T] = \mathbb{E}[Z_0] = 0$. The subtle problem is that, T might not be finite in this case, that the gambler will gamble infinite number of times with non-zero probability. The martingale stopping time theorem shows that, under certain conditions, and in particular when the stopping time is bounded or has bounded expectation, the expected value of the martingale at the stopping time is $\mathbb{E}[Z_0]$.

Theorem 1.13.13 (Martingale Stopping Time Theorem). If $Z_0, Z_1, \dots$ is a martingale with respect to $X_0, X_1, \dots$ and if T is a stopping time for $X_0, X_1, \dots$, then

$$\mathbb{E}[Z_T] = \mathbb{E}[Z_0]$$

whenever one of the following holds:

- The Z_i is bounded, such that there exists constant c for all $i \geq 0, |Z_i| \leq c$.
- T is bounded.
- $\mathbb{E}[T] < \infty$, and there is a constant c such that $\mathbb{E}[|Z_{i+1} - Z_i| \mid X_0, \dots, X_i] \leq c$.

Exercise 1.13.14 (Exercise 13.5 [MU17]). Consider the gambler's ruin example 1.7.51, stopping by either losing ℓ_1 or winning ℓ_2 . Let X_n be the winning in the n^{th} game, Y_n be the winning by the end of the n^{th} game, and $Z_n = Y_n^2 - n$. Show that $Z_1, Z_2, \dots$ is a martingale, and let T be a stopping time, determine $\mathbb{E}[Z_T]$ and $\mathbb{E}[T]$.

Proof. Since $\text{Var}[X_i] = 1$ and $\mathbb{E}[X_i] = 0$, by exercise 1.13.5, $Z_1, Z_2, \dots$ is a martingale.

We have shown in exercise 1.7.68 that $\mathbb{E}[T] = \ell_1 \ell_2$.

The theorem 1.13.13 applies by

$$|Z_{i+1} - Z_i| = |(Y_{i+1} - Y_i)(Y_{i+1} + Y_i) - 1| \leq |Y_{i+1} - Y_i| \cdot |Y_{i+1} + Y_i| + 1 \leq 2 \max(\ell_1, \ell_2) + 1,$$

then $\mathbb{E}[Z_T] = \mathbb{E}[Z_1] = 0$, which means $\mathbb{E}[Y_T^2] = \mathbb{E}[T]$. Since Y_T^2 can either be ℓ_1^2 or ℓ_2^2 , by example 1.7.51¹⁶,

$$\mathbb{E}[Y_T^2] = \ell_1^2 \cdot \frac{\ell_2}{\ell_1 + \ell_2} + \ell_2^2 \cdot \frac{\ell_1}{\ell_1 + \ell_2} = \ell_1 \ell_2,$$

which derives $\mathbb{E}[T]$ in a martingale way. □

¹⁶Now we can show by Y_i is a martingale in example 1.13.2, then $\mathbb{E}[Y_T] = 0$, to derive the winning or losing probabilities.

Exercise 1.13.15 (Exercise 13.6 [MU17]). Consider the gambler's ruin example 1.7.51, stopping by either losing ℓ_1 or winning ℓ_2 , and winning probability $p < 1/2$. Let X_n be the winning in the n^{th} game, and Z_n be the winning by the end of the n^{th} game.

- Show that

$$A_n = \left(\frac{1-p}{p} \right)^{Z_n}$$

is a martingale with mean 1.

- Determine the probability that the player wins ℓ_2 before losing ℓ_1 .
- Show that $B_n = Z_n - (2p - 1)n$ is a martingale with mean 0.
- Let T be the stopping time when the player finishes playing. Determine $\mathbb{E}[Z_T]$ and $\mathbb{E}[T]$.

Proof. We have shown in exercise 1.7.69 that

$$A_n = \left(\frac{1-p}{p} \right)^{Z_n},$$

we have $\mathbb{E}[A_{n+1}] = \mathbb{E}[A_n] = 1$. Since $Z_{n+1} = Z_n + X_{n+1}$,

$$\mathbb{E}[A_{n+1} | X_1, \dots, X_n] = p \left(\frac{1-p}{p} \right)^{Z_n+1} + (1-p) \left(\frac{1-p}{p} \right)^{Z_n} = \left(\frac{1-p}{p} \right)^{Z_n} = A_n,$$

which proves that $A_1, A_2, \dots$ is a martingale by exercise 1.13.4.

With this $A_1, A_2, \dots$ martingale, then let q be the probability of being absorbed into ℓ_2 , and

$$\mathbb{E}[A_T] = q \left(\frac{1-p}{p} \right)^{\ell_2} + (1-q) \left(\frac{1-p}{p} \right)^{-\ell_1} = 1,$$

such that

$$q = \left(1 - \left(\frac{1-p}{p} \right)^{\ell_1} \right) \cdot \left(1 - \left(\frac{1-p}{p} \right)^{\ell_1+\ell_2} \right)^{-1}.$$

Since $\mathbb{E}[X_i] = p - (1-p) = 2p - 1$, then by linearity of expectation,

$$\mathbb{E}[B_{n+1} | X_1, \dots, X_n] = B_n + \mathbb{E}[X_{n+1} - (2p - 1) | X_1, \dots, X_n] = B_n,$$

the second equality holds by the independence of the games. Hence, $B_1, B_2, \dots$ is a martingale by exercise 1.13.4.

Since A_n is bounded, then by theorem 1.13.13, $\mathbb{E}[A_T] = \mathbb{E}[A_1] = 1$. Therefore,

$$\mathbb{E}[Z_T] = -\ell_1 \cdot (1-q) + \ell_2 \cdot q = (\ell_1 + \ell_2) \cdot q - \ell_1.$$

By $\mathbb{E}[B_T] = 0 = \mathbb{E}[Z_T] - (2p - 1)\mathbb{E}[T]$, we have

$$\mathbb{E}[T] = \frac{\mathbb{E}[Z_T]}{2p - 1} = \frac{\ell_1 - (\ell_1 + \ell_2) \cdot q}{1 - 2p}.$$

□

Example 1.13.16 (A Ballot Theorem). Suppose there are 2 candidates running an election, where candidate A obtains a votes, and candidate B obtains $b < a$ votes. The votes come in random order, and can be viewed as sampled from uniformly random permutations. Candidate A always with higher votes has probability $(a - b)/(a + b)$.

Let $n = a + b$, and let S_k be the leading votes of A to B after k votes. Then $S_n = a - b$. For $0 \leq k \leq n - 1$, let

$$X_k = \frac{S_{n-k}}{n - k}.$$

We first show that $X_1, X_2, \dots, X_n$ is a martingale. Since

$$\mathbb{E}[S_{n-k-1} | S_n, \dots, S_{n-k}] = \mathbb{E}[S_{n-k-1} | S_{n-k}] = \frac{n-k+S_{n-k}}{2(n-k)} \cdot (S_{n-k}-1) + \frac{n-k-S_{n-k}}{2(n-k)} \cdot (S_{n-k}+1) = \frac{n-k-1}{n-k} \cdot S_{n-k},$$

where the first equality holds by S_{n-k-1} is only related to S_{n-k} counting backwards, and the second equality holds by the probability of the $(n-k)^{\text{th}}$ vote being A or B, conditioned on S_{n-k} , then

$$\mathbb{E}[X_{k+1} | S_n, \dots, S_{n-k}] = \frac{S_{n-k}}{n-k} = X_k,$$

and

$$\mathbb{E}[X_{k+1} | X_1, \dots, X_k] = \mathbb{E}[\mathbb{E}[X_{k+1} | S_n, \dots, S_{n-k}] | X_1, \dots, X_k] = X_k,$$

which completes the proof of martingale.

Moreover, define a stopping T for the first $X_T = 0$, otherwise $T = n-1$, then since T is bounded, and $X_0, \dots, X_{n-1}$ is also bounded, by theorem 1.13.13,

$$\mathbb{E}[X_T] = \mathbb{E}[X_0] = \frac{S_n}{n} = \frac{a-b}{a+b}.$$

By the stopping time T , the voting is either A and B are drawing, or A leads throughout the n votes. In the first case, $X_T = 0$ as $S_T = 0$, otherwise, $X_T = 1$ as $S_T = n$. Hence,

$$\mathbb{E}[X_T] = \Pr[A \text{ leads throughout}] \cdot 1 + \Pr[A \text{ and B are drawing}] \cdot 0 = \frac{a-b}{a+b},$$


which completes the proof.

Remark 1.13.17. The example 1.13.16 is also called Bertrand's ballot theorem, which can be proved in other ways.

One of the earliest proofs [And87] solved the problem directly, but the “geometric reflection” trick derives from his work. Consider a 2-dimensional lattice, a vote for A moves is a step right, and a vote for B is a step up. Hence, after n moves, the particle moves from $(0,0)$ to (a,b) . An observation is, if A is ahead in the count throughout, then the particle is constantly below the line $y = x$ (except for the initial state with no vote).

Moreover, for any permutation that start with a vote for B, then it has to be at the diagonal $y = x$ at least once. The reflection trick is: we can always reflect a trajectory above the diagonal to the symmetry one below the diagonal $y = x$, such that the move is still a bad move, but the move is beginning with a vote for A to $(1,0)$.

Since all the valid trajectories begin with a vote for A, we just exclude all the bad trajectories beginning with $(1,0)$, and all such bad trajectories can be reflected to the trajectories beginning with $(0,1)$.



Therefore, the number of valid trajectories is

$$\binom{a+b-1}{a-1} - \binom{a+b-1}{b-1} = \frac{a-b}{a+b} \cdot \binom{a+b}{a}.$$

When $a = b + 1$, this deteriorates to lemma 1.7.67, which are Dycks path and Catalans number C_b .

1.13.3 Wald's Equation

Wald's equation is an important corollary of theorem 1.13.13, which handles the expectation of the sum of independent random variables, where the number of random variables being summed is itself a random variable.

Theorem 1.13.18 (Wald's Equation [Wal47]). *Let $X_1, X_2, \dots$ be independent, identically distributed random variables with distribution X , and T be a stopping time for this sequence. If T and X have bounded expectation, then*

$$\mathbb{E} \left[\sum_{i \in [1, T]} X_i \right] = \mathbb{E}[T] \cdot \mathbb{E}[X].$$

Proof. We first show a version for the nonnegative random variables. For $i \geq 1$, let

$$Z_i = \sum_{j \in [1, i]} (X_j - \mathbb{E}[X_j]),$$

then $Z_1, Z_2, \dots$ is a martingale with respect to $X_1, X_2, \dots$, with $\mathbb{E}[Z_i] = 0$ by exercise 1.13.15.

Since $\mathbb{E}[T] < \infty$, and

$$\mathbb{E}[|Z_{i+1} - Z_i| \mid X_1, \dots, X_i] = \mathbb{E}[|X_{i+1} - \mathbb{E}[X]|] \leq 2\mathbb{E}[X], \quad (1.25)$$

then by theorem 1.13.13,

$$\mathbb{E}[Z_T] = \mathbb{E}[Z_1] = 0.$$

Therefore, by linearity of expectation,

$$\mathbb{E}[Z_T] = \mathbb{E} \left[\sum_{i \in [1, T]} (X_i - \mathbb{E}[X]) \right] = \mathbb{E} \left[\sum_{i \in [1, T]} X_i \right] - \mathbb{E}[T] \cdot \mathbb{E}[X],$$

which completes the proof.

In general case, where X is any random variable, the main trick is writing $X = X^+ - X^-$, where

$$X^+ = \max(X, 0), \quad X^- = \max(-X, 0).$$

Hence, write $X_i = X_i^+ - X_i^-$ for all $i \geq 1$, and by the prior nonnegative case,

$$\mathbb{E} \left[\sum_{i \in [1, T]} X_i^+ \right] = \mathbb{E}[T] \cdot \mathbb{E}[X^+], \quad \mathbb{E} \left[\sum_{i \in [1, T]} X_i^- \right] = \mathbb{E}[T] \cdot \mathbb{E}[X^-].$$

By linearity of expectation,

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in [1, T]} X_i \right] &= \mathbb{E} \left[\sum_{i \in [1, T]} X_i^+ - \sum_{i \in [1, T]} X_i^- \right] = \mathbb{E} \left[\sum_{i \in [1, T]} X_i^+ \right] - \mathbb{E} \left[\sum_{i \in [1, T]} X_i^- \right] \\ &= \mathbb{E}[T] \cdot (\mathbb{E}[X^+] - \mathbb{E}[X^-]) \\ &= \mathbb{E}[T] \cdot \mathbb{E}[X], \end{aligned}$$

we conclude the proof for the general case. $\square$

Remark 1.13.19. Since T is a stopping time for $X_1, X_2, \dots$, then event $\{T \geq i\}$ is dependent only on $X_1, \dots, X_{i-1}$ by definition 1.13.12. The mutual independence over X_i ensures that X_i is independent of past $X_1, \dots, X_{i-1}$, hence X_i is independent of $\{T \geq i\}$. The independence is a convenient sufficient condition, that can be used to construct martingale in exercise 1.13.15 and used here, and help simplify the third condition of the theorem 1.13.13 in eq. (1.25).

After theorem 1.13.18, remark 1.13.19, and exercise 1.13.15, for the stopping time in the case of independent random variables, we have an equivalent, yet simpler version of stopping time as follows.

Definition 1.13.20 (Stopping Time for Independent Random Variables). Let $Z_0, Z_1, \dots$ be a sequence of independent random variables. A nonnegative, integer-valued random variable T is a stopping time for the sequence if the event $\{T = n\}$ is independent of $Z_{n+1}, Z_{n+2}, \dots$

Example 1.13.21 (Dice Rolling). Consider a gambling game where a player first rolls a die. If the outcome is X then roll X new standard dice and gain Z , which is the sum of the outcome of the X dice. This is exactly an application of theorem 1.13.18 to derive the expected outcome. By definition 1.13.20 and theorem 1.13.18, let X be the stopping time, and $Y_1, Y_2, \dots$ be i.i.d. random variables for dice rolling scores with distribution Y . Then by theorem 1.13.18,

$$\mathbb{E} \left[\sum_{i \in [1, X]} Y_i \right] = \mathbb{E}[X] \cdot \mathbb{E}[Y] = \left(\frac{1+n}{2} \right)^2,$$

if the die is a n faced fair die.

Example 1.13.22 (Server Message Transmission). Consider n servers communicating using a shared channel, where time is divided in discrete slots. At each time slot, any server that needs to send a packet can transmit it through the channel. If exactly one packet is sent at that time, the transmission is successful, otherwise none are successful. At a time slot, a server transmits a packet with probability $1/n$.

Let $X_1, X_2, \dots$ be the independent 0-1 random variables for a successful packet transmission on the i^{th} time slot with distribution X , then

$$\Pr[X = 1] = \binom{n}{1} \cdot \frac{1}{n} \cdot \left(1 - \frac{1}{n}\right)^{n-1} = \left(1 - \frac{1}{n}\right)^{n-1}.$$

Similar to section 1.5.1, let

$$\varphi_g(n) = (n-1) \ln \left(1 - \frac{1}{n}\right),$$

then

$$\varphi'_g(n) = \ln \left(1 - \frac{1}{n}\right) + \frac{1}{n}.$$

Let $t = 1 - 1/n$, and $n \geq 2$ means $1/2 \leq t < 1$, then $\varphi'_g(n) < 0$. Therefore, when $n \geq 2$, since

$$\lim_{n \rightarrow \infty} \left(1 - \frac{1}{n}\right)^{n-1} = e^{-1},$$

then e^{-1} is a lower bound by

$$p = \Pr[X = 1] = \left(1 - \frac{1}{n}\right)^{n-1} \geq e^{-1}.$$

Let N be the random variable for the number of packets successfully sent until each server has successfully sent at least 1 packet. By the Coupon's Collector's Problem,

$$\mathbb{E}[N] = n \sum_{i \in [1, n]} i^{-1} = nH(n) = n \ln n + cn.$$

Let Y_i be the number of time slots between the $(i-1)^{\text{th}}$ successful transmission and the i^{th} successful transmission. Then $Y_1, Y_2, \dots$ are independent and geometrically distributed with success probability p , therefore $\mathbb{E}[Y_i] = 1/p \leq e$. Moreover, whichever server successfully transmitted a packet is independent of the waiting times Y_i , so the Wald argument applies to this sum. By theorem 1.13.18,

$$\mathbb{E} \left[\sum_{i \in [1, N]} Y_i \right] = \mathbb{E}[N] \cdot \mathbb{E}[Y_1] \leq e(n \ln n + cn).$$

Exercise 1.13.23 (Exercise 13.11 [MU17]). A parking-lot attendant has mixed up n keys for n cars. The n car owners arrive together. The attendant gives each owner a key according to a permutation chosen uniformly at random from all permutations. If an owner receives his key, he takes it and leaves; otherwise, he returns the key to the attendant. The attendant repeats the process with the remaining keys and car owners. This continues until all owners receive the keys to their cars. Let R be the number of rounds until all car owners receive their keys, and let X_i be the number of owners who receive their car keys in the i^{th} round. Show that

$$Y_i = \sum_{j \in [1, i]} (X_j - \mathbb{E}[X_j \mid X_1, \dots, X_{j-1}])$$

is a martingale, and use theorem 1.13.13 to derive $\mathbb{E}[R]$.

Proof. To show that $Y_1, Y_2, \dots$ is a martingale, we have

$$\begin{aligned} \mathbb{E}[Y_i \mid Y_1, \dots, Y_{i-1}] &= Y_{i-1} + \mathbb{E}[X_i - \mathbb{E}[X_i \mid X_1, \dots, X_{i-1}] \mid Y_1, \dots, Y_{i-1}] \\ &= Y_{i-1} + \mathbb{E}[X_i \mid Y_1, \dots, Y_{i-1}] - \mathbb{E}[X_i \mid Y_1, \dots, Y_{i-1}] = Y_{i-1}. \end{aligned}$$

Since both X_{i+1} and $\mathbb{E}[X_{i+1} \mid X_1, \dots, X_i]$ are in $[0, n]$, then by triangle inequality and tower property,

$$\begin{aligned} \mathbb{E}[|Y_{i+1} - Y_i| \mid Y_1, \dots, Y_i] &= \mathbb{E}[|X_{i+1} - \mathbb{E}[X_{i+1} \mid X_1, \dots, X_i]| \mid Y_1, \dots, Y_i] \\ &\leq 2 \mathbb{E}[X_{i+1} \mid Y_1, \dots, Y_i] \leq 2n, \end{aligned}$$

we can apply theorem 1.13.13, that

$$\mathbb{E}[Y_R] = \mathbb{E}[Y_1] = 0.$$

Moreover, by exercise 1.2.4, for $i \geq 1$,

$$\mathbb{E}[X_i \mid X_1, \dots, X_{i-1}] = 1.$$

Hence, by the stopping time R , we have

$$\mathbb{E}[Y_R] = \mathbb{E}\left[\sum_{i \in [1, R]} X_i - \sum_{i \in [1, R]} \mathbb{E}[X_i \mid X_1, \dots, X_{i-1}]\right] = \mathbb{E}\left[\sum_{i \in [1, R]} X_i\right] - \mathbb{E}[R] = 0.$$

Since by stopping time R ,

$$\sum_{i \in [1, R]} X_i = n,$$

we have $\mathbb{E}[R] = n$. □

Exercise 1.13.24 (Exercise 13.12 [MU17]). Alice and Bob play each other in a chess tournament, where the first player to win four games wins the match. The players are evenly matched, so the probability that each player wins each game is $1/2$, independent of other games. The number of minutes for each game is uniformly distributed over the integers in $[30, 60]$, again independent of other games. What is the expected time they spend playing this match?

Proof. Let $X_1, X_2, \dots$ be the time used in the i^{th} game with distribution X , then $\mathbb{E}[X] = 45$.

Let T be the stopping time, and WLOG let Alice win, as Bob wins symmetrically. Since Alice wins the final game, then there are $\binom{3}{0}$ ways to finish in 4 games, $\binom{4}{1}$ ways to finish in 5 games, $\binom{5}{2}$ ways to finish in 6 games, $\binom{6}{3}$ ways to finish in 7 games. The expected number of games is

$$\mathbb{E}[T] = \left(\binom{3}{0} + \binom{4}{1} + \binom{5}{2} + \binom{6}{3}\right)^{-1} \cdot \left(\binom{3}{0} \cdot 4 + \binom{4}{1} \cdot 5 + \binom{5}{2} \cdot 6 + \binom{6}{3} \cdot 7\right) = \frac{32}{5}.$$

By theorem 1.13.18,

$$\mathbb{E}\left[\sum_{i \in [1, T]} X_i\right] = \mathbb{E}[X] \cdot \mathbb{E}[T] = 288.$$

□

Exercise 1.13.25 (Exercise 13.13 [MU17]). Consider the following algorithm for sorting n numbers. Start by choosing one of the n numbers uniformly at random, and place it first. Then choose one of the remaining $n - 1$ numbers uniformly at random, and place it second. If the second number is smaller than the first, start over again from the beginning. Otherwise, next choose one of the remaining $n - 2$ numbers uniformly at random, place it third, and so on. The algorithm starts over from the beginning whenever it finds that the k^{th} item placed is smaller than the $(k - 1)^{\text{th}}$ item. Determine the expected number of times the algorithm tries to place a number, assuming that the input consists of n distinct numbers.

Proof. Let $X_1, X_2, \dots$ be the random variables of the lengths of the incremental sequences drawn from the n distinct numbers, then by the stopping time T , $X_T = n$. Since the success probability of drawing an incremental sequence of length n is $(n!)^{-1}$, then $\mathbb{E}[T] = n!$.

Let $Y_{[1,k]}$ be the random variables for the prefix sequence of length k , then we have

$$\Pr[Y_{[1,k]} = y_{[1,k]}] = \frac{(n-k)!}{n!},$$

as each permutation is uniformly distributed. Hence,

$$\Pr[X \geq k] = \binom{n}{k} \frac{(n-k)!}{n!} = \frac{1}{k!},$$

and the expectation of the length of incremental sequence is

$$\mathbb{E}[X] = \sum_{k \in [1,n]} \Pr[X \geq k] = \sum_{k \in [1,n]} \frac{1}{k!} \leq e,$$

the upper bound can be seen in section 1.5.1.

Since X_i are i.i.d., then we can apply theorem 1.13.18, the expected number of draws is

$$\mathbb{E}\left[\sum_{i \in [1,T]} X_i\right] = \mathbb{E}[T] \cdot \mathbb{E}[X] = n! \sum_{k \in [1,n]} \frac{1}{k!}.$$

□

Exercise 1.13.26 (Exercise 13.14 [MU17]). Suppose we are arranging a chain of n dominoes so that, once we are done, we can have them fall sequentially by knocking down the lead domino. Each time one tries to place a domino in the chain, there is some probability that it falls. In this case, one must start over from the very first domino.

- Call each time placing a domino a trial, succeeding with probability p . Use theorem 1.13.18, find the expected number of trials necessary before the arrangement is ready.
- Suppose instead that one can break the arrangement in k components, each of size n/k , in such a way so that once a component is complete, it will not fall for further dominoes. Find the expected number of trials necessary before the arrangement is ready.

Proof. Let $X_1, X_2, \dots$ be the longest length of a domino arrangement, and T be the stopping time, such that $X_T = n$. Since X_i is geometrically distributed for the number of steps to the first failure, and X_i are i.i.d., then $\mathbb{E}[X_i] = (1-p)^{-1}$. On the other hand, T is also geometrically distributed with success probability p^n , therefore $\mathbb{E}[T] = p^{-n}$.

Hence, by theorem 1.13.18,

$$\mathbb{E}\left[\sum_{i \in [1,T]} X_i\right] = \mathbb{E}[X_1] \cdot \mathbb{E}[T] = \frac{1}{(1-p)p^n}.$$

When there are k segments of sub-arrangements, the expected number of trials to finish a segment is $(1-p)^{-1}p^{-n/k}$, and by linearity of expectation, the expected number of trials to finish all pieces is $k(1-p)^{-1}p^{-n/k}$. □

Exercise 1.13.27 (Exercise 13.15 [MU17]). Use theorem 1.13.18 to derive the following $\mathbb{E}[N]$.

- Let $X_1, X_2, \dots$ be independent exponential random variables with $\mathbb{E}[X_i] = 1$. Given $k \in \mathbb{R}^+$, define

$$N = \min\{n : \sum_{i \in [1, n]} X_i > k\}.$$

- Let $Y_1, Y_2, \dots$ be independent uniform random variables in $(0, 1)$. Given $k \in (0, 1)$, define

$$N = \min\{n : \prod_{i \in [1, n]} Y_i < k\}.$$

Proof. Let $S_m = \sum_{i \in [1, m]} X_i$. Since $\mathbb{E}[X_i] = 1$, by theorem 1.13.18,

$$\mathbb{E}[S_N] = \mathbb{E}[X_1] \cdot \mathbb{E}[N] = \mathbb{E}[N].$$

Moreover, since X_i are exponential random variables, then by the memorylessness property, for any $a \geq 0$,

$$\mathbb{E}[X_i - a \mid X_i > a] = 1.$$

Given that $S_{N-1} \leq k$, and $S_N > k$, we have $X_N > k - S_{N-1}$, then

$$\mathbb{E}[X_N - (k - s) \mid X_N > k - s, S_{N-1} = s] = 1$$

holds for any $0 \leq s \leq k$, therefore $\mathbb{E}[N] = k + 1$.

For $Y_1, Y_2, \dots$, let $Z_i = -\ln Y_i$, then the stopping time is defined as

$$N = \min\{n : \sum_{i \in [1, n]} Z_i > -\ln k\},$$

and $\mathbb{E}[Z_i] = 1$ by

$$\mathbb{E}[Z_i] = \int_0^1 -\ln x \, dx = 1.$$

Let $Q_m = \sum_{i \in [1, m]} Z_i$, by theorem 1.13.18,

$$\mathbb{E}[Q_N] = \mathbb{E}[Z_1] \cdot \mathbb{E}[N] = \mathbb{E}[N].$$

Since $Q_{N-1} \leq -\ln k$, while $Q_N > -\ln k$, we have $Z_N > -\ln k - Q_{N-1}$. Moreover,

$$\Pr[Z_i > a] = \Pr[Y_i < e^{-a}] = e^{-a},$$

as they are the same event. Hence, for all $0 < s \leq -\ln k$,

$$\mathbb{E}[Z_N - (-\ln k - s) \mid Z_N > -\ln k - s, Q_{N-1} = s] = \frac{1}{\Pr[Y_N < k \cdot e^s]} \int_0^{k \cdot e^s} (-\ln x - (-\ln k - s)) \, dx = 1,$$

which proves that $\mathbb{E}[Q_N] = \mathbb{E}[N] = -\ln k + 1$.

Another way of showing is, Z_i is an exponential random variable with probability density function $f(x) = e^{-x}$, hence $\mathbb{E}[Z_i] = 1$. Applying the prior result, the threshold is $-\ln k$, and $\mathbb{E}[N] = -\ln k + 1$. $\square$

1.13.4 Azuma-Hoeffding's Inequality

Azuma-Hoeffding says a martingale with bounded increment per step has total gain $X_t - X_0$ concentrating around 0, even if the steps are not independent.

Theorem 1.13.28 (Azuma-Hoeffding's Inequality [Azu67]). *Let $X_0, X_1, \dots$ be a martingale such that*

$$|X_k - X_{k-1}| \leq c_k,$$

then for all $t \geq 1$ and $\varepsilon > 0$,

$$\Pr[|X_t - X_0| \geq \varepsilon] \leq 2 \exp\left(-\frac{\varepsilon^2}{2 \sum_{i \in [1, t]} c_i^2}\right).$$

Proof. The proof is similar to theorem 1.4.10, using moment-generating functions.

Start by bounding the upper tail. Let $Y_i = X_i - X_{i-1}$, such that $|Y_i| \leq c_i$, and since $X_0, X_1, \dots$ is a martingale,

$$\mathbb{E}[Y_i | X_0, \dots, X_{i-1}] = \mathbb{E}[X_i - X_{i-1} | X_0, \dots, X_{i-1}] = \mathbb{E}[X_i | X_0, \dots, X_{i-1}] - X_{i-1} = 0. \quad (1.26)$$

Write

$$Y_i = -c_i \cdot \frac{1 - Y_i/c_i}{2} + c_i \cdot \frac{1 + Y_i/c_i}{2}.$$

Since $\exp(\lambda Y_i)$ is a convex function, by theorem 1.2.5,

$$\begin{aligned} \exp(\lambda Y_i) &= \exp\left(-c_i \lambda \cdot \frac{1 - Y_i/c_i}{2} + c_i \lambda \cdot \frac{1 + Y_i/c_i}{2}\right) \\ &\leq \frac{1 - Y_i/c_i}{2} \cdot \exp(-c_i \lambda) + \frac{1 + Y_i/c_i}{2} \cdot \exp(c_i \lambda) \\ &= \frac{1}{2} (\exp(\lambda c_i) + \exp(-\lambda c_i)) + \frac{Y_i}{2c_i} (\exp(\lambda c_i) - \exp(-\lambda c_i)). \end{aligned}$$

Now consider

$$\mathbb{E}[\exp(\lambda Y_i) | X_0, \dots, X_{i-1}] \leq \frac{1}{2} (\exp(\lambda c_i) + \exp(-\lambda c_i)). \quad (1.27)$$

By section 1.5.1,

$$\frac{1}{2} (\exp(\lambda c_i) + \exp(-\lambda c_i)) = \sum_{k \geq 0} \frac{(\lambda c_i)^{2k}}{(2k)!}.$$

Since $(2k)! \geq 2^k \cdot k!$,

$$\frac{1}{2} (\exp(\lambda c_i) + \exp(-\lambda c_i)) = \sum_{k \geq 0} \frac{(\lambda c_i)^{2k}}{(2k)!} \leq \sum_{k \geq 0} \frac{(\lambda c_i)^{2k}}{2^k \cdot k!} = \sum_{k \geq 0} \frac{1}{k!} \cdot \left(\frac{(\lambda c_i)^2}{2}\right)^k = \exp\left(\frac{(\lambda c_i)^2}{2}\right). \quad (1.28)$$

To bound the exponential moment used in Markov's inequality, similar to the Chernoff and Hoeffding bounds,

$$\begin{aligned} \mathbb{E}[\exp(\lambda (X_t - X_0))] &= \mathbb{E}\left[\exp\left(\lambda \sum_{i \in [1, t]} Y_i\right)\right] = \mathbb{E}\left[\prod_{i \in [1, t]} \exp(\lambda Y_i)\right] = \mathbb{E}\left[\mathbb{E}\left[\prod_{i \in [1, t]} \exp(\lambda Y_i) \mid X_1, \dots, X_{t-1}\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\exp(\lambda Y_t) \cdot \prod_{i \in [1, t-1]} \exp(\lambda Y_i) \mid X_1, \dots, X_{t-1}\right]\right] \\ &= \mathbb{E}\left[\prod_{i \in [1, t-1]} \exp(\lambda Y_i) \cdot \mathbb{E}[\exp(\lambda Y_t) | X_0, \dots, X_{t-1}]\right] \\ &\leq \exp\left(\frac{\lambda^2 c_t^2}{2}\right) \cdot \mathbb{E}\left[\prod_{i \in [1, t-1]} \exp(\lambda Y_i)\right] \leq \exp\left(\frac{\lambda^2}{2} \sum_{i \in [1, t]} c_i^2\right), \end{aligned} \quad (1.29)$$

where the last inequality is by induction of eq. (1.27) and eq. (1.28). The rest follows the same method as the Chernoff or Hoeffding bounds,

$$\begin{aligned} \Pr[X_t - X_0 \geq \varepsilon] &= \Pr[\exp(\lambda(X_t - X_0)) \geq \exp(\lambda\varepsilon)] \\ &\leq \mathbb{E}[\exp(\lambda(X_t - X_0))] \cdot \exp(-\lambda\varepsilon) \\ &\leq \exp\left(\frac{\lambda^2}{2} \sum_{i \in [1, t]} c_i^2 - \lambda\varepsilon\right) \leq \exp\left(-\frac{\varepsilon^2}{2 \sum_{i \in [1, t]} c_i^2}\right), \end{aligned} \quad (1.30)$$

where the last inequality is by minimizing at $\lambda = \varepsilon / \sum_{i \in [1, t]} c_i^2$.

To bound the lower tail, let $Z_i = X_{i-1} - X_i$, then $\mathbb{E}[Z_i | X_0, \dots, X_{i-1}] = 0$ by the same proof in eq. (1.26), and

$$\mathbb{E}[\exp(\lambda Z_i) | X_0, \dots, X_{i-1}] \leq \exp\left(\frac{(\lambda c_i)^2}{2}\right)$$

by eq. (1.27) and eq. (1.28). We conclude the lower bound by eq. (1.29) and eq. (1.30). $\square$

Theorem 1.13.29 (Generalized Azuma-Hoeffding's Inequality [Azu67]). *Let $X_0, X_1, \dots$ be a martingale such that*

$$B_k \leq X_k - X_{k-1} \leq B_k + d_k,$$

for some constant d_k and random variables B_k that may be functions of $X_0, \dots, X_{k-1}$. Then for all $t \geq 1$ and $\varepsilon > 0$,

$$\Pr[|X_t - X_0| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2}{\sum_{i \in [1, t]} d_i^2}\right).$$

Proof. Start with the upper tail. Let $Y_i = X_i - X_{i-1}$, and $\mathbb{E}[Y_i | X_0, \dots, X_{i-1}] = 0$ by eq. (1.26). Consider bounding the exponential moment as in eq. (1.29):

$$\mathbb{E}[\exp(\lambda(X_t - X_0))] = \mathbb{E}\left[\prod_{i \in [1, t-1]} \exp(\lambda Y_i) \cdot \mathbb{E}[\exp(\lambda Y_t) | X_0, \dots, X_{t-1}]\right].$$

By lemma 1.4.8,

$$\mathbb{E}[\exp(\lambda Y_t) | X_0, \dots, X_{t-1}] \leq \exp\left(\frac{\lambda^2 d_t^2}{8}\right).$$

By induction,

$$\mathbb{E}[\exp(\lambda(X_t - X_0))] \leq \exp\left(\frac{\lambda^2}{8} \sum_{i \in [1, t]} d_i^2\right). \quad (1.31)$$

We conclude by

$$\begin{aligned} \Pr[X_t - X_0 \geq \varepsilon] &= \Pr[\exp(\lambda(X_t - X_0)) \geq \exp(\lambda\varepsilon)] \\ &\leq \mathbb{E}[\exp(\lambda(X_t - X_0))] \cdot \exp(-\lambda\varepsilon) \\ &\leq \exp\left(\frac{\lambda^2}{8} \sum_{i \in [1, t]} d_i^2 - \lambda\varepsilon\right) \leq \exp\left(-\frac{2\varepsilon^2}{\sum_{i \in [1, t]} d_i^2}\right), \end{aligned} \quad (1.32)$$

where the last inequality is by minimizing at $\lambda = 4\varepsilon / \sum_{i \in [1, t]} d_i^2$.

To bound the lower tail, let $Z_i = X_{i-1} - X_i$, then $\mathbb{E}[Z_i | X_0, \dots, X_{i-1}] = 0$ by the same proof in eq. (1.26), and

$$\mathbb{E}[\exp(\lambda Z_i) | X_0, \dots, X_{i-1}] \leq \exp\left(\frac{\lambda^2 d_i^2}{8}\right)$$

by lemma 1.4.8. We conclude the lower bound by eq. (1.31) and eq. (1.32). $\square$

Exercise 1.13.30 (Exercise 13.17 [MU17]). Given a bag with r red balls and g green balls, suppose that we uniformly sample n balls from the bag without replacement. Use Azuma-Hoeffding to show that the number of red balls in the sample concentrates tightly around $nr/(r+g)$.

Proof. Let $X_1, \dots, X_n$ be 0-1 random variables indicating if the i^{th} draw is a red ball, and let $Y = \sum_{i \in [1, n]} X_i$ be the random variable for the total number of red balls drawn. Then

$$Z_0 = \mathbb{E}[Y], \quad Z_k = \mathbb{E}[Y \mid X_1, \dots, X_k]$$

is a Doob martingale. Let $Y_k = \sum_{i \in [1, k]} X_i$, such that $Y_n = Y$. Then

$$\begin{aligned} Z_{k-1} &= \mathbb{E}[Y \mid X_1, \dots, X_{k-1}] = Y_{k-1} + (n - k + 1) \cdot \frac{r - Y_{k-1}}{r + g - k + 1}, \\ Z_k &= \mathbb{E}[Y \mid X_1, \dots, X_k] = Y_k + (n - k) \cdot \frac{r - Y_k}{r + g - k}. \end{aligned}$$

The difference is

$$\begin{aligned} Z_k - Z_{k-1} &= Y_k + (n - k) \cdot \frac{r - Y_k}{r + g - k} - \left(Y_{k-1} + (n - k + 1) \cdot \frac{r - Y_{k-1}}{r + g - k + 1} \right) \\ &= \frac{r + g - n}{r + g - k} \left(X_k + \frac{Y_{k-1} - r}{r + g - k + 1} \right) = \frac{r + g - n}{r + g - k + 1} \left(X_k + \frac{Y_k - r}{r + g - k} \right) \end{aligned}$$

The upper bound is

$$\begin{aligned} Z_k - Z_{k-1} &= \frac{r + g - n}{r + g - k} \left(X_k + \frac{Y_{k-1} - r}{r + g - k + 1} \right) \\ &\leq X_k + \frac{Y_{k-1} - r}{r + g - k + 1} \leq X_k \leq 1. \end{aligned}$$

The lower bound is

$$\begin{aligned} Z_k - Z_{k-1} &= \frac{r + g - n}{r + g - k + 1} \left(X_k + \frac{Y_k - r}{r + g - k} \right) \\ &\geq \frac{r + g - n}{r + g - k + 1} \cdot \frac{Y_k - r}{r + g - k} \geq -\frac{r - Y_k}{r + g - k} \geq -1. \end{aligned}$$

Given that $|Z_k - Z_{k-1}| \leq 1$, and $\mathbb{E}[Y] = n \cdot r/(r+g)$ by linearity of expectation, we conclude with theorem 1.13.28

$$\Pr[|Y - \mathbb{E}[Y]| \geq \varepsilon] \leq 2 \exp\left(-\frac{\varepsilon^2}{2n}\right).$$

□

Remark 1.13.31. Interestingly, if we let $X_0, \dots, X_n$ be the fraction of the remaining red balls in exercise 1.13.30, it is also a martingale. Indeed,

$$\begin{aligned} \mathbb{E}[X_{k+1} \mid X_0, \dots, X_k] &= X_k \cdot \frac{(r + g - k)X_k - 1}{r + g - k - 1} + (1 - X_k) \cdot \frac{(r + g - k)X_k}{r + g - k - 1} \\ &= \frac{-X_k}{r + g - k - 1} + \frac{(r + g - k)X_k}{r + g - k - 1} = X_k. \end{aligned}$$

Moreover, $X_0, \dots, X_n$ is a Doob martingale with respect to itself, and lemma 1.13.9,

$$\mathbb{E}[X_n] = \mathbb{E}[X_0] = \frac{r}{r + g}.$$

The fractional martingale can be viewed as an inverse of example 1.13.11.

McDiarmid's inequality says that, for a function that is not too sensitive, namely changing any one of its inputs will not change its value by much, if the input variables are mutually independent, then the function value depending on the independent random inputs should concentrate around its expectation.

Definition 1.13.32 (Lipschitz Condition [McD89]). A function $f(X_1, \dots, X_n)$ is said to satisfy a Lipschitz condition with bound c if for any i and any set of values $x_1, \dots, x_n$ and y_i ,

$$|f(x_1, \dots, y_i, \dots, x_n) - f(x_1, \dots, x_i, \dots, x_n)| \leq c.$$

That is, changing the value of any single coordinate can change the function value by at most c .

Let

$$Z_0 = \mathbb{E} [f(X_1, \dots, X_n)], \\ Z_k = \mathbb{E} [f(X_1, \dots, X_n) \mid X_1, \dots, X_k].$$

The sequence $Z_0, Z_1, \dots, Z_n$ is a Doob martingale. Moreover, we claim if $X_1, \dots, X_n$ are mutually independent, then there exist random variables $B_1, \dots, B_n$ such that for all $k \in [1, n]$,

$$B_k \leq Z_k - Z_{k-1} \leq B_k + c.$$

We have a counterexample for the dependent case: Let $X_1, \dots, X_n$ be the random variables drawn from the same fair coin toss Y , and $f(X_1, \dots, X_n) = (\sum_{i \in [1, n]} X_i) / n$, with a Lipschitz condition bound $c = 1/n$. Then,

$$Z_0 = \mathbb{E} \left[\frac{1}{n} \sum_{i \in [1, n]} X_i \right] = \frac{1}{2}, \quad Z_1 = \mathbb{E} \left[\frac{1}{n} \sum_{i \in [1, n]} X_i \mid X_1 \right] = Y.$$

Now $Y - 1/2$ has a range much larger than $1/n$, so the claim of existence of $B_1, \dots, B_n$ does not hold.

Another counterexample is as follows. Let $X_1, \dots, X_{k-1}$ be independent random variables drawn from fair coin tosses, and let $X_k, \dots, X_n$ be random variables equal to the same fair coin toss Y . Let $\ell = n - k + 1$, so

$$Z_{k-1} = \mathbb{E} \left[\frac{1}{n} \sum_{i \in [1, n]} X_i \mid X_1, \dots, X_{k-1} \right] = \frac{1}{n} \sum_{i \in [1, k-1]} X_i + \frac{\ell}{2n}, \quad Z_k = \mathbb{E} \left[\frac{1}{n} \sum_{i \in [1, n]} X_i \mid X_1, \dots, X_k \right] = \frac{1}{n} \sum_{i \in [1, k-1]} X_i + \frac{\ell Y}{n}.$$

The difference $\ell/n \cdot (Y - 1/2)$ has a range $\ell/n \geq 1/n$, so the claim fails.

Theorem 1.13.33 (McDiarmid's Inequality [McD89]). Let f be a function on n variables that satisfies the Lipschitz condition in definition 1.13.32 with bound c . Let $X_1, \dots, X_n$ be independent random variables that are in the domain of f . Then

$$\Pr [|f(X_1, \dots, X_n) - \mathbb{E} [f(X_1, \dots, X_n)]| \geq \varepsilon] \leq 2 \exp \left(-\frac{2\varepsilon^2}{nc^2} \right).$$

Proof. Let S_k denote $X_1, \dots, X_k$, simplify $f(X_1, \dots, X_n)$ as $f(\bar{X})$, and let $f_k(\bar{X}, x) = f(X_1, \dots, X_{k-1}, x, X_{k+1}, \dots, X_n)$.

Consider the upper bound and the lower bound of $Z_k - Z_{k-1}$, where $Z_k - Z_{k-1}$ is upper bounded by

$$\sup_x \mathbb{E} [f(\bar{X}) \mid S_{k-1}, X_k = x] - \mathbb{E} [f(\bar{X}) \mid S_{k-1}] = \sup_x ((Z_k \mid X_k = x) - Z_{k-1}),$$

and $Z_k - Z_{k-1}$ is lower bounded by B_k , defined as¹⁷

$$\inf_x \mathbb{E} [f(\bar{X}) \mid S_{k-1}, X_k = x] - \mathbb{E} [f(\bar{X}) \mid S_{k-1}] = \inf_x ((Z_k \mid X_k = x) - Z_{k-1}).$$

¹⁷If $\bar{X}$ here have support (the possible value taken) being finite number of values, we can use max and min to substitute sup and inf.

Then consider the range between the upper and lower bound of $Z_k - Z_{k-1}$:

$$\begin{aligned} \sup_x (Z_k \mid X_k = x) - \inf_x (Z_k \mid X_k = x) &= \sup_x \mathbb{E} \left[f(\bar{X}) \mid S_{k-1}, X_k = x \right] - \inf_x \mathbb{E} \left[f(\bar{X}) \mid S_{k-1}, X_k = x \right] \\ &= \sup_{x,y} \left(\mathbb{E} \left[f(\bar{X}) \mid S_{k-1}, X_k = x \right] - \mathbb{E} \left[f(\bar{X}) \mid S_{k-1}, X_k = y \right] \right) \\ &= \sup_{x,y} \mathbb{E} \left[f_k(\bar{X}, x) - f_k(\bar{X}, y) \mid S_{k-1} \right], \end{aligned} \quad (1.33)$$

where the last equality holds because X_k is independent of $\bar{X} \setminus X_k$; this follows from conditional probability,

$$\begin{aligned} \Pr \left[\bar{X} = (z_1, \dots, z_{k-1}, x, \dots, z_n) \mid S_{k-1} = (z_1, \dots, z_{k-1}), X_k = x \right] &= \\ \Pr \left[\bar{X} \setminus X_k = (z_1, \dots, z_{k-1}, z_{k+1}, \dots, z_n) \mid S_{k-1} = (z_1, \dots, z_{k-1}) \right]. \end{aligned}$$

Since f satisfies the Lipschitz condition with bound c , and $X_1, \dots, X_n$ are independent, eq. (1.33) is at most c . We conclude the proof with theorem 1.13.29. $\square$

Remark 1.13.34. McDiarmid is an instance of Azuma-Hoeffding with a Doob martingale from independent random variables and a function satisfying the Lipschitz condition. The condition to apply Azuma-Hoeffding is to have bounded increments, which are guaranteed by independence and Lipschitz condition bound c . Independence says that revealing X_k does not alter the conditional probability distribution of the unrevealed random variables, as they do not depend on X_k . Therefore, the effect of revealing X_k can be compared as if only the k^{th} coordinate changes, such that the martingale difference is bounded by the same c .

Exercise 1.13.35 (Exercise 13.8 [MU17]). In the bin-packing problem, we are given items with sizes $a_1, \dots, a_n$ with $a_i \in [0, 1]$. The goal is to pack them into minimum number of bins, with each bin being able to hold any collection of items whose total sizes sum to at most 1. Suppose that each of the a_i is chosen independently according to some distribution, let P be the number of bins required in the best packing of the resulting items. Prove that

$$\Pr[|P - \mathbb{E}[P]| \geq \varepsilon] \leq 2 \exp(-2\varepsilon^2/n).$$

Proof. Let $f(a_1, \dots, a_n)$ be the minimum number of bins needed to pack n items. Then the Lipschitz condition bound of f is at most 1, since a bin has capacity at most 1. Starting with $f(a_1, \dots, a_k = 0, \dots, a_n)$ bins, and changing to $a_k = 1$, the worst one can do is to put a_k in a separate bin and preserve the arrangement for the other items. Hence

$$f(a_1, \dots, a_k = 1, \dots, a_n) \leq f(a_1, \dots, a_k = 0, \dots, a_n) + 1.$$

We then conclude the proof with theorem 1.13.33 with $c = 1$. $\square$

Exercise 1.13.36 (Exercise 13.9 [MU17]). Consider an n -cube with $|V| = 2^n$. Let $S \subseteq V$ with $|S| > 0$, and $x \stackrel{R}{\leftarrow} V$. Let $D(x, S)$ be the minimum number of coordinates in which x and y differ over all points $y \in S$. Give a bound on

$$\Pr[|D(x, S) - \mathbb{E}[D(x, S)]| > \varepsilon].$$

Proof. Let $\bar{X} = (X_1, \dots, X_n)$ be independent fair coin tosses, $f(\bar{X}) = D(\bar{X}, S)$, and $f_k(\bar{x}, y) = f(x_1, \dots, x_{k-1}, y, \dots, x_n)$. Let d be the Hamming distance from $\bar{x} \stackrel{R}{\leftarrow} \{0, 1\}^n$ to $\bar{w} \in S$, which is $f(\bar{x})$, and we have the following:

- If $\bar{w}$ is still the closest element in S , then either $w_k = y$ and $f_k(\bar{x}, y) = d - 1$, or $w_k \neq y$ and $f_k(\bar{x}, y) = d + 1$.
- $f_k(\bar{x}, y) = d$, where $y \neq w_k$, but there is an element $\bar{w}' \in S$ that disagrees in $d + 1$ coordinates with $\bar{x}$, and $w'_k = y$.

Alternatively, changing one coordinate either increases or decreases the distance to any element in S by 1. Thus, the Lipschitz condition bound $c = 1$. We conclude with theorem 1.13.33 that

$$\Pr[|D(x, S) - \mathbb{E}[D(x, S)]| > \varepsilon] < 2 \exp(-2\varepsilon^2/n).$$

$\square$

Exercise 1.13.37 (Exercise 13.10 [MU17]). Consider generalizing Hoeffding's bound to independent random variables ranging in $[0, 1]$, and derive a tail bound for $S_n = \sum_{i \in [1, n]} X_i$.

Proof. Consider a Doob martingale over

$$Z_0 = \mathbb{E}[S_n], \quad Z_k = \mathbb{E}[S_n \mid X_1, \dots, X_k],$$

and S_n has Lipschitz condition bound $c = 1$. Immediately by theorem 1.13.33,

$$\Pr[|S_n - \mathbb{E}[S_n]| \geq \varepsilon] \leq 2 \exp(-2\varepsilon^2/n).$$

□

Example 1.13.38 (DNA Pattern Matching). Let $\bar{X} = (X_1, \dots, X_n)$ be a sequence of characters sampled independently and uniformly at random over an alphabet Σ , and let $\mathbf{b} \in \Sigma^k$ be a fixed pattern string of k characters. Let F be the number of occurrences of $\mathbf{b}$ in the random string $\bar{X}$. By linearity of expectation, $\mathbb{E}[F] = (n - k + 1) \cdot s^{-k}$. Let

$$Z_0 = \mathbb{E}[F], \quad Z_i = \mathbb{E}[F \mid X_1, \dots, X_i]$$

be a Doob martingale. Using $|Z_i - Z_{i-1}| \leq k$, theorem 1.13.28 gives

$$\Pr[|F - \mathbb{E}[F]| \geq \varepsilon] \leq 2 \exp\left(-\frac{\varepsilon^2}{2nk^2}\right).$$

Taking a closer look, $|Z_i - Z_{i-1}| \leq k$ is rather wasteful, as F as a function over $\bar{X}$ has Lipschitz condition bound $c = k$. Thus, by theorem 1.13.33 or theorem 1.13.29,

$$\Pr[|F - \mathbb{E}[F]| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2}{nk^2}\right).$$

Exercise 1.13.39 (Exercise 13.16 [MU17]). A subsequence of a string is any string that can be obtained by deleting characters. Consider $\mathbf{x}, \mathbf{y} \stackrel{\text{R}}{\leftarrow} \{0, 1\}^n$, and the longest common subsequence (LCS) of $\mathbf{x}$ and $\mathbf{y}$. Show that the expected length of the LCS is in $(c_1 n, c_2 n)$ where $c_1 > 1/2$ and $c_2 < 1$ when n is sufficiently large¹⁸. Use theorem 1.13.33 to show that the length of the LCS is sharply concentrated around its mean.

Proof. Let L_ℓ be the length of the LCS of two binary strings of length ℓ . Then we have

$$L_\ell + 1 > \mathbb{E}[L_{\ell+1} \mid L_\ell] > L_\ell + \frac{1}{2},$$

as $x_{\ell+1} = y_{\ell+1}$ with probability $1/2$, and $L_{\ell+1} = L_\ell + 1$ with non-zero probability when $x_{\ell+1} \neq y_{\ell+1}$.

Let $Z_i = (X_i, Y_i)$ for the i^{th} bit in $\mathbf{x}$ and $\mathbf{y}$. The Doob martingale is constructed by

$$Z_0 = \mathbb{E}[L_n], \quad Z_k = \mathbb{E}[L_n \mid Z_1, \dots, Z_k].$$

The Lipschitz condition bound is 1, as any common subsequence's length increases or decreases by 1:

- Either the previous LCS remains, no matter what new assignment to the k^{th} coordinate.
- Or the LCS length is unchanged, but altering the k^{th} bits improves another common subsequence.

Hence, we conclude with theorem 1.13.33 by

$$\Pr[|L_n - \mathbb{E}[L_n]| \geq \varepsilon] \leq 2 \exp(-2\varepsilon^2/n).$$

□

¹⁸There are interesting open problems on pushing the bounds closer.

Example 1.13.40 (Balls and Bins). Recall in m balls and n bins model, we let $\{X_i\}_{i \in [1, m]}$ be the random variables for the bin into which the i^{th} ball falls, and let F be the number of empty bins after m bins are thrown. Then

$$Z_0 = \mathbb{E}[F], \quad Z_k = \mathbb{E}[F \mid X_1, \dots, X_k]$$

is a Doob martingale. Moreover, the Lipschitz condition bound is 1, as changing which bin the i^{th} ball lands increases or decreases the number of empty bins after m balls at most by 1. We conclude with theorem 1.13.33 that

$$\Pr[|F - \mathbb{E}[F]| \geq \varepsilon] \leq 2 \exp(-2\varepsilon^2/m),$$

and the expected number of empty bins is $\mathbb{E}[F] = n(1 - 1/n)^m$.

Exercise 1.13.41 (Exercise 13.18 [MU17]). We have shown before in Bloom filter (e.g., exercise 1.5.29) that the fractions of entries that are 0 in a Bloom filter is concentrated around $p' = (1 - 1/n)^{km}$, where m is the number of data items, k is the number of hash functions, and n is the size of the Bloom filter in bits. Derive a similar concentration result using a martingale inequality.

Proof. This is a direct adaptation of example 1.13.40, using a km balls and n bins model. Let $\{X_i\}_{i \in [1, km]}$ be the 0-1 random variables for which bit the hash hits, and Y be the number of 0 bits after m items, then $\mathbb{E}[Y] = p'n$, and

$$Z_0 = \mathbb{E}[Y], \quad Z_k = \mathbb{E}[Y \mid X_1, \dots, X_k]$$

is a Doob martingale, and Y is a function over $\{X_i\}_{i \in [1, km]}$ with Lipschitz condition bound $c = 1$. By theorem 1.13.33,

$$\Pr[|Y - \mathbb{E}[Y]| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2}{km}\right).$$

□

Exercise 1.13.42 (Exercise 13.20 [MU17]). We improve the bound in example 1.13.40 and reuse all defined notations. Let A_i denote the number of bins that are empty after the i^{th} ball is thrown.

- Show that $Z_{i-1} = A_{i-1}(1 - 1/n)^{m-i+1}$.
- Show that, if the i^{th} ball lands in a bin that is empty, $Z_i = (A_{i-1} - 1)(1 - 1/n)^{m-i}$.
- Show that, if the i^{th} ball lands in a bin that is not empty, $Z_i = A_{i-1}(1 - 1/n)^{m-i}$.
- Show that theorem 1.13.29 applies with $d_i = (1 - 1/n)^{m-i}$, and

$$\Pr[|F - \mathbb{E}[F]| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2(2n-1)}{n^2 - \mathbb{E}[F]^2}\right).$$

Proof. The probability of a bin remaining empty after $m - i + 1$ balls is $(1 - 1/n)^{m-i+1}$. By linearity of expectation,

$$Z_{i-1} = \mathbb{E}[F \mid X_1, \dots, X_{i-1}] = A_{i-1}(1 - 1/n)^{m-i+1}.$$

Whether the i^{th} ball lands into an empty bin or not determines the number of empty bins onwards after the i^{th} ball. If the i^{th} ball lands into an empty bin, $A_i = A_{i-1} - 1$; otherwise, $A_i = A_{i-1}$. Still by linearity of expectation,

$$Z_i = \mathbb{E}[F \mid X_1, \dots, X_i] = A_i(1 - 1/n)^{m-i}.$$

Let $B_i = (A_{i-1}/n - 1) \cdot (1 - 1/n)^{m-i}$, then $d_i = (1 - 1/n)^{m-i}$ such that $B_i \leq Z_i - Z_{i-1} \leq B_i + d_i$. By theorem 1.13.29,

$$\sum_{i \in [1, m]} d_i^2 = \sum_{j \in [0, m-1]} \left(1 - \frac{1}{n}\right)^{2j} = \left(1 - \left(1 - \frac{1}{n}\right)^{2m}\right) \cdot \left(1 - \left(1 - \frac{1}{n}\right)^2\right)^{-1} = \frac{n^2 - n^2(1 - 1/n)^{2m}}{2n - 1} = \frac{n^2 - \mathbb{E}[F]^2}{2n - 1},$$

which proves

$$\Pr[|F - \mathbb{E}[F]| \geq \varepsilon] \leq 2 \exp\left(-\frac{2\varepsilon^2(2n-1)}{n^2 - \mathbb{E}[F]^2}\right).$$

□

Remark 1.13.43. The key step is not McDiarmid itself, but showing the bounded increment in the martingale; it can either be directly shown, or by independence and Lipschitz condition bound.

Example 1.13.44 (Chromatic Number [SS87]). Given a random graph G in $G_{n,p}$, the chromatic number $\chi(G)$ is the minimum number of colors needed to color all vertices of the graph, so no adjacent vertices have the same color.

Let $G_i \subseteq G$ be the random subgraph induced by the set of vertices $[1, i]$, so G_i depends on G_{i-1} . Then

$$Z_0 = \mathbb{E}[\chi(G)], \quad Z_k = \mathbb{E}[\chi(G) \mid G_1, \dots, G_k]$$

forms a Doob martingale.

Now consider bounding $Z_k - Z_{k-1}$. For $H = G - v_k$, formed by removing the k^{th} vertex,

$$\chi(H) \leq \chi(G) \leq \chi(H) + 1, \quad (1.34)$$

where the first inequality holds because H has one less vertex than G , and the second inequality holds because adding v_k to H introduces at most one more new color.

Hence, consider G and G' differing only on the edges between v_k and $V \setminus \{v_k\}$. Then $\chi(G), \chi(G') \in [\chi(H), \chi(H)+1]$. By eq. (1.34), it is immediate that

$$\mathbb{E}[\chi(H) \mid G_1, \dots, G_{k-1}] \leq \mathbb{E}[\chi(G) \mid G_1, \dots, G_{k-1}] \leq \mathbb{E}[\chi(H) \mid G_1, \dots, G_{k-1}] + 1.$$

Moreover, since for any G and G' differing on the edges to v_k , $\chi(G), \chi(G') \in [\chi(H), \chi(H) + 1]$, for any G_k ,

$$\mathbb{E}[\chi(H) \mid G_1, \dots, G_{k-1}] \leq \mathbb{E}[\chi(G) \mid G_1, \dots, G_k] \leq \mathbb{E}[\chi(H) \mid G_1, \dots, G_{k-1}] + 1.$$

Thus $d_i = 1$ for theorem 1.13.29, and we conclude that

$$\Pr[|\chi(G) - \mathbb{E}[\chi(G)]| \geq \varepsilon\sqrt{n}] \leq 2\exp(-2\varepsilon^2).$$

Exercise 1.13.45 (Exercise 13.19 [MU17]). Consider a random graph G in $G_{n,N}$ where $N = cn$ for $c > 0$. Let X be the number of isolated vertices, namely vertices of degree 0. Determine $\mathbb{E}[X]$ and show

$$\Pr[|X - \mathbb{E}[X]| \geq 2\varepsilon\sqrt{cn}] \leq 2\exp(-\varepsilon^2/2).$$

Proof. Let X_i be the 0-1 random variable for the i^{th} vertex being unconnected to any of N edges, and $\binom{n}{2}$ is the total number of possible edges. Then there are

$$\binom{\binom{n}{2} - (n-1)}{N} = \binom{\binom{n-1}{2}}{N}$$

graphs isolating the i^{th} vertex, and

$$\Pr[X_i = 1] = \frac{\binom{n-1}{2}}{\binom{n}{2}} \left(\frac{\binom{n-1}{2}}{\binom{n}{2}} \right)^{-1}.$$

By linearity of expectation,

$$\mathbb{E}\left[X = \sum_{i \in [1, n]} X_i\right] = n \left(\frac{\binom{n-1}{2}}{\binom{n}{2}} \right) \left(\frac{\binom{n-1}{2}}{\binom{n}{2}} \right)^{-1}.$$

Let $\{E_i\}_{i \in [1, N]}$ be random variables in $[1, \binom{n}{2}]$ for the location of the edge, and

$$Z_0 = \mathbb{E}[X], \quad Z_k = \mathbb{E}[X \mid E_1, \dots, E_k]$$

is a Doob martingale. An edge can decrease Z_{k-1} by 2 by joining 2 isolated vertices, so $|Z_k - Z_{k-1}| \leq c = 2$, and we conclude with theorem 1.13.28

$$\Pr[|X - \mathbb{E}[X]| \geq 2\varepsilon\sqrt{cn}] \leq 2\exp\left(-\frac{4\varepsilon^2 N}{2 \sum_{i \in [1, N]} c^2}\right) = 2\exp(-\varepsilon^2/2).$$

An alternative way to further bound $Z_k - Z_{k-1}$ follows. Let I_k be the number of isolated vertices removed by the E_k , A_k be the number of isolated vertices after the E_k , such that $A_k = A_{k-1} - I_k$. A vertex remains isolated with probability

$$\Pr[X_i = 1 \mid E_1, \dots, E_k] = \left(\frac{\binom{n-1}{2} - k}{N - k} \right) \left(\frac{\binom{n}{2} - k}{N - k} \right)^{-1} = q_k.$$

Moreover, we have the following iteration

$$q_{k-1} = \frac{\binom{n-1}{2} - (k-1)}{\binom{n}{2} - (k-1)} \cdot q_k.$$

By linearity of expectation, $Z_k = A_k \cdot q_k$. Then

$$Z_k - Z_{k-1} = (A_{k-1} - I_k) \cdot q_k - A_{k-1} \cdot \frac{\binom{n-1}{2} - (k-1)}{\binom{n}{2} - (k-1)} \cdot q_k = q_k \cdot \left(\frac{n-1}{\binom{n}{2} - (k-1)} \cdot A_{k-1} - I_k \right). \quad (1.35)$$

Notice that

$$\mathbb{E}[I_k \mid E_1, \dots, E_{k-1}] = \frac{n-1}{\binom{n}{2} - (k-1)} \cdot A_{k-1},$$

and I_k conditioned on $E_1, \dots, E_{k-1}$ is in $[0, 2]$. Hence eq. (1.35) is of interval length $d_k = 2q_k \leq 2$, and we conclude with theorem 1.13.29

$$\Pr[|X - \mathbb{E}[X]| \geq 2\varepsilon\sqrt{cn}] \leq 2 \exp\left(-\frac{2 \cdot 4\varepsilon^2 N}{\sum_{k \in [1, N]} d_k^2}\right) = 2 \exp(-2\varepsilon^2).$$

A few notes from prior attempts follow. We have

$$\begin{aligned} \mathbb{E}[X \mid E_1, \dots, E_{k-1}] &= \mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 2] \cdot \Pr[I_k = 2 \mid E_1, \dots, E_{k-1}] \\ &\quad + \mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 1] \cdot \Pr[I_k = 1 \mid E_1, \dots, E_{k-1}] \\ &\quad + \mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 0] \cdot \Pr[I_k = 0 \mid E_1, \dots, E_{k-1}]. \end{aligned}$$

Let Y be the number of isolated vertices for a graph in $G_{n, N-1}$. We have

$$\mathbb{E}[X \mid E_1, \dots, E_{k-1}] \leq \mathbb{E}[Y \mid E_1, \dots, E_{k-1}],$$

since the N^{th} edge can still remove isolated vertices by

$$\mathbb{E}[X \mid E_1, \dots, E_{N-1}] \leq \mathbb{E}[Y \mid E_1, \dots, E_{N-1}].$$

Let B_k be the number of edges removing 2 isolated vertices, C_k be the number of edges removing 1 isolated vertex, and D_k be the number of edges removing 0 isolated vertex. For $G_{n, N}$ and $G_{n, N-1}$, the invariant is

$$B_k + C_k + D_k + k = \binom{n}{2}.$$

Then, we have

$$\mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 2] \geq \mathbb{E}[Y \mid E_1, \dots, E_{k-1}] - 2,$$

as $D_k \geq D_{k-1}$, $B_k \leq B_{k-1}$, and $B_{k-1} + C_{k-1} + D_{k-1} = B_k + C_k + D_k + 1$ in this case. Similarly,

$$\mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 1] \geq \mathbb{E}[Y \mid E_1, \dots, E_{k-1}] - 1,$$

as $D_k \geq D_{k-1}$, $B_k \leq B_{k-1}$, and $B_{k-1} + C_{k-1} + D_{k-1} = B_k + C_k + D_k + 1$. Then, we have

$$\mathbb{E}[X \mid E_1, \dots, E_{k-1}, I_k = 0] \leq \mathbb{E}[Y \mid E_1, \dots, E_{k-1}],$$

where $B_k = B_{k-1}$, $C_k = C_{k-1}$, and $D_k = D_{k-1} - 1$. X and Y have $N - k$ edges to go, but X has one fewer edge not decreasing the number of isolated vertices, making it more likely to decrease the number of isolated vertices. $\square$

1.14 Sample Complexity, VC Dimension, and Rademacher Complexity

TODO: ...

1.15 Pairwise Independence and Universal Hash Function

1.15.1 Pairwise Independence

1.15.2 Chebyshev's Inequality for Pairwise Independent Random Variables

1.15.3 Universal Families of Hash Functions

1.15.4 Application in Finding Heavy Hitters in Data Streams

TODO: *after balls and bins, Poisson, and maybe Probabilistic methods.*

1.16 Power Laws and Related Distributions

TODO: ...

1.17 **Balanced Allocations and Cuckoo Hashing**

TODO: *should be after balls bins?*

2 Pseudorandomness

TODO: *Vadhan* [[Vad12](#)] *Pseudorandomness*

3 Probabilistic Method

TODO: *Alon Spencer* [[AS16](#)] *Probabilistic Method*

4 Various Papers

4.1 Leftover Hash Lemma and HILL Entropy

TODO: *after entropy in [MU17], learn LHL [ILL89] and HILL entropy [HILL99]*

4.2 Discrete Gaussian Leftover Hash Lemma

TODO: *should learn from* [AGHS13]

4.3 Schwartz-Zippel by Dana Moshkovitz

The geometric viewpoint in [Mos10] is to foliate $\mathbb{F}_q^n$ by parallel lines, and count their intersections with the zero set of a polynomial. The key is to choose a direction for which no whole line lies in the zero set.

The existence of the direction uses only the following elementary fact.

Lemma 4.3.1 (Nonzero evaluation [Mos10]). *Let $\mathbb{F}_q$ be a finite field with q elements, and $f \in \mathbb{F}_q[X_1, \dots, X_n]$ be nonzero with individual degrees less than q . Then f has a nonzero evaluation in $\mathbb{F}_q^n$.*

Proof. Fixing coordinates successively while keeping the remaining polynomial nonzero is always possible: if every substitution $X_i = a$ for $a \in \mathbb{F}_q$ over a nonzero remaining polynomial P , made P identically zero, then

$$\prod_{a \in \mathbb{F}_q} (X_i - a) = X_i^q - X_i \mid P,$$

contradicting $\deg_{X_i} P < q$. After all coordinates are fixed, the resulting value is nonzero. $\square$

We now restate Schwartz-Zippel [Sch80, Zip79] in Moshkovitz's viewpoint [Mos10].

Lemma 4.3.2 (Schwartz-Zippel [Sch80, Zip79]). *Let $\mathbb{F}_q$ be a finite field with q elements, and $f \in \mathbb{F}_q[X_1, \dots, X_n]$ be nonzero of total degree d , with $n \geq 1$. From uniformly random $\mathbf{x} \xleftarrow{R} \mathbb{F}_q^n$,*

$$\Pr[f(\mathbf{x}) = 0] \leq \frac{d}{q}.$$

Equivalently, as for the zero set of f over $\mathbb{F}_q^n$,

$$\left| \{\mathbf{x} \in \mathbb{F}_q^n : f(\mathbf{x}) = 0\} \right| \leq dq^{n-1}.$$

Proof. The cases $d = 0$ and $d \geq q$ are immediate, so assume $1 \leq d < q$.

Write $f = g + h$, where nonzero g consists of all the degree- d monomials, and thus $\deg h < d$. By lemma 4.3.1, g has a nonzero evaluation over $\mathbb{F}_q^n$, and by g being all degree- d monomials, $g(\mathbf{0}) = 0$. Hence, there exists $\mathbf{y} \in \mathbb{F}_q^n \setminus \{\mathbf{0}\}$ such that $g(\mathbf{y}) \neq 0$.

For any basepoint $\mathbf{x}$, substitution gives $p(t) = f(\mathbf{x} + t\mathbf{y})$ that $\deg_t p = d$, whose degree- d term has coefficient $g(\mathbf{y})$, and has at most d roots.

There are q^{n-1} distinct parallel lines $\mathbf{x} + t\mathbf{y}$, as $|\mathbb{F}_q^n / \text{Span}(\mathbf{y})| = q^{n-1}$, and thus at most dq^{n-1} zero evaluations, as each line contributes at most d roots. $\square$

The same counting pattern appears in the Ligerio argument below: choose a direction along which every parallel line contains few exceptional points, then count across the partition. There the points are received words close to a code; here they are zeros of f .

4.4 Ligero Interleaved Linear Code Testing

For a code $C \subseteq \Sigma^n$ and a vector $\mathbf{v} \in \Sigma^n$,

- denote by $d(\mathbf{v}, C)$ the minimum distance of $\mathbf{v}$ from C , namely the number of positions in which $\mathbf{v}$ differs from the closest codeword in C ,
- denote by $d(V, C)$ the minimum distance between a vector set V and a code C , that

$$d(V, C) = \min_{\mathbf{v} \in V} d(\mathbf{v}, C),$$

- denote by $\Delta(\mathbf{v}, C)$ the set of positions in which $\mathbf{v}$ differs from such a closest codeword,
- denote by $\Delta(V, C)$ the set of positions in which some $\mathbf{v} \in V$ differs from their closest codeword in C , that

$$\Delta(V, C) = \bigcup_{\mathbf{v} \in V} \Delta(\mathbf{v}, C).$$

Definition 4.4.1 (Interleaved Code). Let $L \subset \mathbb{F}^n$ be an $[n, k, d]$ linear code over $\mathbb{F}$. We let L^m denote the $[n, mk, d]$ (interleaved) code over $\mathbb{F}^m$ whose codewords are all $\mathbf{U} \in \mathbb{F}^{m \times n}$, such that each row $\mathbf{u}_i$ of $\mathbf{U}$ satisfies $\mathbf{u}_i \in L$, and let $\mathbf{u}_{*,j}$ denote the j^{th} symbol (column) of $\mathbf{U}$.

A randomized algorithm [AHIV17] to test the membership of $\mathbf{U}$ in L^m follows, where $\mathbf{U}$ is stored in an oracle $\mathcal{O}$:

- $\mathcal{V}$ samples $\mathbf{r} \xleftarrow{\mathbb{R}} \mathbb{F}^m$ and sends over $\mathbf{r}$ to $\mathcal{P}$.
- $\mathcal{P}$ sends back $\mathbf{w} \leftarrow \mathbf{r}^T \mathbf{U}$ to $\mathcal{V}$.
- $\mathcal{V}$ samples a set of positions $Q \xleftarrow{\mathbb{R}} [1, n]^\ell$, and queries $\mathcal{O}$ for symbols $\mathbf{u}_{*,i}$ for all $i \in Q$.
- $\mathcal{V}$ checks if $\mathbf{w}$ is consistent with $\{\mathbf{u}_{*,i}\}_{i \in Q}$ by $\langle \mathbf{r}, \mathbf{u}_{*,i} \rangle = w_i$.

Lemma 4.4.2. Let e be a positive integer such that $e < d/4$. Suppose $d(\mathbf{U}^*, L^m) > e$, then for a random $\mathbf{w}^*$ in $\text{RowSpan}(\mathbf{U}^*)$,

$$\Pr[d(\mathbf{w}^*, L) \leq e] \leq \frac{e+1}{|\mathbb{F}|}.$$

Proof. We introduce the following lemma, stating that if there is an element in the rowspan of $\mathbf{U}^*$ has large distance, then such element is almost everywhere in the rowspan space.

Lemma 4.4.3. Suppose $\mathbf{u}, \mathbf{v} \in \text{RowSpan}(\mathbf{U}^*)$ with $d(\mathbf{u}, L) > 2e$, then over the line $f(\alpha) = \alpha \cdot \mathbf{u} + \mathbf{v}$ where $\alpha \in \mathbb{F}$,

$$\Pr[d(\alpha \cdot \mathbf{u} + \mathbf{v}, L) \leq e] \leq |\mathbb{F}|^{-1}.$$

Proof. On the line $f(\alpha) = \alpha \cdot \mathbf{u} + \mathbf{v}$, suppose there are 2 distinct elements $f(\alpha_0)$ and $f(\alpha_1)$, such that

$$d(\alpha_0 \cdot \mathbf{u} + \mathbf{v}, L) \leq e \text{ and } d(\alpha_1 \cdot \mathbf{u} + \mathbf{v}, L) \leq e.$$

By triangle inequality,

$$d((\alpha_0 - \alpha_1) \cdot \mathbf{u}, L) \leq d(\alpha_0 \cdot \mathbf{u} + \mathbf{v}, L) + d(\alpha_1 \cdot \mathbf{u} + \mathbf{v}, L) \leq 2e,$$

this contradicts against the assumption of $d(\mathbf{u}, L) > 2e$. Hence, at most 1 element on $f(\alpha)$ has distance at most e . $\square$

Corollary 4.4.4. Suppose $\mathbf{u} \in \text{RowSpan}(\mathbf{U}^*)$ has $d(\mathbf{u}, L) > 2e$, then for a random $\mathbf{w}^*$ in $\text{RowSpan}(\mathbf{U}^*)$,

$$\Pr[d(\mathbf{w}^*, L) \leq e] \leq |\mathbb{F}|^{-1}.$$

Proof. We foliate $\text{RowSpan}(\mathbf{U}^*)$ with $|\mathbb{F}|^{m-1}$ parallel lines similar to [Mos10], where each line is defined by $\alpha \cdot \mathbf{u} + \mathbf{v}$. By lemma 4.4.3, each line has at most 1 element of distance at most e , and thus at most $|\mathbb{F}|^{m-1}$ such elements, which completes the proof. $\square$

On the other hand, if every element in $\text{RowSpan}(\mathbf{U}^*)$ has distance to L at most $2e$, fixing a differed position j , we argue that there are at most 1 element over a line, with distance to L is at most e , and the element is not differed at j .

Lemma 4.4.5. Suppose $\mathbf{u}, \mathbf{v} \in \text{RowSpan}(\mathbf{U}^*)$ with $j \in \Delta(\mathbf{u}, L)$, then

$$\Pr \left[d(\alpha \cdot \mathbf{u} + \mathbf{v}, L) \leq e \bigwedge j \notin \Delta(\alpha \cdot \mathbf{u} + \mathbf{v}, L) \right] \leq |\mathbb{F}|^{-1}.$$

Proof. If there are 2 distinct elements on $f(\alpha)$ both have the j^{th} symbol matched, and both distances to L are at most e , then by triangle inequality,

$$\begin{aligned} d(\text{Span}(f(\alpha_0), f(\alpha_1)), L) &= \min_{a,b} d(a \cdot f(\alpha_0) + b \cdot f(\alpha_1), L) \\ &\leq d(\alpha_0 \cdot \mathbf{u} + \mathbf{v}, L) + d(\alpha_1 \cdot \mathbf{u} + \mathbf{v}, L) \\ &\leq 2e < d/2, \end{aligned}$$

all elements in the $\text{Span}(f(\alpha_0), f(\alpha_1))$ are in the unique decoding radius, hence all elements have j^{th} symbol matched. Yet this contradicts $\mathbf{u} \in \text{Span}(f(\alpha_0), f(\alpha_1))$, that $j \in \Delta(\mathbf{u}, L)$. Hence, there are at most 1 such element on the $f(\alpha)$. $\square$

The following corollary is immediate by the space foliation idea in [Mos10].

Corollary 4.4.6. For a random $\mathbf{w}^*$ in $\text{RowSpan}(\mathbf{U}^*)$, and $j \in \Delta(\mathbf{U}^*, L^m)$,

$$\Pr \left[d(\mathbf{w}^*, L) \leq e \bigwedge j \notin \Delta(\mathbf{w}^*, L) \right] \leq |\mathbb{F}|^{-1}.$$

Since $d(\mathbf{U}^*, L^m) > e$, we union bound over at least $e + 1$ differed positions, upper bounded by $(e + 1) \cdot |\mathbb{F}|^{-1}$. $\square$

An improved result of distance preserving of random linear combination over linear code is also in [AHIV17].

Lemma 4.4.7. Let e be a positive integer such that $e < d/3$. Suppose $d(\mathbf{U}^*, L^m) > e$, then for a random $\mathbf{w}^*$ in $\text{RowSpan}(\mathbf{U}^*)$,

$$\Pr[d(\mathbf{w}^*, L) \leq e] \leq \frac{e + 1}{|\mathbb{F}|}.$$

Proof. We start by showing there exists elements in $\text{RowSpan}(\mathbf{U}^*)$ such that the distance to L greater than e .

Lemma 4.4.8. Suppose $d(\mathbf{U}^*, L^m) > e$, and $|\mathbb{F}| > e$, then there exists $\mathbf{w}^* \in \text{RowSpan}(\mathbf{U}^*)$ such that $d(\mathbf{w}^*, L) > e$.

Proof. Suppose all $\mathbf{w}^* \in \text{RowSpan}(\mathbf{U}^*)$ have $d(\mathbf{w}^*, L) \leq e$, then let $\mathbf{u}^* \in \text{RowSpan}(\mathbf{U}^*)$ maximizes the distance to L ,

$$\max_{\mathbf{v}^* \in \text{RowSpan}(\mathbf{U}^*)} d(\mathbf{v}^*, L) = d(\mathbf{u}^*, L) \leq e.$$

There must exist another element $\mathbf{x}^*$ in the rowspan, such that

$$\Delta(\mathbf{x}^*, L) \setminus \Delta(\mathbf{u}^*, L) \neq \emptyset,$$

otherwise $d(\mathbf{u}^*, L) = d(\mathbf{U}^*, L^m) > e$, contradicting to the prior assumption of $d(\mathbf{u}^*, L) \leq e$.

We also want to know, if all $\mathbf{w}^* \in \text{RowSpan}(\mathbf{U}^*)$ are of distance at most e to L , does the nearest codewords follow the same linear combination as the $\mathbf{w}^*$'s (or we cannot analyze by the union of differed positions).

Construct a line $f(\alpha) = \alpha \cdot \mathbf{u}^* + \mathbf{x}^*$, and let $\mathbf{u}, \mathbf{x}, \mathbf{y}$ be the codewords within e distance to $\mathbf{u}^*, \mathbf{x}^*, f(\alpha)$, then

$$d(\alpha \cdot \mathbf{u}^* + \mathbf{x}^*, \alpha \cdot \mathbf{u} + \mathbf{x}) \leq d(\mathbf{u}^*, \mathbf{u}) + d(\mathbf{x}^*, \mathbf{x}) \leq 2e,$$

the first inequality holds by the triangle inequality, and the second holds by the assumption that any element in the rowspan is of distance to L at most e . Then the distance between 2 valid codewords is

$$d(\alpha \cdot \mathbf{u}^* + \mathbf{x}^*, \mathbf{y}) \leq d(\alpha \cdot \mathbf{u}^* + \mathbf{x}^*, f(\alpha)) + d(f(\alpha), \mathbf{y}) \leq 2e + e < d,$$

which leads to a contradiction, as any 2 valid codewords are at least of d distance. Hence, $\mathbf{y} = \alpha \cdot \mathbf{u} + \mathbf{x}$.

Assuming all $\mathbf{w}^* \in \text{RowSpan}(\mathbf{U}^*)$ have $d(\mathbf{w}^*, L) \leq e$, now we know the linear combination of $\mathbf{w}^*$'s is the same as the linear combination of the nearest $\mathbf{w}$ codewords. Let

$$\Delta_{\mathbf{u}^*, \mathbf{x}^*} = \Delta(\mathbf{u}^*, L) \cap \Delta(\mathbf{x}^*, L).$$

Then the probability for any position in $\Delta_{\mathbf{u}^*, \mathbf{x}^*}$ being cancelled on a random α is at most $|\Delta_{\mathbf{u}^*, \mathbf{x}^*}| \cdot |\mathbb{F}|^{-1}$ by union bound. Hence, by

$$|\Delta_{\mathbf{u}^*, \mathbf{x}^*}| < e < |\mathbb{F}|,$$

first inequality holds by $d(\mathbf{u}^*, L) < e$, then the probability that no position in $\Delta_{\mathbf{u}^*, \mathbf{x}^*}$ being cancelled is positive, which proves the existence of $\mathbf{w}^*$ that $d(\mathbf{w}^*, L) > e$. $\square$

Since there exists elements in the rowspan such that the distance to L is greater than e , in order to strengthen the result from lemma 4.4.2, it suffices to show that in an affine subspace of $\mathbb{F}^n$, either all points are of distances within e to L , or almost all are not. This reduces to showing the same in 1-dimensional space.

Lemma 4.4.9. *For any $\mathbf{u}, \mathbf{v} \in \mathbb{F}^n$ constructing a line $\alpha \cdot \mathbf{u} + \mathbf{v}$,*

- *either every points have $d(\alpha \cdot \mathbf{u} + \mathbf{v}, L) \leq e$,*
- *or at most e points have $d(\alpha \cdot \mathbf{u} + \mathbf{v}, L) \leq e$.*

Proof. We show this with 2 vectors of Hamming weight at most e , and see if any element in the 1-dimensional affine subspace can be within distance e to any valid non-zero codeword. Let $\mathbf{u}^*, \mathbf{v}^*$ be 2 vectors of Hamming weights at most e , then they define a line $f(\alpha) = \alpha \cdot \mathbf{u}^* + \mathbf{v}^*$. The lemma can be proved in 2 cases:

- $|\text{supp}(\mathbf{u}^*) \cup \text{supp}(\mathbf{v}^*)| \leq e$. Then any element on the line are within distance e to $\mathbf{0}$.
- $|\text{supp}(\mathbf{u}^*) \cup \text{supp}(\mathbf{v}^*)| > e$. The intersection $\Delta_{\mathbf{u}^*, \mathbf{v}^*}$ has at most $e - 1$ elements. By a similar argument of differed position cancelling in lemma 4.4.8, there are at most $e - 1$ points on the line with overlapped disagreeing positions cancelled, hence there are at most $e - 1$ points within distance e to $\mathbf{0}$.

Suppose there is a non-zero codeword $\mathbf{c}$ such that an $\alpha \cdot \mathbf{u}^* + \mathbf{v}^*$ is within distance e to $\mathbf{c}$, but by triangle inequality,

$$d(\mathbf{0}, \mathbf{c}) \leq d(\mathbf{0}, \alpha \cdot \mathbf{u}^* + \mathbf{v}^*) + d(\alpha \cdot \mathbf{u}^* + \mathbf{v}^*, \mathbf{c}) \leq d(\mathbf{0}, \mathbf{v}^*) + d(\mathbf{0}, \mathbf{u}^*) + d(\alpha \cdot \mathbf{u}^* + \mathbf{v}^*, \mathbf{c}) \leq 3e < d,$$

which is a contradiction. Hence, there is only $\mathbf{0}$ that is within e distance to elements on the line. $\square$

By lemma 4.4.8 and lemma 4.4.9, the proof for lemma 4.4.7 is immediate. $\square$

We conclude with the soundness analysis, when $\mathbf{U}^n$ is far from L^m .

Theorem 4.4.10. *Let $e < d/3$ be a positive integer. Suppose $d(\mathbf{U}^*, L^m) > e$, then for any malicious $\mathcal{P}$ strategy, $\mathcal{V}$ accepts with probability at most*

$$\left(1 - \frac{e+1}{n}\right)^t + \frac{e+1}{|\mathbb{F}|}.$$

Proof. Let $\mathbf{w}^{*T} = \mathbf{r}^T \mathbf{U}^*$, the probability upper bound follows:

$$\begin{aligned} \Pr[\mathcal{V} \text{ accepts}] &\leq \Pr[\mathcal{V} \text{ accepts} \mid d(\mathbf{w}^*, L) > e] + \Pr[d(\mathbf{w}^*, L) \leq e] \\ &\leq \binom{n-e-1}{t} \cdot \binom{n}{t}^{-1} + \frac{e+1}{|\mathbb{F}|} \\ &\leq \left(1 - \frac{e+1}{n}\right)^t + \frac{e+1}{|\mathbb{F}|}, \end{aligned}$$

the first inequality holds by exercise 1.4.14 trick, while the second holds by lemma 4.4.7. $\square$

4.5 Interleaved Linear Code Testing with Logarithmic Randomness

TODO: *Diamond Posen* [DP24]

4.6 Basefold Random Linear Foldable Code and the Distance Proof

Lemma 4.6.1 (Lemma 2 [CF24]). *Fix any generator matrix $\mathbf{G}_{d-1}$ where the encoding of any nonzero message $\mathbf{m} \in \mathbb{F}^{k_{d-1}}$ has fewer than t_{d-1} zeroes. For any subset $S \subseteq [1, n_{d-1}]$,*

- *if $|S| < t_{d-1}$, then $|m_d(S)| < |\mathbb{F}|^{2(t_{d-1}-|S|)}$,*
- *if $|S| \geq t_{d-1}$, then $|m_d(S)| = 1$.*

Proof. When $|S| \geq t_{d-1}$, only $\mathbf{m} = \mathbf{0}$ can make all alphabets over S being 0. Hence, $\dim \ker \left(\mathbf{G}_{d-1,S}^\top \right) = 0$, $|m_d(S)| = 1$.

Let $T \subseteq [1, n_{d-1}] \setminus S$ of size $t_{d-1} - |S|$. We first have

$$\begin{aligned} \text{rank}(\mathbf{G}_{d-1,T \cup S}^\top) &\leq \text{rank} \left(\mathbf{G}_{d-1,S}^\top \right) + \text{rank} \left(\mathbf{G}_{d-1,T}^\top \right) \\ &\leq \text{rank} \left(\mathbf{G}_{d-1,S}^\top \right) + (t_{d-1} - |S|). \end{aligned}$$

By Rank-Nullity,

$$k_{d-1} = \text{rank} \left(\mathbf{G}_{d-1,T \cup S}^\top \right) + \dim \ker \left(\mathbf{G}_{d-1,T \cup S}^\top \right),$$

since the linear code alphabets over $T \cup S$ is $\mathbf{0}$ iff $\mathbf{m} = \mathbf{0}$ ¹⁹, then

$$\dim \ker \left(\mathbf{G}_{d-1,T \cup S}^\top \right) = 0,$$

hence $\text{rank} \left(\mathbf{G}_{d-1,T \cup S}^\top \right) = k_{d-1}$, and $\text{rank} \left(\mathbf{G}_{d-1,S}^\top \right) \geq k_{d-1} - (t_{d-1} - |S|)$. Again by Rank-Nullity,

$$k_{d-1} = \text{rank} \left(\mathbf{G}_{d-1,S}^\top \right) + \dim \ker \left(\mathbf{G}_{d-1,S}^\top \right),$$

and hence,

$$\dim \ker \left(\mathbf{G}_{d-1,S}^\top \right) \leq t_{d-1} - |S|,$$

which completes the proof, that $|m_d(S)| \leq |\mathbb{F}|^{2(t_{d-1}-|S|)}$. □

TODO: basefold stuffs [CF24].

¹⁹or by $t_{d-1} \geq k_{d-1}$ as $n_{d-1} - t_{d-1} + 1$ is the distance of the $(d-1)^{\text{th}}$ level linear code, which is no more than $n_{d-1} - k_{d-1} + 1$ by MDS code.

4.7 Bulletproof

TODO: ... [BBB⁺18]

4.8 Sublinear Lattice ZKP for Arithmetic Circuits

The work [BBC⁺18] achieves sublinear communication complexity in the lattice based ZK argument system for linear relations. More specifically, they can prove knowledge of short pre-images $\{s_i\}_{i \in [1, \ell]}$ over $\mathbb{Z}_q^v$ for a linear relation

$$\mathbf{A}\mathbf{S} \equiv \mathbf{T} \pmod{q},$$

where $\mathbf{A} \in \mathbb{Z}_q^{r \times v}$ and $\mathbf{T} \in \mathbb{Z}_q^{r \times \ell}$ are public, and $\mathbf{S} = [s_1 \mid \dots \mid s_\ell]$. The trick is to use a challenge matrix $\mathbf{C} \xleftarrow{\mathcal{R}} \{0, 1\}^{\ell \times (\lambda+2)}$, such that $\mathbf{S}\mathbf{C}$ is still full of short coefficients, but the number of relations shrink from ℓ to $\lambda + 2$. Applying the same technique in [Lyu09, Lyu12], we achieve zero knowledge.

Ring notation and norms. Let $d = d(\lambda)$ be a power of two, and let $R = \mathbb{Z}[X]/(X^d + 1)$, $R_q = R/qR$, and $R_p = R/pR$. For $f \in R_q$, we always take the centered coefficient representative before measuring its size. If this representative is $f = \sum_{i=0}^{d-1} f_i X^i$, then f_i denotes the coefficient of X^i . Define

$$\|f\|_\infty = \max_{i \in [0, d-1]} |f_i|, \quad \|f\|_2 = \left(\sum_{i=0}^{d-1} f_i^2 \right)^{1/2}.$$

For $\mathbf{z} = (z_1, \dots, z_m) \in R^m$, define

$$\|\mathbf{z}\|_\infty = \max_{j \in [m]} \|z_j\|_\infty, \quad \|\mathbf{z}\|_2 = \left(\sum_{j=1}^m \|z_j\|_2^2 \right)^{1/2}.$$

Commitment Scheme. For a commitment scheme $\Pi_{\text{Com}} = (\text{Setup}, \text{Com})$, let $\mathcal{M}_{\text{Com}}$ denote the message space, and $\mathcal{R}_{\text{Com}}$ denote the randomness space. The setup algorithm outputs a commitment key $\text{ck} \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, and the commitment algorithm takes $\mathbf{m} \in \mathcal{M}_{\text{Com}}$ and $\mathbf{r} \in \mathcal{R}_{\text{Com}}$ as inputs, and outputs $\mathbf{c} \leftarrow \text{Com}_{\text{ck}}(\mathbf{m}; \mathbf{r})$. An opening of $\mathbf{c}$ is a pair $(\mathbf{m}, \mathbf{r})$, and verification recomputes $\text{Com}_{\text{ck}}(\mathbf{m}; \mathbf{r})$.

Definition 4.8.1 (Binding). $\Pi_{\text{Com}} = (\text{Setup}, \text{Com})$ is computationally binding if for every efficient adversary $\mathcal{A}$,

$$\Pr \left[\text{Com}_{\text{ck}}(\mathbf{m}_0; \mathbf{r}_0) = \text{Com}_{\text{ck}}(\mathbf{m}_1; \mathbf{r}_1) \wedge \mathbf{m}_0 \neq \mathbf{m}_1 \mid \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda, 1^\ell), \\ (\mathbf{m}_0, \mathbf{r}_0, \mathbf{m}_1, \mathbf{r}_1) \leftarrow \mathcal{A}(1^\lambda, \text{ck}) \end{array} \right] \leq \text{negl}(\lambda),$$

If the same condition holds for unbounded adversaries, then $\Pi_{\text{Com}} = (\text{Setup}, \text{Com})$ is statistically binding.

Definition 4.8.2 (Hiding). $\Pi_{\text{Com}} = (\text{Setup}, \text{Com})$ is computationally hiding if for every efficient adversary $(\mathcal{A}_0, \mathcal{A}_1)$,

$$\Pr \left[b = b' \mid \begin{array}{l} \text{ck} \leftarrow \text{Setup}(1^\lambda, 1^\ell), \\ (\mathbf{m}_0, \mathbf{m}_1, \text{st}) \leftarrow \mathcal{A}_0(1^\lambda, \text{ck}), \\ b \xleftarrow{\mathcal{R}} \{0, 1\}, \\ \mathbf{r} \leftarrow \mathcal{R}_{\text{Com}}, \\ \mathbf{c}_b \leftarrow \text{Com}_{\text{ck}}(\mathbf{m}_b; \mathbf{r}), \\ b' \leftarrow \mathcal{A}_1(\text{st}, \mathbf{c}_b) \end{array} \right] - \frac{1}{2} \leq \text{negl}(\lambda).$$

If the same condition holds for unbounded adversaries, then $\Pi_{\text{Com}} = (\text{Setup}, \text{Com})$ is statistically hiding.

Module-SIS. The binding property of Ajtai commitments is from the hardness of module-SIS assumption.

Definition 4.8.3 (Module-SIS [LS15]). Let $q = q(\lambda)$, $m = m(\lambda)$, $n = n(\lambda)$, $\beta = \beta(\lambda)$, and $d = d(\lambda)$. We say $\text{MSIS}_{n, m, d, q, \beta}$ holds if for any efficient adversary $\mathcal{A}$ outputting $\mathbf{z} \in R^m$, the following holds

$$\Pr \left[\mathbf{A}\mathbf{z} \equiv \mathbf{0} \pmod{q} \wedge 0 < \|\mathbf{z}\|_\infty \leq \beta \mid \mathbf{A} \xleftarrow{\mathcal{R}} R_q^{n \times m}, \mathbf{z} \leftarrow \mathcal{A}(\mathbf{A}) \right] \leq \text{negl}(\lambda).$$

When $n = 1$, the assumption specializes to the Ring-SIS assumption [PR06, LM06].

Ajtai Commitment. We now recall the Ajtai commitment instantiated from module-SIS assumption.

Definition 4.8.4 (Ajtai Commitment [Ajt96]). Over the above ring, we define the following parameters:

- ℓ : the message dimension,
- k : the number of output ring elements.
- $\mathcal{M}_{\text{Ajtai}}$: the message space R_p^ℓ , where elements of R_p are lifted to R_q by their centered coefficient representatives.
- $\mathcal{R}_{\text{Ajtai}}$: the randomness space $R_q^{2k \log_p q}$.
- $\mathcal{D}_{\text{Ajtai}}$: the subgaussian distribution $\mathcal{D}_\sigma^{2k \log_p q}$ over the randomness space.
- B_{Ajtai} : the binding space defined as follows,

$$B_{\text{Ajtai}} = \left\{ \mathbf{s} = [\mathbf{r}^\top \mid \mathbf{m}^\top]^\top \in R^{\ell+2k \log_p q} : \|\mathbf{s}\|_\infty \leq B \right\}.$$

We construct $\Pi_{\text{Ajtai}} = (\text{Setup}, \text{Com})$ over R_p as follows:

- $\text{Setup}(1^\lambda, 1^\ell)$: Sample $\mathbf{A}_1 \xleftarrow{R} R_q^{k \times 2k \log_p q}$ and $\mathbf{A}_2 \xleftarrow{R} R_q^{k \times \ell}$, and output $\text{ck} = \mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
- $\text{Com}_{\text{ck}}(\mathbf{m}; \mathbf{r})$: On input the message $\mathbf{m} \in R_p^\ell$ and $\mathbf{r} \in R_q^{2k \log_p q}$, output $\mathbf{c} = \mathbf{A}_1 \mathbf{r} + \mathbf{A}_2 \mathbf{m} \in R_q^k$.

An opening is $(\mathbf{m}, \mathbf{r})$, and verification recomputes the commitment.

Lemma 4.8.5. *Ajtai commitment is computationally binding and statistically hiding.*

Proof. Hiding follows from the Leftover Hash Lemma: $(\mathbf{A}_1, \mathbf{A}_1 \mathbf{r})$ is statistically close to uniformly random, and hence the commitment scheme is statistically hiding. Let $n' = k, m' = \ell + 2k \log_p q$, and $2B \leq \beta$, so long as $\text{MSIS}_{n', m', d, q, \beta}$ is hard, the commitment scheme is computationally binding. $\square$

Amortized Linear Relation Proof. We now describe the amortized proof for many linear relations over R_q . The public statement consists of matrices $\mathbf{A} \in R_q^{r \times v}$ and $\mathbf{T} \in R_q^{r \times \ell}$, and the witness is $\mathbf{S} = [\mathbf{s}_1 \mid \dots \mid \mathbf{s}_\ell] \in R^{v \times \ell}$, where each $\mathbf{s}_i$ has ℓ_∞ -norm upper bounded by a bound B_{lin} , satisfying $\mathbf{AS} \equiv \mathbf{T} \pmod{q}$.

By linearity of the Ajtai commitment definition 4.8.4, the commitments to the compressed columns $\mathbf{SC}$ can be obtained by applying the same linear combinations to the commitments of the columns of $\mathbf{S}$.

We begin by describing the “compression lemma” that help establish amortized proof.

Lemma 4.8.6. *Let $\mathbb{G}$ be an abelian group, and $\mathbf{x} = [x_1, \dots, x_\ell]^\top \in \mathbb{G}^\ell$ be nonzero. For $\mathbf{c} \xleftarrow{R} \{0, 1\}^\ell$, $\Pr[\langle \mathbf{c}, \mathbf{x} \rangle = 0] \leq 1/2$. Consequently, for $\mathbf{C} \xleftarrow{R} \{0, 1\}^{\ell \times t}$, $\Pr[\mathbf{x}^\top \mathbf{C} = \mathbf{0}^\top] \leq 2^{-t}$.*

Proof. Pick i with $x_i \neq 0$, and fix all c_j for $j \neq i$. The two possible sums differ by x_i , so at most one of them is zero. The conditional probability is at most $1/2$, and averaging over the fixed bits gives the claim. $\square$

Immediate by lemma 4.8.6, if $\mathbf{AS}' \not\equiv \mathbf{T} \pmod{q}$, for $\mathbf{C} \xleftarrow{R} \{0, 1\}^{\ell \times t}$, $\Pr[\mathbf{AS}'\mathbf{C} \equiv \mathbf{TC} \pmod{q}] \leq 2^{-t}$.

Protocol Construction. The prover commits to each $\mathbf{s}_i$ using the Ajtai commitment definition 4.8.4.

Let C be the challenge set, that is $\{0, 1\}$ when $R = \mathbb{Z}$, otherwise $\{0\} \cup \{\pm X^j\}_{j < d}$ when $R = \mathbb{Z}[X]/(X^d + 1)$.

Commit. For $i \in [1, \ell]$, the prover samples $\mathbf{r}_i \leftarrow \mathcal{D}_{\text{Ajtai}}$ and sends $\mathbf{c}_i = \text{Com}_{\text{ck}}(\mathbf{s}_i; \mathbf{r}_i)$. Let $\mathbf{C}_{\text{com}} = [\mathbf{c}_1 \mid \dots \mid \mathbf{c}_\ell]$.

Challenge. The verifier samples $\mathbf{C} \xleftarrow{R} C^{\ell \times (\lambda+2)}$.

Response. The prover computes $\mathbf{Y} \leftarrow \mathbf{SC}$ and the compressed randomness $\mathbf{R} = [\mathbf{r}_1 \mid \dots \mid \mathbf{r}_\ell] \mathbf{C}$, and sends over $(\mathbf{Y}, \mathbf{R})$.

Verification. The verifier checks $(\mathbf{Y}, \mathbf{R})$ are consistent with the compressed commitments by $\mathbf{C}_{\text{com}} \mathbf{C} = \text{Com}_{\text{ck}}(\mathbf{Y}; \mathbf{R})$.

Completeness is immediate.

Special Knowledge Soundness. We now show the (special) knowledge soundness of an accepting prover.

Here, v is the witness dimension for each relation, $t = \lambda + 2$, and ℓ is the number of relations being batched. Fix a statement $(\mathbf{A}, \mathbf{T})$ and a first message $\mathbf{C}_{\text{Com}}$.

Theorem 4.8.7. *If $\mathcal{P}^*$ responds to a random challenge $\mathbf{C} \xleftarrow{R} \mathcal{C}^{\ell \times t}$, and convinces the verifier with probability $\varepsilon \geq 2^{-\lambda}$, then the extractor should recover a matrix $\mathbf{S} \in R^{v \times \ell}$ such that $\mathbf{AS} \equiv c\mathbf{T} \pmod{q}$ ²⁰, together with extracted randomness, open the commitments in $\mathbf{C}_{\text{Com}}$, and the extracted columns are short. The extractor runs in expected time $\text{poly}(\lambda, 1/\varepsilon)$.*

We start by describing how to extract a column of the witness $(\mathbf{s}_i, \mathbf{r}_i)$, which can be generalized to all columns. For an extractor $\mathcal{E}^{\mathcal{P}^*}$, it rewinds the prover to its first message $\mathbf{C}_{\text{Com}}$ (so as to pin down the prover randomness), together with all rows of the challenge $\mathbf{C} \in \mathcal{C}^{\ell \times t}$. It then resamples only the i^{th} row of $\mathbf{C}$, yielding $\mathbf{C}'$.

We now introduce a helpful lemma via probabilistic method related to the extraction process.

Lemma 4.8.8 (Heavy Row Lemma [Dam10]). *Let $\mathbf{H} \in \{0, 1\}^{n \times m}$ has density at least ε , namely at least εnm entries are 1. We say that a row $a \in [1, n]$ is heavy, if there are more than $\varepsilon m/2$ entries are 1, namely the density in this row is at least $\varepsilon/2$. Then a uniformly random 1-entry of $\mathbf{H}$ lies in a heavy row with probability at least $1/2$.*

Proof. The total number of 1-entries in $\mathbf{H}$ is at least εmn , while the number of 1-entries in non-heavy rows is at most $\varepsilon mn/2$. Hence, the probability of an 1-entry lying in a light row is at most $1/2$, and an 1-entry lying in a heavy row is at least $1/2$. $\square$

Therefore, consider building up a matrix $\mathbf{H}$ over $\{0, 1\}$ as follows. The rows are indexed by the prover randomness and by all challenge rows except the i^{th} row. The columns are indexed by the possible i^{th} row challenges $\alpha_i \in \mathcal{C}^t$. The entry is 1 if $\mathcal{P}^*$ gives an accepting response on the corresponding full challenge matrix, and 0 otherwise.

The density of $\mathbf{H}$ is at least ε as the prover convinces the verifier with probability at least ε . Hence, by lemma 4.8.8, conditioned on a random challenge is accepting, with probability at least $1/2$ the corresponding row of $\mathbf{H}$ is heavy. Thus, after the extractor finds one accepting transcript in a heavy row, resampling only the i^{th} row in $\mathbf{C}$ gives another accepting transcript with probability at least $\varepsilon/2$.

Immediately, this gives a randomized algorithm for the extractor to extract a column of witness $(\mathbf{s}_i, \mathbf{r}_i)$ in expected polynomial time. First, the extractor samples accepting $(\mathbf{C}, \mathbf{Y}, \mathbf{R})$ on the first $\mathbf{C}_{\text{Com}}$. Then, fix the prover randomness and $\mathbf{C}$ except the i^{th} row. Resample the i^{th} row independently up to $N = O(\lambda/\varepsilon)$ times. If one obtains an accepting $(\mathbf{C}', \mathbf{Y}', \mathbf{R}')$ whose i^{th} row differs from that of $\mathbf{C}$, continue with the extraction below. Otherwise, abort and restart.

The runtime analysis is similar to the expectation of a geometrically distributed random variable. Similar to the relaxation in lemma 1.5.36, we have

$$\begin{aligned} \Pr[\text{abort}] &= \Pr[\text{abort} \mid i^{\text{th}} \text{ row of } \mathbf{C} \text{ is heavy for } \mathbf{H}] \cdot \Pr[i^{\text{th}} \text{ row of } \mathbf{C} \text{ is heavy for } \mathbf{H}] \\ &\quad + \Pr[\text{abort} \mid i^{\text{th}} \text{ row of } \mathbf{C} \text{ is light for } \mathbf{H}] \cdot \Pr[i^{\text{th}} \text{ row of } \mathbf{C} \text{ is light for } \mathbf{H}] \\ &\leq \Pr[\text{abort} \mid i^{\text{th}} \text{ row of } \mathbf{C} \text{ is heavy for } \mathbf{H}] + \Pr[i^{\text{th}} \text{ row of } \mathbf{C} \text{ is light for } \mathbf{H}]. \end{aligned}$$

The second term is at most $1/2$ by the lemma 4.8.8. The first term is upper bounded as follows. Since the i^{th} row of $\mathbf{C}$ is heavy on $\mathbf{H}$, the probability of choosing a rejecting challenge is at most $1 - \varepsilon/2$, also excluding the first accepting $\mathbf{C}$, then the rejecting probability is at most $1 - \varepsilon/2 + |\mathcal{C}|^{-t}$. Repeating N times, the first term is at most $(1 - \varepsilon/2 + |\mathcal{C}|^{-t})^N$. Thus, the aborting probability is at most

$$\frac{1}{2} + \left(1 - \frac{\varepsilon}{2} + \frac{1}{|\mathcal{C}|^t}\right)^N.$$

If $\varepsilon > 2^{-\lambda}$, then $1 - \varepsilon/2 + |\mathcal{C}|^{-t} > 1 - \varepsilon/4 > \exp(-\varepsilon/4)$, which leads to the final upper bound being $1/2 + 2^{-\lambda}$. Since each run takes $O(N)$ time, then extractor runs in expected $\text{poly}(\lambda, 1/\varepsilon)$.

²⁰When $R = \mathbb{Z}$, $c = 1$; otherwise when R is a polynomial ring, $c = 2$.

Suppose the extractor obtains two accepting transcripts $(\mathbf{C}, \mathbf{Y}, \mathbf{R})$ and $(\mathbf{C}', \mathbf{Y}', \mathbf{R}')$ whose challenges differ only on the i^{th} row, and differ nontrivially in that row. Write $\alpha_i, \alpha'_i \in C^t$ for the two i^{th} rows and let

$$\Delta \mathbf{C} = \mathbf{C} - \mathbf{C}', \quad \Delta \mathbf{Y} = \mathbf{Y} - \mathbf{Y}', \quad \Delta \mathbf{R} = \mathbf{R} - \mathbf{R}',$$

then the commitment openings satisfy

$$\mathbf{C}_{\text{Com}} \cdot \Delta \mathbf{C} = \text{Com}_{\text{ck}}(\Delta \mathbf{Y}; \Delta \mathbf{R}).$$

Choose a coordinate j such that $\alpha = (\alpha_i - \alpha'_i)_j \in C - C$, then the j^{th} column of the first equation gives

$$\text{Com}_{\text{ck}}(\Delta \mathbf{y}_j; \Delta \mathbf{r}_j) \equiv \mathbf{A}_1 \cdot \Delta \mathbf{r}_j + \mathbf{A}_2 \cdot \Delta \mathbf{y}_j \equiv \alpha \mathbf{c}_i \pmod{q},$$

so $(\Delta \mathbf{y}_j, \Delta \mathbf{r}_j)$ is an opening of the scaled commitment $\alpha \mathbf{c}_i$.

By the challenge-set property, there exists a short $g \in R$ such that $g\alpha = c$. Hence

$$\text{Com}_{\text{ck}}(g \cdot \Delta \mathbf{y}_j; g \cdot \Delta \mathbf{r}_j) \equiv c \mathbf{c}_i \pmod{q}.$$

The extractor therefore obtains the i^{th} scaled opening, and repeating the same row-local procedure for every $i \in [1, \ell]$ yields openings of all columns of $c \mathbf{C}_{\text{Com}}$. Shortness follows from the response norm checks in accepting transcripts; the exact checks and constants are accounted for in the norm-growth analysis.

4.9 A Non-PCP Approach to Succinct Quantum-Safe ZK

The high-level idea of [BLNS20] is twofold. On one hand, it upgrades the amortized linear-relation proof for Ajtai commitments from [BBC⁺18] into a *levelled* SIS commitment, making the commitment object recursively compressing. On the other hand, it gives a lattice analogue of the Bulletproofs folding argument [BBB⁺18]. Concretely, we have:

- levelled SIS commitments, with communication

$$O\left(N^{1/(d+1)} \cdot \left(d^3 \lambda \log^2 N + d\lambda^2\right)\right)$$

for d levels, and hence $\text{polylog}(N)$ communication on $d + 1 = \log N$;

- a lattice analogue of the Bulletproofs folding argument [BBB⁺18], with smaller polylogarithmic proof size but much larger extraction slack in knowledge soundness.

N is the number of committed secret values. Recall the note section 4.8 on [BBC⁺18] that the one-level construction has $\tilde{O}(\sqrt{N})$ communication after balancing the number of commitments and the opening proof size after amortization. We begin by showing how [BLNS20] reduces the communication via a two-level commitment instantiation. As in the previous lattice argument note, we suppress the zero-knowledge layer: the commitment is hiding, and the proof is made zero-knowledge by Gaussian masking with rejection sampling in the Fiat-Shamir-with-aborts style.

Two-Level Commitment Map. The levelling idea starts by committing to commitments. On message $\mathbf{S} \in \mathbb{Z}^{m_1 m_2 \times m_3}$, commitment keys $\mathbf{A}_1 \in \mathbb{Z}_{q_1}^{n \times n m_1}$, $\mathbf{A}_2 \in \mathbb{Z}_{q_2}^{n \times m_2}$, and $q_1 > q_2$. The top commitment is

$$\mathbf{T} = \mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \cdot \mathbf{S} \bmod q_2) \bmod q_1.$$

This is SIS-commitment-like at each level: the lower level compresses blocks of $\mathbf{S}$, and the upper level compresses the lower-level commitments. At this point this is just the commitment map: an honest opening is a matrix $\mathbf{S}$ whose image under the above map is $\mathbf{T}$, and $\mathbf{S}$ is short.

General Levelled Commitments. For d levels, let $N = \prod_{i \in [1, d+1]} m_i$, $m_0 = 1$, and $q_1 > \dots > q_d$. The bottom matrix is $\mathbf{A}_d \in \mathbb{Z}_{q_d}^{n \times m_d}$, and for $i < d$, $\mathbf{A}_i \in \mathbb{Z}_{q_i}^{n \times n m_i}$. Define $M_{i,j} = m_i \cdots m_j$ and $M_i = M_{1,i}$, and define recursively

$$F_{i,j}(\mathbf{S}) = \begin{cases} \mathbf{A}_i \mathbf{S} \bmod q_i, & i = j, \\ F_{i,j-1} \left((\mathbf{I}_{M_{i,j-1}} \otimes \mathbf{A}_j) \cdot \mathbf{S} \bmod q_j \right), & i < j. \end{cases}$$

The statement is $\mathbf{T} \leftarrow F_{1,d}(\mathbf{S})$, where d is the number of levels.

Folding Operations. Similar to the challenges in section 4.8, the challenges are matrices over the challenge set $\mathcal{C}$ the same as the one in section 4.8, with shape $\mathbf{C}_1 \xleftarrow{\mathcal{R}} \mathcal{C}^{m_{d+1} \times \lambda}$, and $\mathbf{C}_i \xleftarrow{\mathcal{R}} \mathcal{C}^{m_{i-1} \lambda \times \lambda}$ for $i \in [2, d]$.

For a matrix $\mathbf{Z}$ whose rows are split into m_i consecutive blocks $\mathbf{Z}_1, \dots, \mathbf{Z}_{m_i}$ of equal size, define the block transpose

$$\text{BT}_i(\mathbf{Z}) = [\mathbf{Z}_1 \mid \dots \mid \mathbf{Z}_{m_i}].$$

We also use the recursive folding map Fold_i from [BLNS20]. The base case is

$$\text{Fold}_0(\mathbf{U}_1; \mathbf{C}_1) = \mathbf{U}_1 \mathbf{C}_1.$$

For $i > 0$, split the input tuple $(\mathbf{U}_1, \dots, \mathbf{U}_{M_i})$ into m_i consecutive sub-tuples, such that for $j \in [1, m_i]$,

$$\tilde{\mathbf{U}}_j = (\mathbf{U}_{(j-1)M_{i-1}+1}, \dots, \mathbf{U}_{jM_{i-1}}).$$

Then

$$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \text{BT}_i \begin{pmatrix} \text{Fold}_{i-1}(\bar{\mathbf{U}}_1; \mathbf{C}_1, \dots, \mathbf{C}_i) \\ \vdots \\ \text{Fold}_{i-1}(\bar{\mathbf{U}}_{m_i}; \mathbf{C}_1, \dots, \mathbf{C}_i) \end{pmatrix} \mathbf{C}_{i+1}.$$

Lemma 4.9.1 (Fold properties [BLNS20]). *Fix i and $\{\mathbf{U}_j\}_{j \in [1, M_i]}$ with compatible dimensions. For challenges $\{\mathbf{C}_j\}_{j \in [1, i+1]}$ with the shapes above:*

- *there exist matrices $\{\mathbf{D}_i\}_{i \in [1, M_i]}$ with $\|\mathbf{D}_i\|_\infty \leq \lambda^i$ such that*

$$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \sum_{t \in [1, M_i]} \mathbf{U}_t \mathbf{D}_t;$$

- *for every compatible matrix $\mathbf{B}$,*

$$\mathbf{B} \cdot \text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \text{Fold}_i(\mathbf{B}\mathbf{U}_1, \dots, \mathbf{B}\mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1});$$

- *if each $\mathbf{U}_t$ is vertically split as $\mathbf{U}_t = [\mathbf{U}_{t,1}^\top \mid \dots \mid \mathbf{U}_{t,r}^\top]^\top$, then*

$$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \begin{bmatrix} \text{Fold}_i(\mathbf{U}_{1,1}, \dots, \mathbf{U}_{M_i,1}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) \\ \vdots \\ \text{Fold}_i(\mathbf{U}_{1,r}, \dots, \mathbf{U}_{M_i,r}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) \end{bmatrix}.$$

Proof. We prove all three parts by induction on i . For $i = 0$, the definition gives $\text{Fold}_0(\mathbf{U}_1; \mathbf{C}_1) = \mathbf{U}_1 \mathbf{C}_1$. Thus the first part holds with $\mathbf{D}_1 = \mathbf{C}_1$, and the second and third parts are immediate.

Assume the claims hold for $i - 1$. Split $(\mathbf{U}_1, \dots, \mathbf{U}_{M_i})$ into subtuples $\bar{\mathbf{U}}_1, \dots, \bar{\mathbf{U}}_{m_i}$ as was defined Fold_i , and write

$$\mathbf{W}_j = \text{Fold}_{i-1}(\bar{\mathbf{U}}_j; \mathbf{C}_1, \dots, \mathbf{C}_i).$$

Also split the rows of $\mathbf{C}_{i+1}$ into m_i blocks $\mathbf{C}_{i+1}^{(1)}, \dots, \mathbf{C}_{i+1}^{(m_i)}$, each of shape $\lambda \times \lambda$. Then

$$\text{Fold}_i(\mathbf{U}_1, \dots, \mathbf{U}_{M_i}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \sum_{j \in [1, m_i]} \mathbf{W}_j \mathbf{C}_{i+1}^{(j)}.$$

By induction, each $\mathbf{W}_j$ is a bounded linear combination of the matrices in $\bar{\mathbf{U}}_j$. Multiplying by $\mathbf{C}_{i+1}^{(j)}$ gives coefficient matrices whose ℓ_∞ norm grows by at most a factor λ , since challenge entries have coefficient norm at most 1. Hence the new coefficients $\mathbf{D}_t$ satisfy $\|\mathbf{D}_t\|_\infty \leq \lambda^i$, proving the first part.

For the second part, using the induction on each subtuple,

$$\mathbf{B}\mathbf{W}_j = \text{Fold}_{i-1}(\mathbf{B}\bar{\mathbf{U}}_j; \mathbf{C}_1, \dots, \mathbf{C}_i).$$

Therefore multiplying Fold_i by $\mathbf{B}$ on the left is the same as replacing every input $\mathbf{U}_t$ by $\mathbf{B}\mathbf{U}_t$ before the recursive fold.

For the third part, since each $\mathbf{U}_t$ splits into r vertical blocks, then by induction, each $\mathbf{W}_j$ splits into the corresponding folded block. Horizontal concatenation followed by multiplication by $\mathbf{C}_{i+1}$ preserves this vertical block decomposition, so the a^{th} vertical output block is

$$\text{Fold}_i(\mathbf{U}_{1,a}, \dots, \mathbf{U}_{M_i,a}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}).$$

This proves the stacking identity and completes the induction. $\square$

To describe the prover's messages that are folded, we write $\mathbf{S}$ as $\mathbf{S}_{i_1, \dots, i_{d-1}}$ for $(i_1, \dots, i_{d-1}) \in [1, m_1] \times \dots \times [1, m_{d-1}]$, as a leaf of $m_d \times m_{d+1}$ elements. In the levelled commitment perspective, the intermediate commitments are viewed as internal nodes of the tree, and for $k \in [1, d-2]$, define intermediate vertices recursively by

$$\mathbf{S}_{i_1, \dots, i_k} = (\mathbf{I}_{m_{k+1}} \otimes \mathbf{A}_{k+2}) \begin{bmatrix} \mathbf{S}_{i_1, \dots, i_k, 1} \\ \vdots \\ \mathbf{S}_{i_1, \dots, i_k, m_{k+1}} \end{bmatrix} \mod q_{k+2},$$

of shape $nm_{k+1} \times m_{d+1}$ elements. Finally,

$$\mathbf{S}_\emptyset = (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \begin{bmatrix} \mathbf{S}_1 \\ \vdots \\ \mathbf{S}_{m_1} \end{bmatrix} \mod q_2,$$

with shape $nm_1 \times m_{d+1}$ elements, so $\mathbf{A}_1 \mathbf{S}_\emptyset \equiv \mathbf{T} \mod q_1$. For the i^{th} response, the prover uses the tuple

$$\mathbf{V}_i = \begin{cases} \mathbf{S}_\emptyset, & i = 1, \\ (\mathbf{S}_\alpha)_{\alpha \in [1, m_1] \times \dots \times [1, m_{i-1}]}, & i > 1, \end{cases}$$

where $\mathbf{V}_i$ is ordered with the first coordinate varying fastest. This is the protocol order used by Fold_{i-1} ; the recursive definition of each intermediate vertex still stacks its children under a fixed prefix.

Amortized Levelled Proof. We now show the 0-1 compression invariant $\mathbf{ASC} \equiv \mathbf{TC} \mod q$ from [BBC⁺18] in the levelled setting of [BLNS20]. The point of lemma 4.9.1 is twofold here: the vertical decomposition lets us treat a folded intermediate commitment as the stack of folded child commitments, while the left multiplication property lets us pull out the next SIS matrix $\mathbf{A}_{i+1}$. Then BT_i only rearranges this vertical stack into the horizontal sibling blocks expected by the next challenge. With this abstraction, we show by induction over the levels of the commitment tree.

- At the root, $\mathbf{A}_1 \mathbf{S}_\emptyset \equiv \mathbf{T} \mod q_1$, and therefore we have $\mathbf{A}_1 \mathbf{Z}_1 \equiv \mathbf{TC}_1 \mod q_1$.
- For the inductive step, write $\mathbf{V}_{i+1} = (\mathbf{V}_{i,1}, \dots, \mathbf{V}_{i,m_i})$, where $\mathbf{V}_{i,j}$ is the tuple of the j^{th} children of the vertices in $\mathbf{V}_i$. By the definition of intermediate commitments, each vertex in $\mathbf{V}_i$ is the vertical stack of the $\mathbf{A}_{i+1}$ -commitments to its children. Using the vertical decomposition and left multiplication properties of lemma 4.9.1, we have

$$\mathbf{Z}_i \equiv \begin{bmatrix} \mathbf{A}_{i+1} \text{Fold}_{i-1}(\mathbf{V}_{i,1}; \mathbf{C}_1, \dots, \mathbf{C}_i) \\ \vdots \\ \mathbf{A}_{i+1} \text{Fold}_{i-1}(\mathbf{V}_{i,m_i}; \mathbf{C}_1, \dots, \mathbf{C}_i) \end{bmatrix} \mod q_{i+1}.$$

Applying BT_i and the next challenge gives

$$\text{BT}_i(\mathbf{Z}_i) \mathbf{C}_{i+1} \equiv \mathbf{A}_{i+1} \text{Fold}_i(\mathbf{V}_{i+1}; \mathbf{C}_1, \dots, \mathbf{C}_{i+1}) = \mathbf{A}_{i+1} \mathbf{Z}_{i+1} \mod q_{i+1}.$$

The block transpose only exposes the m_i sibling blocks of $\mathbf{Z}_i$ in the order needed for the next challenge matrix.

We now show an amortized opening protocol for the levelled commitment, which is public-coin, so all consistency and norm checks can be deferred to the end of the verification.

Commit. The prover computes the levelled commitment $\mathbf{T} \leftarrow F_{1,d}(\mathbf{S})$ and sends $\mathbf{T}$.

i^{th} Challenge. The verifier samples the random challenge matrix $\mathbf{C}_i$ with the shape specified above.

i^{th} Response. After receiving $\mathbf{C}_i$, the prover sends $\mathbf{Z}_i \leftarrow \text{Fold}_{i-1}(\mathbf{V}_i; \mathbf{C}_1, \dots, \mathbf{C}_i)$, of shape $nm_i \times \lambda$ elements.

Verify. After receiving all responses $\mathbf{Z}_1, \dots, \mathbf{Z}_d$, the verifier checks the response are short, and the consistency by

$$\begin{aligned} \mathbf{A}_1 \mathbf{Z}_1 &\equiv \mathbf{TC}_1 \mod q_1, \\ \mathbf{A}_{i+1} \mathbf{Z}_{i+1} &\equiv \text{BT}_i(\mathbf{Z}_i) \mathbf{C}_{i+1} \mod q_{i+1}, \end{aligned}$$

where $i \in [1, d-1]$.

Knowledge soundness and relaxed extraction. The opening argument does not extract an exact opening to the original commitment map, and the source of this relaxation is already illustrated in the one-level vanilla construction of [BBC⁺18]: special soundness subtracts two accepting responses, so the target is not scaled in the plain SIS setting, but the extracted preimage has a larger norm. If each accepted response has norm bounded by B , then the extracted difference is naturally bounded by $2B$.

In the extractor of [BBC⁺18], the extracted opening is represented by bounded integers. However, in a levelled commitment, the extracted preimage is not just an opening witness; it is also meant to be an internal value produced by a lower-level commitment. The same internal value is seen in two ways: the upper-layer consumes an integer representative, as SIS is strong when the integer bound is still small, while the lower-layer only sees the corresponding residue class modulo the lower modulus. The gap between these two views is a multiple of the lower modulus, and these multiples are exactly the carry matrices in F .

For two levels, this definition specializes to a relaxed opening $(\bar{\mathbf{S}}, \mathbf{R})$ satisfying

$$\mathbf{A}_1 \cdot ((\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \bar{\mathbf{S}} \bmod q_2 + q_2 \mathbf{R}) \equiv \mathbf{T} \bmod q_1.$$

The extra $q_2 \mathbf{R}$ term records the missing multiples of q_2 when the lower-level computation is viewed as an integer input to the top commitment. This relaxation is still meaningful when parameters force the extracted $(\bar{\mathbf{S}}, \mathbf{R})$ to stay within SIS-binding bounds.

Although the commitment is tree shaped, extraction is better viewed depth-wise. All vertices at a same level are folded into one response, and the relaxed opening records the consistency between neighboring layers.

We now define the relaxed version of the levelled commitment map, in parallel with the exact map $F_{i,j}$:

$$\tilde{F}_{i,j}(\mathbf{S}; \mathbf{R}_i, \dots, \mathbf{R}_{j-1}) = \begin{cases} \mathbf{A}_i \mathbf{S} \bmod q_i, & i = j, \\ \tilde{F}_{i,j-1} \left((\mathbf{I}_{M_{i,j-1}} \otimes \mathbf{A}_j) \cdot \mathbf{S} \bmod q_j + q_j \mathbf{R}_{j-1}; \mathbf{R}_i, \dots, \mathbf{R}_{j-2} \right), & i < j. \end{cases}$$

A bounded relaxed opening to $\mathbf{T}$ is a tuple $(\bar{\mathbf{S}}, \mathbf{R}_1, \dots, \mathbf{R}_{d-1})$ such that every column of $\bar{\mathbf{S}}$ has norm at most B_S , every column of $\mathbf{R}_i$ has ℓ_∞ norm at most B_R , and

$$\tilde{F}_{1,d}(\bar{\mathbf{S}}; \mathbf{R}_1, \dots, \mathbf{R}_{d-1}) \equiv \mathbf{T} \bmod q_1.$$

An exact opening to $\mathbf{T}$ is the case where all carry matrices $\mathbf{R}$ are zeroed out.

Similar to the theorem 4.8.7 in [BBC⁺18], we define the knowledge soundness for [BLNS20].

Theorem 4.9.2. *If a prover convinces a verifier with probability $\varepsilon > 2^{-\lambda+1} \cdot (4dN)^{2d}$, there is an extractor outputs a bounded relaxed opening, where $B_S = 2^d B$ and $B_{R_k} = 2^k (M_{k-1} m_{d+1} \lambda^{k-1} + 2)$ for $k \in [1, d-1]$. It runs in expected $\text{poly}(\lambda, 1/\varepsilon)$ time.*

Noticing

$$\mathbf{A}_{i+1} \mathbf{Z}_{i+1} \equiv \mathbf{B} \mathbf{T}_i(\mathbf{Z}_i) \mathbf{C}_{i+1} \bmod q_{i+1}$$

is exactly compressing $m_i \times \lambda$ commitments to λ ones as in [BBC⁺18], we can use the same row-local rewinding idea as theorem 4.8.7.

Fix the α^{th} column of $\mathbf{T}$. The desired column-wise relaxed opening has the form

$$\tilde{F}_{1,d}(\bar{\mathbf{s}}_\alpha; \mathbf{r}_1, \dots, \mathbf{r}_{d-1}) \equiv \mathbf{t}_\alpha \bmod q_1,$$

where $\mathbf{r}_i$ is the α^{th} column of $\mathbf{R}_i$ for $i \in [1, d-1]$. Iterating α over $[1, m_{d+1}]$ extracts the full set of relaxed opening.

For now, ignore the runtime of sampling challenges for accepting transcripts, and suppose the extractor obtains two accepting transcripts whose first challenges $\mathbf{C}_1$ and $\mathbf{C}'_1$ differ only in the α^{th} row. The verification equation gives

$$\begin{aligned} \mathbf{A}_1 \mathbf{Z}_1 &\equiv \mathbf{T} \mathbf{C}_1 \bmod q_1, \\ \mathbf{A}_1 \mathbf{Z}'_1 &\equiv \mathbf{T} \mathbf{C}'_1 \bmod q_1. \end{aligned}$$

Choose a coordinate r where the two α^{th} challenge rows differ, and orient the subtraction such that $(\mathbf{C}_1 - \mathbf{C}'_1)_{\alpha,r} = 1$. Taking the r^{th} column after subtracting the two accepting equations gives

$$\mathbf{A}_1(\mathbf{z}_{1,r} - \mathbf{z}'_{1,r}) \equiv \mathbf{t}_\alpha \pmod{q_1}.$$

Thus the extractor obtains a bounded preimage of $\mathbf{t}_\alpha$ under $\mathbf{A}_1$.

In a levelled commitment, this vector is only the upper-layer input; the tree construction is then used to recursively explain this extracted column by lower-level preimages.

For the next transition, write the row blocks of $\mathbf{Z}_1$ as $\mathbf{Z}_{1,1}, \dots, \mathbf{Z}_{1,m_1}$, so that

$$\text{BT}_1(\mathbf{Z}_1) = [\mathbf{Z}_{1,1} \mid \dots \mid \mathbf{Z}_{1,m_1}].$$

The r^{th} column of $\mathbf{Z}_1$ is then spread across the columns $\{r_j = r + (j-1)\lambda\}_{j \in [1, m_1]}$ of $\text{BT}_1(\mathbf{Z}_1)$, and the verification is

$$\mathbf{A}_2 \mathbf{Z}_2 \equiv \text{BT}_1(\mathbf{Z}_1) \mathbf{C}_2 \pmod{q_2}.$$

To extract the j^{th} block of the r^{th} column of $\mathbf{Z}_1$, the extractor therefore rewinds the r_j^{th} row of $\mathbf{C}_2$. If two accepting second challenges $\mathbf{C}_2$ and $\mathbf{C}'_2$ differ only in this row, choose a coordinate s and orient such that $(\mathbf{C}_2 - \mathbf{C}'_2)_{r_j,s} = 1$. Taking the s^{th} column after subtracting the two level-two equations gives

$$\mathbf{A}_2(\mathbf{z}_{2,s} - \mathbf{z}'_{2,s}) \equiv \text{BT}_1(\mathbf{Z}_1)_{r_j} \pmod{q_2}.$$

The right hand side is the j^{th} block of the extracted r^{th} column of $\mathbf{Z}_1$, which is the r_j^{th} column of $\text{BT}_1(\mathbf{Z}_1)$. Iterating over $j \in [1, m_1]$ explains all blocks of that column; these j 's become the child slots in the transcript tree.

The heavy-row argument lemma 4.8.8 from [Dam10] is still the runtime core of the extractor.

Lemma 4.9.3 (Scaled Heavy Row Lemma). *Let $\mathbf{H} \in \{0, 1\}^{n \times m}$ have density at least ε . Fix $\delta \leq \varepsilon$, and call a row δ -heavy if its row-density is at least δ . If a 1-entry of $\mathbf{H}$ is sampled uniformly, then its row is δ -heavy with probability at least $1 - \delta/\varepsilon$. Moreover, conditioned on a δ -heavy row, resampling a column gives another 1-entry with probability at least δ . If the resampled column is required to differ from the original one, this probability is at least $\delta - 1/m$.*

Proof. There are at most δmn 1-entries in light rows, so at least $(\varepsilon - \delta)mn$ 1-entries in heavy rows, and the probability lower bound $1 - \delta/\varepsilon$ holds. The rest follows the same argument in lemma 4.8.8. $\square$

The 0-1 matrices are used in the runtime analysis, serving the same challenge-indexing purpose as in [BBC⁺18], but now used in the extraction tree. Fix a partial transcript prefix $\pi_i = (\mathbf{C}_1, \mathbf{Z}_1, \dots, \mathbf{C}_i, \mathbf{Z}_i)$. First define

$$\mathbf{H}_i^{(\pi_i)}[\chi, \mathbf{C}'_{i+1}] [\mathbf{C}'_{i+2}, \dots, \mathbf{C}'_d] \in \{0, 1\}.$$

The row is indexed by the prover randomness χ and the current challenge $\mathbf{C}'_{i+1}$. The column is indexed by the future challenge choice $\mathbf{C}'_{i+2}, \dots, \mathbf{C}'_d$. The entry is 1 iff these choices give a full accepting transcript extending π_i , such that

$$\mathbf{H}_{i-1}^{\pi_{i-1}}[\chi, \mathbf{C}_i] [\mathbf{C}'_{i+1}, \dots, \mathbf{C}'_d] = \mathbf{H}_i^{(\pi_i)}[\chi, \mathbf{C}'_{i+1}] [\mathbf{C}'_{i+2}, \dots, \mathbf{C}'_d].$$

To apply the heavy-row argument to one row of the next challenge, reindex this same 0-1 matrix. Let ρ be a row index in $\mathbf{C}'_{i+1}$. Write $\mathbf{c}_{i+1,\rho}^T$ for the ρ^{th} row of $\mathbf{C}'_{i+1}$, and $\hat{\mathbf{C}}'_{i+1,\rho}$ for the rows obtained from $\mathbf{C}'_{i+1}$ by deleting $\mathbf{c}_{i+1,\rho}^T$. Define

$$H_i^{(\pi_i, \rho)}[\chi, \hat{\mathbf{C}}'_{i+1,\rho}] [\mathbf{c}_{i+1,\rho}^T, \mathbf{C}'_{i+2}, \dots, \mathbf{C}'_d] = H_i^{(\pi_i)}[\chi, \mathbf{C}'_{i+1}] [\mathbf{C}'_{i+2}, \dots, \mathbf{C}'_d].$$

Thus the rows of $H_i^{(\pi_i, \rho)}$ fix the prover randomness and the non-interesting rows of $\mathbf{C}'_{i+1}$, while its columns are indexed by the ρ^{th} row of $\mathbf{C}'_{i+1}$ together with all later challenges. At the root, $\pi_0 = \emptyset$ and the first rewinding uses the α^{th} row of $\mathbf{C}'_1$. Beyond the root, BT determines how the current extraction target is seen at the next level: it becomes several target rows in $\mathbf{C}'_{i+1}$ spaced by λ , as was previously discussed.

The original one-level use sets $\delta = \varepsilon/2$, and the probability is at least $1/2$ in lemma 4.8.8. For the tree extractor, the same lemma is used with a larger slack K : set the heavy threshold to $\delta = \varepsilon/K$, and an accepting entry lands in a heavy row with probability at least $1 - 1/K$. After landing in such a row, resampling only the target row(s) in the challenge being rewound gives another accepting challenge with probability at least $\varepsilon/K - |C|^{-\lambda}$.

The extractor in [BLNS20] builds a tree of partial transcripts, and the leaves are accepting full transcripts. A state at depth i is written as

$$\text{TC}_i([j_1, \dots, j_i], \mathbf{C}_i, \mathbf{Z}_i, t).$$

Here $[j_1, \dots, j_i]$ is the address of the vertex in the tree, where each j_k records which child slot was taken:

- The block of the current column (because of BT) being followed,
- One of the two rewounded children, left or right.

The arguments $\mathbf{C}_i, \mathbf{Z}_i, t$ are the current challenge, the current response, and the column currently being extracted. Reading the challenges and responses along the addressed path gives the prefix $\pi_i = (\mathbf{C}_1, \mathbf{Z}_1, \dots, \mathbf{C}_i, \mathbf{Z}_i)$ used when rewinding the prover.

The first step is the same operation with empty prefix and target row α in $\mathbf{C}_1$ (namely the α^{th} column of $\mathbf{T}$). Sample two first challenges that differ only in this row, run the prover on both, and choose a coordinate t where the two α^{th} rows differ, oriented so the difference is 1. The two resulting children are depth-1 states storing their own $\mathbf{C}_1, \mathbf{Z}_1$, and both carry the index t . Their edge already gives

$$\mathbf{A}_1(\mathbf{z}_{1,t}^{(0)} - \mathbf{z}_{1,t}^{(1)}) \equiv \mathbf{t}_\alpha \pmod{q_1}.$$

Fix a non-root, non-leaf state $\text{TC}_i([j_1, \dots, j_i], \mathbf{C}_i, \mathbf{Z}_i, t)$. The first task is to determine which rows of the next challenge have to be rewound. The t^{th} column of $\mathbf{Z}_i$ is spread by BT_i across the rows

$$\rho_\ell = t + (\ell - 1)\lambda, \quad \ell \in [1, m_i],$$

of the next challenge $\mathbf{C}_{i+1}$. For each row ρ_ℓ , the algorithm creates a pair of child candidates by sampling two challenges $\tilde{\mathbf{C}}_{i+1}$ and $\tilde{\mathbf{C}}'_{i+1}$ that are equivalent except for the ρ_ℓ^{th} row, and rewinding the prover from the prefix π_i to obtain the two responses $\tilde{\mathbf{Z}}_{i+1}$ and $\tilde{\mathbf{Z}}'_{i+1}$.

For $i < d - 1$, these candidates are the next partial-transcript children. For $i = d - 1$, they must be accepting leaf transcripts instead.

- For the first child, the extractor keeps resampling the ρ_ℓ^{th} row of $\tilde{\mathbf{C}}_{i+1}$ and reruns the prover until the resulting transcript is accepting; it aborts after $\lambda K/\varepsilon$ failed attempts.
- For the second child, it rewinds to the same prefix and repeats the same search for $\tilde{\mathbf{C}}'_{i+1}$; it aborts after $2\lambda K/\varepsilon$ failed attempts.

Choose a coordinate $s_\ell \in [1, \lambda]$ where the two ρ_ℓ^{th} challenge rows differ,

$$(\tilde{\mathbf{C}}_{i+1} - \tilde{\mathbf{C}}'_{i+1})_{\rho_\ell, s_\ell} \neq 0,$$

and orient the pair so the difference is 1. The two responses become the $(\ell, 0)$ and $(\ell, 1)$ children, and both use s_ℓ as the next column to be extracted. The edge created by this local rewind is

$$\mathbf{A}_{i+1}(\tilde{\mathbf{z}}_{i+1, s_\ell} - \tilde{\mathbf{z}}'_{i+1, s_\ell}) \equiv \text{BT}_i(\mathbf{Z}_i)_{\rho_\ell} \pmod{q_{i+1}}.$$

Thus a depth i state with $i \geq 1$ creates $2m_i$ children, and level $i > 0$ has $2^i M_{i-1}$ vertices, as in the paper.

Now we can bound the abort probability in the same order as the paper. Let κ denote the per-step heavy-row fraction parameter, and let the retry parameter in the tree construction be $K = \kappa^{2d}$. Fix a level- $(d-1)$ vertex V , and write

$$(V_0, V_1, \dots, V_{d-2}, V)$$

for the path from the root V_0 to V . Let $r_1, \dots, r_\gamma$ be the indices for which V_{r_i} is a right child. Let $\text{abort}(V)$ be the event that the tree construction aborts while trying to complete this fixed leaf vertex V . Write $h_i(\eta)$ for the event that the row selected by this path is η -heavy in $\mathbf{H}_i$, and write $h'_i(\eta)$ for the corresponding event in the reindexed matrix $\mathbf{H}'_i$. The aborting probability for left and right children is derived differently: for a left child it comes directly from $\mathbf{H}_i$, while for a right child it first comes through the reindexed matrix $\mathbf{H}'_i$ and then through $\mathbf{H}_i$.

The first step is the only boundary case. If V_1 is a right child, then the first challenge was obtained by resampling one target row while keeping the other rows fixed, so the proof first goes through $\mathbf{H}'_0$:

$$\Pr[\text{abort}(V)] \leq \Pr[\text{abort}(V) \mid h'_0(\varepsilon/\kappa)] + \Pr[\neg h'_0(\varepsilon/\kappa)] \leq \Pr[\text{abort}(V) \mid h'_0(\varepsilon/\kappa)] + 1/\kappa.$$

Here $\mathbf{H}'_0$ moves the target row of $\mathbf{C}'_1$ from the row index into the column index: a row of $\mathbf{H}'_0$ is fixed by the non-target rows of $\mathbf{C}'_1$, and is the concatenation of the ordinary $\mathbf{H}_0$ rows obtained by varying that target row. Thus the heavy-row lemma applies to $\mathbf{H}'_0$ after this row-column rearrangement. Conditioned on $h'_0(\varepsilon/\kappa)$, the selected row of $\mathbf{H}'_0$ has density at least ε/κ . Now split this selected row according to the value of the target row. Each piece is one ordinary row of $\mathbf{H}_0$, so another application of the heavy-row lemma gives

$$\Pr[\neg h_0(\varepsilon/\kappa^2) \mid h'_0(\varepsilon/\kappa)] \leq 1/\kappa,$$

thus

$$\Pr[\text{abort}(V) \mid h'_0(\varepsilon/\kappa)] \leq \Pr[\text{abort}(V) \mid h'_0(\varepsilon/\kappa) \wedge h_0(\varepsilon/\kappa^2)] + 1/\kappa.$$

If V_1 is a left child, then the first challenge is sampled as a full challenge, and the bound starts directly from $\mathbf{H}_0$:

$$\Pr[\text{abort}(V)] \leq \Pr[\text{abort}(V) \mid h_0(\varepsilon/\kappa)] + \Pr[\neg h_0(\varepsilon/\kappa)] \leq \Pr[\text{abort}(V) \mid h_0(\varepsilon/\kappa)] + 1/\kappa.$$

After this boundary step, the proof continues in the same way through $h_1, \dots, h_{d-2}$, inserting h'_{r_j-1} immediately before h_{r_j-1} for every right child V_{r_j} on the path. Let E be the conjunction of these heavy-row events, such that the threshold divided by a factor of κ at each peeling step. Then

$$\Pr[\neg E] \leq (d-1+\gamma)/\kappa \leq 2(d-1)/\kappa.$$

On E , the event $h_{d-2}(\varepsilon/\kappa^{d-1+\gamma})$ holds. Concretely, for the prefix π_{d-2} on the path to V , the row $\mathbf{H}_{d-2}^{(\pi_{d-2})}[\chi, \mathbf{C}_{d-1}]$ has density at least $\varepsilon/\kappa^{d-1+\gamma}$. The columns of this row are indexed by the remaining final challenge $\mathbf{C}_d$, so this is the success probability for one attempt in the first final-challenge search. Since $\gamma \leq d-1$ and $K = \kappa^{2d}$, this local success probability is lower-bounded by $\varepsilon/\kappa^{d-1+\gamma} \geq \varepsilon/K$. With $\lambda K/\varepsilon$ attempts, the first leaf abort probability less than $e^{-\lambda}$.

The paired leaf search needs one more heavy-row step. Write ρ for the target row of the final challenge determined by this leaf. After the first final challenge is fixed, the extractor resamples only this row for the second final challenge, so the analysis passes to the row

$$\mathbf{H}_{d-1}^{(\pi_{d-1}, \rho)}[\chi, \hat{\mathbf{C}}_{d, \rho}]$$

whose columns are indexed by the resampled target row. Conditioned on $h'_{d-1}(\varepsilon/\kappa^{d+\gamma})$, this row has density at least $\varepsilon/\kappa^{d+\gamma}$. This event fails with probability at most $1/\kappa$, and, when it holds, the successful resampling probability is at least

$$\varepsilon/\kappa^{d+\gamma} - |\mathbf{C}|^{-\lambda},$$

where the subtraction excludes choosing the same target row again. By the lower bound on ε in theorem 4.9.2, this is still at least $\varepsilon/(2K)$, and the $2\lambda K/\varepsilon$ attempts make the second leaf abort probability less than $e^{-\lambda}$.

Combining the failure of E , the extra possible failure of $h'_{d-1}(\varepsilon/\kappa^{d+\gamma})$, and the two leaf searches, for a fixed leaf vertex we get

$$\Pr[\text{abort}(V)] \leq 2e^{-\lambda} + 2d/\kappa.$$

The paper sets $\kappa = 4dN$. Union bounding over all vertices at level $d-1$, which are $2^{d-1}M_{d-2}$ of them, given that d is a constant, and $2^{d-1}M_{d-2} < N$, then the aborting probability is at most

$$\frac{2N}{e^\lambda} + \frac{2dN}{\kappa} = \frac{2N}{e^\lambda} + \frac{1}{2}.$$

Thus one run of TreeConstruct succeeds with constant probability, up to the negligible $2N/e^\lambda$ term. For a fixed target column α , running $O(\lambda)$ independent copies gives a successful tree except with negligible probability. The extractor needs this for all m_{d+1} columns of $\mathbf{T}$, so it repeats the same procedure for every $\alpha \in [1, m_{d+1}]$. Since $m_{d+1} \leq N$, $N = \text{poly}(\lambda)$, d is constant, and $K = (4dN)^{2d} = \text{poly}(\lambda)$, the total expected running time remains $\text{poly}(\lambda, 1/\varepsilon)$.

Once the tree is built, extraction is just reading off neighboring children. For a vertex whose current target is the t^{th} column, the two children created for the row ρ_ℓ give

$$\mathbf{A}_{i+1} \left(\tilde{\mathbf{z}}_{i+1, s_\ell} - \tilde{\mathbf{z}}'_{i+1, s_\ell} \right) \equiv \text{BT}_i(\mathbf{Z}_i)_{\rho_\ell} \pmod{q_{i+1}}.$$

Collecting these sibling differences over $\ell \in [1, m_i]$ gives the preimage block for the parent target column after the block transition. Repeating down the tree and concatenating the extracted blocks level by level produces the relaxed opening $(\tilde{\mathbf{s}}, \mathbf{r}_1, \dots, \mathbf{r}_{d-1})$ for the chosen column of $\mathbf{T}$; Iterating over $[1, m_{d+1}]$ columns gives the full relaxed opening.

The binding intuition is then about the relaxed relation extracted from the argument. For two levels, write the lifted inputs seen by the top matrix $\mathbf{A}_1$ as

$$\begin{aligned} \tilde{\mathbf{S}}_\emptyset &= (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \tilde{\mathbf{S}} \pmod{q_2 + q_2 \mathbf{R}}, \\ \tilde{\mathbf{S}}'_\emptyset &= (\mathbf{I}_{m_1} \otimes \mathbf{A}_2) \tilde{\mathbf{S}}' \pmod{q_2 + q_2 \mathbf{R}'}. \end{aligned}$$

Both relaxed openings to the same $\mathbf{T}$ satisfy

$$\mathbf{A}_1 \tilde{\mathbf{S}}_\emptyset \equiv \mathbf{A}_1 \tilde{\mathbf{S}}'_\emptyset \equiv \mathbf{T} \pmod{q_1}.$$

Suppose the two bounded relaxed openings $(\tilde{\mathbf{S}}, \mathbf{R})$ and $(\tilde{\mathbf{S}}', \mathbf{R}')$ are distinct, and both open to the same $\mathbf{T}$. If $\mathbf{R} \neq \mathbf{R}'$, then the two integer representations of the upper-layer inputs are different, and hence

$$\mathbf{A}_1(\tilde{\mathbf{S}}_\emptyset - \tilde{\mathbf{S}}'_\emptyset) \equiv \mathbf{0} \pmod{q_1}$$

gives a short SIS relation for $\mathbf{A}_1$. If $\mathbf{R} = \mathbf{R}'$, then $\tilde{\mathbf{S}} \neq \tilde{\mathbf{S}}'$. In this case, either the lower-layer images of $\mathbf{I}_{m_1} \otimes \mathbf{A}_2$ differ modulo q_2 , which again makes the lifted top inputs different and gives SIS for $\mathbf{A}_1$, or their equality modulo q_2 gives a short nonzero SIS solution for $\mathbf{A}_2$. The separated moduli and norm bounds are chosen so that the differences stay inside the SIS-binding range.

Parameters and slack. For constant d , the parameter balancing in the paper gives communication

$$O\left(N^{1/(d+1)} \cdot \left(d^3 \lambda \log^2 N + d \lambda^2\right)\right).$$

If $d+1 = \log N$, this becomes

$$O(\lambda \log^5 N + \lambda^2 \log N).$$

The extractor recovers a relaxed opening: a matrix $\tilde{\mathbf{S}}$ and matrices $\mathbf{R}_1, \dots, \mathbf{R}_{d-1}$ satisfying the relaxed recursive commitment equation. The communication saving therefore comes with knowledge slack, with parameters set so that the slack remains below the relevant SIS bounds.

Bulletproofs-style lattice folding. The second construction asks whether the Bulletproofs [BBB⁺18] folding trick can be made lattice-compatible. The starting statement is the RSIS equation

$$\mathbf{A}\mathbf{s} = \mathbf{t}, \quad \mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2] \in R^{1 \times k}, \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{s}_2 \end{bmatrix} \in R^k,$$

over $R = \mathbb{Z}[X]/(X^n + 1)$. Here the number of integer secrets is $N = kn$. In one folding round, the prover sends

$$\mathbf{l} = \mathbf{A}_1 \mathbf{s}_2, \quad \mathbf{r} = \mathbf{A}_2 \mathbf{s}_1.$$

The verifier samples a monomial challenge $c \xleftarrow{R} \{X^i : i \in \mathbb{Z}_{2n}\}$, and the prover responds with

$$\mathbf{z} = \mathbf{s}_1 + c \mathbf{s}_2.$$

Then

$$(c\mathbf{A}_1 + \mathbf{A}_2)\mathbf{z} = c^2\mathbf{l} + c\mathbf{t} + \mathbf{r},$$

so the verifier can replace the old statement $(\mathbf{A}, \mathbf{t})$ by

$$\mathbf{B} = c\mathbf{A}_1 + \mathbf{A}_2, \quad \mathbf{t}' = c^2\mathbf{l} + c\mathbf{t} + \mathbf{r},$$

with witness $\mathbf{z}$ of half the length. The verifier also checks that $\mathbf{z}$ remains short.

Repeating for d rounds reduces the witness length by 2^d ; setting $d = \log k$ leaves a single ring element and gives logarithmic-length communication.

For one round, extraction rewinds to three accepting transcripts with distinct challenges c_1, c_2, c_3 gives

$$(c_j \mathbf{A}_1 + \mathbf{A}_2) \mathbf{z}_j = c_j^2 \mathbf{l} + c_j \mathbf{t} + \mathbf{r}, \quad j \in [1, 3],$$

for the same $\mathbf{l}, \mathbf{r}, \mathbf{t}$. Choose coefficients $\lambda_1, \lambda_2, \lambda_3 \in R$ by solving the Vandermonde system

$$\begin{bmatrix} c_1^2 & c_2^2 & c_3^2 \\ c_1 & c_2 & c_3 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}.$$

Then the $\mathbf{l}$ and $\mathbf{r}$ terms cancel, and formally

$$\mathbf{A}_1 \sum_{j=1}^3 \lambda_j c_j \mathbf{z}_j + \mathbf{A}_2 \sum_{j=1}^3 \lambda_j \mathbf{z}_j = \mathbf{A} \sum_{j=1}^3 \lambda_j \begin{bmatrix} c_j \mathbf{z}_j \\ \mathbf{z}_j \end{bmatrix} = \mathbf{t}.$$

The challenge set is chosen to consist of monomials X^i . For pairwise distinct monomials, the paper uses the fact that $2/(X^i - X^j)$ has coefficients in $\{-1, 0, 1\}$, and obtains $\|\lambda_j\|_\infty \leq 2n^2$. After clearing this denominator, one obtains a one-round relaxed preimage

$$\bar{\mathbf{z}} = 8 \sum_{j=1}^3 \lambda_j \begin{bmatrix} c_j \mathbf{z}_j \\ \mathbf{z}_j \end{bmatrix}, \quad \mathbf{A} \bar{\mathbf{z}} = 8\mathbf{t},$$

with $\|\bar{\mathbf{z}}\|_\infty \leq O(12n^3p)$ when each accepted response is bounded by $2p$.

This is the sense in which the Bulletproofs-style extraction is approximate: the extractor does not recover a short preimage of the original target $\mathbf{t}$, but a larger preimage of the scaled target $8\mathbf{t}$. Recursing through d folds multiplies both effects. Extraction gives a relaxed solution to

$$\mathbf{A} \bar{\mathbf{z}} = 8^d \mathbf{t}$$

with

$$\|\bar{\mathbf{z}}\|_\infty = O(n^{3d} 12^d p)$$

when $\|\mathbf{s}\|_\infty \leq p$. Thus logarithmic folding buys a much smaller transcript, but pays a large slack: the target is scaled by 8^d and the extracted witness norm grows by roughly $(12n^3)^d$. With repetition to reduce soundness error, the proof size is

$$O\left(\frac{\lambda N \log(2^d p)}{2^d \log n} + \lambda d^2 n\right).$$

Takeaway. Both constructions are non-PCP routes to succinct, plausibly quantum-safe ZK, but neither extracts an exact original witness. Levelled commitments have larger proofs but smaller, more controllable slack. Bulletproofs folding has the better proof-size expression when slack is cheap, but the relaxed equation $\mathbf{A}\bar{\mathbf{z}} = 8^d \mathbf{t}$ and coefficient growth can force much larger parameters in applications where the extracted witness feeds into other cryptographic components. The paper’s main open direction is to keep the algebraic folding/compression while reducing coefficient growth enough to become competitive with generic PCP/Merkle-tree based post-quantum arguments in concrete size.

4.10 Short Invertible Polynomial Ring Elements

This note combines two related papers. The first work [LS18] explains how the factorization pattern of cyclotomic rings modulo q affects the availability of short invertible elements, which is for the challenge spaces in Fiat-Shamir with Abort sigma protocol. The second work [BLS19] uses the fully splitting case to turn a coordinate-wise shortness statement into an algebraic identity via the NTT.

Cyclotomic Polynomial Ring Splitting. Let $R = \mathbb{Z}[X]/(\Phi_\ell(X))$ and $R_q = \mathbb{Z}_q[X]/(\Phi_\ell(X))$. The factorization of Φ_ℓ modulo q controls how much CRT structure is available in R_q . If Φ_ℓ splits into a few irreducible factors, one obtains a partially splitting ring. This is useful for fast multiplication and for constructing large challenge sets whose nonzero differences remain invertible [LS18].

For an integer $m > 1$, write

$$\varphi(m) = m \cdot \prod_{\substack{p \text{ prime} \\ p|m}} \frac{p-1}{p}, \quad \delta(m) = \prod_{\substack{p \text{ prime} \\ p|m}} p,$$

such that $m = \varphi(m)/\varphi(\delta(m)) \times \delta(m)$, and

$$\tau(m) = \begin{cases} m, & m \text{ is odd,} \\ m/2, & m \text{ is even.} \end{cases}$$

For $a \in \mathbb{Z}_m^*$,

$$\text{ord}_m(a) = \min\{k > 0 : a^k \equiv 1 \pmod{m}\}.$$

The paper also writes $s_1(m)$ for the largest singular value of the Vandermonde matrix associated with the primitive m -th roots of unity. We only need φ and ord for the splitting theorem below; s_1 appears in the later bounds showing that small elements are invertible.

The cyclotomic polynomial Φ_m is fully splitting modulo q if it factors into linear terms over $\mathbb{Z}_q$:

$$\Phi_m(X) \equiv \prod_{j \in [1, \varphi(m)]} (X - r_j) \pmod{q},$$

where the $r_j \in \mathbb{Z}_q^*$ are exactly the primitive m^{th} roots of unity modulo q , such that $\text{ord}_q(r_j) = \varphi(m)$. Since $\mathbb{Z}_q^*$ is cyclic of order $q-1$, this happens exactly when $\mathbb{Z}_q$ contains primitive m^{th} roots of unity, namely

$$m \mid q-1 \quad \Leftrightarrow \quad q \equiv 1 \pmod{m}.$$

In this case the CRT decomposes

$$\mathbb{Z}_q[X]/(\Phi_m(X)) \cong \prod_{j \in [1, \varphi(m)]} \mathbb{Z}_q[X]/(X - r_j) \cong \mathbb{Z}_q^{\varphi(m)}.$$

The following theorem describes the characteristics of cyclotomic polynomial ring splitting, and more specifically, in the partial splitting case.

Theorem 4.10.1 (Cyclotomic Partial Splitting [LS18]). *Let*

$$m = \prod_i p_i^{e_i}, \quad z = \prod_i p_i^{f_i},$$

where $e_i \geq 1$ and $1 \leq f_i \leq e_i$. If q is prime, $q \equiv 1 \pmod{z}$, and $\text{ord}_m(q) = m/z$, then for r_j the primitive z^{th} roots of unity,

$$\Phi_m(X) \equiv \prod_{j \in [1, \varphi(z)]} (X^{m/z} - r_j) \pmod{q}$$

and each factor $X^{m/z} - r_j$ is irreducible over $\mathbb{Z}_q$.

The condition $q \equiv 1 \pmod{z}$ means, there is a size z multiplicative subgroup in $\mathbb{Z}_q^*$, such that r_j are z^{th} primitive roots of unity, while $\text{ord}_m(q) = m/z$ means, there is no elements in $\mathbb{Z}_q^*$ that r_j can split to in the directions defined by the prime factors of m/z .

It is better to understand $\text{ord}_m(q)$ in the Frobenius endomorphism aspect of view.

Definition 4.10.2 (Frobenius endomorphism and orbit). Let $\overline{\mathbb{Z}_q}$ be an ambient field containing $\mathbb{Z}_q$ and all the roots of $\Phi_m(X)$. We use it only so that we can discuss roots of Φ_m even when they do not lie in $\mathbb{Z}_q$. The polynomials are still defined over $\mathbb{Z}_q$, meaning their coefficients lie in $\mathbb{Z}_q$; we only evaluate them over $\overline{\mathbb{Z}_q}$. The q -Frobenius map is

$$\text{Endo}_q(\alpha) = \alpha^q.$$

It fixes $\mathbb{Z}_q$ pointwise, since $a^q = a$ for every $a \in \mathbb{Z}_q$. For $\alpha \in \overline{\mathbb{Z}_q}$, its Frobenius orbit is defined as $\alpha, \alpha^q, \alpha^{q^2}, \dots$ and the orbit length is the smallest $t > 0$ such that $\alpha^{q^t} = \alpha$.

If α is a primitive m^{th} root of unity, then the condition $\alpha^{q^t} = \alpha$ for orbit to end is equivalent to $q^t \equiv 1 \pmod{m}$ over the exponent. Hence the Frobenius orbit length of a primitive m^{th} root is $\text{ord}_m(q)$.

We then construct the minimal polynomial from a Frobenius orbit, which is irreducible over $\mathbb{Z}_q$.

Lemma 4.10.3 (Minimal polynomial from Frobenius orbit). Let $\alpha \in \overline{\mathbb{Z}_q}$ have Frobenius orbit length t , then

$$\mu_\alpha(X) = \prod_{i \in [0, t-1]} (X - \alpha^{q^i})$$

is the minimal polynomial for α over $\mathbb{Z}_q$. In particular, $\mu_\alpha(X)$ is irreducible over $\mathbb{Z}_q$ and has coefficients in $\mathbb{Z}_q$, and $\deg(\mu_\alpha) = t$.

Proof. This product lies in $\mathbb{Z}_q[X]$ as the q -Frobenius map cyclically permutes the roots in the product, which means

$$\prod_{i \in [0, t-1]} (X - (\alpha^{q^i})^q) = \prod_{i \in [1, t]} (X - \alpha^{q^i}) = \mu_\alpha(X),$$

and all the binomial terms are over $\mathbb{Z}_q$ satisfying $b = b^q$, then each coefficient $c^q = c$, therefore each $c \in \mathbb{Z}_q$.

Let $g(X) \in \mathbb{Z}_q[X]$ satisfy $g(\alpha) = 0$. By prior discussion, $g(\alpha^q) = 0$ as each coefficient over $\mathbb{Z}_q$ has $c = c^q$. Hence, $g(X)$ vanishes over the Frobenius orbit, so $\mu_\alpha | g$, and μ_α is the minimal polynomial for α over $\mathbb{Z}_q$.

A minimal polynomials over a field are irreducible, or otherwise contradicting the minimality of the degree. $\square$

If ζ is a m^{th} root of unity, $\zeta^{m/z}$ is a z^{th} root of unity. The primitiveness is preserved, say ζ is a primitive m^{th} root of unity, then $\zeta^{m/z}$ is a primitive z^{th} root of unity.

With $\text{ord}_m(q) = m/z$, for a primitive m^{th} root of unity in $\overline{\mathbb{Z}_q}$, the Frobenius orbit has length m/z . By $q \equiv 1 \pmod{z}$, we know the z^{th} roots of unity are over $\mathbb{Z}_q$. Hence, for each primitive m^{th} root of unity ζ , we have

$$(\zeta^{q^i})^{m/z} = (\zeta^{m/z})^{q^i} = r^{q^i} = r,$$

where $\zeta^{m/z} = r$, and r is a primitive z^{th} root of unity. Therefore, the m/z long Frobenius orbit vanishes over $X^{m/z} - r_j$ that is irreducible over $\mathbb{Z}_q$ for all $j \in [1, \varphi(z)]$.

The choice of z controls the splitting pattern: the ring has $\varphi(z)$ CRT slots, each tied to an extension field of degree m/z over $\mathbb{Z}_q$. The fully splitting case is when $z = m$, where $\text{ord}_m(q) = 1$, namely $q \equiv 1 \pmod{m}$, then

$$\Phi_m(X) \equiv \prod_{j \in [1, \varphi(m)]} (X - r_j) \pmod{q}.$$

By CRT decomposition, evaluation at these roots gives the NTT isomorphism. Let $n = \varphi(m)$, then

$$R_q \cong \mathbb{Z}_q^n, \quad s \mapsto \widehat{s} = (s(r_1), \dots, s(r_n)),$$

and multiplication in R_q becomes coordinate-wise multiplication after applying the NTT. In the power-of-two case $R = \mathbb{Z}[X]/(X^n + 1) = \mathbb{Z}[X]/(\Phi_{2n}(X))$, this condition becomes $q \equiv 1 \pmod{2n}$, so $X^n + 1$ has all of its roots in $\mathbb{Z}_q$.

Splitting and Short Invertability. Let $R_m = \mathbb{Z}[X]/(\Phi_m(X))$ and let $\omega_1, \dots, \omega_{\varphi(m)}$ be the primitive m^{th} roots of unity defined over complex field. The short-invertibility question in [LS18] is whether a nonzero short element of $R_{m,q}$ can vanish in one of the CRT coordinates induced by the above splitting. Suppose

$$\Phi_m(X) \equiv \prod_j f_j(X) \pmod{q},$$

where the f_j are irreducible over $\mathbb{Z}_q$, then

$$R_{m,q} \cong \prod_j \mathbb{Z}_q[X]/(f_j(X)),$$

and $y \in R_{m,q}$ is invertible when $y \not\equiv 0 \pmod{f_j(X)}$ for every irreducible f_j . Hence, non-invertibility means that y lands in the kernel of at least one CRT projection, and [LS18] views such kernels as ideal lattices under the canonical embedding. For $\mathbf{y} = \sum_{i \in [0, \varphi(m)-1]} y_i X^i \in R_m$, the canonical embedding is the complex evaluation vector

$$\sigma(\mathbf{y}) = (\mathbf{y}(\omega_1), \dots, \mathbf{y}(\omega_{\varphi(m)})) \in \mathbb{C}^{\varphi(m)}, \quad (4.1)$$

and the embedding norm is

$$\|\mathbf{y}\|_e = \|\sigma(\mathbf{y})\|_2. \quad (4.2)$$

Let V_m be the associated Vandermonde matrix

$$\mathbf{V}_m = \begin{pmatrix} 1 & \omega_1 & \cdots & \omega_1^{\varphi(m)-1} \\ \vdots & \vdots & & \vdots \\ 1 & \omega_{\varphi(m)} & \cdots & \omega_{\varphi(m)}^{\varphi(m)-1} \end{pmatrix}.$$

The parameter $s_1(m)$ defined as

$$s_1(m) = \max_{\mathbf{u} \in \mathbb{C}^{\varphi(m)} \setminus \{0\}} \frac{\|\mathbf{V}_m \mathbf{u}\|_2}{\|\mathbf{u}\|_2}$$

is the largest singular value of V_m , which is used to compare coefficient size with embedding size:

$$\|\mathbf{y}\|_e = \|\mathbf{V}_m \mathbf{y}\|_2 \leq s_1(m) \|\mathbf{y}\|_2, \quad (4.3)$$

where $\mathbf{y} = (y_0, \dots, y_{\varphi(m)-1})$ is the coefficient vector of $\mathbf{y}$.

We now construct the main result of [LS18] from the lower bound of embedding norm of an ideal lattice.

Lemma 4.10.4 (Embedding lower bound for ideal lattices [PR07, LS18]). *Let Λ be an ideal lattice in R_m , for every nonzero $\mathbf{y} \in \Lambda$ satisfies*

$$\|\mathbf{y}\|_e \geq \sqrt{\varphi(m)} \cdot \det(\Lambda)^{1/\varphi(m)}.$$

Applying eq. (4.3), we get

$$\|\mathbf{y}\|_2 \geq \lambda_1(\Lambda) \geq \frac{\sqrt{\varphi(m)}}{s_1(m)} \cdot \det(\Lambda)^{1/\varphi(m)}.$$

Consider $\Lambda = \{\mathbf{x} \in R_m : \mathbf{x} \equiv 0 \pmod{(f_i, q)}\}$, and the next result is immediate by lemma 4.10.4. First, Λ itself is an additive subgroup of $\mathbb{Z}^{\varphi(m)}$. Second, for any $a \in \Lambda$ and $b \in R_m$, we have $ab \in \Lambda$. Hence, Λ is an ideal lattice.

Lemma 4.10.5 (Polynomial ring ℓ_2 lower bound [LS18]). *Suppose $\Phi_m(X)$ splits modulo q into d distinct irreducible factors of degree $\varphi(m)/d$. If $\mathbf{y} \in R_m$ is nonzero and $\mathbf{y} \equiv 0 \pmod{(f_j(X), q)}$ for one irreducible factor f_j , then*

$$\|\mathbf{y}\|_2 \geq \frac{\sqrt{\varphi(m)}}{s_1(m)} q^{1/d}.$$

Since $\det(\Lambda) = |\mathbb{Z}^{\varphi(m)}/\Lambda| = q^{\varphi(m)/d}$, then lemma 4.10.5 is immediate.

If a nonzero $\mathbf{y}$ vanishes in one CRT coordinate, then $\mathbf{y}$ lies in Λ , and lemma 4.10.5 forces $\|\mathbf{y}\|_2$ to be large. Thus any nonzero $\mathbf{y}$ whose coefficient norm is below this lower bound cannot vanish in any CRT coordinate.

Theorem 4.10.6 (Short invertibility in partially splitting rings [LS18]). *Under the partial-splitting setup*

$$\Phi_m(X) \equiv \prod_{j \in [1, \varphi(z)]} (X^{m/z} - r_j) \pmod{q},$$

any nonzero $\mathbf{y} \in R_{m,q}$ satisfying either

$$\|\mathbf{y}\|_\infty < \frac{1}{s_1(z)} q^{1/\varphi(z)}$$

or

$$\|\mathbf{y}\|_2 < \frac{\sqrt{\varphi(m)}}{s_1(m)} q^{1/\varphi(z)}$$

is invertible in $R_{m,q}$.

Informally, short nonzero elements do not wrap around to zero in any NTT coordinate.

4.11 Lattice-Based Exact Argument of Knowledge

From [LS18] to [BLS19]. The role of splitting is different in [LS18] and [BLS19]. In [LS18], the NTT coordinates are used to argue that short nonzero elements cannot vanish in any coordinate, and hence are invertible. In [BLS19], we instead use the fully splitting case to expose the witness coordinates themselves: after the NTT, the ternary condition on $\widehat{\mathbf{s}}$ is equivalent to the single ring identity $\mathbf{s}(\mathbf{s} - \mathbf{1})(\mathbf{s} - \mathbf{2}) = \mathbf{0}$. The remaining issue is no longer ring splitting, but how to enforce this identity together with $\mathbf{A}\widehat{\mathbf{s}} = \overline{\mathbf{u}}$ inside an exact proof of knowledge. In our case, $R_q = \mathbb{Z}_q[X]/(\Phi_\ell)$ is a power-of-2 fully splitting cyclotomic ring by $q \equiv 1 \pmod{\ell}$, $\Phi_\ell = X^n + 1$, and $n = \varphi(\ell)$, so $\ell = 2n$.

The NTT Trick for Shortness. The starting exact relation is to prove knowledge of a short vector $\widehat{\mathbf{s}} \in \mathbb{Z}_q^n$ such that

$$\mathbf{A}\widehat{\mathbf{s}} = \overline{\mathbf{u}} \pmod{q}.$$

We focus on $\widehat{\mathbf{s}} \in \{0, 1, 2\}^n$. For $\mathbf{s} \in R_q$ whose NTT representation is $\widehat{\mathbf{s}}$, the ternary constraint is equivalent to

$$\widehat{\mathbf{s}} \in \{0, 1, 2\}^n \Leftrightarrow \widehat{\mathbf{s}} \circ (\widehat{\mathbf{s}} - \widehat{\mathbf{1}}) \circ (\widehat{\mathbf{s}} - \widehat{\mathbf{2}}) = \widehat{\mathbf{0}} \Leftrightarrow \mathbf{s}(\mathbf{s} - \mathbf{1})(\mathbf{s} - \mathbf{2}) = \mathbf{0} \in R_q, \quad (4.4)$$

and the “shortness” check is turned into an algebraic identity over a fully splitting ring. This is how exact shortness is proved without falling back to a coordinate-wise combinatorial argument.

First Attempt From Prior Approximate Arguments of Knowledge. We start from the Commit-and-Open Sigma-Protocol, which originated from [Lyu12], and was explicitly constructed in [BDL⁺18] for the approximate argument of knowledge for the commitment opening.

Since we cannot challenge $\mathbf{A}\widehat{\mathbf{s}}$ with short invertible ring elements, as $\widehat{\mathbf{s}} \in \mathbb{Z}_q^n$ is incompatible with the challenge applied through the polynomial product. Going down the path, for the knowledge soundness of the Sigma-Protocol, the challenge has to be from $\mathbb{Z}_q$ for a sufficiently large challenge set. Yet, given that the new challenge set is $\mathbb{Z}_q$, the norm of $c \cdot \widehat{\mathbf{s}}$ can be arbitrary, and the masking message $\overline{\mathbf{r}} \in \mathbb{Z}_q^n$ has to be uniformly random. Again, this violates the commitment binding in the Commit-and-Open Sigma-Protocol, where $\mathbf{A}$ as the Ajtai commitment has to bind $\overline{\mathbf{r}}$ to pin down the masking against the random oracle.

Looking closely, in previous approximate arguments of knowledge [BBC⁺18, BDL⁺18], $\mathbf{A}$ simultaneously takes over the role of the statement, and the binding commitment that binds the masking message and the last opening. The dual roles enforce that the masking and opening have to be short, or otherwise Ajtai will not bind.

In [BLS19], the role of binding commitment is lifted from $\mathbf{A}$, and moved to a lattice-based commitment scheme that binds arbitrary messages, such that the canonical representations $\mathbf{s}$ and $\mathbf{r} = \text{NTT}^{-1}(\overline{\mathbf{r}})$ are committed and pinned to the random oracle. In this way, $\mathbf{A}$ is treated as more of a linear system of the relation to be proved, rather than a commitment scheme that instantiates a Commit-and-Open Sigma-Protocol. The lifting of the binding role allows the challenge c to be uniformly random over $\mathbb{Z}_q$, and the masking $\overline{\mathbf{r}}$ for $c \cdot \widehat{\mathbf{s}}$ to be uniformly random over $\mathbb{Z}_q^n$.

Why [BDL⁺18] Commitment. Ajtai commitment [Ajt96] has been great as a compressing commitment, when the message is short. The binding and extractability hold by the SIS hardness assumption. However in our scenario, the NTT representation is short for the message to be committed, but not the canonical form of the message itself. Hence, we need a commitment that lets us manipulate arbitrary ring messages algebraically and homomorphically, while preserving hiding, binding, and extractability.

Now consider a commitment satisfying the requirements above, though it is not compressing like Ajtai commitment. The vanilla BDLOP commitment [BDL⁺18] is constructed as follows. For $\beta > 0$, write

$$S_\beta = \{\mathbf{a} \in R : \|\mathbf{a}\|_\infty \leq \beta\}. \quad (4.5)$$

The public parameter is a pair of matrices:

$$\begin{aligned} \mathbf{A}_1 &= [\mathbf{I}_n \quad \mathbf{A}'_1] \in R_q^{n \times k}, & \mathbf{A}'_1 &\stackrel{\mathbf{R}}{\leftarrow} R_q^{n \times (k-n)}, \\ \mathbf{A}_2 &= [\mathbf{0}_{\ell_{\text{Com}} \times n} \quad \mathbf{I}_{\ell_{\text{Com}}} \quad \mathbf{A}'_2] \in R_q^{\ell_{\text{Com}} \times k}, & \mathbf{A}'_2 &\stackrel{\mathbf{R}}{\leftarrow} R_q^{\ell_{\text{Com}} \times (k-n-\ell_{\text{Com}})}. \end{aligned}$$

Equivalently, the commitment map is defined by the vertical stack

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \in R_q^{(n+\ell_{\text{com}}) \times k}.$$

On message $\bar{\mathbf{m}} \in R_q^{\ell_{\text{com}}}$, sample short randomness $\bar{\mathbf{r}} \xleftarrow{R} S_\beta^k$ and output

$$\bar{\mathbf{c}} = \begin{bmatrix} \bar{\mathbf{c}}_1 \\ \bar{\mathbf{c}}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0}^n \\ \bar{\mathbf{m}} \end{bmatrix}.$$

Thus the message is added only to the lower coordinates, while both coordinates are still linear functions of the same short randomness $\bar{\mathbf{r}}$. The relaxed opening form used in the proof system allows an extracted scalar $\mathbf{f}$ from the challenge set:

$$\mathbf{f} \cdot \bar{\mathbf{c}} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix} \bar{\mathbf{r}} + \mathbf{f} \cdot \begin{bmatrix} \mathbf{0}^n \\ \bar{\mathbf{m}} \end{bmatrix}, \quad (4.6)$$

where the honest opening has $\mathbf{f} = \mathbf{1}$ and $\bar{\mathbf{r}}$ remains short. This identity-block form is what makes opening proofs compact: the prover can prove knowledge of a short $\bar{\mathbf{r}}$ for the top part and recover the message from the lower part.

Definition 4.11.1 ($\text{SKS}_{n,k,\gamma}^2$). Given $\mathbf{A}' \xleftarrow{R} R_q^{n \times (k-n)}$, find a nonzero vector $\bar{\mathbf{y}} = (\mathbf{y}_1, \dots, \mathbf{y}_k)^\top \in R^k$ such that

$$\begin{bmatrix} \mathbf{I}_n & \mathbf{A}' \end{bmatrix} \bar{\mathbf{y}} = \mathbf{0} \in R_q^n, \quad \|\mathbf{y}_i\|_2 \leq \gamma.$$

This is exactly the Module-SIS hardness assumption where $\mathbf{A}$ is in Hermite Normal Form.

Definition 4.11.2 ($\text{DKS}_{n,k,\beta}^\infty$). Given $\mathbf{A}' \xleftarrow{R} R_q^{n \times (k-n)}$, distinguish

$$(\mathbf{A}', [\mathbf{I}_n \quad \mathbf{A}'] \bar{\mathbf{y}}), \quad \bar{\mathbf{y}} \xleftarrow{R} S_\beta^k,$$

from $(\mathbf{A}', \bar{\mathbf{u}})$ for uniform $\bar{\mathbf{u}} \xleftarrow{R} R_q^n$. This is equivalent to the Module-LWE hardness assumption, when the number of samples is limited, namely for a limited number of MLWE samples sharing the same secret.

The hiding reduction is to $\text{DKS}_{n+\ell_{\text{com}},k,\beta}^\infty$: if $\mathbf{A}\bar{\mathbf{r}}$ is computationally (or statistically) indistinguishable from uniform, then adding $(\mathbf{0}, \bar{\mathbf{m}})$ does not reveal $\bar{\mathbf{m}}$. The binding reduction is to $\text{SKS}_{n,k,\gamma}^2$ for the top block $\mathbf{A}_1$: two different short relaxed openings give a nonzero vector of the form $\mathbf{f}'\bar{\mathbf{r}} - \bar{\mathbf{f}}'\bar{\mathbf{r}}$ with $\mathbf{A}_1(\mathbf{f}'\bar{\mathbf{r}} - \bar{\mathbf{f}}'\bar{\mathbf{r}}) = \mathbf{0}$.

BDLOP Specialized in [BLS19]. Now back to the [BLS19]. The challenge set C is defined by

$$C = \{\mathbf{0}\} \cup \{X^i\}_{i \in [0, \ell-1]} = \{\mathbf{0}\} \cup \{\pm X^i\}_{i \in [0, n-1]}, \quad (4.7)$$

such that $|C| = \ell + 1 = 2n + 1$. The challenge-difference set is $\bar{C} := (C - C) \setminus \{\mathbf{0}\}$, and every element in $\bar{C}$ is invertible: Since R_q is fully splitting, then

$$R_q \cong \prod_{j \in \mathbb{Z}_\ell^*} \mathbb{Z}_q[X]/(X - \xi^j),$$

where ξ is the ℓ^{th} primitive root of unity. Each $X^\alpha - X^\beta \in \bar{C}$ has $X^\alpha - X^\beta \equiv \xi^{\alpha j} - \xi^{\beta j} \pmod{(X - \xi^j)}$ for all $j \in \mathbb{Z}_\ell^*$, and $\xi^{\alpha j} - \xi^{\beta j}$ is zeroed out iff $\alpha \equiv \beta \pmod{\ell}$. Hence, all of the CRT slots are nonzero, and they are invertible. Differences involving $\mathbf{0}$ are of form $\pm X^\alpha$, which are units in R_q and are invertible.

Moreover, for every $\bar{\mathbf{f}} = \mathbf{f}_1 - \mathbf{f}_2 \in \bar{C}$ and $\mathbf{a} \in R$, multiplication by $\bar{\mathbf{f}}$ increases the coefficient ℓ_2 -norm by at most a factor of 2:

$$\|\bar{\mathbf{f}}\mathbf{a}\|_2 \leq \|\mathbf{f}_1\mathbf{a}\|_2 + \|\mathbf{f}_2\mathbf{a}\|_2 \leq 2\|\mathbf{a}\|_2. \quad (4.8)$$

One thing to note: the ℓ_2 -norm in [BLS19] is the canonical embedding ℓ_2 -norm, while the ℓ_2 -norm is the coefficient ℓ_2 -norm. The canonical embedding ℓ_2 -norm is the ℓ_2 -norm for the eq. (4.1), which was shown in [LS18]. In our R power-of-two cyclotomic ring, we have

$$\begin{aligned}\|\mathbf{a}\|_e^2 &= \sum_{j \in \mathbb{Z}_\ell^*} |\mathbf{a}(\omega^j)|^2 = \sum_{j \in \mathbb{Z}_\ell^*} \mathbf{a}(\omega^j) \overline{\mathbf{a}(\omega^j)} = \sum_{j \in \mathbb{Z}_\ell^*} \mathbf{a}(\omega^j) \mathbf{a}(\omega^{-j}) \\ &= n \|\mathbf{a}\|_2^2 + \sum_{i \neq k \in [0, n-1]} a_i a_k \sum_{j \in \mathbb{Z}_\ell^*} \omega^{j(i-k)} = n \|\mathbf{a}\|_2^2,\end{aligned}$$

where ω here is the primitive ℓ^{th} complex root of unity, namely $\omega = \exp(2\pi i/\ell)$, as was written in eq. (4.2).

In the [BLS19] specialization for [BDL⁺18], RLWE and RSIS are needed for hiding and binding.

Definition 4.11.3 (RSIS $_{k,B}$). An efficient algorithm $\mathcal{A}$ has advantage ε in solving RSIS $_{k,B}$ if

$$\Pr \left[\|\bar{\mathbf{s}}\|_2 \leq B \wedge [\mathbf{1}, \bar{\mathbf{a}}^\top] \bar{\mathbf{s}} = \mathbf{0} \wedge \bar{\mathbf{s}} \neq \bar{\mathbf{0}} \mid \bar{\mathbf{a}} \xleftarrow{R} R_q^k, \bar{\mathbf{s}} \leftarrow \mathcal{A}(\bar{\mathbf{a}}) \right] \geq \varepsilon,$$

where $\bar{\mathbf{s}} \in R^{k+1}$ and the zero-equation is evaluated in R_q , and $\|\bar{\mathbf{s}}\|_2$ is the coefficient ℓ_2 -norm.

Fix a distribution χ over $\mathbb{Z}$. We write χ^n for the distribution over R , where a sample is obtained by sampling n coefficients independently from χ , and χ^{kn} denotes the corresponding distribution over R^k . In the instantiation of [BLS19], the distribution χ is over $\{-1, 0, 1\}$, outputting ± 1 with probability $5/16$ each, and 0 with probability $3/8$.

Definition 4.11.4 (RLWE $_m$). Let $m \geq 1$. An efficient algorithm $\mathcal{A}$ has advantage ε in solving RLWE $_m$ if

$$\left| \Pr \left[b = 1 \mid \begin{array}{l} \bar{\mathbf{a}} \xleftarrow{R} R_q^m, \\ \mathbf{s} \leftarrow \chi^n, \\ \bar{\mathbf{e}} \leftarrow \chi^{mn}, \\ b \leftarrow \mathcal{A}(\bar{\mathbf{a}}, \bar{\mathbf{a}}\mathbf{s} + \bar{\mathbf{e}}) \end{array} \right] - \Pr \left[b = 1 \mid \begin{array}{l} \bar{\mathbf{a}} \xleftarrow{R} R_q^m, \\ \bar{\mathbf{u}} \xleftarrow{R} R_q^m, \\ b \leftarrow \mathcal{A}(\bar{\mathbf{a}}, \bar{\mathbf{u}}) \end{array} \right] \right| \geq \varepsilon.$$

The RLWE $_m$ problem is hard if every efficient algorithm has only negligible advantage.

The BDLOP commitment is specialized to commit to $\bar{\mathbf{m}} = (\mathbf{m}_2, \mathbf{m}_3, \mathbf{m}_4, \mathbf{m}_5)^\top \in R_q^4$, and the public parameter is a matrix $\mathbf{B} \in R_q^{5 \times 6}$ of the form

$$\mathbf{B} = \begin{bmatrix} \mathbf{1} & \mathbf{b}_{1,2} & \mathbf{b}_{1,3} & \mathbf{b}_{1,4} & \mathbf{b}_{1,5} & \mathbf{b}_{1,6} \\ \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{b}_{2,6} \\ \mathbf{0} & \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{0} & \mathbf{b}_{3,6} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{1} & \mathbf{0} & \mathbf{b}_{4,6} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{1} & \mathbf{b}_{5,6} \end{bmatrix}, \quad \mathbf{b}_{i,j} \xleftarrow{R} R_q.$$

This is the same identity-block idea as vanilla BDLOP, but with one top coordinate and four message coordinates. To commit to $\bar{\mathbf{m}}$, sample short randomness $\bar{\mathbf{r}} \leftarrow \chi^{6n}$ from the error distribution and output

$$\bar{\mathbf{t}} = \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}} \end{bmatrix} \in R_q^5.$$

For $\sigma > 0$, let $\rho_\sigma(\bar{\mathbf{u}}) := \exp(-\|\bar{\mathbf{u}}\|_2^2/(2\sigma^2))$, and $\rho_\sigma(S) := \sum_{\bar{\mathbf{x}} \in S} \rho_\sigma(\bar{\mathbf{x}})$. The discrete Gaussian centered at $\bar{\mathbf{v}} \in R^k$ is

$$D_{\bar{\mathbf{v}}, \sigma}^{kn}(\bar{\mathbf{z}}) = \rho_\sigma(\bar{\mathbf{z}} - \bar{\mathbf{v}}) \rho_\sigma^{-1}(R^k). \quad (4.9)$$

When $\bar{\mathbf{v}} = \bar{\mathbf{0}}$, we write D_σ^{kn} .

The relaxed opening relation from extractor is

$$\mathbf{f} \cdot \bar{\mathbf{t}} = \mathbf{B}\bar{\mathbf{r}} + \mathbf{f} \cdot \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}} \end{bmatrix},$$

where $\mathbf{f}$ is either from the challenge-difference set $\bar{C}$, or $\mathbf{f} = \mathbf{1}$ for an honest opening. $\bar{\mathbf{r}}$ needs to be short, written as $\|\bar{\mathbf{r}}\|_2 \leq 2B$, with $B = \sigma\sqrt{12n}$.

In [BLS19] the hiding is reduced to the RLWE₅ hardness, and the binding is reduced to the RSIS_{5,8B} hardness, where $8B$ derives from the ℓ_2 -norm of $\mathbf{f}_1\bar{\mathbf{r}}_0 - \mathbf{f}_0\bar{\mathbf{r}}_1$ in eq. (4.10), recall eq. (4.8) for factor 2 and $\|\bar{\mathbf{r}}\|_2 \leq 2B$.

Short Invertible Elements Alone Won't Cut It. A short invertible element itself cannot give knowledge with relaxed norm bounds, not even the exact knowledge. Its role is more of, so long as the elements in the challenge-difference $\bar{C}$ are invertible, we can reduce binding to the hardness assumptions like SIS.

After the first definition of binding by definition 4.8.1, we define the property again as follows.

Definition 4.11.5 (Binding). A commitment scheme is ε -binding if for all efficient $\mathcal{A}$,

$$\Pr \left[\mathbf{m}_0 \neq \mathbf{m}_1 \wedge \text{Open}_{\text{pp}}(\mathbf{m}_0, \mathbf{c}, \mathbf{r}_0) = 1 \wedge \text{Open}_{\text{pp}}(\mathbf{m}_1, \mathbf{c}, \mathbf{r}_1) = 1 \mid \begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda), \\ (\mathbf{m}_0, \mathbf{r}_0, \mathbf{m}_1, \mathbf{r}_1, \mathbf{c}) \leftarrow \mathcal{A}(1^\lambda, \text{pp}) \end{array} \right] < \varepsilon.$$

For [BDL⁺18], the $(\mathbf{f}, \bar{\mathbf{m}}, \bar{\mathbf{r}})$ is a relaxed opening proof. Supposing we have a pair of relaxed openings satisfying

$$\begin{aligned} \mathbf{f}_0 \cdot \bar{\mathbf{t}} &= \mathbf{B}\bar{\mathbf{r}}_0 + \mathbf{f}_0 \cdot \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}}_0 \end{bmatrix}, \\ \mathbf{f}_1 \cdot \bar{\mathbf{t}} &= \mathbf{B}\bar{\mathbf{r}}_1 + \mathbf{f}_1 \cdot \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}}_1 \end{bmatrix} \end{aligned}$$

with $\mathbf{m}_0 \neq \mathbf{m}_1$, then cancelling out $\bar{\mathbf{t}}$ term, we have

$$\mathbf{B}(\mathbf{f}_1 \cdot \bar{\mathbf{r}}_0 - \mathbf{f}_0 \cdot \bar{\mathbf{r}}_1) + \mathbf{f}_0 \cdot \mathbf{f}_1 \cdot \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_1 \end{bmatrix} = \bar{\mathbf{0}}, \quad (4.10)$$

separating by the vertical stack

$$\begin{aligned} \mathbf{A}_1(\mathbf{f}_1 \cdot \bar{\mathbf{r}}_0 - \mathbf{f}_0 \cdot \bar{\mathbf{r}}_1) &= \mathbf{0}, \\ \mathbf{A}_2(\mathbf{f}_1 \cdot \bar{\mathbf{r}}_0 - \mathbf{f}_0 \cdot \bar{\mathbf{r}}_1) &= \mathbf{f}_0 \cdot \mathbf{f}_1 \cdot (\bar{\mathbf{m}}_1 - \bar{\mathbf{m}}_0). \end{aligned}$$

It is that $\mathbf{f}_0 \cdot \mathbf{f}_1$ being invertible, ensuring $\mathbf{f}_0 \cdot \mathbf{f}_1 \cdot (\bar{\mathbf{m}}_0 - \bar{\mathbf{m}}_1) \neq \bar{\mathbf{0}}$, or otherwise the product with a non-invertible ring element a may be 0. Given that $\mathbf{f}_1 \cdot \bar{\mathbf{r}}_0 - \mathbf{f}_0 \cdot \bar{\mathbf{r}}_1 \neq \bar{\mathbf{0}}$, reducing to the SIS hardness of $\mathbf{A}_1(\mathbf{f}_1 \cdot \bar{\mathbf{r}}_0 - \mathbf{f}_0 \cdot \bar{\mathbf{r}}_1)$.

For [Ajt96], recall definition 4.8.4, write $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$ as the commitment key, $\mathbf{s} = [\mathbf{r}^\top \mid \mathbf{m}^\top]^\top$ as the message and randomness, and $\mathbf{c} \leftarrow \mathbf{A}\mathbf{s} = \mathbf{A}_1\mathbf{r} + \mathbf{A}_2\mathbf{m}$ as the Ajtai commitment. Deriving from the same Sigma-Protocol in [Lyu12], a relaxed opening proof $(\mathbf{f}, \bar{\mathbf{m}}, \bar{\mathbf{r}})$ from extraction must satisfy

$$\mathbf{f} \cdot \bar{\mathbf{c}} = \mathbf{A}_1\bar{\mathbf{r}} + \mathbf{A}_2\bar{\mathbf{m}},$$

where the honest opening has $\mathbf{f} = \mathbf{1}$.

If the same commitment is opened with $\mathbf{s}' = [\mathbf{r}'^\top \mid \mathbf{m}'^\top]^\top$, then for $\mathbf{s} = [\bar{\mathbf{r}}^\top \mid \bar{\mathbf{m}}^\top]^\top$ with $\mathbf{s} \cdot \mathbf{f}^{-1} \neq \mathbf{s}'$, we immediately reduce binding to SIS hardness assumption; or otherwise without invertibility the existence of an SIS solution when $\mathbf{f} \cdot \mathbf{s}' - \mathbf{s} = \bar{\mathbf{0}}$ is hard to reason (say an entry of $\mathbf{s}$ is 0 while the same entry in $\mathbf{s}'$ is non-zero but annihilated). Again, the preservation of non-zeros by invertible elements helps with the hardness reduction.

To sum up, the short invertible elements in the relaxed opening ensure the binding reduction, but do not directly give knowledge soundness.

What if $\mathbf{A}$ is Weak. For [BLS19], the relation to prove is $\mathbf{A}\hat{\mathbf{s}} \equiv \bar{\mathbf{u}} \pmod{q}$, where each entry in $\hat{\mathbf{s}}$ is within $\{0, 1, 2\}$. In the prior Commit-and-Open Sigma-Protocol attempt based on previous approximate arguments of knowledge, we lifted the role of the binding commitment from $\mathbf{A}$. My ill-contrived question is now: if $\mathbf{A}$ is not a strong Ajtai instance, can we still have zero knowledge for arbitrary $\mathbf{A}$, unlike in previous works such as [BBC⁺18, BLNS20, BDL⁺18]?

I started by considering an extreme case where $\mathbf{A}$ is an identity matrix. Then $\bar{\mathbf{u}} = \hat{\mathbf{s}}$, and the random masking $\bar{\mathbf{r}}$ is sent in the clear. An honest-but-curious verifier can remove $\bar{\mathbf{r}}$ from $\bar{\mathbf{z}}$, which is the third move of the Sigma-Protocol proposed by modifying the original opening protocol in [BDL⁺18] (ignoring the BDLOP commitment for now), and derive the witness $\hat{\mathbf{s}}$.

On the other hand, such an accepting transcript is totally simulatable! Still ignoring the BDLOP part,

- an accepting transcript contains the first message $\mathbf{A}\bar{\mathbf{r}}$ of the masking through the $\mathbf{A}$ linear system, the second challenge c , which is uniformly random over $\mathbb{Z}_q$, and the third response $c \cdot \hat{\mathbf{s}} + \bar{\mathbf{r}}$;
- a simulated transcript first samples the third response $\bar{\mathbf{z}}$ uniformly at random and then samples the challenge $c' \xleftarrow{\mathbb{R}} \mathbb{Z}_q$; the first masking through $\mathbf{A}$ can be simulated by $\mathbf{A}\bar{\mathbf{z}} - c' \cdot \bar{\mathbf{u}}$.

Both the second and third messages are uniformly distributed. The first masking $\mathbf{A}\bar{\mathbf{r}}$ in a real accepting transcript is uniform over $\text{im}(\mathbf{A})$ (uniformly random over the whole codomain if $\mathbf{A}$ has full row rank), which is determined by the distribution of $\bar{\mathbf{r}}$, while in the simulated transcript, the distribution of $\mathbf{A}\bar{\mathbf{z}} - c' \cdot \bar{\mathbf{u}}$ is determined by c' and $\bar{\mathbf{z}}$, which are both uniformly random, and $\bar{\mathbf{u}} \in \text{im}(\mathbf{A})$ as HVZK is defined for a valid statement. Then $\mathbf{A}\bar{\mathbf{z}} - c' \cdot \bar{\mathbf{u}} = \mathbf{A}(\bar{\mathbf{z}} - c' \cdot \bar{\mathbf{s}}')$, where $\mathbf{A}\bar{\mathbf{s}}' \equiv \bar{\mathbf{u}} \pmod{q}$, so the first masking is also simulatable as part of the joint transcript distribution. Generalizing from $\mathbf{A} = \mathbf{I}_n$, simulatability holds regardless of $\mathbf{A}$.

So what exactly does HVZK really say? Literally, HVZK says that an honestly generated accepting transcript, where the prover holds the witness and the statement is public, is computationally or statistically close to a simulated transcript generated using only the statement. Now let T_{real} be the honestly generated accepting transcript, T_{sim} be the simulated transcript, $\mathcal{A}$ be an efficient algorithm attempting to recover the witness from the transcript, and

$$P_{\text{real}} = \Pr[R(x, \mathcal{A}(x, T_{\text{real}})) = 1], \quad P_{\text{sim}} = \Pr[R(x, \mathcal{A}(x, T_{\text{sim}})) = 1].$$

$R(x, \mathcal{A}(x, T))$ can be computed efficiently because membership in the NP relation R can be checked efficiently and $\mathcal{A}$ is efficient. If HVZK is computational, then let $R(x, \mathcal{A}(x, \cdot))$ be an efficient distinguisher, hence

$$|P_{\text{real}} - P_{\text{sim}}| \leq \text{negl}(\lambda). \quad (4.11)$$

If HVZK is statistical, then eq. (4.11) follows from the statistical distance being negligible. Let $\mathcal{B}$ be an algorithm that samples T_{sim} from the simulator and runs $\mathcal{A}(x, T_{\text{sim}})$. Since the simulator is efficient, $\mathcal{B}$ is efficient. Then by eq. (4.11),

$$\Pr[R(x, \mathcal{B}(x)) = 1] - \text{negl}(\lambda) \leq P_{\text{real}} \leq \Pr[R(x, \mathcal{B}(x)) = 1] + \text{negl}(\lambda),$$

or equivalently

$$P_{\text{real}} - \text{negl}(\lambda) \leq \Pr[R(x, \mathcal{B}(x)) = 1] \leq P_{\text{real}} + \text{negl}(\lambda).$$

In this way, HVZK can be interpreted as, it does not give one more advantage in finding the witness for the statement, with or without an honestly generated accepting transcript. But ZK is irrelevant to $\mathbf{A}$ being a hard instance or not. If the problem is easy, then the problem is in P , and there is no witness to hide.

Thus, any efficient witness-recovery advantage obtained from the honest verifier's view could already be obtained from the public statement alone. If the relation is hard, no such non-negligible advantage can exist, and the *witness hiding* against an honest verifier follows from HVZK and the hardness of the relation²¹. See also [Dam10].

In our particular setting, if $(\mathbf{A}, \bar{\mathbf{u}})$ is sampled from a hard ISIS instance distribution, then finding $\bar{\mathbf{r}}$ from $\mathbf{A}\bar{\mathbf{r}}$ in the transcript for short $\hat{\mathbf{s}}'$ is at least as hard as solving the ISIS instance; namely, efficient recovery of the masking value $\bar{\mathbf{r}}$ would give an efficient solver for the ISIS relation.

Henceforth, we let the statement $(\mathbf{A}, \bar{\mathbf{u}})$ be sampled from a hard ISIS instance distribution for witness privacy.

²¹This is an honest-verifier analogue of the witness-hiding interpretation introduced by [FS90]. Their formal definition is distributional and applies to potentially malicious verifiers, whereas the present conclusion follows only for the honest verifier from HVZK.

After the First Sigma-Protocol. Recall the construction of the first Sigma-Protocol attempt without BDLOP:

- The prover sends out $\bar{w} \leftarrow A\bar{r}$, where $\bar{r}$ is uniformly random.
- The verifier challenges the prover by sending $c \xleftarrow{R} \mathbb{Z}_q$.
- The prover replies with $c \cdot \hat{s} + \bar{r}$.

It is definitely flawed, a malicious prover who does not know s can send out the first masking $\bar{w}'$ at random, and on verifier replies with $c \xleftarrow{R} \mathbb{Z}_q$, find an arbitrary image $\bar{z}'$ such that $A\bar{z}' = \bar{w}' + c \cdot \bar{u}$, where $A\hat{s} \equiv \bar{u} \pmod{q}$ in statement.

Starting from the approximate argument similar to [BDL⁺18, Lyu12]. For HVZK, namely the accepting transcript is simulatable, the first message for the masking randomness should be simulatable in the joint distribution of the transcript, or otherwise supposing the accepting transcript whose first message being raw masking randomness is efficiently simulatable, then such efficient simulator can be used to solve the statement, which is a hard ISIS instance, and it does not make sense. By contradiction, the first message to be sent out should be “hiding”.

On the other hand, one could have just sent over the raw masking randomness to the RO, and let the RO decide which challenge to send back. But due to the goal of reaching ZK, one has to use a (cryptographic) object in between the masking randomness and the RO. Intuitively, one should pin the masking randomness against the RO, such that the challenge is decided and fixed after the RO absorbs the first message, and the represented masking randomness.

Since the RO itself cannot guarantee what value the first message represents, we need the “binding” property, which says it is computationally/statistically hard to find different preimages that lead to the same message. With such property, an efficient knowledge extractor can extract (relaxed) knowledge of the witness from a pair of forked transcripts from the first message, or otherwise a malicious prover can provide “branch-dependent” decompositions, say $z_0 = y_0 + c_0 \cdot s_0$ and $z_1 = y_1 + c_1 \cdot s_1$, so rewinding does not give valid witness. Binding removes the possibility to equivocate, and makes it possible for extraction from rewinding.

With the property of binding and hiding, we introduce the notion of commitment.

In the vanilla approximate arguments such as [BDL⁺18, Lyu12], stopping at the third move is sufficient. It has to be either breaking the SIS hardness assumption, with probability at most $\text{negl}(\lambda)$, or the malicious prover prepares a forged commitment $\bar{t}$ to the masking randomness, and a forged response short $\bar{z}$, such that the challenge c happens to satisfy $A\bar{z} = \bar{t} + c \cdot \bar{u}$, with probability at most $|C|^{-1}$. That sums up to be the knowledge soundness, and it suffices to stop at the third move as the binding property allows for rewinding for knowledge extraction.

For the current protocol [BLS19], the same argument no longer applies. The masking randomness is not pinned to the RO, even if A stands for an Ajtai commitment, it will not bind for a message with unbounded length. Moreover, on the third response, in the previous protocols [BDL⁺18, Lyu12], it suffices to use binding of short messages to stop there, while in [BLS19], the norm is arbitrary, so we will resort to proving the linear relation $\bar{r} + c \cdot \hat{s} = \hat{z}$, let alone proving the shortness of $\hat{s}$.

Quickly recap what we have here, so after the first Sigma-Protocol, we transformed the problem into proving the linear relation $\hat{z} = c \cdot \hat{s} + \bar{r}$, where $\hat{z}$ is in the third move prover response, while $\bar{r}$ is the masking randomness, and $\hat{s}$ is the witness. Finally, we also need to prove the shortness of $\hat{s}$ in the statement.

Proving Linear Relation over Elements in [BDL⁺18]. Deriving from the protocol for linear relation in [BDL⁺18], we consider an adaptation here in the RSIS and RLWE settings to prove the linear relation of the committed elements.

Let $B \in R_q^{k \times (k+1)}$ be a BDLOP commitment public parameter, and $\bar{m} \in R_q^{k-1}$ be the message, then

$$\bar{c} = B\bar{r} + \begin{bmatrix} 0 \\ \bar{m} \end{bmatrix} \in R_q^k$$

is the commitment to $\bar{m}$, where $\bar{r} \leftarrow \chi^{(k+1)n}$.

Let $f(x_1, \dots, x_{k-1})$ be a linear function of $\{x_i\}_{i \in [1, k-1]}$ whose total degree is at most 1, $\bar{b}_i$ be the i^{th} row of B , and

$\bar{\mathbf{b}}_f = f(\bar{\mathbf{b}}_2, \dots, \bar{\mathbf{b}}_k)$. Then Π_{lin} is defined for statement

$$f(\mathbf{m}_1, \dots, \mathbf{m}_{k-1}) = \mathbf{w}.$$

- The prover samples $\bar{\mathbf{r}}' \leftarrow D_\sigma^{(k+1)n}$, and sends over

$$\bar{\mathbf{w}} = \begin{bmatrix} \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_f^\top \end{bmatrix} \bar{\mathbf{r}}'.$$

- The verifier samples $\mathbf{f} \xleftarrow{\mathcal{R}} C$, where C is the short challenge set, and sends over $\mathbf{f}$.
- The prover runs reject sampling, and outputs $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{r}}' + \mathbf{f} \cdot \bar{\mathbf{r}}$.
- The verifier checks $\bar{\mathbf{z}}$ being short, and the following equality holds

$$\bar{\mathbf{w}} + \mathbf{f} \cdot \begin{bmatrix} \mathbf{c}_1 \\ f(\mathbf{c}_2, \dots, \mathbf{c}_k) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_f^\top \end{bmatrix} \bar{\mathbf{z}} + \mathbf{f} \cdot \begin{bmatrix} \mathbf{0} \\ \mathbf{w} \end{bmatrix}.$$

Moreover, the linear relation scheme Π_{lin} can be extended to multiple linear functions.

Piece Things Together Without Shortness Check. We now piece Π_{lin} and the first Sigma-Protocol attempt together for the construction of the argument system just to prove the reduced statement $\hat{\mathbf{z}} = \bar{\mathbf{r}} + c \cdot \hat{\mathbf{s}}$ without the norm check.

Here we let $\hat{\mathbf{y}} = \text{NTT}(\mathbf{y})$ be the masking randomness, which used to be represented by $\bar{\mathbf{r}}$. We use $\mathbf{B} \in R_q^{3 \times 4}$ for the BDLOP commitment public parameter, such that we commit to $\mathbf{s}$ and $\mathbf{y}$.

- The prover sends out $\bar{\mathbf{w}} \leftarrow \mathbf{A}\hat{\mathbf{y}}$, where $\hat{\mathbf{y}}$ is uniformly random. Moreover, the prover samples $\bar{\mathbf{r}} \leftarrow \chi^{4n}$, commits

$$\bar{\mathbf{c}} \leftarrow \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0} \\ \mathbf{y} \\ \mathbf{s} \end{bmatrix},$$

and sends $\bar{\mathbf{c}}$ over to the verifier.

- The verifier challenges the prover by sending $c \xleftarrow{\mathcal{R}} \mathbb{Z}_q$.
- The prover replies with $\mathbf{z} \leftarrow c \cdot \mathbf{s} + \mathbf{y}$, and starts Π_{lin} by:
 - The prover samples $\bar{\mathbf{y}}' \leftarrow D_\sigma^{4n}$, let $f(\mathbf{y}, \mathbf{s}) = c \cdot \mathbf{s} + \mathbf{y}$, and sends over

$$\bar{\mathbf{w}}' = \begin{bmatrix} \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_f^\top \end{bmatrix} \bar{\mathbf{y}}'.$$

- The verifier samples $\mathbf{f} \xleftarrow{\mathcal{R}} C$, where C is the short challenge set, and sends over $\mathbf{f}$.
 - The prover runs reject sampling, and outputs $\bar{\mathbf{z}}' \leftarrow \bar{\mathbf{y}}' + \mathbf{f} \cdot \bar{\mathbf{r}}$.
- The verifier checks that $\|\bar{\mathbf{z}}'\|_2 \leq B$ and that the following equality holds:

$$\bar{\mathbf{w}}' + \mathbf{f} \cdot \begin{bmatrix} \mathbf{c}_1 \\ f(\mathbf{c}_2, \mathbf{c}_3) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_f^\top \end{bmatrix} \bar{\mathbf{z}}' + \mathbf{f} \cdot \begin{bmatrix} \mathbf{0} \\ \mathbf{z} \end{bmatrix},$$

and $\bar{\mathbf{w}} + c \cdot \bar{\mathbf{u}} = \mathbf{A}\hat{\mathbf{z}}$ holds.

Shortness Check Over the Sigma-Protocol Response. To adapt the previous NTT shortness check in eq. (4.4) for $\hat{\mathbf{z}}$, we have

$$(\mathbf{z} - 2c)(\mathbf{z} - c)\mathbf{z} = \mathbf{y}^3 + 3(\mathbf{s} - 1)\mathbf{y}^2 \cdot c + (3\mathbf{s}^2 - 6\mathbf{s} + 2)\mathbf{y} \cdot c^2 + (\mathbf{s} - 1)(\mathbf{s} - 2)\mathbf{s} \cdot c^3.$$

To prove the degree-3 evaluation over c , the trivial way is to use Π_{lin} to prove the commitments to $\mathbf{y}^3$, $3(\mathbf{s} - 1)\mathbf{y}^2$, and $(3\mathbf{s}^2 - 6\mathbf{s} + 2)\mathbf{y}$ satisfy the linear relation over coefficients $1, c, c^2$, then the public parameter should be $R_q^{6 \times 7}$.

One optimization in [BLS19] is to reuse the commitment to $\mathbf{s}$, such that the linear relation to be done is over $1, c$, while c^2 should be multiplying 0, and

$$\begin{aligned} (\mathbf{z} - 2c)(\mathbf{z} - c)\mathbf{s} &= \mathbf{s}\mathbf{y}^2 + \mathbf{y}\mathbf{s}(2\mathbf{s} - 3) \cdot c + (\mathbf{s} - 2)(\mathbf{s} - 1)\mathbf{s} \cdot c^2 \\ &= \mathbf{y}(2\mathbf{s} - 3) \cdot \mathbf{z} - \mathbf{y}^2(\mathbf{s} - 3) + (\mathbf{s} - 2)(\mathbf{s} - 1)\mathbf{s} \cdot c^2. \end{aligned}$$

Either way is fine, but this allows for committing to only 4 elements, and the public parameter is decreased to $R_q^{5 \times 6}$.

Final Construction. We now present the final construction, and we use $\mathbf{B} \in R_q^{5 \times 6}$ for the BDLOP public parameter, such that we commit to $\mathbf{y}$, $\mathbf{s}$, $\mathbf{y}(2\mathbf{s} - 3)$, and $\mathbf{y}^2(\mathbf{s} - 3)$.

- The prover sends out $\bar{\mathbf{w}} \leftarrow \mathbf{A}\hat{\mathbf{y}}$, where $\hat{\mathbf{y}}$ is uniformly at random. Moreover, the prover samples $\bar{\mathbf{r}} \leftarrow \chi^{6n}$, commits

$$\bar{\mathbf{c}} \leftarrow \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} 0 \\ \mathbf{y} \\ \mathbf{s} \\ \mathbf{y}(2\mathbf{s} - 3) \\ \mathbf{y}^2(\mathbf{s} - 3) \end{bmatrix},$$

and sends $\bar{\mathbf{c}}$ over to the verifier.

- The verifier challenges the prover by sending $c \xleftarrow{R} \mathbb{Z}_q$.
- The prover replies with $\mathbf{z} \leftarrow c \cdot \mathbf{s} + \mathbf{y}$, and starts Π_{lin} by:
 - The prover samples $\bar{\mathbf{y}}' \leftarrow D_{\sigma}^{6n}$, let

$$\begin{aligned} f_0(\mathbf{y}, \mathbf{s}, \mathbf{m}_4, \mathbf{m}_5) &= c \cdot \mathbf{s} + \mathbf{y}, \\ f_1(\mathbf{y}, \mathbf{s}, \mathbf{m}_4, \mathbf{m}_5) &= (\mathbf{z} - 2c)(\mathbf{z} - c) \cdot \mathbf{s} - \mathbf{m}_4 \cdot \mathbf{z} + \mathbf{m}_5, \end{aligned}$$

and sends over

$$\bar{\mathbf{w}}' = \begin{bmatrix} \bar{\mathbf{b}}_1^T \\ \bar{\mathbf{b}}^T \\ \bar{\mathbf{b}}_{f_0}^T \\ \bar{\mathbf{b}}_{f_1}^T \end{bmatrix} \bar{\mathbf{y}}'.$$

- The verifier samples $\mathbf{f} \xleftarrow{R} C$, where C is the short challenge set, and sends over $\mathbf{f}$.
- The prover runs reject sampling, and outputs $\bar{\mathbf{z}}' \leftarrow \bar{\mathbf{y}}' + \mathbf{f} \cdot \bar{\mathbf{r}}$.
- The verifier checks $\bar{\mathbf{z}}'$ being short, the following equality holds

$$\bar{\mathbf{w}}' + \mathbf{f} \cdot \begin{bmatrix} \mathbf{c}_1 \\ f_0(\mathbf{c}_2, \dots, \mathbf{c}_5) \\ f_1(\mathbf{c}_2, \dots, \mathbf{c}_5) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{b}}_1^T \\ \bar{\mathbf{b}}^T \\ \bar{\mathbf{b}}_{f_0}^T \\ \bar{\mathbf{b}}_{f_1}^T \end{bmatrix} \bar{\mathbf{z}}' + \mathbf{f} \cdot \begin{bmatrix} \mathbf{0} \\ \mathbf{z} \\ \mathbf{0} \end{bmatrix},$$

and $\bar{\mathbf{w}} + c \cdot \bar{\mathbf{u}} = \mathbf{A}\hat{\mathbf{z}}$ holds.

Knowledge Extraction. We let ε be the convincing probability of the prover, so the prover generates an accepting transcript in expected $1/\varepsilon$ steps. The knowledge extraction process is similar to that in [BLNS20], where one needs distinct challenges $c_1, c_2, c_3 \in \mathbb{Z}_q$ because there is a degree-2 polynomial over c in the shortness check, and the last step is an adaptation of the BDLOP proof of a linear relation, for each c_i , by [BDL⁺18], we need 2 distinct $\mathbf{f}_{i,1}, \mathbf{f}_{i,2} \in C$ for a relaxed opening, and therefore in total we need 6 accepting transcripts forked in 3 by 2.

We begin by bounding the expected runtime of the extractor, where the “heavy-row” argument [Dam10] is used in a similar way as [BBC⁺18, BLNS20]. If the probability of convincing is ε , consider all transcripts indexed by the first challenge $c \in \mathbb{Z}_q$ and the second challenge $\mathbf{f} \in C$, in a 0-1 matrix.

Start with the first pair of accepting transcripts decided by c_1 . A “heavy row” indexed by challenge c is a row with entries being 1 of fraction at least $\varepsilon/2$, then by the “heavy-row” argument in [Dam10], the first accepting transcript is on a heavy row with probability at least $1/2$. Conditioned that the first accepting transcript is on a heavy row, the second accepting transcript is obtained with probability at least $\varepsilon/2 - |C|^{-1}$, as $\mathbf{f}_{1,1} \neq \mathbf{f}_{1,2}$. Hence, the expected runtime for the first pair is

$$\frac{1}{\varepsilon} + \frac{2}{\varepsilon - 2|C|^{-1}}.$$

The second pair of accepting transcripts is similar: The first of them is obtained with probability at least $\varepsilon - 1/q$ as $c_2 \neq c_1$, then let the heavy row with entries being 1 of fraction at least $\varepsilon/2 - 1/(2q)$, then the first of them lands on a heavy row with probability at least $1/2$, and conditioned that, the second of them is obtained with probability at least $\varepsilon/2 - 1/(2q) - |C|^{-1}$. Hence, the expected runtime for the second pair is

$$\frac{1}{\varepsilon - q^{-1}} + \frac{2}{\varepsilon - q^{-1} - 2|C|^{-1}}.$$

Hence, conditioned on all three pairs being on heavy rows, occurring with probability at least $1/8$, the expected runtime is at most

$$T = \frac{9}{\varepsilon - 2q^{-1} - 2|C|^{-1}},$$

and we cap the runtime of the extractor at $2T$. By Markov’s inequality, finding 3 pairs of accepting transcripts in $2T$ steps succeeds with probability at least $1/16$ (conditioned that all are on heavy rows, the probability is at least $1/2$ by Markov’s inequality, then $1/16$ is immediate). Restarting after failure thus gives expected runtime at most $32T$ ²².

For $i \in [1, 3]$, let $\mathbf{f}'_i = \mathbf{f}_{i,1} - \mathbf{f}_{i,2}$, and $\bar{\mathbf{z}}''_i = \bar{\mathbf{z}}'_{i,1} - \bar{\mathbf{z}}'_{i,2}$ be the relaxed-opening randomness obtained by taking the difference of the two masked responses. The difference $\mathbf{f}'_i$ is invertible over R_q by $\bar{C}$, as $\mathbf{f}_{i,1}, \mathbf{f}_{i,2} \in C$ are distinct. The two transcripts have the same $c_i, \mathbf{z}_i$, and three masking terms. Subtracting their three verification equations cancels the common masking terms and gives a branch-specific pair of linear functions. Namely, let

$$\begin{aligned} g_i &:= (\mathbf{z}_i - c_i)(\mathbf{z}_i - 2c_i), \\ f_{0,i}(\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5) &:= \mathbf{x}_2 + c_i \mathbf{x}_3, \\ f_{1,i}(\mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5) &:= g_i \mathbf{x}_3 - \mathbf{z}_i \mathbf{x}_4 + \mathbf{x}_5, \end{aligned}$$

with the corresponding rows

$$\begin{aligned} \bar{\mathbf{b}}_{f_{0,i}}^T &:= \bar{\mathbf{b}}_2^T + c_i \bar{\mathbf{b}}_3^T, \\ \bar{\mathbf{b}}_{f_{1,i}}^T &:= g_i \bar{\mathbf{b}}_3^T - \mathbf{z}_i \bar{\mathbf{b}}_4^T + \bar{\mathbf{b}}_5^T. \end{aligned}$$

The subtraction is then written in the same matrix form as the masked BDLOP opening:

$$\mathbf{f}'_i \begin{bmatrix} \mathbf{c}_1 \\ f_{0,i}(\mathbf{c}_2, \dots, \mathbf{c}_5) \\ f_{1,i}(\mathbf{c}_2, \dots, \mathbf{c}_5) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{b}}_1^T \\ \bar{\mathbf{b}}_{f_{0,i}}^T \\ \bar{\mathbf{b}}_{f_{1,i}}^T \end{bmatrix} \bar{\mathbf{z}}''_i + \mathbf{f}'_i \begin{bmatrix} \mathbf{0} \\ \mathbf{z}_i \\ \mathbf{0} \end{bmatrix}.$$

²²[BLS19] states $16T$, apparently omitting the factor 2 from the timeout.

Define the extracted messages by

$$\begin{aligned} \mathbf{y}_i^* &= \mathbf{c}_2 - (\mathbf{f}'_i)^{-1} \bar{\mathbf{b}}_2^\top \bar{\mathbf{z}}_i'', & \mathbf{s}_i^* &= \mathbf{c}_3 - (\mathbf{f}'_i)^{-1} \bar{\mathbf{b}}_3^\top \bar{\mathbf{z}}_i'', \\ \mathbf{m}_{4,i}^* &= \mathbf{c}_4 - (\mathbf{f}'_i)^{-1} \bar{\mathbf{b}}_4^\top \bar{\mathbf{z}}_i'', & \mathbf{m}_{5,i}^* &= \mathbf{c}_5 - (\mathbf{f}'_i)^{-1} \bar{\mathbf{b}}_5^\top \bar{\mathbf{z}}_i''. \end{aligned}$$

The first row of this equation and these definitions give the relaxed opening

$$\mathbf{f}'_i \bar{\mathbf{c}} = \mathbf{B} \bar{\mathbf{z}}_i'' + \mathbf{f}'_i \begin{bmatrix} \mathbf{0} \\ \mathbf{y}_i^* \\ \mathbf{s}_i^* \\ \mathbf{m}_{4,i}^* \\ \mathbf{m}_{5,i}^* \end{bmatrix}, \quad \|\bar{\mathbf{z}}_i''\|_2 \leq 2B.$$

If two outer branches yield different messages, the two relaxed openings break binding and give an $\text{RSIS}_{5,8B}$ solution. Otherwise, all three branches open to fixed messages $\mathbf{y}^*, \mathbf{s}^*, \mathbf{m}_4^*, \mathbf{m}_5^*$ that are independent of c_i .

The second row of the matrix equation then gives

$$\mathbf{z}_i = f_{0,i}(\mathbf{y}^*, \mathbf{s}^*, \mathbf{m}_4^*, \mathbf{m}_5^*) = \mathbf{y}^* + c_i \mathbf{s}^* \quad \text{for } i \in [1, 3], \quad (4.12)$$

while the third row gives

$$f_{1,i}(\mathbf{y}^*, \mathbf{s}^*, \mathbf{m}_4^*, \mathbf{m}_5^*) = g_i \mathbf{s}^* - \mathbf{z}_i \mathbf{m}_4^* + \mathbf{m}_5^* = 0.$$

Substituting eq. (4.12) into the definition of g_i gives

$$g_i = (\mathbf{y}^* + c_i(\mathbf{s}^* - 1))(\mathbf{y}^* + c_i(\mathbf{s}^* - 2)),$$

substituting this expression for g_i , together with the expression for $\mathbf{z}_i$ in eq. (4.12), into the third-row equation yields a degree-2 polynomial in c_i whose leading coefficient is

$$\mathbf{s}^*(\mathbf{s}^* - 1)(\mathbf{s}^* - 2).$$

The polynomial vanishes at the three distinct points c_1, c_2, c_3 . The corresponding Vandermonde matrix is invertible over $\mathbb{Z}_q$, hence also over R_q , and therefore

$$\mathbf{s}^*(\mathbf{s}^* - 1)(\mathbf{s}^* - 2) = 0.$$

Applying the NTT shows that $\widehat{\mathbf{s}}^* \in \{0, 1, 2\}^n$.

Finally, subtracting the outer verification equations for two distinct challenges gives

$$\mathbf{A}(\widehat{\mathbf{z}}_1 - \widehat{\mathbf{z}}_2) = (c_1 - c_2)\bar{\mathbf{u}}.$$

By eq. (4.12), $\widehat{\mathbf{z}}_1 - \widehat{\mathbf{z}}_2 = (c_1 - c_2)\widehat{\mathbf{s}}^*$. Since $c_1 - c_2$ is invertible in $\mathbb{Z}_q$, we conclude that

$$\mathbf{A}\widehat{\mathbf{s}}^* = \bar{\mathbf{u}}.$$

Therefore, the extractor outputs an exact witness $\widehat{\mathbf{s}}^* \in \{0, 1, 2\}^n$, or it produces an $\text{RSIS}_{5,8B}$ solution.

Knowledge Soundness. Following [BG93, Dam10], the knowledge error κ can be viewed as an upper bound on the acceptance probability that does not yet force extractable knowledge. Informally, it is the best convincing probability that the analysis permits for a prover that does not know a correct witness. If the prover goes beyond this threshold, the extractor can take a witness out of its behavior. This interpretation does not assert that a prover without a witness can actually attain κ : the bound need not be tight, and a particular prover may do strictly worse.

For comparison, in a flat protocol with a uniform challenge set of size N , if any r distinct accepting branches from the same prefix suffice for extraction, then at most $r - 1$ branches may remain non-extractable, giving $\kappa = (r - 1)/N$.

Thus two-branch extraction gives $1/N$, while three-branch extraction gives $2/N$. This is only the flat intuition: the extractor above requires a nested 3-by-2 challenge tree.

For the current protocol, let

$$\kappa := \frac{2}{q} + \frac{2}{|C|}.$$

When obtaining the third pair of transcripts, excluding the two previously used outer challenges loses at most $2/q$ of the acceptance probability. On a heavy row, the acceptance probability for the second challenge is therefore at least one half of $\varepsilon - 2/q$. Excluding the previously used second challenge then leaves probability at least

$$\frac{\varepsilon}{2} - \frac{1}{q} - \frac{1}{|C|} = \frac{1}{2}(\varepsilon - \kappa)$$

for obtaining the required inner fork. Hence $\varepsilon > \kappa$ makes every denominator in the extraction above positive, and the heavy-row argument of [Dam10] finds the nested challenge tree in expected time $\text{poly}(\lambda)/(\varepsilon - \kappa)$.

The extractor then outputs either an exact witness $\widehat{\mathbf{s}}^* \in \{0, 1, 2\}^n$ satisfying $\mathbf{A}\widehat{\mathbf{s}}^* = \overline{\mathbf{u}}$, or an $\text{RSIS}_{5,8B}$ solution. Under the $\text{RSIS}_{5,8B}$ hardness assumption, the latter event is negligible. This realizes the Bellare-Goldreich interpretation: acceptance beyond κ by a non-negligible gap yields extractable knowledge, with the inverse gap governing the expected extraction time.

Rejection Sampling and Honest-Prover Bounds. All norms in this paragraph are coefficient ℓ_2 -norms. Let $\bar{\mathbf{v}} := \mathbf{f}\bar{\mathbf{r}}$, where $\bar{\mathbf{r}} \leftarrow \chi^{6n}$ and $\mathbf{f} \xleftarrow{\mathbf{R}} C$. Since $\mathbf{f} \in C$ defined in eq. (4.7), it does not increase the coefficient norm. Moreover, $\|\bar{\mathbf{r}}\|_2^2$ is a sum of $6n$ independent Bernoulli trials with success probability $5/8$. Hence, for $0 < \delta \leq 1$,

$$\Pr[\|\bar{\mathbf{v}}\|_2 \leq T_{\text{norm}}] \geq 1 - \exp\left(-\frac{5\delta^2 n}{4}\right), \quad T_{\text{norm}} := \sqrt{(1 + \delta)\frac{15n}{4}}$$

by Chernoff's bound. Choose δ so that the failure probability is at most 2^{-101} , and choose $\sigma \geq 5T_{\text{norm}}$.

For $\bar{\mathbf{y}}' \leftarrow D_{\sigma}^{6n}$, the uncorrected response $\bar{\mathbf{z}}' = \bar{\mathbf{y}}' + \bar{\mathbf{v}}$ is distributed according to $D_{\bar{\mathbf{v}},\sigma}^{6n}$, centered at $\bar{\mathbf{v}}$ rather than $\bar{\mathbf{0}}$. The prover accepts this response with probability

$$\min\left\{\frac{1}{12} \cdot \frac{\rho_{\sigma}(\bar{\mathbf{z}}')}{\rho_{\sigma}(\bar{\mathbf{z}}' - \bar{\mathbf{v}})}, 1\right\}.$$

By the rejection-sampling lemma of [Lyu12, BLS19], acceptance occurs with probability at least $1/12 - 2^{-104}$. Conditioned on acceptance, the joint distribution of $(\bar{\mathbf{v}}, \bar{\mathbf{z}}')$ is within statistical distance 2^{-100} of independently sampling $\bar{\mathbf{v}}$ as above and $\bar{\mathbf{z}}' \leftarrow D_{\sigma}^{6n}$. Thus the accepted response is statistically close to a zero-centered Gaussian independent of the shift $\bar{\mathbf{v}} = \mathbf{f}\bar{\mathbf{r}}$.

Finally, by [Ban93], over $\bar{\mathbf{z}}' \leftarrow D_{\sigma}^{6n}$, the Gaussian tail bound gives

$$\Pr\left[\|\bar{\mathbf{z}}'\|_2 \leq \sigma\sqrt{12n}\right] \geq 1 - (2/e)^{3n/2}.$$

We therefore set $B := \sigma\sqrt{12n}$. An accepted response fails the verifier's norm check with probability at most

$$2^{-100} + (2/e)^{3n/2}.$$

Upon rejection, the prover aborts and restarts with fresh randomness, so an honest execution succeeds with probability approximately $1/12$ per attempt.

4.12 Practical Product Proof from Lattice Commitments

Polynomial Ring Elements Don't Split Into Many Factors. One contribution of [ALS20], building on [LS18], is the observation that polynomial ring elements cannot split into many factors for a short invertible challenge set. More specifically, consider the particular setting of a power-of-two ring from theorem 4.10.1 and theorem 4.10.6.

Lemma 4.12.1 (Power-of-two Subgroup Order [LS18]). *Let $4 \leq k \leq n$ be both power of two, and $a \equiv 1 + k \pmod{2k}$, then $\text{ord}_n(a) = n/k$, namely the cyclic subgroup $\langle a \rangle \subset \mathbb{Z}_n^*$ generated by a has order n/k .*

Proof. The $k = n$ case is trivial, as $a \equiv 1 + n \pmod{2n}$. Henceforth assume $k < n$.

Since $a \equiv 1 + k \pmod{2k}$, we write $a = 1 + a_1 k$ where a_1 is odd. As a result,

$$a^2 = 1 + 2ka_1 + k^2 a_1^2 = 1 + (2k)a_1(1 + ka_1/2) = 1 + (2k)a_2,$$

Since $k \geq 4$ is a power of two, $1 + ka_1/2$ is odd, and hence so is a_2 . By $a^2 \equiv 1 + 2k \pmod{4k}$, we have

$$a^{n/k} \equiv 1 + n \pmod{2n},$$

and moreover, $a^{n/(2k)} \equiv 1 + n/2 \pmod{n}$. Therefore, $\langle a \rangle \subset \mathbb{Z}_n^*$ has order n/k . $\square$

Then theorem 4.10.1 can be updated as follows.

Corollary 4.12.2. *Let $m = 2d$ and $z = 2\ell$ both be powers of two, with $4 \leq z \leq m$. If q is prime with $q \equiv 1 + z \pmod{2z}$, then for distinct primitive z^{th} roots of unity r_j ,*

$$\Phi_m(X) = X^d + 1 \equiv \prod_{j \in [1, \ell]} (X^{d/\ell} - r_j) \pmod{q}$$

and each factor $X^{d/\ell} - r_j$ is irreducible over $\mathbb{Z}_q$.

We also update theorem 4.10.6 as follows.

Corollary 4.12.3. *Following corollary 4.12.2, any $\mathbf{y} \in R_{m,q}$ satisfying*

$$0 < \|\mathbf{y}\|_\infty < \ell^{-1/2} \cdot q^{1/\ell} = \frac{1}{s_1(z)} q^{1/\varphi(z)}$$

is invertible in $R_{m,q}$.

Proof. By [ACX21, LPR13], $s_1(z) \leq \sqrt{\tau(z)} = \sqrt{\ell}$, with equality here. The rest follows from $\varphi(z) = \ell$. $\square$

For pairwise difference of distinct challenges being invertible (to reduce binding to the openings to SIS hardness), we need $q^{1/\ell} \cdot \ell^{-1/2} > 2$, and eventually $q > (2\sqrt{\ell})^\ell$. The following table shows that the polynomial ring cannot split into many irreducible factors without making the modulus too large for computational efficiency.

Number of factors ℓ	Required modulus q
4	$q > 2^8$
8	$q > 2^{20}$
16	$q > 2^{48}$
32	$q > 2^{112}$
64	$q > 2^{256}$

Table 4.12.1: Required modulus size for different numbers of irreducible factors.

This gives rise to the question behind [ALS20]: what if the challenge set is not globally invertible, but the probability that a pairwise difference is noninvertible is low, such as $\text{poly}(q^{-1})$? The goal is a good challenge set even when the ring splits into a large number of irreducible factors.

Prime Splittings and Automorphisms. To describe the splitting in corollary 4.12.2 better, we point out that R_q has a group of automorphisms $\text{Aut}(R_q)$ that is isomorphic to $\mathbb{Z}_{2d}^*$. We write $\text{Aut}(R_q)$ for the cyclotomic automorphism subgroup consisting of the maps $f(X) \mapsto f(X^i)$ for $i \in \mathbb{Z}_{2d}^*$. This group is isomorphic to $\mathbb{Z}_{2d}^*$:

$$i \mapsto \sigma_i : \mathbb{Z}_{2d}^* \rightarrow \text{Aut}(R_q),$$

where

$$\begin{aligned} f(X) &= \sum_{j \in [0, d-1]} a_j X^j, \\ \sigma_i(f(X)) &= \sum_{j \in [0, d-1]} a_j \sigma_i(X^j) = \sum_{j \in [0, d-1]} a_j X^{ij} = f(X^i). \end{aligned} \quad (4.13)$$

The group $\text{Aut}(R_q)$ acts transitively on prime ideals $(X^{d/\ell} - \xi)$ in R_q , and every σ_i factors through field isomorphisms

$$\mathbb{Z}_q[X]/(X^{d/\ell} - \xi) \cong R_q/(X^{d/\ell} - \xi) \rightarrow \mathbb{Z}_q[X]/(X^{d/\ell} - \xi^{i^{-1}}) \cong R_q/(\sigma_i(X^{d/\ell} - \xi)).$$

Concretely, by automorphism, we have

$$\sigma_i(\langle X^{d/\ell} - \xi \rangle) = \langle X^{id/\ell} - \xi \rangle = \langle X^{d/\ell} - \xi^{i^{-1}} \rangle$$

in R_q , then for $\mathbf{f} \in R_q$,

$$\sigma_i(\mathbf{f} \bmod (X^{d/\ell} - \xi)) = \sigma_i(\mathbf{f}) \bmod (X^{d/\ell} - \xi^{i^{-1}}).$$

The multiplicative subgroup $\langle 2\ell + 1 \rangle \subseteq \mathbb{Z}_{2d}^*$ has order d/ℓ by corollary 4.12.2, and any $i \in \langle 2\ell + 1 \rangle$ stabilizes the prime ideal $(X^{d/\ell} - \xi)$, namely

$$\sigma_i(X^{d/\ell} - \xi) = X^{d/\ell} - \xi,$$

as ξ is a primitive $2\ell^{\text{th}}$ root of unity, then $i \in \langle 2\ell + 1 \rangle \subseteq \mathbb{Z}_{2d}^*$ has $\xi^{i^{-1}} = \xi$. The quotient group $\mathbb{Z}_{2d}^*/\langle 2\ell + 1 \rangle \cong \mathbb{Z}_{2\ell}^*$ has order ℓ , and since elements in $\langle 2\ell + 1 \rangle$ stabilize the prime ideal, we can index the ℓ prime ideals by

$$(X^d + 1) = \prod_{i \in \mathbb{Z}_{2d}^*/\langle 2\ell + 1 \rangle} (X^{d/\ell} - \xi^i). \quad (4.14)$$

Moreover, for $k \mid \ell$ with $k < \ell$, $\langle 2\ell/k + 1 \rangle \subseteq \mathbb{Z}_{2d}^*$ has order kd/ℓ , and thus $\langle 2\ell/k + 1 \rangle/\langle 2\ell + 1 \rangle$ has order k . Then

$$(X^{kd/\ell} - \xi^k) = \prod_{i \in \langle 2\ell/k + 1 \rangle/\langle 2\ell + 1 \rangle} (X^{d/\ell} - \xi^i), \quad (4.15)$$

and therefore

$$(X^d + 1) = \prod_{i \in \mathbb{Z}_{2d}^*/\langle 2\ell/k + 1 \rangle} (X^{kd/\ell} - \xi^{ik}) = \prod_{i \in \mathbb{Z}_{2d}^*/\langle 2\ell/k + 1 \rangle} \prod_{j \in \langle 2\ell/k + 1 \rangle/\langle 2\ell + 1 \rangle} (X^{d/\ell} - \xi^{ij}).$$

Another way of showing it is, let $\sigma = \sigma_{2\ell/k+1}$, since $\{(2\ell/k + 1)^i \bmod 2d\}_{i \in [0, k-1]}$ gives $\langle 2\ell/k + 1 \rangle/\langle 2\ell + 1 \rangle$, then

$$(X^{kd/\ell} - \xi^k) = \prod_{i \in [0, k-1]} \sigma^i(X^{d/\ell} - \xi). \quad (4.16)$$

By $\mathbb{Z}_{2d}^*/\langle 2\ell/k + 1 \rangle \cong \mathbb{Z}_{2\ell/k}^*$, then

$$(X^d + 1) = \prod_{i \in \mathbb{Z}_{2\ell/k}^*} \prod_{j \in [0, k-1]} \sigma^j(X^{d/\ell} - \xi^i),$$

where the prime ideals are indexed by $\mathbb{Z}_{2\ell/k}^* \times [0, k-1]$.

Distribution in the NTT Coefficients. Consider the distribution for the coefficients in $R_q = \mathbb{Z}_q[X]/(X^d + 1)$, where each coefficient is over $\{-1, 0, 1\}$, 0 has probability p , and ± 1 has probability $(1 - p)/2$ each. Such distribution C is the distribution for the challenge set.

Lemma 4.12.4 (Distributional Invariance across NTT Factors [ALS20]). *Let $\mathbf{x} \leftarrow C$ where the coefficients are i.i.d. Then $\mathbf{x} \bmod (X^{d/\ell} - \xi^i)$ is identically distributed as $\mathbf{x} \bmod (X^{d/\ell} - \xi^j)$ for all $i, j \in \mathbb{Z}_{2\ell}^*$.*

Proof. First, for any automorphism $\sigma \in \text{Aut}(R_q)$, if $\mathbf{x} \leftarrow C$, then $\mathbf{x}$ is identically distributed as $\sigma(\mathbf{x})$, as the coefficients of $\sigma(\mathbf{x})$ are also i.i.d. over $\{-1, 0, 1\}$ with the same distribution as C .

Suppose $(X^{d/\ell} - \xi^i)$ is a prime ideal in $\mathbb{Z}_q[X]/(X^d + 1)$, then as previously discussed,

$$R_q / (X^{d/\ell} - \xi^i) \cong R_q / (X^{d/\ell} - \xi^j)$$

as ξ has an order 2ℓ , and each NTT component is isomorphic to each other under an automorphism in $\text{Aut}(R_q)$, then

$$\sigma(\mathbf{x} \bmod (X^{d/\ell} - \xi^i)) = \sigma(\mathbf{x}) \bmod (X^{d/\ell} - \xi^j)$$

if $\sigma \in \text{Aut}(R_q)$ induces the isomorphism. Thus, the prime-ideal case for distributional invariance across NTT factors is concluded.

If $(X^{d/\ell} - \xi^i)$ is not irreducible, then neither is $(X^{d/\ell} - \xi^j)$, and both split into the same number of irreducible factors. By the existence of an automorphism inducing an isomorphism from one irreducible factor to the other, we conclude the proof. $\square$

Let $\alpha = d/\ell$. We reorder the coefficients of $\mathbf{c} \leftarrow C$ as follows:

$$\begin{aligned} \mathbf{c} = & \begin{array}{cccc} c_0 & + c_\alpha X^\alpha & + \cdots & + c_{d-\alpha} X^{d-\alpha} \\ c_1 X & + c_{\alpha+1} X^{\alpha+1} & + \cdots & + c_{d-\alpha+1} X^{d-\alpha+1} \\ & & \cdots & \\ c_{\alpha-1} X^{\alpha-1} & + c_{2\alpha-1} X^{2\alpha-1} & + \cdots & + c_{d-1} X^{d-1}. \end{array} \end{aligned}$$

Hence, the d/ℓ coefficients of $\mathbf{c} \bmod (X^\alpha - \xi)$ are i.i.d., which can be shown as

$$\mathbf{c} \equiv \begin{bmatrix} 1 & X & \cdots & X^{\alpha-1} \end{bmatrix} \begin{bmatrix} c_0 & + & c_\alpha \xi & + \cdots & + & c_{d-\alpha} \xi^{\ell-1} \\ c_1 & + & c_{\alpha+1} \xi & + \cdots & + & c_{d-\alpha+1} \xi^{\ell-1} \\ & & \cdots & & & \\ c_{\alpha-1} & + & c_{2\alpha-1} \xi & + \cdots & + & c_{d-1} \xi^{\ell-1} \end{bmatrix} \bmod (X^\alpha - \xi), \quad (4.17)$$

where each row, namely each coefficient of $\mathbf{c} \bmod (X^\alpha - \xi)$, is represented by the random variable

$$Y = \sum_{i \in [0, \ell-1]} \xi^i X_i,$$

and the X_i are i.i.d. according to the $\{-1, 0, 1\}$ distribution. We now present the lemma for the distribution of Y .

Lemma 4.12.5 (Projected Challenge Anti-Concentration [ALS20]). *For all $x \in \mathbb{Z}_q$,*

$$\Pr[Y = x] \leq \frac{1}{q} + \frac{1}{q} \sum_{j \in \mathbb{Z}_q^*} \prod_{k \in [0, \ell-1]} \left| p + (1-p) \cos(2\pi j \xi^k / q) \right|.$$

Proof. Let $P : \mathbb{Z}_q \rightarrow [0, 1]$ be the probability mass function. It is a convolution of the following distributions:

$$\mu_k(0) = p, \quad \mu_k(\xi^k) = \mu_k(-\xi^k) = (1-p)/2$$

for $k \in [0, \ell - 1]$, as $Y = \sum_{i \in [0, \ell - 1]} \xi^i X_i$. Hence, by Fourier analysis,

$$\hat{\mu}_k(j) = \sum_{\gamma \in \mathbb{Z}_q} \mu_k(\gamma) \exp(-2\pi i j \gamma / q) = p + (1 - p) \cos(2\pi j \xi^k / q).$$

By convolution,

$$\hat{P}(j) = \prod_{k \in [0, \ell - 1]} \hat{\mu}_k(j) = \prod_{k \in [0, \ell - 1]} \left(p + (1 - p) \cos(2\pi j \xi^k / q) \right).$$

By inverse Fourier transform,

$$\begin{aligned} P(x) &= \frac{1}{q} \sum_{j \in \mathbb{Z}_q} \hat{P}(j) \exp(2\pi i j x / q) = \frac{1}{q} + \frac{1}{q} \sum_{j \in \mathbb{Z}_q^*} \hat{P}(j) \exp(2\pi i j x / q) \\ &\leq \frac{1}{q} + \frac{1}{q} \sum_{j \in \mathbb{Z}_q^*} |\hat{P}(j)| = \frac{1}{q} + \frac{1}{q} \sum_{j \in \mathbb{Z}_q^*} \prod_{k \in [0, \ell - 1]} \left| p + (1 - p) \cos(2\pi j \xi^k / q) \right|. \end{aligned}$$

□

By taking advantage of symmetry, we can reduce the number of terms needed to compute the probability upper bound by a factor of 2ℓ .

Lemma 4.12.6. *For all $x \in \mathbb{Z}_q$,*

$$\Pr[Y = x] \leq \frac{1}{q} + \frac{2\ell}{q} \sum_{j \in \mathbb{Z}_q^* / \langle \xi \rangle} \prod_{k \in [0, \ell - 1]} \left| p + (1 - p) \cos(2\pi j \xi^k / q) \right|. \quad (4.18)$$

Proof. Since $\langle \xi \rangle \subset \mathbb{Z}_q^*$ has order 2ℓ , then $\{1, \xi, \dots, \xi^{\ell-1}\} = \langle \xi \rangle / \pm 1$. Supposing $g \in \mathbb{Z}_q^*$ has order $q - 1$ that generates $\mathbb{Z}_q^*$, and $\xi = g^{(q-1)/(2\ell)}$, then the cosets of $\langle \xi \rangle$ can be represented by $\mathbb{Z}_q^* / \langle \xi \rangle = \{1, \dots, g^{(q-1)/(2\ell)-1}\}$. Hence, multiplying any ξ^m to $\{j, j\xi, \dots, j\xi^{\ell-1}\}$ permutes these elements up to sign. Since cosine is even, then $\hat{P}(j) = \hat{P}(j\xi^m)$. □

Each row can be represented in Horner's polynomial evaluation form

$$x_0 + \xi(x_1 + \xi(x_2 + \dots + \xi(x_{\ell-2} + \xi x_{\ell-1}))),$$

which can be transformed into a Markov chain: $Z_0 = 0, Z_{i+1} = \xi Z_i + X_i$ where X_i follows the $\{-1, 0, 1\}$ distribution.

Module-SIS and Module-LWE. We recall the assumptions underlying the binding and hiding properties of the commitment scheme in [ALS20].

Definition 4.12.7 (MSIS $_{n,m,B}$). Let $n, m \geq 1$ and $0 < B < q$. An efficient $\mathcal{A}$ has advantage ε in solving MSIS $_{n,m,B}$ if

$$\Pr \left[0 < \|\bar{\mathbf{x}}\|_2 \leq B \wedge \mathbf{A}\bar{\mathbf{x}} = \bar{\mathbf{0}} \mid \mathbf{A} \xleftarrow{\mathcal{R}} R_q^{n \times m}, \bar{\mathbf{x}} \leftarrow \mathcal{A}(\mathbf{A}) \right] \geq \varepsilon,$$

where $\bar{\mathbf{x}} \in R_q^m$, the equation is over R_q , and the norm is the coefficient ℓ_2 -norm of the whole vector, using centered representatives.

Fix a coefficient distribution χ over $\mathbb{Z}$. As before, χ^{kd} denotes independent sampling of all coefficients in R^k .

Definition 4.12.8 (MLWE $_{n,m,\chi}$). Let $n, m \geq 1$. A PPT algorithm $\mathcal{A}$ has advantage ε in solving MLWE $_{n,m,\chi}$ if

$$\left| \Pr \left[\mathcal{A}(\mathbf{A}, \mathbf{A}\bar{\mathbf{s}} + \bar{\mathbf{e}}) = 1 \mid \begin{array}{l} \mathbf{A} \xleftarrow{\mathcal{R}} R_q^{m \times n}, \\ \bar{\mathbf{s}} \leftarrow \chi^{nd}, \\ \bar{\mathbf{e}} \leftarrow \chi^{md} \end{array} \right] - \Pr \left[\mathcal{A}(\mathbf{A}, \bar{\mathbf{u}}) = 1 \mid \begin{array}{l} \mathbf{A} \xleftarrow{\mathcal{R}} R_q^{m \times n}, \\ \bar{\mathbf{u}} \xleftarrow{\mathcal{R}} R_q^m \end{array} \right] \right| \geq \varepsilon.$$

All samples within each experiment are independent, and the arithmetic is over R_q .

Following [ALS20], we omit m and write $\text{MSIS}_{n,B}$ and $\text{MLWE}_{n,\chi}$. The BDLOP commitment matrix follows, where μ and λ denote the module ranks for MSIS and MLWE, respectively.

For L committed ring elements, put $N = \mu + \lambda + L$. The block structure used in [BLS19] generalizes to

$$\mathbf{B} = \begin{bmatrix} \mathbf{I}_\mu & \mathbf{U} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L \end{bmatrix} \begin{bmatrix} \mathbf{I}_\mu & \mathbf{0}^{\mu \times L} & \mathbf{V} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L & \mathbf{W} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_\mu & \mathbf{U} & \mathbf{V} + \mathbf{U}\mathbf{W} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L & \mathbf{W} \end{bmatrix} \in R_q^{(\mu+L) \times N}, \quad (4.19)$$

where $\mathbf{U} \xleftarrow{R} R_q^{\mu \times L}$, $\mathbf{V} \xleftarrow{R} R_q^{\mu \times \lambda}$, and $\mathbf{W} \xleftarrow{R} R_q^{L \times \lambda}$. The instance in [BLS19] has $\mu = \lambda = 1$ and $L = 4$, and for the instance in [ALS20], the binding and hiding relies on $\text{MSIS}_{\mu,8\kappa\beta}$ and $\text{MLWE}_{\lambda,\chi}$.

In [ALS20], the public matrix is sampled uniformly with the same dimensions by $\mathbf{B} \xleftarrow{R} R_q^{(\mu+L) \times N}$ (for readability). For the hiding property, supposing the first $\mu + L$ columns form an invertible $\mathbf{C} \in R_q^{(\mu+L) \times (\mu+L)}$, such that

$$\mathbf{B} = [\mathbf{C} \quad \mathbf{D}] = \mathbf{C} [\mathbf{I}_{\mu+L} \quad \mathbf{C}^{-1}\mathbf{D}] = \mathbf{C}\mathbf{E},$$

then multiplying $\mathbf{E}$ with $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ yields a MLWE sample $\bar{\mathbf{g}}$, and $\bar{\mathbf{g}}$ multiplying with invertible $\mathbf{C}$ makes the result computationally indistinguishable to a uniform sample, so as to mask the message $\bar{\mathbf{m}}$. The formal argument is the knapsack-MLWE duality, see [EZZ⁺19, Appendix C]. The binding is still reduced to $\text{MSIS}_{\mu,8\kappa\beta}$ in the same reduction.

Another way of seeing is

$$\mathbf{B}' = \mathbf{R} \begin{bmatrix} \mathbf{I}_\mu & \mathbf{U} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L \end{bmatrix} \begin{bmatrix} \mathbf{I}_\mu & \mathbf{0}^{\mu \times L} & \mathbf{V} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L & \mathbf{W} \end{bmatrix} = \mathbf{R} \begin{bmatrix} \mathbf{I}_\mu & \mathbf{U} & \mathbf{V} + \mathbf{U}\mathbf{W} \\ \mathbf{0}^{L \times \mu} & \mathbf{I}_L & \mathbf{W} \end{bmatrix},$$

where $\mathbf{R}$ is uniform over invertible $R_q^{(\mu+L) \times (\mu+L)}$, then hiding and binding are still reduced to $\text{MLWE}_{\lambda,\chi}$ and $\text{MSIS}_{\mu,8\kappa\beta}$.

To commit to $\bar{\mathbf{m}} = (\mathbf{m}_1, \dots, \mathbf{m}_L)^\top \in R_q^L$, sample $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ and output

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \\ \vdots \\ \mathbf{t}_L \end{bmatrix} = \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_1 \\ \vdots \\ \mathbf{m}_L \end{bmatrix}, \quad \bar{\mathbf{t}}_0 = \mathbf{B}_0\bar{\mathbf{r}}, \quad \mathbf{t}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{r}} \rangle + \mathbf{m}_i.$$

Challenge Set without Difference Invertibility. An immediate application of the result on the NTT slot coefficient distribution is a distribution over C under which the coefficients are sampled i.i.d. over $\{0, 1\}$.²³ Hence, a difference of independently sampled distinct elements is non-invertible with probability at most $\ell p^{d/\ell}$ by union bound.

It is helpful to start from the BDLOP construction, as [ALS20] improves on [BDL⁺18] with a challenge set without the property of *globally invertible pairwise differences*. Regarding the commitment

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_0 \\ -\bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m} \end{bmatrix},$$

in the proof of opening, the prover commits to a short masking vector $\bar{\mathbf{y}}$ from a discrete Gaussian distribution, sends $\bar{\mathbf{w}} = \mathbf{B}_0\bar{\mathbf{y}}$, the verifier challenges with $\mathbf{c} \leftarrow C$, the prover replies with $\bar{\mathbf{z}} = \bar{\mathbf{y}} + \mathbf{c} \cdot \bar{\mathbf{r}}$, and rejection sampling is used to make $\bar{\mathbf{z}}$ statistically close to a discrete Gaussian sample centered at $\mathbf{0}$. The verifier checks $\|\bar{\mathbf{z}}\|_2 \leq \beta$ and $\mathbf{B}_0\bar{\mathbf{z}} = \bar{\mathbf{w}} + \mathbf{c} \cdot \bar{\mathbf{t}}_0$.

If C has this property, the best [BDL⁺18] can extract from the proof of opening is the relaxed opening eq. (4.6).

- On one hand, one cannot guarantee that $\mathbf{f}^{-1} \cdot \bar{\mathbf{r}}$ is short after cancelling out $\mathbf{f}$ from $\bar{\mathbf{r}}$, which can be arbitrarily long.
- On the other hand, the binding property should prevent a single commitment from admitting relaxed openings to distinct messages; otherwise, the opening proof would be meaningless.

The binding reduction to the SIS is achieved by the pairwise difference being invertible, see eq. (4.10).

²³ A bit flawed construction, elaborate later.

Hence, the best [BDL⁺18] can do is the relaxed opening, and uniquely interpolating $\bar{\mathbf{r}}^*$ and $\bar{\mathbf{y}}^*$ out of a pair of forked accepting transcripts. But the scheme cannot guarantee the shortness of $\bar{\mathbf{r}}^*$ and $\bar{\mathbf{y}}^*$.

In [ALS20], the challenge set is not pairwise difference invertible, so an extractor can obtain a “relaxed opening” from a pair of forked accepting transcripts, but binding is not immediately reducible to SIS, and we cannot uniquely interpolate the prover’s replies $\bar{\mathbf{z}}$ and $\bar{\mathbf{z}}'$ to obtain $\bar{\mathbf{y}}^*$ and $\bar{\mathbf{r}}^*$, from only a pair of forked accepting transcripts

$$\bar{\mathbf{z}} = \bar{\mathbf{y}}^* + \mathbf{c} \cdot \bar{\mathbf{r}}^*, \quad \bar{\mathbf{z}}' = \bar{\mathbf{y}}^* + \mathbf{c}' \cdot \bar{\mathbf{r}}^*. \quad (4.20)$$

But we can restore the interpolation for $\bar{\mathbf{y}}^*$ and $\bar{\mathbf{r}}^*$, by piecing together local interpolations over each prime ideal splitting R_q . More specifically, though $\Delta\mathbf{c} = \mathbf{c} - \mathbf{c}'$ is not invertible, not all of the CRT slots are zeroed out. Suppose that

$$R_q = \mathbb{Z}_q[X]/(X^d + 1) \cong \mathbb{Z}_q[X]/(\varphi_1) \times \cdots \times \mathbb{Z}_q[X]/(\varphi_\ell),$$

where $\varphi_1, \dots, \varphi_\ell$ are irreducible polynomial factors, and supposing $\Delta\mathbf{c} \not\equiv 0 \pmod{\varphi_i}$, then interpolation over the i^{th} CRT slot gives

$$\bar{z}_i \equiv \bar{y}_i^* + \mathbf{c}_i \cdot \bar{r}_i^* \pmod{\varphi_i}, \quad \bar{z}'_i \equiv \bar{y}_i^* + \mathbf{c}'_i \cdot \bar{r}_i^* \pmod{\varphi_i},$$

and by piecing together the interpolated $\bar{r}_i^*$ and $\bar{y}_i^*$ over each CRT slot $i \in [1, \ell]$, we have a pair $(\bar{\mathbf{r}}^*, \bar{\mathbf{y}}^*)$.

Moreover, we will show in the following, for this pair of global interpolations $(\bar{\mathbf{y}}^*, \bar{\mathbf{r}}^*)$ pieced out CRT slot-wise, on every accepting transcript prefixed by $\bar{\mathbf{w}}$, the prover must reply in the degree-1 affine form in eq. (4.20).

Lemma 4.12.9. *If we have, for $i \in [1, \ell]$, a pair of accepting transcripts $(\bar{\mathbf{w}}, \mathbf{c}_i, \bar{z}_i)$ and $(\bar{\mathbf{w}}, \mathbf{c}'_i, \bar{z}'_i)$ for the same commitment, with $\Delta\mathbf{c}_i := \mathbf{c}_i - \mathbf{c}'_i \not\equiv 0 \pmod{\varphi_i}$, then every accepting transcript for that commitment with the same $\bar{\mathbf{w}}$ must have that $\bar{\mathbf{z}} = \bar{\mathbf{y}}^* + \mathbf{c} \cdot \bar{\mathbf{r}}^*$ where $(\bar{\mathbf{y}}^*, \bar{\mathbf{r}}^*)$ are fixed independently of $\mathbf{c}$, or otherwise we have a $\text{MSIS}_{\mu, 8\kappa\beta}$ solution for $\mathbf{B}_0$, where κ is a bound for the ℓ_1 -norm of the challenges. Moreover, $\mathbf{B}_0 \bar{\mathbf{y}}^* = \bar{\mathbf{w}}$ and $\mathbf{B}_0 \bar{\mathbf{r}}^* = \bar{\mathbf{t}}_0$.*

Proof. Let $\bar{\mathbf{y}}^{*'} be such that $\bar{\mathbf{z}} = \bar{\mathbf{y}}^{*'} + \mathbf{c} \cdot \bar{\mathbf{r}}^*$. Then fix an arbitrary $i \in [1, \ell]$,$

$$\mathbf{B}_0 (\bar{z}_i - \bar{z}'_i) = \Delta\mathbf{c}_i \cdot \bar{\mathbf{t}}_0, \quad \mathbf{B}_0 (\bar{\mathbf{z}} - \bar{z}_i) = (\mathbf{c} - \mathbf{c}_i) \bar{\mathbf{t}}_0.$$

Cross-multiplying $(\mathbf{c} - \mathbf{c}_i)$ and $\Delta\mathbf{c}_i$, then subtracting them, should yield either a $\text{MSIS}_{\mu, 8\kappa\beta}$ solution for $\mathbf{B}_0$, where the norm bound follows from

$$\|\Delta\mathbf{c}_i\|_1 \cdot \|\bar{\mathbf{z}} - \bar{z}_i\|_2 + \|\mathbf{c} - \mathbf{c}_i\|_1 \cdot \|\bar{z}_i - \bar{z}'_i\|_2 \leq (2\kappa) \cdot (2\beta) + (2\kappa) \cdot (2\beta) = 8\kappa\beta,$$

or it yields

$$(\mathbf{c} - \mathbf{c}_i) \cdot (\bar{z}_i - \bar{z}'_i) = \Delta\mathbf{c}_i \cdot (\bar{\mathbf{z}} - \bar{z}_i).$$

Since $\Delta\mathbf{c}_i \not\equiv 0 \pmod{\varphi_i}$, then

$$(\mathbf{c} - \mathbf{c}_i) \cdot \Delta\mathbf{c}_i \cdot \bar{\mathbf{r}}^* \equiv \Delta\mathbf{c}_i \cdot (\bar{\mathbf{y}}^{*'} - \bar{\mathbf{y}}^* + (\mathbf{c} - \mathbf{c}_i) \cdot \bar{\mathbf{r}}^*) \pmod{\varphi_i},$$

namely $\Delta\mathbf{c}_i \cdot (\bar{\mathbf{y}}^{*'} - \bar{\mathbf{y}}^*) \equiv 0 \pmod{\varphi_i}$. Hence, $\bar{\mathbf{y}}^{*'} \equiv \bar{\mathbf{y}}^* \pmod{\varphi_i}$. If it holds for all $i \in [1, \ell]$, then $\bar{\mathbf{y}}^{*'} = \bar{\mathbf{y}}^*$.

The statement for $\mathbf{B}_0 \bar{\mathbf{y}}^* = \bar{\mathbf{w}}$ and $\mathbf{B}_0 \bar{\mathbf{r}}^* = \bar{\mathbf{t}}_0$ is immediate by CRT slot-wise local interpolation. $\square$

Lemma 4.12.9 is helpful, as it explains that the prover has to reply in the degree-1 form w.r.t. the challenge, and it gives rise to the following augmentation of the relaxed opening, allowing for binding reduction to SIS even if we don’t rely on the pairwise difference invertibility property.

Definition 4.12.10 (Weak Opening). A weak opening for the commitment $\bar{\mathbf{t}}^\top = [\bar{\mathbf{t}}_0^\top \mid \mathbf{t}_1^\top]$ consists of $(\Delta\mathbf{c}_i)_{i \in [1, \ell]} \in R_q^\ell$, $\bar{\mathbf{r}}^* \in R_q^N$ and $\mathbf{m}^* \in R_q$ such that for each $\Delta\mathbf{c}_i \in R_q$, $\Delta\mathbf{c}_i \not\equiv 0 \pmod{(\varphi_i)}$, $\|\Delta\mathbf{c}_i\|_1 \leq 2\kappa$, $\|\Delta\mathbf{c}_i \cdot \bar{\mathbf{r}}^*\|_2 \leq 2\beta$, and

$$\mathbf{B}_0 \bar{\mathbf{r}}^* = \bar{\mathbf{t}}_0, \quad \langle \bar{\mathbf{b}}_1, \bar{\mathbf{r}}^* \rangle + \mathbf{m}^* = \mathbf{t}_1.$$

The key point is to supply ℓ pairwise differences of challenges, such that each CRT slot is nonzero, and eventually this allows for binding reduction to SIS, which shows as follows.

Lemma 4.12.11. *The commitment scheme is binding w.r.t. the weak openings if $\text{MSIS}_{\mu, 8\kappa\beta}$ is hard, namely, valid weak openings*

$$(\{\Delta \mathbf{c}_i\}, \bar{\mathbf{r}}, \mathbf{m}^*), \quad (\{\Delta \mathbf{c}'_i\}, \bar{\mathbf{r}}', \mathbf{m}^{*'})$$

with $\mathbf{m}^* \neq \mathbf{m}^{*'}$ give an $\text{MSIS}_{\mu, 8\kappa\beta}$ solution for $\mathbf{B}_0$.

Proof. If $\mathbf{m}^* \neq \mathbf{m}^{*'}$, then $\bar{\mathbf{r}}^* \neq \bar{\mathbf{r}}^{*'}$, and there must be at least one $i \in [1, \ell]$ such that $\bar{\mathbf{r}}^* - \bar{\mathbf{r}}^{*'} \not\equiv \bar{\mathbf{0}} \pmod{(\varphi_i)}$. Then

$$\Delta \mathbf{c}_i \cdot (\Delta \mathbf{c}'_i \cdot \bar{\mathbf{r}}^*) - \Delta \mathbf{c}'_i \cdot (\Delta \mathbf{c}_i \cdot \bar{\mathbf{r}}^{*'}) = \Delta \mathbf{c}_i \cdot \Delta \mathbf{c}'_i \cdot (\bar{\mathbf{r}}^* - \bar{\mathbf{r}}^{*'}) \not\equiv \bar{\mathbf{0}} \pmod{\varphi_i}, \quad (4.21)$$

which gives a $\text{MSIS}_{\mu, 8\kappa\beta}$ solution for $\mathbf{B}_0$. The $8\kappa\beta$ derives from $\|\Delta \mathbf{c}_i\|_1 \leq 2\kappa$ and $\|\Delta \mathbf{c}_i \cdot \bar{\mathbf{r}}^*\|_2 \leq 2\beta$ in definition 4.12.10. $\square$

Looking back, we can reverse engineer the motivation from [BDL⁺18] as follows. Starting from a challenge set that is not globally pairwise difference invertible, then to obtain an augmentation of the relaxed opening, we should collect ℓ pairs of forked accepting transcripts starting from $\bar{\mathbf{w}}$ the commitment to the masking vector, such that each CRT slot is once nonzero, and we piece out a global interpolation over R_q by local interpolations over $R_q/(\varphi_i)$. What is different from [BDL⁺18] in [ALS20] is that we need to interpolate many times, while in [BDL⁺18] one interpolation is sufficient, as C in [BDL⁺18] is globally pairwise difference invertible. To ensure that the linearity of the prover reply $\mathbf{z}$ against the short challenge set elements is still preserved, the lemma 4.12.9 shows that the affine form is respected for all accepting transcripts beginning with $\bar{\mathbf{w}}$ for the commitment. A good trick in lemma 4.12.9 is, on a reply that does not have the affine form, move all the eccentric parts down to the degree 0 constant $\bar{\mathbf{y}}^{*'}$, breaking SIS.

With lemma 4.12.9, the form of definition 4.12.10 makes sense, as lemma 4.12.9 rules out the non-linearity in each “relaxed opening” in the weak opening, so $\bar{\mathbf{y}}^*$ is always cancelled out. Finally, the lemma 4.12.11 concludes that the weak opening binds against the commitment by SIS hardness.

All of this effort is necessary because C is not pairwise difference invertible, so extraction cannot be done by one pair of forked accepting transcripts, but rather interpolations over multiple pairs. The weak opening enables binding reduction to SIS by providing $\Delta \mathbf{c}_i$ whose i^{th} CRT slot is non-zero, and lemma 4.12.9 rules out all other forms of extracted openings by showing the uniqueness of the line.

The analysis in [ALS20] has some interesting reasoning for a larger challenge set: Suppose the challenge set C is globally pairwise difference invertible, and R_q is full splitting. Fix a CRT slot, there are q distinct values on that CRT slot, and for the challenge set C , by the pigeonhole principle, there are at most q challenges in C , or otherwise there exists collision on the CRT slot, then the pairwise difference is not always invertible. Hence, the soundness is at best $1/q$. Even when R_q splits into a few low-degree factors, we cannot quantify the exact size of such globally pairwise difference invertible C .

Now I should explain why section 4.12 does not work out of the box. Indeed, if $\mathbf{c}, \mathbf{c}' \xleftarrow{R} \{0, 1\}^d$, then the coefficients of $\mathbf{c} - \mathbf{c}'$ are i.i.d. with $P(\pm 1) = 1/4$ and $P(0) = 1/2$. Yet it does not mean the anti-concentration bound p in eq. (4.18) can be immediately applied here, whose coefficient 0 probability is $1/2$, such that the knowledge soundness is $p^{d/\ell}$. Looking closely, let X, X' be the random variables for the j^{th} coefficient in the i^{th} CRT slot of $\mathbf{c}, \mathbf{c}'$, then

$$\sum_{x \in \mathbb{Z}_q} \Pr[X = x \wedge X' = x] = \sum_{x \in \mathbb{Z}_q} \Pr[X = x]^2 = \Pr[X = X'] = p.$$

The last equality holds from eq. (4.18), where the inequality is equality when $x = 0$ for the distribution of $X - X'$. Hence

$$\max_{x \in \mathbb{Z}_q} \Pr[X = x]^2 \leq \sum_{x \in \mathbb{Z}_q} \Pr[X = x]^2 = p,$$

and hence the knowledge soundness induced by the challenge set over $\{0, 1\}^d$ is no longer $p^{d/\ell}$ but rather $p^{d/(2\ell)}$, as the maximum point probability of a coefficient in the i^{th} CRT slot of $\mathbf{c}$ is at most the square root of the anti-concentration bound formed for the $\{-1, 0, 1\}$ symmetric distribution.

BDLOP with Ring Automorphism. Now we construct a BDLOP in the ring automorphism settings. By changing the challenge set distribution, say [LS18], to the distribution defined previously over $\{-1, 0, 1\}^d$, where distribution for each coefficient is i.i.d., $P(\pm 1) = 1/4$ and $P(0) = 1/2$, then let p be the maximum probability in eq. (4.18), then the knowledge soundness should be $p^{d/\ell}$. More specifically, if the prover convinces the verifier with probability $\varepsilon > p^{d/\ell}$, we obtain the first accepting transcript on challenge $\mathbf{c}$ in expected $1/\varepsilon$ time, then for another challenge $\mathbf{c}'$,

$$\begin{aligned} \varepsilon &= \Pr[\mathcal{P} \text{ succeeds}] = \Pr[\mathcal{P} \text{ succeeds} \mid \mathbf{c}' \not\equiv \mathbf{c} \pmod{(\varphi_i)}] \cdot \Pr[\mathbf{c}' \not\equiv \mathbf{c} \pmod{(\varphi_i)}] \\ &\quad + \Pr[\mathcal{P} \text{ succeeds} \mid \mathbf{c}' \equiv \mathbf{c} \pmod{(\varphi_i)}] \cdot \Pr[\mathbf{c}' \equiv \mathbf{c} \pmod{(\varphi_i)}] \\ &\leq \Pr[\mathcal{P} \text{ succeeds} \mid \mathbf{c}' \not\equiv \mathbf{c} \pmod{(\varphi_i)}] + \Pr[\mathbf{c}' \equiv \mathbf{c} \pmod{(\varphi_i)}], \end{aligned}$$

the convincing probability on $\mathbf{c}' \not\equiv \mathbf{c} \pmod{\varphi_i}$ is at least $\varepsilon - p^{d/\ell}$, and the extractor runs in expected polynomial time if $\varepsilon - p^{d/\ell}$ is $\text{poly}^{-1}(\lambda)$.

If $p^{d/\ell}$ is not negligible, we can run k copies of the protocol in parallel, and reduce the soundness down to $p^{kd/\ell}$. To achieve this in a 3-move Sigma protocol, they challenge with $\mathbf{c}$ and its automorphisms. I suppose the observation stems from: for a $\mathbf{c}$ sampled over $\{-1, 0, 1\}^d$, a coefficient in a coarse ideal $X^{kd/\ell} - \xi^{jk}$ takes on a value with probability at most p by eq. (4.18)²⁴. Recall that $\sigma = \sigma_{2\ell/k+1}$ stabilizes the ideal $(X^{kd/\ell} - \xi^{jk})$, such that

$$(X^{kd/\ell} - \xi^{jk}) = \prod_{i \in [0, k-1]} \sigma^i (X^{d/\ell} - \xi^j) = \prod_{i \in \langle 2\ell/k+1 \rangle / \langle 2\ell+1 \rangle} (X^{d/\ell} - \xi^{ij}) \quad (4.22)$$

by eqs. (4.15) and (4.16), such that if $\mathbf{c} - \mathbf{c}' \not\equiv 0 \pmod{(X^{kd/\ell} - \xi^{jk})}$, then any of $\sigma^i(\mathbf{c} - \mathbf{c}') \not\equiv 0 \pmod{(X^{kd/\ell} - \xi^{jk})}$, and any of the CRT slot under the coarse CRT factor is at least once non-zero in the k automorphisms.

The construction is then immediate. Sample k random masking elements and commit to them, then the verifier replies with the challenge $\mathbf{c}$, and the prover replies with $\bar{\mathbf{z}}_i = \bar{\mathbf{y}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{r}}$ under rejection sampling. If the prover runs in unit time and convinces the verifier with probability $\varepsilon > p^{kd/\ell}$, then the extraction takes

$$\frac{1}{\varepsilon} + \frac{\ell}{k} \cdot \frac{1}{\varepsilon - p^{kd/\ell}}$$

expected time, as sampling

- The first accepting transcript with challenge $\mathbf{c}$ takes expected $1/\varepsilon$ time.
- An accepting transcript with challenge $\mathbf{c}_j$ with $\mathbf{c} - \mathbf{c}_j \not\equiv 0 \pmod{(X^{kd/\ell} - \xi^{jk})}$ takes expected $(\varepsilon - p^{kd/\ell})^{-1}$ time.

Strawman Product Proof. We begin with a rudimentary product proof using random masking elements. For

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_1 \\ \bar{\mathbf{t}}_2 \\ \bar{\mathbf{t}}_3 \end{bmatrix} = \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_1 \\ \mathbf{m}_2 \\ \mathbf{m}_3 \end{bmatrix}, \quad (4.23)$$

the relation to be proved is $\mathbf{m}_1 \cdot \mathbf{m}_2 = \mathbf{m}_3$. The strawman solution in [ALS20] is by

$$\begin{bmatrix} \bar{\mathbf{t}}'_0 \\ \bar{\mathbf{t}}'_1 \\ \bar{\mathbf{t}}'_2 \\ \bar{\mathbf{t}}'_3 \\ \bar{\mathbf{t}}'_4 \\ \bar{\mathbf{t}}'_5 \end{bmatrix} = \mathbf{B}'\bar{\mathbf{r}}' + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{a}_1 \\ \mathbf{a}_2 \\ \mathbf{a}_3 \\ \mathbf{a}_1\mathbf{m}_2 + \mathbf{a}_2\mathbf{m}_1 - \mathbf{a}_3 \\ \mathbf{a}_1\mathbf{a}_2 \end{bmatrix}, \quad (4.24)$$

²⁴Here the p is not defined for $X^{d/\ell} - \xi^j$ but for $X^{kd/\ell} - \xi^{kj}$.

that is committing to random masking polynomials $\mathbf{a}_i$ and garbage terms. The invariant is the masked relation

$$(\mathbf{a}_1 + \mathbf{x}\mathbf{m}_1)(\mathbf{a}_2 + \mathbf{x}\mathbf{m}_2) - \mathbf{x}(\mathbf{a}_3 + \mathbf{x}\mathbf{m}_3) = \mathbf{a}_1\mathbf{a}_2 + \mathbf{x}(\mathbf{m}_1\mathbf{a}_2 + \mathbf{m}_2\mathbf{a}_1 - \mathbf{a}_3) + \mathbf{x}^2(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3) \quad (4.25)$$

after a challenge $\mathbf{x}$ from the verifier. More specifically, for the masked openings

$$\mathbf{f}_i = \mathbf{a}_i + \mathbf{x}\mathbf{m}_i,$$

we can prove that $\mathbf{t}'_i + \mathbf{x}\mathbf{t}_i$ commits to $\mathbf{f}_i$. The verifier can then check eq. (4.25) by seeing if $\mathbf{t}'_5 + \mathbf{x}\mathbf{t}'_4$ commits to $\mathbf{f}_1\mathbf{f}_2 - \mathbf{x}\mathbf{f}_3$

$$\mathbf{f}_1\mathbf{f}_2 - \mathbf{x}\mathbf{f}_3 = \mathbf{a}_1\mathbf{a}_2 + \mathbf{x}(\mathbf{m}_1\mathbf{a}_2 + \mathbf{m}_2\mathbf{a}_1 - \mathbf{a}_3) + \mathbf{x}^2(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3),$$

such that the quadratic relation is checked linearly over committed elements, and the desired product relation is the vanishing of its degree-2 coefficient. Essentially, the masking technique is used in [BLS19, YAZ⁺19], where a product proof appears inside a shortness proof.

We now put things together conceptually as follows. Suppose the prover knows an opening $(\mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_3, \bar{\mathbf{r}})$ for the commitment in eq. (4.23):

- The prover commits to the masked terms and garbage terms in eq. (4.24).
- The verifier replies with a uniformly random challenge $\mathbf{x}$.
- The prover replies with masked openings $\mathbf{f}_i = \mathbf{a}_i + \mathbf{x}\mathbf{m}_i$.

Moreover, it commits to the masking randomness $\bar{\mathbf{y}}$ and $\bar{\mathbf{y}}'$ by

$$\bar{\mathbf{w}} = \mathbf{B}_0\bar{\mathbf{y}}, \quad \bar{\mathbf{w}}' = \mathbf{B}'_0\bar{\mathbf{y}}',$$

and sends over $\mathbf{v}_i = \mathbf{x} \cdot \langle \bar{\mathbf{b}}_i, \bar{\mathbf{y}} \rangle + \langle \bar{\mathbf{b}}'_i, \bar{\mathbf{y}}' \rangle$, and $\mathbf{u} = \mathbf{x} \cdot \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle + \langle \bar{\mathbf{b}}'_5, \bar{\mathbf{y}}' \rangle$ ²⁵.

- The verifier replies with a short challenge $\mathbf{c}$.
- The prover runs rejection sampling, and replies

$$\bar{\mathbf{z}} = \bar{\mathbf{y}} + \mathbf{c} \cdot \bar{\mathbf{r}}, \quad \bar{\mathbf{z}}' = \bar{\mathbf{y}}' + \mathbf{c} \cdot \bar{\mathbf{r}}'.$$

- The verifier checks $\bar{\mathbf{z}}$ and $\bar{\mathbf{z}}'$ are short,

$$\mathbf{B}_0\bar{\mathbf{z}} = \bar{\mathbf{w}} + \mathbf{c} \cdot \bar{\mathbf{t}}_0,$$

$$\mathbf{B}'_0\bar{\mathbf{z}}' = \bar{\mathbf{w}}' + \mathbf{c} \cdot \bar{\mathbf{t}}'_0,$$

$$\mathbf{v}_i + \mathbf{c}(\mathbf{t}'_i + \mathbf{x}\mathbf{t}_i) = \mathbf{x} \cdot \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}} \rangle + \langle \bar{\mathbf{b}}'_i, \bar{\mathbf{z}}' \rangle + \mathbf{c}\mathbf{f}_i,$$

and checks the quadratic masked relation eq. (4.25) by ensuring $\gamma = \mathbf{x}\mathbf{t}'_4 + \mathbf{t}'_5 - (\mathbf{f}_1\mathbf{f}_2 - \mathbf{x}\mathbf{f}_3)$ commits to 0 by

$$\mathbf{u} + \mathbf{c}\gamma = \mathbf{x} \cdot \langle \bar{\mathbf{b}}_4, \bar{\mathbf{z}} \rangle + \langle \bar{\mathbf{b}}'_5, \bar{\mathbf{z}}' \rangle.$$

Regardless of the proof size, there is a catch regarding the knowledge soundness: $\mathbf{x}$ should be pairwise difference invertible, which falls back the prior discussion of the challenge set construction²⁶, and reusing the same masks for different $\mathbf{x}$ breaks zero knowledge.

²⁵Recall the proving linear relation over elements in BDLOP abstracted in note for [ALS20].

²⁶Suppose the polynomial ring splits completely. The knowledge soundness will be at best $1/q$ by the pigeonhole principle.

Computationally Indistinguishable Masked Opening. An observation to the prior product proof from masked opening is, the values $\mathbf{f}_i = \mathbf{a}_i + \mathbf{x}\mathbf{m}_i$ provide unconditional statistical hiding. We have

$$\mathbf{f}'_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}} \rangle - \mathbf{c}\mathbf{t}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{y}} \rangle - \mathbf{c}\mathbf{m}_i,$$

such that $\langle \bar{\mathbf{b}}_i, \bar{\mathbf{y}} \rangle$ is the computationally indistinguishable mask, and the masked relation in eq. (4.25) converts to

$$\mathbf{f}'_1 \mathbf{f}'_2 + \mathbf{c}\mathbf{f}'_3 = \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle + \mathbf{c} \left(\langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle - \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \mathbf{m}_2 - \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle \mathbf{m}_1 \right) + \mathbf{c}^2 (\mathbf{m}_1 \mathbf{m}_2 - \mathbf{m}_3). \quad (4.26)$$

With one garbage term $\langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle - \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \mathbf{m}_2 - \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle \mathbf{m}_1$ to commit to (and sending out claim $\langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle + \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle$), the new construction can be more efficient. Hence, the masked openings are inherent in the Σ -protocol.

The protocol so far works as follows:

- The prover commits to the masking randomness $\bar{\mathbf{w}} = \mathbf{B}_0 \bar{\mathbf{y}}$, the message and the garbage term

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \\ \mathbf{t}_4 \end{bmatrix} = \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_1 \\ \mathbf{m}_2 \\ \mathbf{m}_3 \\ \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle - \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \mathbf{m}_2 - \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle \mathbf{m}_1 \end{bmatrix},$$

and derives the claim $\mathbf{v} = \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}} \rangle \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle + \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle$.

- The verifier challenges with $\mathbf{c}$.
- The prover runs rejection sampling and replies with $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{y}} + \mathbf{c}\bar{\mathbf{r}}$.
- The verifier checks that $\bar{\mathbf{z}}$ is short and that

$$\begin{aligned} \bar{\mathbf{w}} + \mathbf{c}\bar{\mathbf{t}}_0 &= \mathbf{B}_0 \bar{\mathbf{z}}, \\ \mathbf{f}_1 \mathbf{f}_2 + \mathbf{c}\mathbf{f}_3 + \mathbf{f}_4 &= \mathbf{v}, \end{aligned}$$

where $\mathbf{f}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}} \rangle - \mathbf{c}\mathbf{t}_i$ is the new computationally indistinguishable masked opening.

Soundness Amplification from Automorphism Repetition, Recap. After the strawman construction of the product proof, we pointed out the knowledge soundness issue from uniformly random $\mathbf{x}$, and the same issue remains in the computationally indistinguishable masked opening, which is challenged by $\mathbf{c}$.

The prior structured parallel repetition using automorphisms becomes handy, and it is a good time to recap here. Before the “automorphism repetition”, the malicious prover can cheat by, say, if $\bar{\mathbf{r}}$ is correct except at the i^{th} CRT slot, then there must exist $\bar{\delta}$ and $\bar{\Delta} = \mathbf{B}_0 \bar{\delta}$ satisfying

$$\bar{\delta} \not\equiv \bar{\mathbf{0}} \pmod{(X^{d/\ell} - \xi^i)}, \quad \bar{\delta} \equiv \bar{\mathbf{0}} \pmod{(X^{d/\ell} - \xi^j)}$$

for $j \neq i$, and $\bar{\Delta}$ is non-zero only at the i^{th} CRT slot, such that $\mathbf{B}_0 \bar{\mathbf{r}}_0 = \bar{\mathbf{t}}_0$ for $\bar{\mathbf{r}} = \bar{\mathbf{r}}_0 + \bar{\delta}$, and hence $\mathbf{B}_0 \bar{\mathbf{r}} = \bar{\mathbf{t}}_0 + \bar{\Delta}$.

Supposing the prover guessed the challenge's i^{th} CRT lane is $\mathbf{s} \in \mathbb{Z}_q[X]/(X^{d/\ell} - \xi^i)$, let $\mathbf{c}_s$ be the lifted R_q element of $\mathbf{s}$ such that it is only non-zero at the i^{th} CRT lane, and the forging follows:

- The prover samples $\bar{\mathbf{y}}$ the masking randomness, and sends out $\bar{\mathbf{w}} \leftarrow \mathbf{B}_0 \bar{\mathbf{y}} + \mathbf{c}_s \bar{\Delta}$.
- The verifier challenges with an element $\mathbf{c}$ such that $\mathbf{c} \equiv \mathbf{c}_s \pmod{(X^{d/\ell} - \xi^i)}$.
- The prover runs rejection sampling and replies with $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{y}} + \mathbf{c}\bar{\mathbf{r}}$.
- The forged proof still goes through by

$$\bar{\mathbf{w}} + \mathbf{c}\bar{\mathbf{t}}_0 = \mathbf{B}_0 \bar{\mathbf{y}} + \mathbf{c}_s \bar{\Delta} + \mathbf{c}(\mathbf{B}_0 \bar{\mathbf{r}} - \bar{\Delta}) = \mathbf{B}_0(\bar{\mathbf{y}} + \mathbf{c}\bar{\mathbf{r}}) + \bar{\Delta}(\mathbf{c}_s - \mathbf{c}) = \mathbf{B}_0 \bar{\mathbf{z}}.$$

The trick is to compensate by the right amount of $\bar{\Delta}$ if the prover guesses the relevant part of $\mathbf{c}$ correctly, as $\bar{\mathbf{t}}_0$ has a deficit of $\bar{\Delta}$ relative to $\mathbf{B}_0\bar{\mathbf{r}}$.

Moving to the “automorphism repetition”, for a CRT slot determined by $(X^{d/\ell} - \xi^j)$, instead of being tested by a single challenge $\mathbf{c} \bmod (X^{d/\ell} - \xi^j)$, it is tested by k challenges deriving from

$$\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj}) \cong \left(\mathbf{c} \bmod (X^{d/\ell} - \xi^j), \sigma(\mathbf{c}) \bmod (X^{d/\ell} - \xi^j), \dots, \sigma^{k-1}(\mathbf{c}) \bmod (X^{d/\ell} - \xi^j) \right). \quad (4.27)$$

Suppose any one of the small CRT slots under $(X^{kd/\ell} - \xi^{kj})$ ²⁷ is incorrect in $\bar{\mathbf{r}}$, the malicious prover needs to guess $\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj})$, and forge across the k repetitions at the corresponding slots with the same algorithm. The guessing probability is at most the knowledge-soundness bound $p^{kd/\ell}$.

Automorphism Repetition in Product Proof. We can view the proof of $\mathbf{m}_1\mathbf{m}_2 = \mathbf{m}_3$ as proving

$$\mathbf{m}_1\mathbf{m}_2 \equiv \mathbf{m}_3 \bmod (X^{d/\ell} - \xi^j)$$

for each $j \in \mathbb{Z}_{2d}^*/\langle 2\ell + 1 \rangle$, simultaneously for ℓ CRT slots. With the same motivation, the vanilla product proof testing eq. (4.26) only suffices for non-negligible soundness, as a CRT slot of the message is tested with a single CRT slot of the challenge, instead of being tested k times by eq. (4.27).

To build up the intuition, consider the vanilla product proof for eq. (4.26), and suppose $\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3$ is nonzero only at the j^{th} CRT slot. For simplicity, we write $\mathbf{a}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{y}} \rangle$, and the malicious prover guesses by adding an element $\mathbf{s}$ to the claim $\mathbf{v}$, such that $\mathbf{s}$ is only nonzero at the j^{th} CRT slot, then

$$\mathbf{f}_1\mathbf{f}_2 + \mathbf{c}\mathbf{f}_3 + \mathbf{f}_4 - \mathbf{v} = (\mathbf{a}_1\mathbf{a}_2 + \mathbf{a}_4 - \mathbf{v}) + \mathbf{c}(\mathbf{a}_3 - \mathbf{a}_1\mathbf{m}_2 - \mathbf{a}_2\mathbf{m}_1 - \mathbf{m}_4) + \mathbf{c}^2(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3) \equiv \mathbf{0} \bmod (X^{d/\ell} - \xi^j) \quad (4.28)$$

has at most 2 roots for the j^{th} CRT slot, and thus the guessing probability is at most $2p^{d/\ell}$ by a union bound.

Now the genuine question is how to implement “automorphism repetition” for the product proof using eq. (4.26). Consider a product proof utilizing the “automorphism parallelism” that uses the challenge $\mathbf{c}$ and its automorphisms but separates the claims and commitments to the garbage terms as follows:

- The prover commits to the masking randomness $\bar{\mathbf{w}}_i = \mathbf{B}_0\bar{\mathbf{y}}_i$ for $i \in [0, k-1]$, the message commitment

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \end{bmatrix} = \mathbf{B}\bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0} \\ \mathbf{m}_1 \\ \mathbf{m}_2 \\ \mathbf{m}_3 \end{bmatrix},$$

and the garbage terms and the claims for $i \in [0, k-1]$ by

$$\begin{aligned} \mathbf{t}_{4,i} &= \langle \bar{\mathbf{b}}_4, \bar{\mathbf{r}} \rangle + \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}}_i \rangle - \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_i \rangle \mathbf{m}_2 - \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_i \rangle \mathbf{m}_1, \\ \mathbf{v}_i &= \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}}_i \rangle + \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_i \rangle \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_i \rangle. \end{aligned}$$

- The verifier challenges with $\mathbf{c}$.
- The prover runs rejection sampling and replies with $\bar{\mathbf{z}}_i \leftarrow \bar{\mathbf{y}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{r}}$.
- The verifier checks that $\bar{\mathbf{z}}_i$ is short and that

$$\begin{aligned} \bar{\mathbf{w}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{t}}_0 &= \mathbf{B}_0\bar{\mathbf{z}}_i, \\ \mathbf{f}_{1,i}\mathbf{f}_{2,i} + \sigma^i(\mathbf{c}) \cdot \mathbf{f}_{3,i} + \mathbf{f}_{4,i} &= \mathbf{v}_i, \end{aligned}$$

where, for $a \in [1, 3]$,

$$\mathbf{f}_{a,j} = \langle \bar{\mathbf{b}}_a, \bar{\mathbf{z}}_j \rangle - \sigma^j(\mathbf{c}) \cdot \mathbf{t}_{a,j}, \quad \mathbf{f}_{4,j} = \langle \bar{\mathbf{b}}_4, \bar{\mathbf{z}}_j \rangle - \sigma^j(\mathbf{c}) \cdot \mathbf{t}_{4,j}.$$

²⁷See eq. (4.22).

Again, let $\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3$ be nonzero only at one of the fine CRT slots under $(X^{kd/\ell} - \xi^{kj})$. By eq. (4.28), for one forgery of the claim $\mathbf{v}$, the quadratic constraint admits at most 2 roots modulo the fine CRT slot, then by eq. (4.27), forging the claims $\mathbf{v}_0, \dots, \mathbf{v}_{k-1}$ for the nonzero CRT slot gives at most 2^k roots modulo $(X^{kd/\ell} - \xi^{kj})$, namely, in each repetition, a forgery on the CRT slot admits 2 roots for $\sigma^i(\mathbf{c}) \bmod (X^{d/\ell} - \xi^j)$. Hence, at most 2^k values of $\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj})$ can make the k CRT slots across k repetitions vanish simultaneously.

Note that in the degree-2 setting, the malicious prover can take advantage of the higher degree: it does not have to guess the exact $\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj})$, but instead can guess among 2^k values. The resulting guessing probability is at most $2^k p^{kd/\ell}$.

Automorphism and Communication Complexity. The repetition over automorphisms provides a concrete benefit. In fact, for the prior construction utilizing the automorphism repetition, where $\{\mathbf{t}_{4,i}\}_{i \in [0, k-1]}$ commits to the garbage terms, and $\{\mathbf{v}_i\}_{i \in [0, k-1]}$ are the claims, using a set of independent challenges still suffices, and the independence makes it easier to reason about the soundness.

This is surely better than the strawman construction regarding the concrete proof size, yet the structural nature of the automorphism is absent. Recalling the automorphism in eq. (4.13), consider an inverse σ_i^{-1} such that

$$\sigma_i^{-1}(\sigma_i(\mathbf{f}_1) \cdot \sigma_i(\mathbf{f}_2)) = \sigma_i^{-1}(\sigma_i(\mathbf{f}_1 \cdot \mathbf{f}_2)) = \mathbf{f}_1 \cdot \mathbf{f}_2,$$

then for eq. (4.28) and eq. (4.26), we can “inverse rotate” from

$$\mathbf{f}_{1,i}\mathbf{f}_{2,i} + \sigma^i(\mathbf{c}) \cdot \mathbf{f}_{3,i} = \mathbf{a}_{1,i}\mathbf{a}_{2,i} + \sigma^i(\mathbf{c})(\mathbf{a}_{3,i} - \mathbf{a}_{1,i}\mathbf{m}_2 - \mathbf{a}_{2,i}\mathbf{m}_1) + \sigma^i(\mathbf{c})^2(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3)$$

to

$$\sigma^{-i}(\mathbf{f}_{1,i}\mathbf{f}_{2,i}) + \mathbf{c} \cdot \sigma^{-i}(\mathbf{f}_{3,i}) = \sigma^{-i}(\mathbf{a}_{1,i}\mathbf{a}_{2,i}) + \mathbf{c}\sigma^{-i}(\mathbf{a}_{3,i} - \mathbf{a}_{1,i}\mathbf{m}_2 - \mathbf{a}_{2,i}\mathbf{m}_1) + \mathbf{c}^2\sigma^{-i}(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3),$$

such that for $i \in [0, k-1]$, the $\mathbf{c}$ automorphisms are aligned back to the original $\mathbf{c}$. This structure allows us to combine k claims and garbage commitments into one through a random linear combination.

Let $\alpha_0, \dots, \alpha_{k-1}$ be sampled uniformly over R_q . Then the garbage-term commitment and the claim are

$$\mathbf{t}_4 = \langle \bar{\mathbf{b}}_4, \bar{\mathbf{r}} \rangle + \sum_{i \in [0, k-1]} \alpha_i \sigma^{-i}(\mathbf{a}_{3,i} - \mathbf{a}_{1,i}\mathbf{m}_2 - \mathbf{a}_{2,i}\mathbf{m}_1), \quad \mathbf{v} = \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}}_0 \rangle + \sum_{i \in [0, k-1]} \alpha_i \sigma^{-i}(\mathbf{a}_{1,i}\mathbf{a}_{2,i}), \quad (4.29)$$

for the combined statement

$$\sum_{i \in [0, k-1]} \alpha_i \cdot \sigma^{-i}(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3) = \mathbf{0}. \quad (4.30)$$

The final protocol follows:

- The prover commits to the masking randomness $\bar{\mathbf{w}}_i = \mathbf{B}_0 \bar{\mathbf{y}}_i$ for $i \in [0, k-1]$, the message commitment

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \\ \mathbf{t}_2 \\ \mathbf{t}_3 \end{bmatrix} = \mathbf{B} \bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_1 \\ \mathbf{m}_2 \\ \mathbf{m}_3 \end{bmatrix}.$$

- The verifier replies with the uniform randomness $\alpha_0, \dots, \alpha_{k-1}$.
- The prover derives $\mathbf{t}_4$ and $\mathbf{v}$ in eq. (4.29) and replies.
- The verifier challenges with $\mathbf{c}$.
- The prover runs rejection sampling and replies with $\bar{\mathbf{z}}_i \leftarrow \bar{\mathbf{y}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{r}}$.
- The verifier checks that $\bar{\mathbf{z}}_i$ is short, that $\bar{\mathbf{w}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{t}}_0 = \mathbf{B}_0 \bar{\mathbf{z}}_i$, and

$$\mathbf{f}_{4,0} + \sum_{i \in [0, k-1]} \alpha_i \sigma^{-i}(\mathbf{f}_{1,i}\mathbf{f}_{2,i} + \sigma^i(\mathbf{c}) \cdot \mathbf{f}_{3,i}) = \mathbf{f}_{4,0} + \sum_{i \in [0, k-1]} \alpha_i \sigma^{-i}(\mathbf{f}_{1,i}\mathbf{f}_{2,i}) + \mathbf{c} \sum_{i \in [0, k-1]} \alpha_i \sigma^{-i}(\mathbf{f}_{3,i}) = \mathbf{v},$$

where $\mathbf{f}_{i,j} = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}}_j \rangle - \sigma^j(\mathbf{c}) \cdot \mathbf{t}_i$.

Now supposing a CRT slot is nonzero for $\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3$, then the automorphisms

$$\sigma^{-i}(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3) \bmod \sigma^{-i}(X^{d/\ell} - \xi^j)$$

are all nonzero. For the combined statement eq. (4.30), we analyze the number of CRT slots under

$$\sum_{i \in [0, k-1]} \alpha_i \cdot \sigma^{-i}(\mathbf{m}_1\mathbf{m}_2 - \mathbf{m}_3) \bmod (X^{kd/\ell} - \xi^{kj})$$

are nonzero, and analyze the guessing probability of a malicious prover.

- If the k CRT slots are all nonzero, then it follows exactly from my version of the product proof without utilizing automorphism well, namely the guessing probability is at most $2^k p^{kd/\ell}$, as $\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj})$ takes on at most 2^k values.
- If i of the k CRT slots are zero, then on a set of i slots being zeroed out by the randomness, $\mathbf{c} \bmod (X^{kd/\ell} - \xi^{kj})$ takes on at most $q^{id/\ell} 2^{k-i}$ values.

Hence, the guessing probability is at most

$$\sum_{i \in [0, k]} \binom{k}{i} \left(\frac{1}{q^{d/\ell}} \right)^i \left(1 - \frac{1}{q^{d/\ell}} \right)^{k-i} 2^{k-i} q^{id/\ell} p^{kd/\ell} \leq 3^k p^{kd/\ell}. \quad (4.31)$$

Moreover, to batch multiple product relations, it suffices to have

$$\sum_j \sum_{i \in [0, k-1]} \alpha_{i,j} \cdot \sigma^{-i}(\mathbf{m}_{1,j}\mathbf{m}_{2,j} - \mathbf{m}_{3,j}) = \mathbf{0},$$

namely, to take a random linear combination with $\{\{\alpha_{i,j}\}_{i \in [0, k-1]}\}_j$ over R_q .

The observation is, if there exists nonzero $\mathbf{m}_{1,j}\mathbf{m}_{2,j} - \mathbf{m}_{3,j}$, under the chosen coarse slot containing nonzero error, for each of the k CRT slots, random linear combination over independent $\alpha_{i,j} \bmod \varphi$ vanishes with probability at most $q^{-d/\ell}$ by Schwartz-Zippel, and the randomness of each slot is independent, so the soundness in eq. (4.31) does not change.

4.13 Practical Exact Proofs from Lattices

We now apply the product proof in [ALS20] from the computationally indistinguishable masked opening to [BLS19]. A baseline protocol for the exact proof follows:

- The prover sends out $\bar{w} \leftarrow A\hat{y}$ and $\bar{\mathbf{w}} \leftarrow \mathbf{B}_0\bar{\mathbf{y}}$, where $\mathbf{y} \xleftarrow{R} R_q$ and $\bar{\mathbf{y}} \leftarrow D_\sigma^{Nd}$, and samples $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ to commit

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_1 \\ \bar{\mathbf{t}}_2 \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_2^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{y} \\ \mathbf{s} \end{bmatrix}.$$

- The verifier challenges by sending $c \xleftarrow{R} \mathbb{Z}_q$.
- The prover replies with $\mathbf{z} \leftarrow c \cdot \mathbf{s} + \mathbf{y}$, and starts Π_{lin} and the shortness check in [ALS20] by:
 - The prover derives $\bar{\mathbf{b}}_{f_0} := \bar{\mathbf{b}}_1 + c \cdot \bar{\mathbf{b}}_2$, replies $\mathbf{w}' \leftarrow \langle \bar{\mathbf{b}}_{f_0}, \bar{\mathbf{y}} \rangle$, then it commits

$$\begin{bmatrix} \bar{\mathbf{t}}_3 \\ \bar{\mathbf{t}}_4 \end{bmatrix} \leftarrow \begin{bmatrix} \bar{\mathbf{b}}_3^\top \\ \bar{\mathbf{b}}_4^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle - 3\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^2 \mathbf{s} \\ \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle (3\mathbf{s}^2 - 1) \end{bmatrix}, \quad (4.32)$$

and replies $\bar{\mathbf{t}}_3, \bar{\mathbf{t}}_4$, and the claim $\mathbf{v} \leftarrow \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^3 + \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle$.

- The verifier samples $\mathbf{c}$ over the $\{-1, 0, 1\}$ challenge set C , with each coefficient has $P(\pm 1) = 1/4$, $P(0) = 1/2$.
- The prover runs rejection sampling, and outputs $\bar{\mathbf{z}}' \leftarrow \bar{\mathbf{y}} + \mathbf{c} \cdot \bar{\mathbf{r}}$.
- The verifier checks by ensuring $\bar{\mathbf{z}}'$ being short, $A\hat{\mathbf{z}} = \bar{w} + c \cdot \bar{u}$,

$$\begin{bmatrix} \bar{\mathbf{w}} \\ \mathbf{w}' \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_1 + c \cdot \bar{\mathbf{t}}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_{f_0}^\top \end{bmatrix} \bar{\mathbf{z}}' + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{z} \end{bmatrix},$$

and the shortness check from the product proof in [ALS20] by

$$\mathbf{f}_2 (\mathbf{f}_2 - \mathbf{c}) (\mathbf{f}_2 + \mathbf{c}) + \mathbf{f}_3 + \mathbf{c}\mathbf{f}_4 = \mathbf{v}, \quad (4.33)$$

where $\mathbf{f}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}}' \rangle - \mathbf{c}\mathbf{t}_i$ for $i \in [2, 4]$.

The 2 garbage terms defined in eq. (4.32) derive from the degree-3 shortness check over $\{-1, 0, 1\}$ from the product proof over the masked opening, similar to the one in [ALS20]:

$$\begin{aligned} \mathbf{f}_2 (\mathbf{f}_2 - \mathbf{c}) (\mathbf{f}_2 + \mathbf{c}) &= (\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle - \mathbf{c}\mathbf{s}) (\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle - \mathbf{c}(\mathbf{s} - 1)) (\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle - \mathbf{c}(\mathbf{s} + 1)) \\ &= \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^3 - 3\mathbf{s}\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^2 \mathbf{c} + \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle (3\mathbf{s}^2 - 1) \mathbf{c}^2 - (\mathbf{s}^3 - \mathbf{s}) \mathbf{c}^3. \end{aligned}$$

Hence, when $\mathbf{f}_2$ is the masked opening for $\mathbf{s}$, we let the garbage terms defined in eq. (4.32), such that

$$\begin{aligned} \mathbf{f}_4 &= \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle - \mathbf{c}\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle (3\mathbf{s}^2 - 1), \\ \mathbf{f}_3 &= \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle - \mathbf{c} (\langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle - 3\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^2 \mathbf{s}), \end{aligned}$$

so $\mathbf{f}_3$ can cancel out the computationally indistinguishable masking $\langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle$ in $\mathbf{f}_4$, and the claim $\mathbf{v}$ is

$$\mathbf{f}_2 (\mathbf{f}_2 - \mathbf{c}) (\mathbf{f}_2 + \mathbf{c}) + \mathbf{f}_3 + \mathbf{c}\mathbf{f}_4 = \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^3 + \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle - (\mathbf{s}^3 - \mathbf{s}) \mathbf{c}^3 = \mathbf{v} - (\mathbf{s}^3 - \mathbf{s}) \mathbf{c}^3,$$

where a satisfying $\mathbf{s}$ should be vanishing in $\mathbf{s}^3 - \mathbf{s}$.

The current baseline protocol deviates from [BLS19], as c was needed for both the random linear combination to embed $\mathbf{s}$ to $\mathbf{z}$ (later proved by the linear relation proof), and the shortness check from the product proof, now c only serves the first purpose.

Now we move forward and amplify the soundness for the baseline protocol via automorphism repetition, such that the BDLOP commits to k uniform maskings, 1 message, and 2 garbage terms for shortness check. The masked openings for the garbage terms, and the combined claim, are formed as

$$\begin{aligned} \mathbf{f}_{k+3} &= \langle \bar{\mathbf{b}}_{k+3}, \bar{\mathbf{y}}_1 \rangle - \mathbf{c} \sum_{j \in [0, k-1]} \alpha_{j+1} \sigma^{-j} \left(\langle \bar{\mathbf{b}}_{k+1}, \bar{\mathbf{y}}_{j+1} \rangle (3\mathbf{s}^2 - \mathbf{1}) \right) = \langle \bar{\mathbf{b}}_{k+3}, \bar{\mathbf{y}}_1 \rangle - \mathbf{c} \mathbf{m}_{k+3}, \\ \mathbf{f}_{k+2} &= \langle \bar{\mathbf{b}}_{k+2}, \bar{\mathbf{y}}_1 \rangle - \mathbf{c} \left(\langle \bar{\mathbf{b}}_{k+3}, \bar{\mathbf{y}}_1 \rangle - 3 \sum_{j \in [0, k-1]} \alpha_{j+1} \sigma^{-j} \left(\langle \bar{\mathbf{b}}_{k+1}, \bar{\mathbf{y}}_{j+1} \rangle^2 \mathbf{s} \right) \right) = \langle \bar{\mathbf{b}}_{k+2}, \bar{\mathbf{y}}_1 \rangle - \mathbf{c} \mathbf{m}_{k+2}, \\ \mathbf{v} &= \langle \bar{\mathbf{b}}_{k+2}, \bar{\mathbf{y}}_1 \rangle + \sum_{j \in [0, k-1]} \alpha_{j+1} \sigma^{-j} \left(\langle \bar{\mathbf{b}}_{k+1}, \bar{\mathbf{y}}_{j+1} \rangle^3 \right), \end{aligned} \quad (4.34)$$

such that the random linear combination yields the statement and the invariant for combined claim

$$\begin{aligned} \mathbf{0} &= \sum_{j \in [0, k-1]} \alpha_{j+1} \sigma^{-j} (\mathbf{s}^3 - \mathbf{s}), \\ \mathbf{v} &= \sum_{j \in [0, k-1]} \alpha_{j+1} \sigma^{-j} (\mathbf{f}_{k+1, j+1} (\mathbf{f}_{k+1, j+1} - \sigma^j(\mathbf{c})) (\mathbf{f}_{k+1, j+1} + \sigma^j(\mathbf{c}))) + \mathbf{f}_{k+2} + \mathbf{c} \mathbf{f}_{k+3}. \end{aligned} \quad (4.35)$$

The amplified baseline protocol follows:

- The prover sends out $\bar{w}_i \leftarrow A \hat{\mathbf{y}}_i$ and $\bar{\mathbf{w}}_i \leftarrow \mathbf{B}_0 \bar{\mathbf{y}}_i$, where $\mathbf{y}_i \xleftarrow{\mathbf{R}} R_q$ and $\bar{\mathbf{y}}_i \leftarrow D_\sigma^{Nd}$, and samples $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ to commit

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \\ \vdots \\ \mathbf{t}_k \\ \mathbf{t}_{k+1} \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \\ \vdots \\ \bar{\mathbf{b}}_k^\top \\ \bar{\mathbf{b}}_{k+1}^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_k \\ \mathbf{s} \end{bmatrix}.$$

- The verifier challenges by sending $c_1, \dots, c_k \xleftarrow{\mathbf{R}} \mathbb{Z}_q$ and $\alpha_1, \dots, \alpha_k \xleftarrow{\mathbf{R}} R_q$.
- The prover replies with $\mathbf{z}_i \leftarrow c_i \cdot \mathbf{s} + \mathbf{y}_i$ for $i \in [1, k]$, and starts Π_{lin} and the shortness check in [ALS20] by:
 - The prover derives $\bar{\mathbf{b}}_{f_i} := \bar{\mathbf{b}}_i + c_i \cdot \bar{\mathbf{b}}_{k+1}$, replies $\mathbf{w}'_i \leftarrow \langle \bar{\mathbf{b}}_{f_i}, \bar{\mathbf{y}}_i \rangle$, then it commits

$$\begin{bmatrix} \mathbf{t}_{k+2} \\ \mathbf{t}_{k+3} \end{bmatrix} \leftarrow \begin{bmatrix} \bar{\mathbf{b}}_{k+2}^\top \\ \bar{\mathbf{b}}_{k+3}^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \mathbf{m}_{k+2} \\ \mathbf{m}_{k+3} \end{bmatrix},$$

and replies with $\mathbf{t}_{k+2}, \mathbf{t}_{k+3}$, and the claim $\mathbf{v}$ in eq. (4.34),

- The verifier samples $\mathbf{c}$ over the $\{-1, 0, 1\}$ challenge set $\mathcal{C}$, with each coefficient has $P(\pm 1) = 1/4, P(0) = 1/2$.
- The prover runs rejection sampling, and outputs $\bar{\mathbf{z}}'_i \leftarrow \bar{\mathbf{y}}_i + \sigma^{i-1}(\mathbf{c}) \cdot \bar{\mathbf{r}}$ for $i \in [1, k]$.
- The verifier checks by ensuring $\bar{\mathbf{z}}'_i$ being short, $A \hat{\mathbf{z}}_i = \bar{w}_i + c_i \cdot \bar{u}$ for $i \in [1, k]$,

$$\begin{bmatrix} \bar{\mathbf{w}}_i \\ \mathbf{w}'_i \end{bmatrix} + \sigma^{i-1}(\mathbf{c}) \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_i + c_i \cdot \mathbf{t}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_{f_i}^\top \end{bmatrix} \bar{\mathbf{z}}'_i + \sigma^{i-1}(\mathbf{c}) \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{z}_i \end{bmatrix},$$

and the shortness check from the product proof in [ALS20] by checking the second equality eq. (4.35), where

$$\mathbf{f}_{k+1, j+1} = \langle \bar{\mathbf{b}}_{k+1}, \bar{\mathbf{z}}'_{j+1} \rangle - \sigma^j(\mathbf{c}) \mathbf{t}_{k+1},$$

for $j \in [0, k-1]$, and $\mathbf{f}_j = \langle \bar{\mathbf{b}}_j, \bar{\mathbf{z}}'_1 \rangle - \mathbf{c} \mathbf{t}_j$ for $j \in \{k+2, k+3\}$.

But now the baseline protocol augmented with automorphism repetition is more like the prior attempt to amplify soundness for product proof with automorphism repetition. The shortness proof itself needs 2 garbage terms, while the linearity test for each repetition needs k separate masking terms. A question arises that how can BDLOP commits to only a constant number of messages, while we can still amplify the soundness.

Linear Relation Testing via NTT Averaging. We consider an interactive protocol testing the matrix multiplication statement $A\hat{\mathbf{s}} = \bar{\mathbf{u}}$ without $\mathbf{y} \xleftarrow{R} R_q$: Consider a linear prover, which is bound to a witness chosen before the challenge, and must return the verifier-specified linear evaluation of the witness, then

- The verifier challenges by $\bar{\mathbf{y}} \xleftarrow{R} \mathbb{Z}_q^d$.
- The prover replies with $\mathbf{z} \leftarrow \bar{\mathbf{y}}^\top A\hat{\mathbf{s}}$.
- The verifier checks $\mathbf{z} = \langle \bar{\mathbf{y}}, \bar{\mathbf{u}} \rangle$.

By Schwartz-Zippel, if $A\hat{\mathbf{s}} - \bar{\mathbf{u}} \neq \bar{\mathbf{0}}$, over the randomness of $\bar{\mathbf{y}}$, the probability of $\bar{\mathbf{y}}^\top (A\hat{\mathbf{s}} - \bar{\mathbf{u}}) = 0$ is q^{-1} .

Put it in the polynomial ring computation over the NTT terms, we first consider the lemma in [ENS20].

Lemma 4.13.1 (NTT Averaging Lemma). *Let R_q be defined as eq. (4.14), $\mathbf{f} \in R_q$, and $\hat{\mathbf{f}}_i = \mathbf{f} \bmod (X^{d/\ell} - \xi^i)$ for $i \in \mathbb{Z}_{2\ell}^*$, then when we lift $\hat{\mathbf{f}}_i$ to R_q , we have*

$$\frac{1}{\ell} \sum_{i \in \mathbb{Z}_{2\ell}^*} \hat{\mathbf{f}}_i = \sum_{j \in [0, d/\ell - 1]} f_j X^j.$$

Proof. Recall the eq. (4.17), then noticing

$$\sum_{i \in [0, \ell - 1]} \xi^{(2i+1)v} = \xi^v \frac{\xi^{2\ell v} - 1}{\xi^{2v} - 1} = 0$$

for any $v \in [1, \ell - 1]$, then all the ξ^j in eq. (4.17) with degree $j > 0$ are filtered out in the averaging summation of $\hat{\mathbf{f}}_i$, while the constant coefficients $\{f_i\}_{i \in [0, d/\ell - 1]}$ are scaled up by ℓ . $\square$

Consider $\text{NTT}^{-1}(\bar{\mathbf{y}}^\top A) \in R_q$, then by lemma 4.13.1, $d^{-1} \langle \bar{\mathbf{y}}^\top A, \hat{\mathbf{s}} \rangle$ is in the constant coefficient of $\text{NTT}^{-1}(\bar{\mathbf{y}}^\top A) \cdot \mathbf{s} \in R_q$.

- The prover sends out $\bar{\mathbf{w}} \leftarrow \mathbf{B}_0 \bar{\mathbf{y}}$, where $\bar{\mathbf{y}} \leftarrow D_\sigma^{Nd}$, samples $\mathbf{g} \xleftarrow{R} \{R_q : g_0 = 0\}$ and $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$, to commit

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_1 \\ \bar{\mathbf{t}}_2 \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_2^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{g} \\ \mathbf{s} \end{bmatrix}.$$

- The verifier challenges by sending $\bar{\mathbf{y}} \xleftarrow{R} \mathbb{Z}_q^d$.
- The prover replies with $\mathbf{z} \leftarrow \text{NTT}^{-1}(d\bar{\mathbf{y}}^\top A) \cdot \mathbf{s} + \mathbf{g}$, and starts Π_{lin} and the shortness check in [ALS20] by:
 - The prover derives $\bar{\mathbf{b}}_{f_0} := \bar{\mathbf{b}}_1 + \text{NTT}^{-1}(d\bar{\mathbf{y}}^\top A) \cdot \bar{\mathbf{b}}_2$, replies $\mathbf{w}' \leftarrow \langle \bar{\mathbf{b}}_{f_0}, \bar{\mathbf{y}} \rangle$, then it commits

$$\begin{bmatrix} \bar{\mathbf{t}}_3 \\ \bar{\mathbf{t}}_4 \end{bmatrix} \leftarrow \begin{bmatrix} \bar{\mathbf{b}}_3^\top \\ \bar{\mathbf{b}}_4^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}} \rangle - 3\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^2 \mathbf{s} \\ \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle (3\mathbf{s}^2 - 1) \end{bmatrix},$$

and replies $\bar{\mathbf{t}}_3, \bar{\mathbf{t}}_4$, and the claim $\mathbf{v} \leftarrow \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}} \rangle^3 + \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}} \rangle$.

- The verifier samples $\mathbf{c}$ over the $\{-1, 0, 1\}$ challenge set C , with each coefficient has $P(\pm 1) = 1/4$, $P(0) = 1/2$.

- The prover runs rejection sampling, and outputs $\bar{\mathbf{z}}' \leftarrow \bar{\mathbf{y}} + \mathbf{c} \cdot \bar{\mathbf{r}}$.
- The verifier checks by ensuring $\bar{\mathbf{z}}'$ being short, $z_0 = \langle \bar{\mathbf{y}}, \bar{\mathbf{u}} \rangle$,

$$\begin{bmatrix} \bar{\mathbf{w}} \\ \bar{\mathbf{w}}' \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 + \text{NTT}^{-1}(d\bar{\mathbf{y}}^\top A) \cdot \mathbf{t}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_{f_0}^\top \end{bmatrix} \bar{\mathbf{z}}' + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{z} \end{bmatrix},$$

and the shortness check from the product proof in [ALS20] by eq. (4.33).

Coefficient Filtering via Automorphisms. The prior construction from NTT averaging suffices as a single repetition, but for soundness amplification by automorphism repetition, the masking term over $\{R_q : g_0 = 0\}$ makes it admissible for only a single repetition.

We now want a way to filter out coefficients over R_q via linear operations and automorphisms, or otherwise the new linearity test is no better than the [BLS19] one.

We start by an observation that can filter out all the odd powered coefficients. Since

$$\sigma_{d+1}(X) = X^{d+1} \equiv -X \pmod{(X^d + 1)},$$

then

$$\mathbf{f} + \sigma_{d+1}(\mathbf{f}) = 2 \left(f_0 + f_2 X^2 + \dots + f_{d-2} X^{d-2} \right).$$

Then, in $\log k$ rounds, the automorphisms $\sigma_{d+1}, \sigma_{d/2+1}, \dots, \sigma_{2d/k+1}$ can yield

$$k \left(f_0 + f_k X^k + \dots + f_{d-k} X^{d-k} \right).$$

Moreover, after r round for $r \in [1, \log_2 k]$, all the automorphisms used are indexed by $\{\sigma_{i \cdot 2d/2^r + 1}\}_{i \in [0, 2^r - 1]}$.

Lemma 4.13.2 (Automorphism Filtering Lemma). *For $k \leq d$ and $k \mid d$, then for any $\mathbf{f} \in R_q$,*

$$F_k(\mathbf{f}) = \sum_{i \in [0, k-1]} \sigma_{i \cdot 2d/k+1}(\mathbf{f}) = k \left(f_0 + f_k X^k + \dots + f_{d-k} X^{d-k} \right).$$

Moreover, when $k < d$,

$$F_k(\mathbf{f}) = \sum_{i \in [0, k-1]} \sigma_{2d/k+1}^i(\mathbf{f}). \quad (4.36)$$

Now with lemma 4.13.2, each repetition can be mapped to coefficients of degree in $i \in [0, k-1]$, by multiplying terms X^i , then the masking term can be relaxed to $\mathbf{g} \xleftarrow{R} \{R_q : g_i = 0, i \in [0, k-1]\}$, the responses masked by $\mathbf{g}$ is,

$$\mathbf{z} \leftarrow \mathbf{g} + \sum_{i \in [0, k-1]} k^{-1} X^i F_k \left(\text{NTT}^{-1} \left(d\bar{\mathbf{y}}_i^\top A \right) \cdot \mathbf{s} \right), \quad (4.37)$$

which can be derived from the commitment

$$\mathbf{t} = \mathbf{t}_1 + \sum_{i \in [0, k-1]} k^{-1} X^i F_k \left(\text{NTT}^{-1} \left(d\bar{\mathbf{y}}_i^\top A \right) \cdot \mathbf{t}_2 \right) = \langle \bar{\mathbf{b}}_1, \bar{\mathbf{r}} \rangle + \sum_{i \in [0, k-1]} k^{-1} X^i F_k \left(\text{NTT}^{-1} \left(d\bar{\mathbf{y}}_i^\top A \right) \cdot \langle \bar{\mathbf{b}}_2, \bar{\mathbf{r}} \rangle \right) + \mathbf{z}, \quad (4.38)$$

where $\mathbf{t}_1$ commits to $\mathbf{g}$, and $\mathbf{t}_2$ commits to $\mathbf{s}$.

Yet, there are automorphism terms in eq. (4.38), and the prior vanilla Π_{lin} no longer suffices. We start by stepping back and thinking up a way to prove linear relation over 2 BDLOP commitments from different commitment keys (assuming they are of same size). Supposing they commit to $\mathbf{m}_0$ and $\mathbf{m}_1$, by sampling $\bar{\mathbf{r}}_0, \bar{\mathbf{r}}_1 \leftarrow \chi^{Nd}$, and

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_1 \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{r}}_0 + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_0 \end{bmatrix}, \quad \begin{bmatrix} \bar{\mathbf{t}}_2 \\ \bar{\mathbf{t}}_3 \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_2 \\ \bar{\mathbf{b}}_3^\top \end{bmatrix} \bar{\mathbf{r}}_1 + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{m}_1 \end{bmatrix},$$

and $\mathbf{m}_2 = \mathbf{m}_0 + \mathbf{m}_1$ is the public statement. The protocol works as follows:

- The prover samples $\bar{\mathbf{y}}_0, \bar{\mathbf{y}}_1 \leftarrow D_\sigma^{Nd}$, and outputs $\bar{\mathbf{w}}_0 \leftarrow \mathbf{B}_0 \bar{\mathbf{y}}_0$, $\bar{\mathbf{w}}_1 \leftarrow \mathbf{B}_2 \bar{\mathbf{y}}_1$, and $\mathbf{w}_2 \leftarrow \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_0 \rangle + \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}}_1 \rangle$.
- The verifier challenges with $\mathbf{c}$ over the challenge set $\mathcal{C}$, where each coefficient has $P(\pm 1) = 1/4$, $P(0) = 1/2$.
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}'_0 \leftarrow \bar{\mathbf{y}}_0 + \mathbf{c} \cdot \bar{\mathbf{r}}_0$ and $\bar{\mathbf{z}}'_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{r}}_1$.
- The verifier checks $\bar{\mathbf{z}}'_0$ and $\bar{\mathbf{z}}'_1$ are short, and

$$\begin{bmatrix} \mathbf{B}_0 \\ \mathbf{0} \\ \bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{z}}'_0 + \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_2 \\ \bar{\mathbf{b}}_3^\top \end{bmatrix} \bar{\mathbf{z}}'_1 + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{0}} \\ \mathbf{m}_2 \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{w}}_0 \\ \bar{\mathbf{w}}_1 \\ \mathbf{w}_2 \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \bar{\mathbf{t}}_2 \\ \mathbf{t}_1 + \mathbf{t}_3 \end{bmatrix}$$

The soundness is at most p by eq. (4.18). Automorphism repetition can be applied to have soundness to at most p^k .

With this protocol, we can actually prove the public statement of automorphism filtering $F_k(\mathbf{s})$ for committed $\mathbf{s}$. The intuition was to view $\sigma(\mathbf{B}_0)$ and $\sigma(\bar{\mathbf{b}}_1)$ as different commitment keys, where $\sigma = \sigma_{2d/k+1}$ for some $k < d$ and $k|d$, and the commitment is formed from $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ by

$$\begin{bmatrix} \bar{\mathbf{t}}_0 \\ \mathbf{t}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0} \\ \mathbf{s} \end{bmatrix}.$$

- The prover samples $\bar{\mathbf{y}}_i \leftarrow D_\sigma^{Nd}$, outputs $\bar{\mathbf{w}}_i \leftarrow \mathbf{B}_0 \bar{\mathbf{y}}_i$ for $i \in [0, k-1]$, and

$$\mathbf{w}_k \leftarrow \sum_{i \in [0, k-1]} \sigma^{-i} \left(\langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_i \rangle \right). \quad (4.39)$$

- The verifier challenges with $\mathbf{c}$ over the challenge set $\mathcal{C}$, where each coefficient has $P(\pm 1) = 1/4$, $P(0) = 1/2$.
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}'_i \leftarrow \bar{\mathbf{y}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{r}}$ for $i \in [0, k-1]$.
- The verifier checks $\bar{\mathbf{z}}'_i$ are short for $i \in [0, k-1]$, and

$$\begin{bmatrix} \mathbf{B}_0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{z}}'_0 + \begin{bmatrix} \mathbf{0} \\ \sigma^{-1}(\mathbf{B}_0) \\ \vdots \\ \mathbf{0} \\ \sigma^{-1}(\bar{\mathbf{b}}_1^\top) \end{bmatrix} \sigma^{-1}(\bar{\mathbf{z}}'_1) + \cdots + \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \sigma^{-(k-1)}(\mathbf{B}_0) \\ \sigma^{-(k-1)}(\bar{\mathbf{b}}_1^\top) \end{bmatrix} \sigma^{-(k-1)}(\bar{\mathbf{z}}'_{k-1}) + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \vdots \\ \vdots \\ \bar{\mathbf{0}} \\ F_k(\mathbf{s}) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{w}}_0 \\ \sigma^{-1}(\bar{\mathbf{w}}_1) \\ \vdots \\ \sigma^{-(k-1)}(\bar{\mathbf{w}}_{k-1}) \\ \mathbf{w}_k \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \sigma^{-1}(\bar{\mathbf{t}}_0) \\ \vdots \\ \sigma^{-(k-1)}(\bar{\mathbf{t}}_0) \\ F_k(\mathbf{t}_1) \end{bmatrix}. \quad (4.40)$$

The soundness is at most p by eq. (4.18). Now take a closer look: The prior protocol uses $\mathbf{c}$ to test the automorphism filtering via the linear relation test over distinct BDLOP commitment keys in eq. (4.40). For automorphism repetition, think concretely from $\bar{\mathbf{z}}'_1 \leftarrow \bar{\mathbf{y}}_1 + \sigma(\mathbf{c}) \cdot \bar{\mathbf{r}}$, then the next $\mathbf{w}_{k,1}$ to send before $\mathbf{c}$ is

$$\mathbf{w}_{k,1} \leftarrow \sum_{i \in [0, k-1]} \sigma^{-i} \left(\langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_{i+1 \bmod k} \rangle \right),$$

and eq. (4.40) becomes

$$\begin{bmatrix} \mathbf{B}_0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \bar{\mathbf{b}}_1^\top \end{bmatrix} \bar{\mathbf{z}}'_1 + \begin{bmatrix} \mathbf{0} \\ \sigma^{-1}(\mathbf{B}_0) \\ \vdots \\ \mathbf{0} \\ \sigma^{-1}(\bar{\mathbf{b}}_1^\top) \end{bmatrix} \sigma^{-1}(\bar{\mathbf{z}}'_2) + \cdots + \begin{bmatrix} \mathbf{0} \\ \vdots \\ \mathbf{0} \\ \sigma^{-(k-1)}(\mathbf{B}_0) \\ \sigma^{-(k-1)}(\bar{\mathbf{b}}_1^\top) \end{bmatrix} \sigma^{-(k-1)}(\bar{\mathbf{z}}'_0) + \sigma(\mathbf{c}) \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \vdots \\ \vdots \\ \bar{\mathbf{0}} \\ F_k(\mathbf{s}) \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{w}}_1 \\ \sigma^{-1}(\bar{\mathbf{w}}_2) \\ \vdots \\ \sigma^{-(k-1)}(\bar{\mathbf{w}}_0) \\ \mathbf{w}_{k,1} \end{bmatrix} + \sigma(\mathbf{c}) \cdot \begin{bmatrix} \bar{\mathbf{t}}_0 \\ \sigma^{-1}(\bar{\mathbf{t}}_0) \\ \vdots \\ \sigma^{-(k-1)}(\bar{\mathbf{t}}_0) \\ F_k(\mathbf{t}_1) \end{bmatrix}.$$

By induction,

$$\mathbf{w}_{k,j} \leftarrow \sum_{i \in [0, k-1]} \sigma^{-i} \left(\langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_{i+j \bmod k} \rangle \right) = \sigma^j(\mathbf{w}_k). \quad (4.41)$$

Supposing a slot under the slots grouped by $(X^k - \xi^k)$ is wrong, then the adversary needs to guess $\mathbf{c} \bmod (X^k - \xi^k)$ to convince the verifier, with the same strategy for the affine forging in weak opening in the note for [ALS20], by fixing the pre-challenge $\bar{\mathbf{w}}$ for the relevant error term, yielding the soundness of at most p^k , where p is eq. (4.18) for the guessing probability of each coefficient in coarse $(X^k - \xi^k)$.

The rest is to prove the automorphism filtering of each $\text{NTT}^{-1}(d\bar{\gamma}_j^T A)\mathbf{s}$ for $j \in [0, k-1]$ by stacking them up over X^j , normalizing by k^{-1} , masking by $\mathbf{g}$, and eventually eq. (4.37). To carry on the automorphism filtering test derived from the linearity test, each $\mathbf{w}_{k,j}$ derived from eq. (4.39) is indexed by $j \in [0, k-1]$ for each $\bar{\gamma}_j$ by

$$\mathbf{w}_{k,j} \leftarrow \sum_{i \in [0, k-1]} \sigma^{-i} \left(\langle \text{NTT}^{-1}(d\bar{\gamma}_j^T A) \cdot \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_i \rangle \right).$$

Correspondingly, the verifier can derive

$$\mathbf{u}_{k,j} \leftarrow \sum_{i \in [0, k-1]} \sigma^{-i} \left(\langle \text{NTT}^{-1}(d\bar{\gamma}_j^T A) \cdot \bar{\mathbf{b}}_2, \bar{\mathbf{z}}'_i \rangle \right),$$

and on the third round prover response, by eq. (4.41), for each $i \in [0, k-1]$ automorphic challenges $\sigma^i(\mathbf{c})$,

$$\mathbf{w}_{f_i} \leftarrow \langle \bar{\mathbf{b}}_1, \bar{\mathbf{y}}_i \rangle + \sum_{j \in [0, k-1]} k^{-1} X^j \sigma^i(\mathbf{w}_{k,j}), \quad (4.42)$$

and the verifier can derive correspondingly

$$\mathbf{u}_{f_i} \leftarrow \langle \bar{\mathbf{b}}_1, \bar{\mathbf{z}}'_i \rangle + \sum_{j \in [0, k-1]} k^{-1} X^j \sigma^i(\mathbf{u}_{k,j}), \quad (4.43)$$

such that for each $i \in [0, k-1]$ automorphic challenges $\sigma^i(\mathbf{c})$,

$$\underbrace{\mathbf{w}_{f_i} + \sigma^i(\mathbf{c}) \cdot \left(\mathbf{t}_1 + \sum_{j \in [0, k-1]} k^{-1} X^j F_k \left(\text{NTT}^{-1}(d\bar{\gamma}_j^T A) \cdot \mathbf{t}_2 \right) \right)}_{\mathbf{t}} = \underbrace{\mathbf{u}_{f_i} + \sigma^i(\mathbf{c}) \cdot \left(\mathbf{g} + \sum_{j \in [0, k-1]} k^{-1} X^j F_k \left(\text{NTT}^{-1}(d\bar{\gamma}_j^T A) \cdot \mathbf{s} \right) \right)}_{\mathbf{z}}, \quad (4.44)$$

where $\mathbf{t}, \mathbf{z}$ were in eq. (4.38) and eq. (4.37).

Again, now the masked openings for the garbage terms, and the combined claim, are formed as

$$\begin{aligned} \mathbf{f}_4 &= \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}}_0 \rangle - \mathbf{c} \sum_{j \in [0, k-1]} \alpha_j \sigma^{-j} \left(\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_j \rangle (3s^2 - 1) \right) = \langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}}_0 \rangle - \mathbf{c} \mathbf{m}_4, \\ \mathbf{f}_3 &= \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}}_0 \rangle - \mathbf{c} \left(\langle \bar{\mathbf{b}}_4, \bar{\mathbf{y}}_0 \rangle - 3 \sum_{j \in [0, k-1]} \alpha_j \sigma^{-j} \left(\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_j \rangle^2 s \right) \right) = \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}}_0 \rangle - \mathbf{c} \mathbf{m}_3, \\ \mathbf{v} &= \langle \bar{\mathbf{b}}_3, \bar{\mathbf{y}}_0 \rangle + \sum_{j \in [0, k-1]} \alpha_j \sigma^{-j} \left(\langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_j \rangle^3 \right), \end{aligned} \quad (4.45)$$

such that the random linear combination yields the statement and the invariant for combined claim

$$\begin{aligned} \mathbf{0} &= \sum_{j \in [0, k-1]} \alpha_j \sigma^{-j} (s^3 - s), \\ \mathbf{v} &= \sum_{j \in [0, k-1]} \alpha_j \sigma^{-j} (\mathbf{f}_{2,j} (\mathbf{f}_{2,j} - \sigma^j(\mathbf{c})) (\mathbf{f}_{2,j} + \sigma^j(\mathbf{c}))) + \mathbf{f}_3 + \mathbf{c} \mathbf{f}_4, \end{aligned} \quad (4.46)$$

where $\mathbf{f}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{z}}'_0 \rangle - \mathbf{c} \mathbf{t}_i = \langle \bar{\mathbf{b}}_i, \bar{\mathbf{y}}_0 \rangle - \mathbf{c} \mathbf{m}_i$ for $i \in \{3, 4\}$, each $\mathbf{f}_{2,j}$ for $j \in [0, k-1]$ is

$$\mathbf{f}_{2,j} = \langle \bar{\mathbf{b}}_2, \bar{\mathbf{z}}'_j \rangle - \sigma^j(\mathbf{c}) \mathbf{t}_2 = \langle \bar{\mathbf{b}}_2, \bar{\mathbf{y}}_j \rangle - \sigma^j(\mathbf{c}) \mathbf{s},$$

and the second equality holds if the prover is honest.

The protocol for a single matrix multiplication claim follows:

- The prover sends $\bar{\mathbf{w}}_i \leftarrow \mathbf{B}_0 \bar{\mathbf{y}}_i$, where $\bar{\mathbf{y}}_i \leftarrow D_\sigma^{Nd}$, samples $\mathbf{g} \xleftarrow{R} \{R_q : g_j = 0, j \in [0, k-1]\}$ and $\bar{\mathbf{r}} \leftarrow \chi^{Nd}$ to commit

$$\begin{bmatrix} \mathbf{t}_0 \\ \mathbf{t}_1 \\ \mathbf{t}_2 \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{B}_0 \\ \bar{\mathbf{b}}_1^\top \\ \bar{\mathbf{b}}_2^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \mathbf{0} \\ \mathbf{g} \\ \mathbf{s} \end{bmatrix}.$$

- The verifier challenges by sending $\bar{\gamma}_0, \dots, \bar{\gamma}_{k-1} \xleftarrow{R} \mathbb{Z}_q^d$ and $\alpha_0, \dots, \alpha_{k-1} \xleftarrow{R} R_q$.
- The prover replies $\mathbf{z}$ in eq. (4.37), and starts Π_{lin} and the shortness check in [ALS20] by:
 - The prover replies each $\mathbf{w}_{f_i}$ for $i \in [0, k-1]$ in eq. (4.42), then it commits

$$\begin{bmatrix} \mathbf{t}_3 \\ \mathbf{t}_4 \end{bmatrix} \leftarrow \begin{bmatrix} \bar{\mathbf{b}}_3^\top \\ \bar{\mathbf{b}}_4^\top \end{bmatrix} \bar{\mathbf{r}} + \begin{bmatrix} \mathbf{m}_3 \\ \mathbf{m}_4 \end{bmatrix},$$

and replies with $\mathbf{t}_3, \mathbf{t}_4$, and the claim $\mathbf{v}$ in eq. (4.45),

- The verifier samples $\mathbf{c}$ over the $\{-1, 0, 1\}$ challenge set C , with each coefficient has $P(\pm 1) = 1/4$, $P(0) = 1/2$.
- The prover runs rejection sampling, and outputs $\bar{\mathbf{z}}'_i \leftarrow \bar{\mathbf{y}}_i + \sigma^i(\mathbf{c}) \cdot \bar{\mathbf{r}}$ for $i \in [0, k-1]$.
- The verifier checks $\bar{\mathbf{z}}'_i$ being short, $z_j = \langle \bar{\gamma}_j, \bar{\mathbf{u}} \rangle$ for $j \in [0, k-1]$.

It then derives $\mathbf{u}_{f_i}$ for $i \in [0, k-1]$ in eq. (4.43), $\mathbf{t}$ in eq. (4.38), checks

$$\begin{aligned} \bar{\mathbf{w}}_i + \sigma^i(\mathbf{c}) \cdot \mathbf{t}_0 &= \mathbf{B}_0 \bar{\mathbf{z}}'_i, \\ \mathbf{w}_{f_i} + \sigma^i(\mathbf{c}) \cdot \mathbf{t} &= \mathbf{u}_{f_i} + \sigma^i(\mathbf{c}) \cdot \mathbf{z} \end{aligned}$$

for $i \in [0, k-1]$, which is eq. (4.44) and eq. (4.40),

Finally, it checks the shortness of $\hat{\mathbf{s}}$ by the second equality eq. (4.46).

Cost Comparison Against [BLS19]. We compare the published noninteractive constructions of [BLS19, ENS20] by counting transmitted elements and their coefficient widths.

Let n be the total witness length, $A \in \mathbb{Z}_q^{m \times n}$, d be the R_q degree, and $b_q = \lceil \log_2 q \rceil$. We retain N for the opening-vector width and μ for the top commitment rank. Counts exclude the public statement and commitment key. The following are the papers' historical parameter choices for approximately 128-bit security.

Parameter	BLS19	ENS20
Matrix dimensions $m \times n$	1024×2048	1024×2048
Coefficient width b_q	32	32
Ring degree d	2048	128
Witness polynomials n/d	1	16
Repetitions	$t = 4, t' = 3$	$k = 4$
Commitment ranks (λ, μ)	(1, 1)	(10, 9)

In particular, ENS separates the witness length from the ring degree: Their example uses 16 witness polynomials.

BLS: transmitted objects. In the optimized Fiat–Shamir protocol of [BLS19, Algorithms 2–3 and Section 3.4], each of the t commitment vectors has five ring elements, with one shared across repetitions. There are t masked witness replies and tt' short opening vectors of width 6.

Component	Count	Bits per coefficient or choice
Commitments	$1 + 4t$ ring elements	b_q
Masked witness replies	t ring elements	b_q
Short opening replies	$6tt'$ ring elements	b_z
Scalar challenges	t scalars	b_q
Monomial challenges	tt' choices	$\lceil \log_2 \ell \rceil$

Here $d = n$ and $\ell = 2d$ for the power-of-two benchmark. The proof size in bits is therefore

$$S_{\text{BLS}} = (1 + 5t)db_q + 6tt'db_z + t(b_q + t'\lceil \log_2 \ell \rceil). \quad (4.47)$$

For $\sigma = 919.13$, the short-coefficient width is $b_z = \lceil \log_2(12\sigma) \rceil = 14$. Substituting the benchmark parameters gives

Component	Calculation in bytes	Size
Full part	$21 \cdot 2048 \cdot 32/8 = 172,032$	168 KiB
Short part	$72 \cdot 2048 \cdot 14/8 = 258,048$	252 KiB
Challenges	$4(32 + 3 \cdot 12)/8 = 34$	34 bytes
Total	430,114	420.0332 KiB

ENS: transmitted objects. For n/d witness polynomials, the committed message count and opening width are

$$L = n/d + 3, \quad N = \lambda + \mu + L.$$

The additional 3 messages are the $\mathbf{g}$ linear-test mask, and the two shortness garbage terms.

Besides the $\mu + L$ commitment elements, ENS transmits one masked linear-test polynomial, called h in the paper. For our single-polynomial notation it is the publicly centered response

$$h = \mathbf{z} - \sum_{j=0}^{k-1} X^j \langle \bar{\gamma}_j, \bar{u} \rangle,$$

so its first k coefficients are known to be zero. After reconstructing the first messages, the remaining short replies contain kN ring elements. With $b_z = \lceil \log_2(2\delta_1) \rceil$, [ENS20, Section 4.2, Equation (11)] gives the basic bit count

$$\begin{aligned} S_{\text{ENS}}^{\text{basic}} &= (\mu + n/d + 4)db_q + k(\lambda + \mu + n/d + 3)db_z \\ &= [n + (\mu + 4)d]b_q + k[n + (\lambda + \mu + 3)d]b_z. \end{aligned} \quad (4.48)$$

This leading count omits the 256-bit challenge seed and the saving from the k known coefficients of h .

For bookkeeping, insert $\delta_1 = 2^{18}$ from [ENS20, Appendix B.1], so $b_z = 19$, $L = 19$, and $N = 38$:

Component	Calculation in bytes	Size
Full part	$29 \cdot 128 \cdot 32/8 = 14,848$	14.5 KiB
Short part	$152 \cdot 128 \cdot 19/8 = 46,208$	45.125 KiB
Basic total	61,056	59.625 KiB

Although ENS has more short polynomials (152 versus 72), it has 19456 short coefficients, compared to BLS's 147456. ENS and BLS contain 29 and 21 full ring elements, respectively, or 3712 and 43008 coefficients.

4.14 Lattice Proofs for Exact Euclidean Norm Bounds and Quadratic Relations

One of the motivations of [LNP22] is that, the prior works [BLS19, ENS20] for $A\hat{\mathbf{s}} = \bar{\mathbf{u}}$ restrict each entry of $\hat{\mathbf{s}}$ being over $\{-1, 0, 1\}$, and thus committing the $\mathbf{m} = \text{NTT}^{-1}(\hat{\mathbf{s}})$, which can be arbitrarily long, can be expensive with [BDL⁺18]. If we want to work with $\hat{\mathbf{s}}$ with larger ℓ_∞ -norm, then more garbage terms are needed in the commitment.

The question they asked was: Is NTT packing all that necessary, as both concrete efficiency and communication complexity can be improved if no NTT packing is involved. In [LNP22], the encoding over CRT slots [BLS19, ALS20, ENS20] is not used, messages are committed in coefficient form. In this case, R_q is not necessarily fully splitting.

Challenge Set. There are many ways to construct challenge sets for the R_q :

- In [LS18], the corollary 4.12.3 gives a short challenge set, such that all nonzero pairwise differences are invertible, but by corollaries 4.12.2 and 4.12.3, the modulus grows with the number of splitting factors and the norm bound of the set of short elements, and some computation is in table 4.12.1;
- In [BLS19], the eq. (4.7) suffices;
- In [ALS20], the anti-concentration lemma 4.12.5 gives the max point probability of a CRT slot, and the probability for pairwise difference being non-invertible is upper bounded by a union bound. The extractor can extract CRT slot by slot, so it is not required to be globally pairwise difference invertible.

In [LNP22], they found another set of elements in R_q that is invertible.

Lemma 4.14.1. *By corollary 4.12.2, let $p \equiv 5 \pmod{8}$ be a prime, then any nonzero $\mathbf{c} \in R_p$ such that $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$ is invertible.*

Proof. By corollary 4.12.2, if $p \equiv 5 \pmod{8}$ is a prime, then

$$R_p = \mathbb{Z}_p[X]/(X^d + 1) \cong \mathbb{Z}_p[X]/(X^{d/2} - r) \times \mathbb{Z}_p[X]/(X^{d/2} + r),$$

r being the 4th primitive root of unity, splitting 2 irreducible factors. Since $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$, then satisfying elements are

$$\mathbf{c} = c_0 + \sum_{i \in [1, d/2-1]} c_i (X^i - X^{d-i}) = c_0 + \sum_{i \in [1, d/2-1]} c_i X^i - \sum_{i \in [1, d/2-1]} c_i X^{d-i},$$

as any $X^i - X^{d-i}$ satisfies $\sigma_{-1}(X^i - X^{d-i}) = X^i - X^{d-i}$. Hence, for each factor $(X^{d/2} \pm r)$, we have

$$\mathbf{c} \equiv c_0 + \sum_{i \in [1, d/2-1]} (c_i \pm r c_{d/2-i}) X^i \pmod{(X^{d/2} \pm r)}.$$

We claim at least one of $c_i - r c_{d/2-i}$ and $c_{d/2-i} - r c_i$ is nonzero, given that at least one of $c_i, c_{d/2-i}$ is nonzero: Supposing $c_i = r c_{d/2-i}$, then $(1 - r^2) c_{d/2-i} \neq 0$. The case for $c_i + r c_{d/2-i}$ and $c_{d/2-i} + r c_i$ is the same idea.

Hence, a nonzero $\mathbf{c}$ is nonzero at both factors if $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$. □

The lemma 4.14.1 can be applied to the R_q , where $q = \prod q_i$, and $q_1 < \dots < q_k$, with all $q_i \equiv 5 \pmod{8}$ are primes, then $R_q = \prod R_{q_i}$, and thus each R_{q_i} splits into 2 irreducible factors. Hence, any nonzero $\mathbf{c} \in R_q$ satisfying $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$ and $0 < \|\mathbf{c}\|_\infty < q_1$, $\mathbf{c}$ is invertible.

The short invertible set of elements from lemma 4.14.1 shares a similar vibe to [LS18]. The similarity is generally “If it is just too short, it will not wrap around in any of the CRT slots”.

- The current one finds a pattern of invertibility $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$ that applies to any R_{q_i} splitting into 2 irreducible factors in the above setting, then any nonzero $\mathbf{c} \in R_q$ such that $\sigma_{-1}(\mathbf{c}) = \mathbf{c}$, that is shorter than q_1 in ℓ_∞ -norm is invertible, as the coefficient representations within $[-q_1 + 1, q_1 - 1]$ can remain the same for all $[-q_i + 1, q_i - 1]$ where $q_i \geq q_1$.
- The [LS18] one is, considering $R_q \cong \prod \mathbb{Z}_q[X]/(f_i(X))$, then for the ideal lattice $\Lambda = \{\mathbf{x} \in R_m : \mathbf{x} \equiv 0 \pmod{(f_i, q)}\}$, if nonzero $\mathbf{x} \in R_q$ is shorter than $\lambda_1(\Lambda)$ in ℓ_2 -norm, then $\mathbf{x}$ cannot belong to the CRT kernel lattice Λ .

Now we move on and consider the norm growth when multiplying with challenge set elements, namely we want to bound $\|\mathbf{c}\bar{\mathbf{r}}\|_2$. The main idea of the next lemma is to upper bound the operator norm of $\text{rot}(\mathbf{c}) \in \mathbb{Z}^{d \times d}$, such that

$$\|\mathbf{c}\bar{\mathbf{r}}\| \leq \left(\max_{\|\mathbf{x}\|_2=1} \|\mathbf{c}\mathbf{x}\|_2 \right) \|\bar{\mathbf{r}}\|_2 = \|\text{rot}(\mathbf{c})\|_2 \|\bar{\mathbf{r}}\|_2.$$

Lemma 4.14.2. *Let $\bar{\mathbf{r}} \in R^\ell$ and $\mathbf{c} \in R$, then for any power-of-two κ , we have*

$$\|\mathbf{c}\bar{\mathbf{r}}\|_2 \leq \|\sigma_{-1}(\mathbf{c}^\kappa) \mathbf{c}^\kappa\|_1^{1/(2\kappa)} \cdot \|\bar{\mathbf{r}}\|_2.$$

Proof. We first write $\mathbf{C} = \text{rot}(\mathbf{c}) \in \mathbb{Z}^{d \times d}$, then

$$\|\mathbf{C}\|_2^2 = \max_{\|\bar{\mathbf{x}}\|_2=1} \|\mathbf{C}\bar{\mathbf{x}}\|_2^2 = \max_{\|\bar{\mathbf{x}}\|_2=1} \bar{\mathbf{x}}^\top \mathbf{C}^\top \mathbf{C} \bar{\mathbf{x}}, \quad (4.49)$$

where the last equality holds by squaring up and summing up all entries in $\mathbf{C}\bar{\mathbf{x}}$, then by eq. (4.49), $\mathbf{C}^\top \mathbf{C}$ is positive semidefinite. By the spectral theorem, $\mathbf{C}^\top \mathbf{C}$ has an orthonormal eigenbasis, that eigenvectors $\{\bar{\mathbf{v}}_i\}_{i \in [1,d]}$ are mutually perpendicular and of ℓ_2 -norm being 1. Hence, for $\bar{\mathbf{x}}$ with $\|\bar{\mathbf{x}}\|_2 = 1$,

$$\mathbf{C}^\top \mathbf{C} \bar{\mathbf{x}} = \mathbf{C}^\top \mathbf{C} \sum_{i \in [1,d]} \alpha_i \bar{\mathbf{v}}_i = \sum_{i \in [1,d]} \alpha_i \lambda_i \bar{\mathbf{v}}_i, \quad (4.50)$$

and by the $\|\bar{\mathbf{x}}\|_2 = 1$,

$$1 = \left\| \sum_{i \in [1,d]} \alpha_i \bar{\mathbf{v}}_i \right\|_2^2 = \sum_{i \in [1,d]} \alpha_i^2 + \sum_{i \neq j \in [1,d]} \alpha_i \alpha_j \langle \bar{\mathbf{v}}_i, \bar{\mathbf{v}}_j \rangle = \sum_{i \in [1,d]} \alpha_i^2. \quad (4.51)$$

Therefore, conditioned that eq. (4.51), we have

$$\|\mathbf{C}^\top \mathbf{C}\|_2 = \max_{\|\bar{\mathbf{x}}\|=1} \|\mathbf{C}^\top \mathbf{C} \bar{\mathbf{x}}\|_2 = \max_{\{\alpha_i\}_{i \in [1,d]}} \left\| \sum_{i \in [1,d]} \alpha_i \lambda_i \bar{\mathbf{v}}_i \right\|_2 = \max_{\{\alpha_i\}_{i \in [1,d]}} \left(\sum_{i \in [1,d]} \alpha_i^2 \lambda_i^2 \right)^{1/2} = \max_{i \in [1,d]} |\lambda_i|,$$

where the third equality holds by eq. (4.51). Moreover, continuing from eq. (4.50), we finish eq. (4.49) by

$$\|\mathbf{C}\|_2^2 = \max_{\{\alpha_i\}_{i \in [1,d]}} \sum_{i \in [1,d]} \alpha_i^2 \lambda_i^2 = \max_{i \in [1,d]} \lambda_i^2 = \|\mathbf{C}^\top \mathbf{C}\|_2,$$

where the last equality holds by the positive semidefiniteness of $\mathbf{C}^\top \mathbf{C}$, such that each λ_i is nonnegative. Continuing on, since $\mathbf{A} = \mathbf{C}^\top \mathbf{C}$ is positive semidefinite, then $\mathbf{A}^\top \mathbf{A} = \mathbf{A}^2$ as $\mathbf{A}$ is diagonal symmetric, then

$$\|\mathbf{C}\|_2^{2\kappa} = \|\mathbf{C}^\top \mathbf{C}\|_2^\kappa = \|(\mathbf{C}^\top \mathbf{C})^\kappa\|_2.$$

For $\mathbf{u}, \mathbf{v} \in R$, another observation for $\mathbf{u}\mathbf{v} = \text{rot}(\mathbf{u})\bar{\mathbf{v}}$ is,

$$\text{rot}(\mathbf{u})\bar{\mathbf{v}} = \sum_{i \in [0,d-1]} u_i \text{rot}(X^i) \bar{\mathbf{v}},$$

and therefore by triangle inequality,

$$\|\text{rot}(\mathbf{u})\bar{\mathbf{v}}\|_2 \leq \sum_{i \in [0,d-1]} |u_i| \|\text{rot}(X^i) \bar{\mathbf{v}}\|_2 = \|\mathbf{u}\|_1 \cdot \|\bar{\mathbf{v}}\|_2. \quad (4.52)$$

Hence, for $\mathbf{v}$ with $\|\mathbf{v}\|_2 = 1$ and coefficients over $\mathbb{R}$, by operator norm eq. (4.49), $\|\text{rot}(\mathbf{u})\|_2 \leq \|\mathbf{u}\|_1$. The rest is immediate by the fact that $\mathbf{C}^\top = \text{rot}(\sigma_{-1}(\mathbf{c}))$, then

$$(\mathbf{C}^\top \mathbf{C})^2 = \text{rot}(\sigma_{-1}(\sigma_{-1}(\mathbf{c})\mathbf{c})) \text{rot}(\sigma_{-1}(\mathbf{c})\mathbf{c}) = \text{rot}(\sigma_{-1}(\mathbf{c})^2 \mathbf{c}^2) = \text{rot}(\sigma_{-1}(\mathbf{c}^2) \mathbf{c}^2),$$

and by induction,

$$\|\mathbf{C}\|_2^{2\kappa} = \|\mathbf{C}^\top \mathbf{C}\|_2^\kappa = \|(\mathbf{C}^\top \mathbf{C})^\kappa\|_2 = \|\text{rot}(\sigma_{-1}(\mathbf{c}^\kappa) \mathbf{c}^\kappa)\|_2 \leq \|\sigma_{-1}(\mathbf{c}^\kappa) \mathbf{c}^\kappa\|_1.$$

□

Moreover, by eq. (4.52), we have $\|\mathbf{u}\mathbf{v}\|_1 \leq \|\mathbf{u}\|_1 \|\mathbf{v}\|_1$, then $\|\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa\|_1^{1/(2\kappa)}$ is monotonically nonincreasing by

$$\|\sigma_{-1}(\mathbf{c}^{2\kappa})\mathbf{c}^{2\kappa}\|_1^{1/(4\kappa)} = \|(\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa)^2\|_1^{1/(4\kappa)} \leq \|\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa\|_1^{1/(2\kappa)}.$$

On the other hand, by Cauchy-Schwartz, $\|\mathbf{v}\|_1 \leq \sqrt{d} \|\mathbf{v}\|_2$, then

$$\|\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa\|_1^{1/(2\kappa)} \leq d^{1/(4\kappa)} \|\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa\|_2^{1/(2\kappa)} \leq d^{1/(4\kappa)} \|(\mathbf{C}^\top \mathbf{C})^\kappa\|_2^{1/(2\kappa)} = d^{1/(4\kappa)} \|\mathbf{C}\|_2,$$

where the second inequality holds by letting the rotational matrix multiplying against $\bar{\mathbf{e}}_0 = [1, 0, \dots, 0]^\top$.

Therefore,

$$\|\mathbf{C}\|_2 \leq \|\sigma_{-1}(\mathbf{c}^{2\kappa})\mathbf{c}^{2\kappa}\|_1^{1/(4\kappa)} \leq \|\sigma_{-1}(\mathbf{c}^\kappa)\mathbf{c}^\kappa\|_1^{1/(2\kappa)} \leq d^{1/(4\kappa)} \|\mathbf{C}\|_2,$$

and the upper bound by lemma 4.14.2 is monotonically nonincreasing, and is approaching the operator norm $\|\mathbf{C}\|_2$.

Given lemma 4.14.2 as an estimation for the ℓ_2 operator norm of $\mathbf{c} \in R$, we construct the challenge set C defined over the elements specified by lemma 4.14.1, that for a power-of-two k ,

$$C := \{\mathbf{c} \in S_\kappa^\sigma : \|\sigma_{-1}(\mathbf{c}^k)\mathbf{c}^k\|_1^{1/(2k)} \leq \eta\}, \quad (4.53)$$

where $S_\kappa^\sigma = \{\mathbf{c} \in S_\kappa : \sigma(\mathbf{c}) = \mathbf{c}\}$ for $\sigma \in \text{Aut}(R_q)$, and S_κ is defined in eq. (4.5), that $\kappa \leq (q_1 - 1)/2$. Eventually, the goal is to make the ℓ_∞ -norm of challenges to be bounded by $(q_1 - 1)/2$, and estimate $\|\text{rot}(\mathbf{c})\|_2 \leq \eta$, such that $\|\mathbf{c}\bar{\mathbf{r}}\|_2 \leq \eta \|\bar{\mathbf{r}}\|_2$.

Rejection Sampling and Parameters. We haven't get to the extended MLWE from [LNS21], so we use the vanilla rejection sampling from [Lyu12] throughout this note. In [LNP22], the masking randomness for [BDL⁺18] and [Ajt96] is sampled over a Gaussian distribution, and the next lemma help describe the concentration property of a sample.

Lemma 4.14.3 (Gaussian Tail Bound [Ban93]). *Let $\bar{\mathbf{z}} \leftarrow D_s^{md}$, then for $t > 1$,*

$$\Pr \left[\|\bar{\mathbf{z}}\|_2 > t \cdot s \sqrt{md} \right] < \left(t \exp \left(\frac{1-t^2}{2} \right) \right)^{md}.$$

In [LNP22], the lemma 4.14.3 is parameterized with $t = \sqrt{2}$, then

$$\Pr \left[\|\bar{\mathbf{z}}\|_2 > s \sqrt{2md} \right] < (2/e)^{md/2}.$$

We continue on describing the rejection sampling from the same algorithm in [Lyu12], but parameterized differently from [BLS19]. The key idea is using the outputting probability threshold by [Lyu12] for $\bar{\mathbf{z}} \leftarrow \bar{\mathbf{y}} + \bar{\mathbf{v}}$ where $\bar{\mathbf{y}} \leftarrow D_s^{md}$

$$\min \left(\frac{1}{M} \cdot \frac{D_s^{md}(\bar{\mathbf{z}})}{D_{\bar{\mathbf{v}},s}^{md}(\bar{\mathbf{z}})}, 1 \right) = \min \left(\frac{1}{M} \cdot \frac{\rho_s(\bar{\mathbf{z}})}{\rho_s(\bar{\mathbf{z}} - \bar{\mathbf{v}})}, 1 \right), \quad (4.54)$$

where ρ_s is defined in eq. (4.9), and since we are working over Gaussian distributions, the equality holds. Moreover, over $\bar{\mathbf{z}} \leftarrow D_s^{md}$,

$$\Pr \left[D_s^{md}(\bar{\mathbf{z}}) \leq M \cdot D_{\bar{\mathbf{v}},s}^{md}(\bar{\mathbf{z}}) \right] \geq 1 - 2^{-128}, \quad (4.55)$$

then the threshold eq. (4.54) uses the left term with overwhelming probability: The M is chosen by $s := \gamma \|\bar{\mathbf{v}}\|_2$, then

$$\frac{\rho_s(\bar{\mathbf{z}})}{\rho_s(\bar{\mathbf{z}} - \bar{\mathbf{v}})} = \exp \left(\frac{-2\langle \bar{\mathbf{z}}, \bar{\mathbf{v}} \rangle + \|\bar{\mathbf{v}}\|_2^2}{2s^2} \right) \leq \exp \left(\frac{28s \|\bar{\mathbf{v}}\|_2 + \|\bar{\mathbf{v}}\|_2^2}{2s^2} \right) = \exp \left(\frac{14}{\gamma} + \frac{1}{2\gamma^2} \right) = M, \quad (4.56)$$

(if $\|\bar{\mathbf{v}}\|_2 \leq T$, then $s := \gamma T$), and the second inequality holds with overwhelming probability by Lemma 4.3 of [Lyu12]:

$$\Pr [|\langle \bar{\mathbf{z}}, \bar{\mathbf{v}} \rangle| \geq 14s \|\bar{\mathbf{v}}\|_2] \leq 2 \exp(-98) \leq 2^{-128},$$

thus eq. (4.55) holds, the accepting probability is at least $(1 - 2^{-128})/M$, and conditioned on the sample's acceptance, the statistical distance of the real distribution is within 2^{-128} against the simulated distribution.

ABDLOP Commitment Scheme. Consider combining Ajtai [Ajt96] and BDLOP [BDL⁺18] together, then we have ABDLOP commitment [LNP22]. Recall eq. (4.19), where the commitment key is over $R_q^{(\mu+L) \times N}$:

- Binding is reduced to MSIS $_{\mu, N, B}$.
- Hiding is reduced to MLWE $_{\lambda, \mu+L, \chi}$, where commitment randomness is drawn from distribution χ^{Nd} .

Onwards, we use m_1 to denote the length of the Ajtai message, and m_2 to denote the length of the BDLOP randomness. The combination is done by stacking an Ajtai commitment to the SIS part in the BDLOP commitment:

$$\begin{bmatrix} \bar{\mathbf{t}}_A \\ \bar{\mathbf{t}}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{m}} \end{bmatrix} \in R_q^{\mu+L},$$

such that binding is reduced to MSIS $_{\mu, m_1+m_2, B}$, while hiding is reduced to MLWE $_{\lambda, \mu+L, \chi}$.

Assume $\|\bar{\mathbf{s}}_1\|_2 \leq \alpha$ and $\|\bar{\mathbf{s}}_2\|_\infty \leq \nu$. By $\|\text{rot}(\mathbf{c})\|_2 \leq \eta$ from lemma 4.14.2, the ℓ_2 -norm bound for the centering shifts are $T_1 := \eta\alpha$ and $T_2 := \eta\nu\sqrt{m_2d}$. Choose $\gamma_1, \gamma_2 > 0$, and thus $s_i := \gamma_i T_i$ and $M_i := \exp(14/\gamma_i + 1/(2\gamma_i^2))$ in eq. (4.56).

Combining the proof of opening for Ajtai [Ajt96] and [BDL⁺18], a three move ABDLOP proof of opening follows:

- The prover samples $\bar{\mathbf{y}}_1 \leftarrow D_{s_1}^{m_1d}$ and $\bar{\mathbf{y}}_2 \leftarrow D_{s_2}^{m_2d}$, and sends over $\bar{\mathbf{w}} \leftarrow \mathbf{A}_1 \bar{\mathbf{y}}_1 + \mathbf{A}_2 \bar{\mathbf{y}}_2$.
- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{s}}_1$ and $\bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{y}}_2 + \mathbf{c} \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_1\|_2 \leq s_1\sqrt{2m_1d}$, $\|\bar{\mathbf{z}}_2\|_2 \leq s_2\sqrt{2m_2d}$, and

$$\bar{\mathbf{w}} + \mathbf{c} \cdot \bar{\mathbf{t}}_A = \mathbf{A}_1 \bar{\mathbf{z}}_1 + \mathbf{A}_2 \bar{\mathbf{z}}_2.$$

For MSIS hardness reduction for binding, the concrete parameter instantiated is $B = 8\eta\sqrt{s_1^2(2m_1d) + s_2^2(2m_2d)}$. Similar to eqs. (4.10) and (4.21), supposing we have a pair of accepting $(\bar{\mathbf{w}}, \bar{\mathbf{v}}, \mathbf{c}, \bar{\mathbf{z}}_1, \bar{\mathbf{z}}_2)$ and $(\bar{\mathbf{w}}, \bar{\mathbf{v}}, \mathbf{c}', \bar{\mathbf{z}}'_1, \bar{\mathbf{z}}'_2)$, then let $\Delta\mathbf{c} := \mathbf{c} - \mathbf{c}'$, $\Delta\bar{\mathbf{z}}_1 := \bar{\mathbf{z}}_1 - \bar{\mathbf{z}}'_1$, and $\Delta\bar{\mathbf{z}}_2 := \bar{\mathbf{z}}_2 - \bar{\mathbf{z}}'_2$, an extracted relaxed opening has $(\Delta\mathbf{c}, \Delta\bar{\mathbf{z}}_1, \Delta\bar{\mathbf{z}}_2, \bar{\mathbf{m}})$, such that

$$\Delta\mathbf{c} \begin{bmatrix} \bar{\mathbf{t}}_A \\ \bar{\mathbf{t}}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \Delta\bar{\mathbf{z}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \end{bmatrix} \Delta\bar{\mathbf{z}}_2 + \Delta\mathbf{c} \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{m}} \end{bmatrix}, \quad (4.57)$$

where the term $\Delta\mathbf{c} \cdot \bar{\mathbf{m}} = \Delta\mathbf{c} \cdot \bar{\mathbf{t}}_B - \mathbf{B} \Delta\bar{\mathbf{z}}_2$. Given relaxed openings $(\Delta\mathbf{c}, \Delta\bar{\mathbf{z}}_1, \Delta\bar{\mathbf{z}}_2, \bar{\mathbf{m}})$ and $(\Delta\mathbf{c}', \Delta\bar{\mathbf{z}}'_1, \Delta\bar{\mathbf{z}}'_2, \bar{\mathbf{m}}')$, that at least one of $\Delta\mathbf{c}' \Delta\bar{\mathbf{z}}_1 \neq \Delta\mathbf{c} \Delta\bar{\mathbf{z}}'_1$ and $\bar{\mathbf{m}} \neq \bar{\mathbf{m}}'$ is true, then

$$\bar{\mathbf{0}} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \left(\Delta\mathbf{c}' \cdot \begin{bmatrix} \Delta\bar{\mathbf{z}}_1 \\ \Delta\bar{\mathbf{z}}_2 \end{bmatrix} - \Delta\mathbf{c} \cdot \begin{bmatrix} \Delta\bar{\mathbf{z}}'_1 \\ \Delta\bar{\mathbf{z}}'_2 \end{bmatrix} \right) + \Delta\mathbf{c} \Delta\mathbf{c}' \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{m}} - \bar{\mathbf{m}}' \end{bmatrix},$$

which gives a solution to the MSIS instance $[\mathbf{A}_1 \mid \mathbf{A}_2] \in R_q^{\mu \times (m_1+m_2)}$ in either case. Since for $\bar{\boldsymbol{\beta}} := [\Delta\bar{\mathbf{z}}_1^\top \mid \Delta\bar{\mathbf{z}}_2^\top]^\top$,

$$\|\bar{\boldsymbol{\beta}}\|_2 \leq 2\sqrt{s_1^2(2m_1d) + s_2^2(2m_2d)},$$

and by lemma 4.14.2, $\|\text{rot}(\mathbf{c})\|_2 \leq \eta$ for a $\mathbf{c} \in C$, thus $\|\text{rot}(\Delta\mathbf{c})\|_2 \leq 2\eta$, which concludes the proof for B derivation.

Now consider a linear relation proof over R_q . Consider instantiating from just Ajtai commitment for $\mathbf{R}\bar{\mathbf{s}}_1 = \bar{\mathbf{u}}$.

- The prover samples $\bar{\mathbf{y}}_1 \leftarrow D_{s_1}^{m_1d}$ and sends over $\bar{\mathbf{w}} \leftarrow \mathbf{A}_1 \bar{\mathbf{y}}_1$ and $\bar{\mathbf{v}} \leftarrow \mathbf{R}\bar{\mathbf{y}}_1$.
- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{s}}_1$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_1\|_2 \leq s_1\sqrt{2m_1d}$ and

$$\begin{aligned} \bar{\mathbf{w}} + \mathbf{c} \cdot \bar{\mathbf{t}} &= \mathbf{A}_1 \bar{\mathbf{z}}_1, \\ \bar{\mathbf{v}} + \mathbf{c} \cdot \bar{\mathbf{u}} &= \mathbf{R}\bar{\mathbf{z}}_1. \end{aligned}$$

Continuing on, consider the BDLOP commitment for $\mathbf{R}\bar{\mathbf{m}} = \bar{\mathbf{u}}$, which is basically Π_{lin} .

- The prover samples $\bar{\mathbf{y}}_2 \leftarrow D_{s_2}^{m_2d}$ and sends over $\bar{\mathbf{w}} \leftarrow \mathbf{A}_2\bar{\mathbf{y}}_2$ and $\bar{\mathbf{v}} \leftarrow \mathbf{R}\mathbf{B}\bar{\mathbf{y}}_2$.
- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{y}}_2 + \mathbf{c} \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_2\|_2 \leq s_2\sqrt{2m_2d}$ and

$$\begin{bmatrix} \bar{\mathbf{w}} \\ \bar{\mathbf{v}} \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_A \\ \mathbf{R}\bar{\mathbf{t}}_B \end{bmatrix} = \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{R}\mathbf{B} \end{bmatrix} \bar{\mathbf{z}}_2 + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{u}} \end{bmatrix}.$$

Hence, we can somehow merge them together for the statement $\mathbf{R}_1\bar{\mathbf{s}}_1 + \mathbf{R}_2\bar{\mathbf{m}} = \bar{\mathbf{u}}$ by

- The prover samples $\bar{\mathbf{y}}_1 \leftarrow D_{s_1}^{m_1d}$, $\bar{\mathbf{y}}_2 \leftarrow D_{s_2}^{m_2d}$ and sends over $\bar{\mathbf{w}} \leftarrow \mathbf{A}_1\bar{\mathbf{y}}_1 + \mathbf{A}_2\bar{\mathbf{y}}_2$ and $\bar{\mathbf{v}} \leftarrow \mathbf{R}_2\mathbf{B}\bar{\mathbf{y}}_2 - \mathbf{R}_1\bar{\mathbf{y}}_1$.
- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{s}}_1$ and $\bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{y}}_2 + \mathbf{c} \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_1\|_2 \leq s_1\sqrt{2m_1d}$, $\|\bar{\mathbf{z}}_2\|_2 \leq s_2\sqrt{2m_2d}$, and

$$\begin{bmatrix} \bar{\mathbf{w}} \\ \bar{\mathbf{v}} \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_A \\ \mathbf{R}_2\bar{\mathbf{t}}_B \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{R}_1\bar{\mathbf{z}}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{z}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{R}_2\mathbf{B} \end{bmatrix} \bar{\mathbf{z}}_2 + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{u}} \end{bmatrix}.$$

Hence, a relaxed opening $(\Delta\mathbf{c}, \Delta\bar{\mathbf{z}}_1, \Delta\bar{\mathbf{z}}_2, \bar{\mathbf{m}})$ satisfies eq. (4.57), $\|\Delta\bar{\mathbf{z}}_1\|_2 \leq 2s_1\sqrt{2m_1d}$, $\|\Delta\bar{\mathbf{z}}_2\|_2 \leq 2s_2\sqrt{2m_2d}$, and

$$\mathbf{R}_1\Delta\bar{\mathbf{z}}_1 + \mathbf{R}_2(\Delta\mathbf{c} \cdot \bar{\mathbf{m}}) = \Delta\mathbf{c} \cdot \bar{\mathbf{u}}.$$

Now from the linear relation proof over R_q , consider the linear relation proof over $\mathbb{Z}_q$. Consider the inner product performed over the coefficients of polynomial ring via automorphism σ_{-1} , that the constant coefficient of $\sigma_{-1}(\mathbf{a})\mathbf{b}$ is the inner product $\langle \bar{\mathbf{a}}, \bar{\mathbf{b}} \rangle$ of coefficients in $\mathbf{a} = \sum a_i X^i$ and $\mathbf{b} = \sum b_i X^i$, allowing for encoding the $\mathbb{Z}_q$ linear relation over the coefficients: $\langle \sigma_{-1}(\bar{\mathbf{r}}_1), \bar{\mathbf{s}}_1 \rangle + \langle \sigma_{-1}(\bar{\mathbf{r}}_2), \bar{\mathbf{m}} \rangle = \bar{\mathbf{u}}$. Augmenting the BDLOP message for maskings like [ENS20] by

$$\begin{bmatrix} \bar{\mathbf{t}}_A \\ \bar{\mathbf{t}}_B \\ \bar{\mathbf{t}}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \\ \mathbf{B}_g \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \bar{\mathbf{0}} \\ \bar{\mathbf{m}} \\ \bar{\mathbf{g}} \end{bmatrix} \in R_q^{\mu+L+\rho},$$

where $\bar{\mathbf{g}} \in R_q^\rho$ are the masking terms. Start by proving with the Π_{lin} for R_q in ABDLOP, with a single repetition:

- The prover commits $\mathbf{g} \xleftarrow{R} \{R_q : g_0 = 0\}$ by $\mathbf{t}_g \leftarrow \langle \bar{\mathbf{b}}_g, \bar{\mathbf{s}}_2 \rangle + \mathbf{g}$, and sends $\mathbf{t}_g$ over.
- The verifier challenges by sampling $\gamma \xleftarrow{R} \mathbb{Z}_q$.
- The prover replies $\mathbf{u}' \leftarrow \mathbf{g} + \gamma \cdot \bar{\mathbf{u}}$, and starts the Π_{lin} with ABDLOP for $\langle \gamma\sigma_{-1}(\bar{\mathbf{r}}_1), \bar{\mathbf{s}}_1 \rangle + \langle \gamma\sigma_{-1}(\bar{\mathbf{r}}_2), \bar{\mathbf{m}} \rangle + \mathbf{g} = \mathbf{u}'$ by
 - The prover samples $\bar{\mathbf{y}}_1 \leftarrow D_{s_1}^{m_1d}$, $\bar{\mathbf{y}}_2 \leftarrow D_{s_2}^{m_2d}$ and replies

$$\begin{aligned} \bar{\mathbf{w}} &\leftarrow \mathbf{A}_1\bar{\mathbf{y}}_1 + \mathbf{A}_2\bar{\mathbf{y}}_2, \\ \mathbf{v} &\leftarrow \left(\gamma\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B} + \bar{\mathbf{b}}_g^\top \right) \bar{\mathbf{y}}_2 - \gamma\sigma_{-1}(\bar{\mathbf{r}}_1)^\top \bar{\mathbf{y}}_1. \end{aligned}$$

- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{s}}_1$ and $\bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{y}}_2 + \mathbf{c} \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_1\|_2 \leq s_1\sqrt{2m_1d}$, $\|\bar{\mathbf{z}}_2\|_2 \leq s_2\sqrt{2m_2d}$,

$$\begin{bmatrix} \bar{\mathbf{w}} \\ \mathbf{v} \end{bmatrix} + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{t}}_A \\ \gamma\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B} + \bar{\mathbf{t}}_g \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{0}} \\ \gamma\sigma_{-1}(\bar{\mathbf{r}}_1)^\top \bar{\mathbf{z}}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{z}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \gamma\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B} + \bar{\mathbf{b}}_g^\top \end{bmatrix} \bar{\mathbf{z}}_2 + \mathbf{c} \cdot \begin{bmatrix} \bar{\mathbf{0}} \\ \mathbf{u}' \end{bmatrix},$$

and the constant term satisfies $u'_0 = \gamma u_0$.

For multiple linear relations, apply random linear combination over $\mathbb{Z}_q$ under the same mask $\mathbf{g}$, with soundness q_1^{-1} , say any CRT slot under the split by $\mathbb{Z}_q = \prod \mathbb{Z}_{q_i}$ is wrong, the uniformly random challenges that are independently sampled can fix up that slot with probability at most q_1^{-1} by Schwartz-Zippel. The random linear combination over multiple linear relations is similar to the style in [ALS20] for multiple product relations.

When repetitions are greater than once, a straightforward way is to make ρ exactly the number of repetitions. But a closer observation would be, given that R_q at least has 2 automorphisms, we can use the automorphism filtering from [ENS20], such that a masking can hold 2 repetitions by $\mathbf{g} \xleftarrow{R} \{R_q : g_0 = g_1 = 0\}$.

We show by 1 linear relation with 2 γ challenges setting:

- The prover commits $\mathbf{g} \xleftarrow{R} \{R_q : g_0 = g_1 = 0\}$ by $\mathbf{t}_g \leftarrow \langle \bar{\mathbf{b}}_g, \bar{\mathbf{s}}_2 \rangle + \mathbf{g}$, and sends $\mathbf{t}_g$ over.
- The verifier challenges by sampling $\gamma_0, \gamma_1 \xleftarrow{R} \mathbb{Z}_q$.
- The prover replies

$$\mathbf{u}' \leftarrow \mathbf{g} + \frac{\gamma_0 + \gamma_1 X}{2} \cdot F_2(\mathbf{u}),$$

and starts the Π_{lin} with ABDLOP for

$$\mathbf{g} + \frac{\gamma_0 + \gamma_1 X}{2} \cdot F_2(\langle \sigma_{-1}(\bar{\mathbf{r}}_1), \bar{\mathbf{s}}_1 \rangle + \langle \sigma_{-1}(\bar{\mathbf{r}}_2), \bar{\mathbf{m}} \rangle),$$

where F_2 is defined in eq. (4.36), by

- The prover samples $\bar{\mathbf{y}}_{1,0}, \bar{\mathbf{y}}_{1,1} \leftarrow D_{s_1}^{m_1 d}, \bar{\mathbf{y}}_{2,0}, \bar{\mathbf{y}}_{2,1} \leftarrow D_{s_2}^{m_2 d}$, computes

$$\alpha \leftarrow 2^{-1} \sum_{i \in [0,1]} \sigma_{d+1}^{-i} \left(\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B} \bar{\mathbf{y}}_{2,i} - \sigma_{-1}(\bar{\mathbf{r}}_1)^\top \bar{\mathbf{y}}_{1,i} \right),$$

and replies

$$\begin{aligned} \bar{\mathbf{w}}_i &\leftarrow \mathbf{A}_1 \bar{\mathbf{y}}_{1,i} + \mathbf{A}_2 \bar{\mathbf{y}}_{2,i}, \\ \mathbf{v}_i &\leftarrow (\gamma_0 + \gamma_1 X) \sigma_{d+1}^i(\alpha) + \bar{\mathbf{b}}_g^\top \bar{\mathbf{y}}_{2,i}, \end{aligned}$$

where σ_{d+1} is due to lemma 4.14.1 on R_q splitting.

- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_{1,i} \leftarrow \bar{\mathbf{y}}_{1,i} + \sigma_{d+1}^i(\mathbf{c}) \cdot \bar{\mathbf{s}}_1$ and $\bar{\mathbf{z}}_{2,i} \leftarrow \bar{\mathbf{y}}_{2,i} + \sigma_{d+1}^i(\mathbf{c}) \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_{1,i}\|_2 \leq s_1 \sqrt{2m_1 d}$, $\|\bar{\mathbf{z}}_{2,i}\|_2 \leq s_2 \sqrt{2m_2 d}$, automorphism filtering checks for Π_{lin} :

$$\begin{aligned} &\begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ -\frac{\gamma_0 + \gamma_1 X}{2} \sigma_{-1}(\bar{\mathbf{r}}_1)^\top \end{bmatrix} \bar{\mathbf{z}}_{1,i} + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{0} \\ \frac{\gamma_0 + \gamma_1 X}{2} \sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B} + \bar{\mathbf{b}}_g^\top \end{bmatrix} \bar{\mathbf{z}}_{2,i} \\ &+ \begin{bmatrix} \mathbf{0} \\ \sigma_{d+1}^{-1}(\mathbf{A}_1) \\ -\frac{\gamma_0 + \gamma_1 X}{2} \sigma_{d+1}^{-1}(\sigma_{-1}(\bar{\mathbf{r}}_1)^\top) \end{bmatrix} \sigma_{d+1}^{-1}(\bar{\mathbf{z}}_{1,1-i}) + \begin{bmatrix} \mathbf{0} \\ \sigma_{d+1}^{-1}(\mathbf{A}_2) \\ \frac{\gamma_0 + \gamma_1 X}{2} \sigma_{d+1}^{-1}(\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \mathbf{B}) \end{bmatrix} \sigma_{d+1}^{-1}(\bar{\mathbf{z}}_{2,1-i}) + \sigma_{d+1}^i(\mathbf{c}) \cdot \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{u}' \end{bmatrix} \\ &= \begin{bmatrix} \bar{\mathbf{w}}_i \\ \sigma_{d+1}^{-1}(\bar{\mathbf{w}}_{1-i}) \\ \mathbf{v}_i \end{bmatrix} + \sigma_{d+1}^i(\mathbf{c}) \begin{bmatrix} \bar{\mathbf{t}}_A \\ \sigma_{d+1}^{-1}(\bar{\mathbf{t}}_A) \\ \frac{\gamma_0 + \gamma_1 X}{2} F_2(\sigma_{-1}(\bar{\mathbf{r}}_2)^\top \bar{\mathbf{t}}_B) + \mathbf{t}_g \end{bmatrix} \end{aligned}$$

holds for each $i \in [0, 1]$, and the constant term satisfies $u'_i = \gamma_i u_0$.

Quadratic Relation Proof. Now also consider for quadratic relation proof over ABDLOP, extending from [ALS20]. Considering the quadratic relation proof over BDLOP by the computationally indistinguishable masking in eq. (4.26), the quadratic relation can be extended to arbitrary pair of subset sums of the BDLOP messages by

$$\begin{aligned}\bar{\mathbf{f}}^T \mathbf{R} \bar{\mathbf{f}} &= (\bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2)^T \mathbf{R} (\bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2) \\ &= \mathbf{c}^2 \bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{m}} - \mathbf{c} (\bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{m}} + \bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{y}}) + \bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{y}},\end{aligned}\quad (4.58)$$

where $\bar{\mathbf{f}} = \bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2$ the computationally indistinguishable masked opening, and $\bar{\mathbf{y}} = \mathbf{B}\bar{\mathbf{y}}_2$. Continuing on, consider the linear terms and the constant terms, we have the relation

$$\bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{m}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{m}} + \mathbf{r}_0 = \mathbf{0}, \quad (4.59)$$

the corresponding masked opening should give

$$\begin{aligned}\bar{\mathbf{f}}^T \mathbf{R} \bar{\mathbf{f}} + \mathbf{c} \bar{\mathbf{r}}_1^T \bar{\mathbf{f}} + \mathbf{c}^2 \mathbf{r}_0 &= (\bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2)^T \mathbf{R} (\bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2) + \mathbf{c} \bar{\mathbf{r}}_1^T (\bar{\mathbf{c}}\bar{\mathbf{m}} - \mathbf{B}\bar{\mathbf{y}}_2) + \mathbf{c}^2 \mathbf{r}_0 \\ &= \mathbf{c}^2 (\bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{m}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{m}} + \mathbf{r}_0) - \mathbf{c} (\bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{m}} + \bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{y}}) + \bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{y}},\end{aligned}\quad (4.60)$$

such that the degree-1 term $\mathbf{g}_0 = -(\bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{m}} + \bar{\mathbf{m}}^T \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{y}})$ is the additional garbage term to commit.

Now we consider the Ajtai part. Instead of the computationally indistinguishable opening $\bar{\mathbf{f}} = \bar{\mathbf{c}}\bar{\mathbf{m}} - \bar{\mathbf{y}}$, we have statistically indistinguishable opening $\bar{\mathbf{z}} = \mathbf{y}_1 + \bar{\mathbf{c}}\bar{\mathbf{s}}_1$. Hence, the computation structure in eqs. (4.58) and (4.60) remains if we combine the masks by

$$\bar{\mathbf{y}} = \begin{bmatrix} -\bar{\mathbf{y}}_1 \\ \mathbf{B}\bar{\mathbf{y}}_2 \end{bmatrix}, \quad \bar{\mathbf{s}} = \begin{bmatrix} \bar{\mathbf{s}}_1 \\ \bar{\mathbf{m}} \end{bmatrix}, \quad (4.61)$$

and $\bar{\mathbf{f}} = \mathbf{c} \cdot \bar{\mathbf{s}} - \bar{\mathbf{y}}$, then eq. (4.60) becomes

$$\bar{\mathbf{f}}^T \mathbf{R} \bar{\mathbf{f}} + \mathbf{c} \bar{\mathbf{r}}_1^T \bar{\mathbf{f}} + \mathbf{c}^2 \mathbf{r}_0 = \mathbf{c}^2 (\bar{\mathbf{s}}^T \mathbf{R} \bar{\mathbf{s}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{s}} + \mathbf{r}_0) - \mathbf{c} (\bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{s}} + \bar{\mathbf{s}}^T \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{y}}) + \bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{y}},$$

such that the degree-1 term $\mathbf{g}_0 = -(\bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{s}} + \bar{\mathbf{s}}^T \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{r}}_1^T \bar{\mathbf{y}})$ is the additional garbage term to commit.

Hence, a prototype for quadratic relation over ABDLOP commitment scheme follows:

- The prover samples $\bar{\mathbf{y}}_1 \leftarrow D_{s_1}^{m_1 d}$, $\bar{\mathbf{y}}_2 \leftarrow D_{s_2}^{m_2 d}$, commits to $\mathbf{g}_0$ by $\mathbf{t}_g \leftarrow \bar{\mathbf{b}}_g^T \bar{\mathbf{s}}_2 + \mathbf{g}_0$, and replies

$$\begin{aligned}\bar{\mathbf{w}} &\leftarrow \mathbf{A}_1 \bar{\mathbf{y}}_1 + \mathbf{A}_2 \bar{\mathbf{y}}_2, \\ \mathbf{v} &\leftarrow \bar{\mathbf{y}}^T \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{b}}_g^T \bar{\mathbf{y}}_2.\end{aligned}$$

- The verifier challenges by $\mathbf{c} \leftarrow C$ defined in eq. (4.53).
- The prover runs rejection sampling, and replies $\bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{y}}_1 + \mathbf{c} \cdot \bar{\mathbf{s}}_1$ and $\bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{y}}_2 + \mathbf{c} \cdot \bar{\mathbf{s}}_2$.
- The verifier checks shortness $\|\bar{\mathbf{z}}_1\|_2 \leq s_1 \sqrt{2m_1 d}$, $\|\bar{\mathbf{z}}_2\|_2 \leq s_2 \sqrt{2m_2 d}$, derives $\mathbf{f}_g \leftarrow \bar{\mathbf{b}}_g^T \bar{\mathbf{z}}_2 - \mathbf{c} \mathbf{t}_g$ and

$$\bar{\mathbf{f}} \leftarrow \begin{bmatrix} \bar{\mathbf{z}}_1 \\ \mathbf{c} \bar{\mathbf{t}}_B - \mathbf{B} \bar{\mathbf{z}}_2 \end{bmatrix},$$

and checks by

$$\begin{aligned}\mathbf{A}_1 \bar{\mathbf{z}}_1 + \mathbf{A}_2 \bar{\mathbf{z}}_2 &= \bar{\mathbf{w}} + \mathbf{c} \bar{\mathbf{t}}_A \\ \bar{\mathbf{f}}^T \mathbf{R} \bar{\mathbf{f}} + \mathbf{c} \bar{\mathbf{r}}_1^T \bar{\mathbf{f}} + \mathbf{c}^2 \mathbf{r}_0 + \mathbf{f}_g &= \mathbf{v}.\end{aligned}$$

Additionally, noticing the challenges from C has $\sigma(\mathbf{c}) = \mathbf{c}$, then the quadratic relation can be extended over all $\bar{\mathbf{s}}$'s automorphisms, that the message got expanded to

$$\bar{\mathbf{s}} = \left[\begin{array}{c} \{\sigma^i(\bar{\mathbf{s}}_1)\}_{i \in [0, k-1]} \\ \{\sigma^i(\bar{\mathbf{m}})\}_{i \in [0, k-1]} \end{array} \right], \quad (4.62)$$

as each opening has $\sigma^j(\bar{\mathbf{z}}_i) = \sigma^j(\bar{\mathbf{y}}_i) + \mathbf{c} \cdot \sigma^j(\bar{\mathbf{s}}_i)$ by self-automorphism $\mathbf{c} = \sigma(\mathbf{c})$, which is similar to the technique for automorphism filtering in eq. (4.40). Then eq. (4.61) can be changed to

$$\bar{\mathbf{y}} = \left[\begin{array}{c} \{\sigma^i(-\bar{\mathbf{y}}_1)\}_{i \in [0, k-1]} \\ \{\sigma^i(\mathbf{B}\bar{\mathbf{y}}_2)\}_{i \in [0, k-1]} \end{array} \right], \quad \bar{\mathbf{s}} = \left[\begin{array}{c} \{\sigma^i(\bar{\mathbf{s}}_1)\}_{i \in [0, k-1]} \\ \{\sigma^i(\bar{\mathbf{m}})\}_{i \in [0, k-1]} \end{array} \right]. \quad (4.63)$$

If the prover is honest, the convincing probability is approximately

$$\frac{1}{M_1 M_2} = \exp \left(\frac{14}{\gamma_1} + \frac{14}{\gamma_2} + \frac{1}{2\gamma_1^2} + \frac{1}{2\gamma_2^2} \right)^{-1}$$

by eq. (4.55), as we are just using [Lyu12] rejection sampling.

The knowledge soundness analysis is similar to the heavy row argument lemma 4.8.8, but the extractor behavior is described and developed in [ACK21]. We begin by the expected draws for negative hypergeometric distribution.

Lemma 4.14.4. *Given a bin with N balls of which M are marked. Then the attempts to draw uniformly at random from the bin without replacement until $k \leq M$ marked balls are drawn is distributed according to the negative hypergeometric distribution. The expected number of draws is $k(N+1)/(M+1)$.*

Proof. The sequence of draws without replacement can be viewed as a permutation of N balls, and we are interested in the number of draws up to the first k marked ones. For an unmarked ball, it can be in one of the $M+1$ positions separated by M marked ones, and there are k of them before the k^{th} marked one. By the linearity of expectation, the expected number of draws is

$$k + (N - M) \cdot \frac{k}{M + 1} = \frac{k(N + 1)}{M + 1}.$$

□

An immediate application is, let $\mathbf{H} \in \{0, 1\}^{R \times N}$ be the binary matrix for R rows indexed by the prover's randomness, and the N columns indexed by verifier challenges, then let the i^{th} row has accepting dencity ε_i , supposing the first draw over the row is accepting, then drawing for the next 2 acceptance is expected to be $2/\varepsilon_i$ by lemma 4.14.4. Supposing the extractor aborts if the first draw is rejected, or otherwise it samples uniformly at random the remaining entries of the row without replacement until three acceptances or the row is exhausted, and supposing there are at least 3 acceptance on the row, then

$$\mathbb{E}[Q \mid i^{\text{th}} \text{ row}] = 1 \cdot (1 - \varepsilon_i) + (1 + 2/\varepsilon_i) \cdot \varepsilon_i = 3,$$

and supposing there are less than 3 acceptance on the row, then

$$\mathbb{E}[Q \mid i^{\text{th}} \text{ row}] = 1 \cdot (1 - \varepsilon_i) + N \cdot \varepsilon_i \leq 3.$$

Averaging over the prover randomness, we thus have

$$\mathbb{E}[Q] = \mathbb{E}[\mathbb{E}[Q \mid i^{\text{th}} \text{ row}]] \leq 3, \quad (4.64)$$

where the success probability is

$$\sum_{k \in [3, N]} \frac{k}{N} \cdot \delta_k = \sum_{k \in [1, N]} \frac{k}{N} \cdot \delta_k - \frac{\delta_1}{N} - \frac{2\delta_2}{N} = \varepsilon - \frac{\delta_1 + 2\delta_2}{N}, \quad (4.65)$$

where δ_k is the fraction of rows that have k acceptances, ε is the fraction of acceptance in $\mathbf{H}$, and eq. (4.65) is at least $\varepsilon - 3/N$ by a coarse analysis, and $\varepsilon - 2/N$ by a finer one. Hence, if checks on an entry of $\mathbf{H}$ takes T , and $\varepsilon > 2/N$, then extracting three accepting transcripts by restarting over fresh prover randomness, is at most expected $3T/(\varepsilon - 2/N)$.

Now, in expected $3T/(\varepsilon - 2/N)$ time, we have 3 accepting transcripts forked from the same prover randomness, hence the same first prover message, by $(\bar{\mathbf{w}}, \mathbf{v}, \mathbf{t}_g, \bar{\mathbf{z}}_{1,i}, \bar{\mathbf{z}}_{2,i})$ for $i \in [0, 2]$.

Let $\Delta \mathbf{c} \leftarrow \mathbf{c}_0 - \mathbf{c}_1$, $\Delta \bar{\mathbf{z}}_1 \leftarrow \bar{\mathbf{z}}_{1,0} - \bar{\mathbf{z}}_{1,1}$, $\Delta \bar{\mathbf{z}}_2 \leftarrow \bar{\mathbf{z}}_{2,0} - \bar{\mathbf{z}}_{2,1}$, and $\bar{\mathbf{s}}_1 \leftarrow \Delta \bar{\mathbf{z}}_1 \cdot (\Delta \mathbf{c})^{-1}$, $\bar{\mathbf{s}}_2 \leftarrow \Delta \bar{\mathbf{z}}_2 \cdot (\Delta \mathbf{c})^{-1}$. Further, we have $\bar{\mathbf{m}} \leftarrow \bar{\mathbf{t}}_B - \mathbf{B} \bar{\mathbf{s}}_2$ and $\mathbf{g}_0 \leftarrow \mathbf{t}_g - \bar{\mathbf{b}}_g^\top \bar{\mathbf{s}}_2$, then

$$\begin{bmatrix} \bar{\mathbf{t}}_A \\ \bar{\mathbf{t}}_B \\ \bar{\mathbf{t}}_g \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bar{\mathbf{s}}_1 + \begin{bmatrix} \mathbf{A}_2 \\ \mathbf{B} \\ \bar{\mathbf{b}}_g^\top \end{bmatrix} \bar{\mathbf{s}}_2 + \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{m}} \\ \mathbf{g}_0 \end{bmatrix},$$

and $\|\Delta \mathbf{c}\|_\infty \leq 2\kappa$ by eq. (4.53), $\|\Delta \bar{\mathbf{z}}_i\|_2 \leq 2s_i \sqrt{2m_i d}$ for $i \in [1, 2]$.

Continuing on, we can have $\bar{\mathbf{y}}_1 \leftarrow \bar{\mathbf{z}}_{1,0} - \mathbf{c}_0 \cdot \bar{\mathbf{s}}_1$, $\bar{\mathbf{y}}_2 \leftarrow \bar{\mathbf{z}}_{2,0} - \mathbf{c}_0 \cdot \bar{\mathbf{s}}_2$, and under the same interpolation,

$$\begin{aligned} \bar{\mathbf{z}}_{1,0} - \mathbf{c}_0 \cdot \bar{\mathbf{s}}_1 &= \bar{\mathbf{z}}_{1,1} - \mathbf{c}_1 \cdot \bar{\mathbf{s}}_1, \\ \bar{\mathbf{z}}_{2,0} - \mathbf{c}_0 \cdot \bar{\mathbf{s}}_2 &= \bar{\mathbf{z}}_{2,1} - \mathbf{c}_1 \cdot \bar{\mathbf{s}}_2. \end{aligned}$$

Applying the derived $\bar{\mathbf{s}}_1$, $\bar{\mathbf{s}}_2$, and $\mathbf{c}_2$, against $\bar{\mathbf{z}}_{1,2}$ and $\bar{\mathbf{z}}_{2,2}$, we have either $\bar{\mathbf{y}}_1$ and $\bar{\mathbf{y}}_2$ being the masking randomness, or we find a $\text{MSIS}_{n, m_1+m_2, B}$ solution for $[\mathbf{A}_1 \mathbf{A}_2]$, where B was defined near eq. (4.57), that for $i \in [1, 2]$,

$$\bar{\mathbf{u}}_i \leftarrow \Delta \mathbf{c} (\bar{\mathbf{z}}_{i,2} - \mathbf{c}_2 \bar{\mathbf{s}}_i - \bar{\mathbf{y}}_i) = \Delta \mathbf{c} (\bar{\mathbf{z}}_{i,2} - \bar{\mathbf{z}}_{i,0}) + (\mathbf{c}_0 - \mathbf{c}_2) \Delta \bar{\mathbf{z}}_i$$

solves $[\mathbf{A}_1 \mathbf{A}_2]$ by $\mathbf{A}_1 \bar{\mathbf{u}}_1 + \mathbf{A}_2 \bar{\mathbf{u}}_2 = \mathbf{0}$.

Onwards assuming the MSIS is respected, then derive eq. (4.63), $\mathbf{g}'_0 = -(\bar{\mathbf{y}}^\top \mathbf{R} \bar{\mathbf{s}} + \bar{\mathbf{s}}^\top \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{r}}_1^\top \bar{\mathbf{y}})$, such that

$$\begin{bmatrix} 1 & \mathbf{c}_0 & \mathbf{c}_0^2 \\ 1 & \mathbf{c}_1 & \mathbf{c}_1^2 \\ 1 & \mathbf{c}_2 & \mathbf{c}_2^2 \end{bmatrix} \begin{bmatrix} \bar{\mathbf{y}}^\top \mathbf{R} \bar{\mathbf{y}} + \bar{\mathbf{b}}_g^\top \bar{\mathbf{y}}_2 - \mathbf{v} \\ \mathbf{g}'_0 - \mathbf{g}_0 \\ \bar{\mathbf{s}}^\top \mathbf{R} \bar{\mathbf{s}} + \bar{\mathbf{r}}_1^\top \bar{\mathbf{s}} + \mathbf{r}_0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

must hold for the extracted variables. Since pairwise challenge differences are invertible, the Vandermonde matrix is invertible; hence every coefficient is zero, including the claimed quadratic relation.

Many Quadratic Relations Proof. To instantiate the proof for multiple quadratic relations, we sample uniformly random $\mu_1, \dots, \mu_N \xleftarrow{R} R_q$ and apply linear combination over quadratic relations defined by eq. (4.59).

For soundness analysis, it is similar to the linear combination for multiple linear relations. Recall the soundness error q_1^{-1} derived by $\gamma_1, \dots, \gamma_\ell \xleftarrow{R} \mathbb{Z}_q$, it was due to $\mathbb{Z}_q \cong \prod \mathbb{Z}_{q_i}$, then the random linear combination over a CRT slot being fixed to 0 has probability at most q_1^{-1} . Now consider random linear combination over $\mu_1, \dots, \mu_N \xleftarrow{R} R_q$, over some fixed $\mathbf{e}_1, \dots, \mathbf{e}_N \in R_q$ that are not all zero, then by

$$R_q \cong \prod \left(\mathbb{F}_{q_i^{d/2}} \times \mathbb{F}_{q_i^{d/2}} \right),$$

supposing $q_i \equiv 5 \pmod{8}$ makes $R_{q_i} \cong \mathbb{F}_{q_i^{d/2}} \times \mathbb{F}_{q_i^{d/2}}$, then fixing a nonzero CRT slot has probability at most $q_1^{-d/2}$.

Hence, starting from the extractor that extracts for a single quadratic relation, the knowledge soundness here is $2/N + q_1^{-d/2}$, where the additional term is from some relation is not satisfying.

Consider using the extractor for one quadratic relation. After the first extraction, we have an opening satisfying

$$\bar{\mathbf{s}}^\top \left(\sum \mu_i \mathbf{R}_i \right) \bar{\mathbf{s}} + \left(\sum \mu_i \bar{\mathbf{r}}_{1,i}^\top \right) \bar{\mathbf{s}} + \sum \mu_i \mathbf{r}_{0,i} = 0.$$

Yet, we cannot rule out the case that some subordinate quadratic relation is not satisfying.

Again, we design an extractor for multiple quadratic relations, utilizing the extractor for one quadratic relation, under the trick by [ACK21] utilizing lemma 4.14.4. For a fixed opening violating some underlying quadratic relations, it may succeed with probability at most $q_1^{-d/2}$ over random linear combination by prior discussion, without breaking MSIS.

Here, the opening is fixed independently of the sampled $\bar{\mu}$ for random linear combination, and a malicious prover capable of producing different openings for the same commitment may instead, choose an opening according to $\bar{\mu}$. Hence, in this case, $q_1^{-d/2}$ does not directly bound the probability of obtaining some invalid opening in one extraction.

We therefore seek a second successful extraction for the same commitment, forked from independently sampled $\bar{\mu}$. If the two extracted relaxed openings differ, the binding reduction yields an MSIS solution; otherwise, the same openings may still violate some underlying quadratic relations. Hence, a second extraction tells malicious prover being lucky from the case of finding an MSIS solution.

The new extractor invokes the extractor for one quadratic relation, if the first extraction failed, abort immediately; otherwise, restart extraction over independently sampled $\bar{\mu}$ until a success. Let p be the probability of a successful extraction for one quadratic relation, E be the event that the first extraction is successful, and X be the number of invocations for the extractor for one quadratic relation, then the expected number of invocations is

$$\mathbb{E}[X] = \mathbb{E}[X \mid E] \Pr[E] + \mathbb{E}[X \mid \neg E] \Pr[\neg E] = \left(1 + \frac{1}{p}\right) \cdot p + 1 \cdot (1 - p) = 2.$$

Hence, by eq. (4.64), the expected runtime of the extractor for one quadratic relation is at most $3T$, then the expected runtime of the extractor for multiple quadratic relations is at most $6T$.

For the success probability of such extractor for multiple quadratic relations, there are three disjoint events:

- E_0 : Both openings are the same, and all underlying quadratic relations are satisfied.
- E_1 : Both openings are the same, but some underlying quadratic relations are violated.
- E_2 : The openings are distinct, or an MSIS solution on either extraction.

A prover capable of breaking binding may choose an invalid opening by $\bar{\mu}$ in the first extraction, yet return the same opening in the second extraction. Such an execution does not give an MSIS solution and belongs to E_1 . We therefore need to bound $\Pr[E_1]$ even for such provers. By the construction of the extractor for multiple quadratic relations,

$$\Pr[E_0 \vee E_1 \vee E_2] = p \geq \varepsilon - \frac{2}{N}, \quad (4.66)$$

where ε is the prover's convincing probability, and eq. (4.66) holds by, on the first successful extraction, the second extraction keeps rebooting up to another successful extraction. Hence,

$$\Pr[E_0 \vee E_2] = p - \Pr[E_1] \geq \varepsilon - \frac{2}{N} - \Pr[E_1].$$

Let X_0, X_1 be the random variables for the output of the first and second extraction, and h_x be the probability of an extraction gives relaxed opening x , then $h_{\perp} = 1 - p \leq 1 - \varepsilon + 2/N$, and

$$\Pr[X_1 = x \mid X_0 = x] = \sum_{r \geq 0} (1-p)^r h_x = \frac{h_x}{p}.$$

Therefore, summing overall openings that violates some underlying quadratic relations, we have

$$\Pr[E_1] = \sum_x \Pr[X_1 = x \mid X_0 = x] \Pr[X_0 = x] = \sum_x \frac{h_x}{p} \cdot h_x \leq \frac{q_1^{-d/2}}{p} \sum_x h_x \leq q_1^{-d/2},$$

the third inequality holds by $h_x \leq q_1^{-d/2}$ for a fixed invalid opening x , and the last inequality holds by the successful extraction probability, and therefore

$$\Pr[E_0 \vee E_2] \geq \varepsilon - \frac{2}{N} - q_1^{-d/2}.$$

Vanishing Constant Coefficient Proof. We denote the protocol for proving multiple quadratic relations by Π_{many}^2 . We consider proving multiple quadratic relations, whose evaluations

$$\bar{\mathbf{s}}^T \mathbf{R}_i \bar{\mathbf{s}} + \bar{\mathbf{r}}_{1,i}^T \bar{\mathbf{s}} + \mathbf{r}_{0,i} \quad (4.67)$$

over $\bar{\mathbf{s}}$ defined in eq. (4.62), for $i \in [1, M]$, have constant coefficients being 0. Write $F_i(\bar{\mathbf{s}})$ for eq. (4.67), and extend the BDLOP commitment key by $\mathbf{B}_g \in R_q^{\lambda \times m_2}$. The protocol proceeds as follows:

- The prover samples $\bar{\mathbf{g}} \xleftarrow{R} \{\mathbf{a} \in R_q : a_0 = 0\}^\lambda$, commits $\bar{\mathbf{t}}_g \leftarrow \mathbf{B}_g \bar{\mathbf{s}}_2 + \bar{\mathbf{g}}$, and sends $\bar{\mathbf{t}}_g$ to the verifier
- The verifier replies $\Gamma \xleftarrow{R} \mathbb{Z}_q^{\lambda \times M}$.
- The prover computes

$$\begin{bmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_\lambda \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{g}_1 \\ \vdots \\ \mathbf{g}_\lambda \end{bmatrix} + \Gamma \begin{bmatrix} F_1(\bar{\mathbf{s}}) \\ \vdots \\ F_M(\bar{\mathbf{s}}) \end{bmatrix},$$

and sends $\mathbf{h}_1, \dots, \mathbf{h}_\lambda$.

- Invoke Π_{many}^2 to prove

$$\mathbf{g}_j + \sum_{i=1}^M \gamma_{j,i} F_i(\bar{\mathbf{s}}) - \mathbf{h}_j = \mathbf{0}$$

for $j \in [1, \lambda]$.

- The verifier accepts if Π_{many}^2 accepts and $h_{j,0} = 0$ for $j \in [1, \lambda]$.

Additional quadratic relations over R_q over the original committed witness can be included in the same invocation of Π_{many}^2 , which combines them with the above well-formedness equations.

Similar to the analysis for the multiple quadratic relation proof, the knowledge soundness here is $2/N + q_1^{-d/2} + q_1^{-\lambda}$. The extractor rewinds at the Γ challenge, and invokes the extractor for multiple quadratic relations. By the [ACK21] trick, the expected number of invocations to the extractor for multiple quadratic relations is 2, and thus the expected runtime is at most $12T$. With the same analysis, splitting into 3 disjoint events, the probability of 2 openings being distinct is at most $q_1^{-\lambda}$, completing the soundness analysis.

Proving Norm Bounds. We begin by a lemmas from [BL17, LNS21] as a motivation. Recall the binomial distribution from fair coin, where Bin_κ is a distribution from

$$\sum_{i \in [1, \kappa]} (X_i - Y_i),$$

where X_i and Y_i are fair coins over $\{0, 1\}$. We begin with the prototype in [BL17].

Lemma 4.14.5 (ℓ_∞ Lower-Tail Bound for Binary Compression [BL17]). *For all $\bar{\mathbf{s}} \in \mathbb{Z}_p^k$, $\bar{\mathbf{c}}' \in \{0, 1\}^k$, and $\mathbf{y} \in \mathbb{Z}_p$, over the randomness of $\bar{\mathbf{c}} \xleftarrow{R} \{0, 1\}^k$,*

$$\Pr \left[\left\| \langle \bar{\mathbf{s}}, \bar{\mathbf{c}} - \bar{\mathbf{c}}' \rangle + \mathbf{y} \right\|_\infty < \frac{1}{2} \|\bar{\mathbf{s}}\|_\infty \right] \leq \frac{1}{2}.$$

Proof. Let $F(\bar{\mathbf{c}}) := \langle \bar{\mathbf{s}}, \bar{\mathbf{c}} - \bar{\mathbf{c}}' \rangle + \mathbf{y}$, and $\|\bar{\mathbf{s}}\|_\infty = |s_i|$, then let $\bar{\mathbf{c}}''$ has c_i bit flipped, and

$$F(\bar{\mathbf{c}}) - F(\bar{\mathbf{c}}'') = (2c_i - 1)s_i,$$

and hence $\|F(\bar{\mathbf{c}}) - F(\bar{\mathbf{c}}'')\|_\infty = \|\bar{\mathbf{s}}\|_\infty$. By triangle inequality, for all $\bar{\mathbf{c}} \in \{0, 1\}^k$,

$$\|\bar{\mathbf{s}}\|_\infty \leq \|F(\bar{\mathbf{c}})\|_\infty + \|F(\bar{\mathbf{c}}'')\|_\infty,$$

such that at least one of $\|F(\bar{\mathbf{c}})\|_\infty$ and $\|F(\bar{\mathbf{c}}'')\|_\infty$ is at least $\|\bar{\mathbf{s}}\|_\infty / 2$, completing the proof. $\square$

With lemma 4.14.5, the following result is immediate.

Lemma 4.14.6 (ℓ_∞ Lower-Tail Bound for Binomial Compression [BL17, LNS21]). *For all $\bar{s} \in \mathbb{Z}_p^k$ and $y \in \mathbb{Z}_p$, over the randomness of $\bar{c} \xleftarrow{R} \text{Bin}_1^k$,*

$$\Pr \left[\left| \langle \bar{s}, \bar{c} \rangle + y \right|_\infty < \frac{1}{2} \|\bar{s}\|_\infty \right] \leq \frac{1}{2}.$$

Proof. By lemma 4.14.5, we have the claim for the fixed $\bar{c}'$ version of claim. By $\bar{c}' \xleftarrow{R} \{0, 1\}^k$,

$$\Pr \left[\left| \langle \bar{s}, \bar{c}'' \rangle + y \right|_\infty < \frac{1}{2} \|\bar{s}\|_\infty \right] = \mathbb{E} \left[\Pr \left[\left| \langle \bar{s}, \bar{c} - \bar{c}' \rangle + y \right|_\infty < \frac{1}{2} \|\bar{s}\|_\infty \mid \bar{c}' \right] \right] \leq \frac{1}{2},$$

where $\bar{c}'' \leftarrow \text{Bin}_1^k$, and the inequality holds by lemma 4.14.5, as any probability in the expectation is at most 1/2. $\square$

The range by lemma 4.14.6 gives a probability bound for the ℓ_∞ -norm in $[\|\bar{s}\|_\infty/2, k\|\bar{s}\|_\infty]$, and by lemma 1.4.8, we can show without y that

$$\Pr \left[|\langle \bar{c}, \bar{s} \rangle| \geq C\sqrt{k} \|\bar{s}\|_\infty \right] \leq 2 \exp(-C^2/2),$$

and given a concretely sufficiently large k , we can use a large constant C . The form of $y + \langle \bar{s}, \bar{c} \rangle$ is useful for the sigma protocol, yet the range of the ℓ_∞ -norm has a factor of $O(\sqrt{k})$, which is too large.

Johnson-Lindenstrauss Lemma [JL84], conceptually shows the existence of a linear map $\mathbf{V} \in \mathbb{R}^{\ell \times m}$ sampled over continuous Gaussian distribution, such that for k distinct points in m -dimensional vector space be linearly embedded into ℓ -dimensional space, where $\ell = O(\log k)$, and their distances over ℓ_2 -norm are preserved approximately.

Lemma 4.14.7 (Johnson–Lindenstrauss [JL84]). *Let $S \subseteq \mathbb{R}^m$ with $|S| = k \geq 2$, and $0 < \varepsilon < 1$. There exists $\ell = O(\varepsilon^{-2} \log k)$ and a linear map $\mathbf{V} \in \mathbb{R}^{\ell \times m}$ such that, for all $\bar{x}, \bar{y} \in S$,*

$$(1 - \varepsilon) \|\bar{x} - \bar{y}\|_2^2 \leq \|\mathbf{V}\bar{x} - \mathbf{V}\bar{y}\|_2^2 \leq (1 + \varepsilon) \|\bar{x} - \bar{y}\|_2^2.$$

In [JL84], the lemma 4.14.7 shows the existence of such linear map, and in [IM98] we have a randomized algorithm over Gaussian matrix, such that with high probability, the linear map satisfies the conditions in lemma 4.14.7.

Lemma 4.14.8 (Gaussian Randomized Johnson–Lindenstrauss [IM98]). *Let $S \subseteq \mathbb{R}^m$ with $|S| = k \geq 2$, and let $\varepsilon, \delta \in (0, 1)$. There exists a constant $C > 0$ such that the following holds.*

Choose a positive integer

$$\ell \geq C\varepsilon^{-2} \log(k/\delta),$$

and sample $\mathbf{V} \in \mathbb{R}^{\ell \times m}$ by independently drawing each entry from $\mathcal{N}(0, 1/\ell)$.

Then, with probability at least $1 - \delta$ over $\mathbf{V}$, simultaneously for all $\bar{x}, \bar{y} \in S$,

$$(1 - \varepsilon) \|\bar{x} - \bar{y}\|_2^2 \leq \|\mathbf{V}\bar{x} - \mathbf{V}\bar{y}\|_2^2 \leq (1 + \varepsilon) \|\bar{x} - \bar{y}\|_2^2.$$

Achlioptas [Ach03] shows that the same asymptotic dimension bound and distance-preservation guarantee can be obtained using independent random signs.

Lemma 4.14.9 (Sign Randomized Johnson–Lindenstrauss [Ach03]). *The conclusion of lemma 4.14.8 remains valid for entries $V_{i,j} = X_{i,j}/\sqrt{\ell}$, where the $X_{i,j}$ are sampled independently and uniformly over ± 1 .*

In our scenario, we don't need to handle k points in this case, but we do need to preserve the pairwise difference in Euclidean space. More specifically, we want the distance between a fixed vector and zero preserved approximately. The fixed vector is the witness in the Sigma protocol. We therefore study how closely its projected norm follows its original norm, with high probability.

Gentry et al. [GHL22] use an integer projection from $\mathbf{R} \leftarrow \text{Bin}_1^{256 \times m}$ to reduce the dimension of the vector whose shortness is to be proved. To obtain concrete norm estimates, they compare this projection with a Gaussian matrix $\mathbf{G} \in \mathbb{R}^{256 \times m}$ whose entries are sampled independently from $\mathcal{N}(0, 1/2)$, matching the variance of Bin_1 .

For a fixed nonzero vector $\bar{w} \in \mathbb{Z}^m$,

$$\frac{2 \|\mathbf{G}\bar{w}\|_2^2}{\|\bar{w}\|_2^2} \sim \chi_{256}^2.$$

To see this, consider the output of row i :

$$Y_i := (\mathbf{G}\bar{w})_i = \sum_{j=1}^m G_{i,j} w_j.$$

A fixed linear combination of independent Gaussian variables is Gaussian. Since $\mathbb{E}[G_{i,j}] = 0$ and $\text{Var}[G_{i,j}] = 1/2$,

$$Y_i \sim \mathcal{N}\left(0, \frac{1}{2} \sum_{j \in [1,m]} w_j^2\right) = \mathcal{N}\left(0, \frac{\|\bar{w}\|_2^2}{2}\right).$$

Moreover, the Y_i are independent, since they use independently sampled rows of $\mathbf{G}$.

Thus, normalizing each output by its standard deviation gives independent standard normal variables

$$Z_i := \frac{\sqrt{2}}{\|\bar{w}\|_2} Y_i \sim \mathcal{N}(0, 1).$$

In particular,

$$\|\mathbf{G}\bar{w}\|_2^2 = \sum_{i \in [1,256]} Y_i^2 = \frac{\|\bar{w}\|_2^2}{2} \sum_{i \in [1,256]} Z_i^2.$$

The factor $\|\bar{w}\|_2^2$ therefore cancels:

$$\frac{2 \|\mathbf{G}\bar{w}\|_2^2}{\|\bar{w}\|_2^2} = \sum_{i=1}^{256} Z_i^2 \sim \chi_{256}^2.$$

Here, the 256 degrees of freedom count the independent standard normal variables being squared.

Let $X \sim \chi_{256}^2$. Numerically evaluating its tails gives

$$\Pr[X < 60] \approx 2^{-130.98} < 2^{-128}, \quad \Pr[X > 674] \approx 2^{-128.30} < 2^{-128}.$$

Consequently, the normalization above yields the Gaussian bounds in [GHL22, Fact 3.1]:

$$\begin{aligned} \Pr_{\mathbf{G}} \left[\|\mathbf{G}\bar{w}\|_2^2 < 30 \|\bar{w}\|_2^2 \right] &= \Pr[X < 60] < 2^{-128}, \\ \Pr_{\mathbf{G}} \left[\|\mathbf{G}\bar{w}\|_2^2 > 337 \|\bar{w}\|_2^2 \right] &= \Pr[X > 674] < 2^{-128}. \end{aligned} \tag{4.68}$$

The constants 30 and 337 arise by dividing the chi-square thresholds 60 and 674 by 2.

Gentry et al. [GHL22, Corollary 3.2] then heuristically apply these same bounds to the projection over binomial distribution $\mathbf{R} \leftarrow \text{Bin}_1^{256 \times m}$, whose entries have the same variance as those of $\mathbf{G}$. The Gaussian bounds in eq. (4.68) are exact probability statements; their transfer to the binomial projection is heuristic.

More generally, replacing the Gaussian variance $1/2$ by $\kappa/2$ gives

$$\frac{2 \|\mathbf{G}\bar{w}\|_2^2}{\kappa \|\bar{w}\|_2^2} \sim \chi_{256}^2.$$

Under the same Gaussian substitution heuristic, for $\mathbf{R} \leftarrow \text{Bin}_{\kappa}^{256 \times m}$ and $X \sim \chi_{256}^2$, LNP [LNP22, Lemma 2.8] obtains

$$\begin{aligned} \Pr_{\mathbf{R}} \left[\|\mathbf{R}\bar{w}\|_2^2 < 13\kappa \|\bar{w}\|_2^2 \right] &\approx \Pr[X < 26] \leq 2^{-256}, \\ \Pr_{\mathbf{R}} \left[\|\mathbf{R}\bar{w}\|_2^2 > 337\kappa \|\bar{w}\|_2^2 \right] &\approx \Pr[X > 674] \leq 2^{-128}. \end{aligned}$$

Here, $\approx$ denotes the heuristic comparison. The smaller lower-tail threshold gives the stronger 2^{-256} estimate needed in the subsequent argument.

Lemma 4.14.10 (ℓ_2 Lower-Tail Bound for Binomial Compression [LNP22]). *Fix $m, p \in \mathbb{Z}^+$ and a bound $0 < b \leq p/(41m)$. Let $\bar{w} \in \mathbb{Z}_p^m$ satisfy $\|\bar{w}\|_2 \geq b$, and $\bar{y} \in \mathbb{Z}_p^{256}$ is arbitrary. Norms of residue vectors use centered representatives.*

For $\mathbf{R} \leftarrow \text{Bin}_1^{256 \times m}$ with independent entries, over the randomness of $\mathbf{R}$,

$$\Pr \left[\|\mathbf{R}\bar{w} + \bar{y} \bmod p\|_2 < 4b \right] < 2^{-128}.$$

Proof. We begin by separating out the case where $\|\bar{w}\|_{\infty} \geq p/(4m)$. By the given bound condition,

$$4b \leq \frac{4p}{41m} \leq \frac{p}{8m}.$$

By lemma 4.14.6, for any row $i \in [1, 256]$, over the randomness of i^{th} row $\bar{r}_i^{\top}$,

$$\Pr \left[|\langle \bar{r}_i, \bar{w} \rangle + y_i| < \frac{1}{2} \|\bar{w}\|_{\infty} \right] \leq \frac{1}{2}.$$

All rows must satisfy such event to have the ℓ_2 -norm of the projected vector be at most $4b$, of at most 2^{-256} probability.

Now consider the complementary case $\|\bar{w}\|_{\infty} < p/(4m)$. Consider all the computation over $\mathbb{Z}$ rather than $\mathbb{Z}_p$, and let $X_i = \langle \bar{r}_i, \bar{w} \rangle$, then we have $|X_i| \leq \|\bar{w}\|_1 < p/4$, and thus each $X_i + y_i$ has $|X_i + y_i| < 3p/4$.

With a centered representation, a long $|X_i + y_i| > p/2$ has wrapped representation $p/4 < |X_i + y_i \bmod p| < p/2$, yet we have $4b \leq p/(8m) < p/4$, and thus with $\|\bar{w}\|_{\infty} < p/(4m)$,

$$\Pr \left[\|\mathbf{R}\bar{w} + \bar{y} \bmod p\|_2 < 4b \right] = \Pr \left[\|\mathbf{R}\bar{w} + \bar{y}\|_2 < 4b \right],$$

as the settings with rows of large ℓ_{∞} -norm wrapping into $(p/4, p/2)$ are not counted anyways.

Continuing on, we use the MGF trick in Chernoff's bound, by

$$\begin{aligned} \Pr \left[\|\mathbf{R}\bar{w} + \bar{y}\|_2 < 4b \right] &= \Pr \left[\sum_{i \in [1, 256]} (X_i + y_i)^2 < 16b^2 \right] \leq \exp(16sb^2) \mathbb{E} \left[\prod_{i \in [1, 256]} \exp(-s(X_i + y_i)^2) \right] \\ &= \exp(16sb^2) \prod_{i \in [1, 256]} \mathbb{E} \left[\exp(-s(X_i + y_i)^2) \right] \end{aligned}$$

for any $s \geq 0$. We observe that, for a binomial random variable $B \sim \text{Bin}_1$,

$$\mathbb{E} \left[\exp(iuB) \right] = \frac{1}{4} (\exp(-iu) + \exp(iu)) + \frac{1}{2} = \frac{1}{2} (\cos(u) + 1) = \cos^2(u/2).$$

With this result, we can introduce the linear shift y such that

$$\begin{aligned} \mathbb{E} \left[\cos(u(X_i + y_i)) \right] &= \frac{1}{2} (\mathbb{E} \left[\exp(ui(X_i + y_i)) \right] + \mathbb{E} \left[\exp(-ui(X_i + y_i)) \right]) = \cos(uy_i) \prod_{j \in [1, m]} \cos^2(uw_j/2) \\ &\leq \prod_{j \in [1, m]} \cos^2(uw_j/2) = \mathbb{E} [\cos(uX_i)]. \end{aligned}$$

Another fact we are using is, let $G \sim \mathcal{N}(0, 1)$ and any $t \in \mathbb{R}$, we have $\mathbb{E}[\cos(tG)] = \exp(-t^2/2)$. Taking independent G of X_i , we substitute $u = \sqrt{2s}G$ into the preceding inequality and average over G . Since cosine is bounded, we can exchange the order of expectations and obtain

$$\begin{aligned} \mathbb{E}_{X_i} [\exp(-s(X_i + y_i)^2)] &= \mathbb{E}_G \mathbb{E}_{X_i} [\cos(\sqrt{2s}G(X_i + y_i))] \\ &\leq \mathbb{E}_G \mathbb{E}_{X_i} [\cos(\sqrt{2s}GX_i)] = \mathbb{E}_{X_i} [\exp(-sX_i^2)]. \end{aligned}$$

Hence, the shift y_i is removed from the upper bound.

Substituting this into the Chernoff bound, and using that the X_i have the same distribution, gives

$$\Pr [\|\mathbf{R}\bar{w} + \bar{y}\|_2 < 4b] \leq \exp(16sb^2) \prod_{i \in [1, 256]} \mathbb{E} [\exp(-sX_i^2)]$$

for every $s \geq 0$. It remains to bound the expectation for a single row and then choose s .

The next goal is to upper bound $\mathbb{E} [\exp(-sX_i^2)]$ such that the X_i can be taken out, and it only matters with $\|\bar{w}\|_2$.

TODO: *continue on?*

□

TODO: *Stuffs for [LNP22].*

4.15 Practical Sublinear Proofs for R1CS from Lattices

TODO: *stuff from* [NS22].

4.16 **LaBRADOR**

TODO: ... [BS23]

4.17 Greyhound

TODO: ... [NS24]

4.18 Lattice-Based Succinct Arguments for NP with Polylog Verification

TODO: ... [BCS23]

References

- [Ach03] Dimitris Achlioptas. Database-friendly random projections: Johnson–Lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, 2003.
- [ACK21] Thomas Attema, Ronald Cramer, and Lisa Kohl. A compressed Σ -protocol theory for lattices. In *CRYPTO*, pages 549–579, 2021.
- [ACX21] Thomas Attema, Ronald Cramer, and Chaoping Xing. A note on short invertible ring elements and applications to cyclotomic and trinomial number fields. *Mathematical Cryptology*, 1(1):45–70, 2021.
- [AGHS13] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete gaussian leftover hash lemma over infinite domains. In *ASIACRYPT*, pages 97–116, 2013.
- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Liger: Lightweight sublinear arguments without a trusted setup. In *ACM CCS*, pages 2087–2104, 2017.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *STOC*, pages 99–108, 1996.
- [ALS20] Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In *CRYPTO*, pages 470–499, 2020.
- [And87] Désiré André. Solution directe du problème résolu par m. bertrand. *Comptes Rendus de l'Académie des Sciences, Paris*, 105:436–437, 1887.
- [AS16] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. Wiley, 4th edition, 2016.
- [Azu67] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, 19(3):357–367, 1967.
- [Ban93] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE S&P*, pages 315–334, 2018.
- [BBC⁺18] Carsten Baum, Jonathan Bootle, Andrea Cerulli, Rafaël del Pino, Jens Groth, and Vadim Lyubashevsky. Sub-linear lattice-based zero-knowledge arguments for arithmetic circuits. In *CRYPTO*, pages 669–699, 2018.
- [BCS23] Jonathan Bootle, Alessandro Chiesa, and Katerina Sotiraki. Lattice-based succinct arguments for NP with polylogarithmic-time verification. In *CRYPTO*, pages 227–251, 2023.
- [BD97] Russ Bubley and Martin E. Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *FOCS*, pages 223–231, 1997.
- [BDL⁺18] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, pages 368–385, 2018.
- [Bec91] József Beck. An algorithmic approach to the lovász local lemma. i. *Random Struct. Algorithms*, 2(4):343–365, 1991.
- [BG93] Mihir Bellare and Oded Goldreich. On defining proofs of knowledge. In *CRYPTO*, pages 390–420, 1993.
- [BL17] Carsten Baum and Vadim Lyubashevsky. Simple amortized proofs of shortness for linear relations over polynomial rings. Cryptology ePrint Archive, Paper 2017/759, 2017.

- [BLNS20] Jonathan Bootle, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. A Non-PCP approach to succinct Quantum-Safe Zero-Knowledge. In *CRYPTO*, pages 441–469, 2020.
- [BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993.
- [BLS19] Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In *CRYPTO*, pages 176–202, 2019.
- [BS23] Ward Beullens and Gregor Seiler. LaBRADOR: Compact proofs for R1CS from Module-SIS. In *CRYPTO*, pages 518–548, 2023.
- [CF24] Hadas Zeilberger Binyi Chen and Ben Fisch. Basefold: Efficient multilinear polynomial commitment schemes from foldable codes. In *CRYPTO*, pages 138–169, 2024.
- [Cho94] Kwok Pui Choi. On the medians of gamma distributions and an equation of ramanujan. *Proceedings of the American Mathematical Society*, 121(1):245–251, 1994.
- [Dam10] Ivan Damgård. On Σ -protocols. Lecture notes, Cryptologic Protocol Theory, 2010. CPT 2010, v.2.
- [Dob40] Wolfgang Doblin. Éléments d’une théorie générale des chaînes simples constantes de Markoff. *Annales scientifiques de l’École Normale Supérieure*, 57:61–111, 1940.
- [Doo40] Joseph L. Doob. Regularity properties of certain families of chance variables. *Transactions of the American Mathematical Society*, 47(3):455–486, 1940.
- [DP24] Benjamin E. Diamond and Jim Posen. Proximity testing with logarithmic randomness. *IACR Communications in Cryptology*, 1(1), 2024.
- [EL75] Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In András Hajnal, Richard Rado, and Vera T. Sós, editors, *Infinite and Finite Sets*, volume 10 of *Colloquia Mathematica Societatis János Bolyai*, pages 609–627. North-Holland, Amsterdam, 1975.
- [Eli72] Peter Elias. The efficient construction of an unbiased random sequence. *The Annals of Mathematical Statistics*, 43(3):865–870, 1972.
- [ENS20] Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In *ASIACRYPT*, pages 259–288, 2020.
- [Ete81] Nasrollah Etemadi. An elementary proof of the strong law of large numbers. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 55(1):119–122, 1981.
- [EZX⁺19] Muhammed F. Esgin, Raymond K. Zhao, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. MatRiCT: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In *ACM CCS*, pages 567–584, 2019.
- [Fel68a] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, 1968.
- [Fel68b] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 2. John Wiley & Sons, 1968.
- [FS90] Uriel Feige and Adi Shamir. Witness indistinguishable and witness hiding protocols. In *STOC*, pages 416–426, 1990.
- [GHL22] Craig Gentry, Shai Halevi, and Vadim Lyubashevsky. Practical non-interactive publicly verifiable secret sharing with thousands of parties. In *EUROCRYPT*, pages 458–487, 2022.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions. In *STOC*, pages 12–24, 1989.
- [IM98] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *STOC*, pages 604–613, 1998.
- [Jer95] Mark Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, 1995.
- [JL84] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in Modern Analysis and Probability*, volume 26 of *Contemporary Mathematics*, pages 189–206. American Mathematical Society, 1984.
- [JS89] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [JVV86] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.*, 43:169–188, 1986.
- [KL51] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [Kra49] Leon G. Kraft. A device for quantizing, grouping, and coding amplitude-modulated pulses. M.s. thesis, Massachusetts Institute of Technology, Cambridge, MA, 1949.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP*, pages 144–155, 2006.
- [LMR94] F. Thomson Leighton, Bruce M. Maggs, and Satish B. Rao. Packet routing and job-shop scheduling in $O(\text{congestion} + \text{dilation})$ steps. *Combinatorica*, 14(2):167–186, 1994.
- [LNP22] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In *CRYPTO*, pages 71–101, 2022.
- [LNS21] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In *PKC*, pages 215–241, 2021.
- [Lov79] László Lovász. On the shannon capacity of a graph. *IEEE Transactions on Information Theory*, 25(1):1–7, 1979.
- [LP17] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI, 2 edition, 2017. With contributions by Elizabeth L. Wilmer.
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. In *EUROCRYPT*, pages 35–54, 2013.
- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.
- [LS18] Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In *EUROCRYPT*, pages 204–224, 2018.
- [Lub66] David Lubell. A short proof of sperner’s lemma. *Journal of Combinatorial Theory*, 1:299, 1966.

- [Lyu09] Vadim Lyubashevsky. Fiat–Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
- [Mat88] Peter C. Matthews. Covering problems for brownian motion on spheres. *Annals of Probability*, 16(1):189–199, 1988.
- [McD89] Colin McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics, 1989*, volume 141 of *London Mathematical Society Lecture Note Series*, pages 148–188. Springer, 1989.
- [McM56] Brockway McMillan. Two inequalities implied by unique decipherability. *IRE Transactions on Information Theory*, 2(4):115–116, 1956.
- [Mos09] Robin A. Moser. A constructive proof of the lovász local lemma. In *STOC*, pages 343–350, 2009.
- [Mos10] Dana Moshkovitz. An alternative proof of the schwartz–zippel lemma. *Electron. Colloquium Comput. Complex.*, 2010:96, 2010.
- [Mos12] Robin A. Moser. *Exact Algorithms for Constraint Satisfaction Problems*. PhD thesis, ETH Zurich, 2012.
- [MT10] Robin A. Moser and Gábor Tardos. A constructive proof of the general lovász local lemma. *J. ACM*, 57(2), 2010.
- [MU17] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, Cambridge, UK, 2nd edition, 2017.
- [NS22] Ngoc Khanh Nguyen and Gregor Seiler. Practical sublinear proofs for R1CS from lattices. In *CRYPTO*, pages 133–162, 2022.
- [NS24] Ngoc Khanh Nguyen and Gregor Seiler. Greyhound: Fast polynomial commitments from lattices. In *CRYPTO*, pages 243–275, 2024.
- [Pap91] Christos H. Papadimitriou. On selecting a satisfying truth assignment. In *FOCS*, pages 163–169, 1991.
- [Per92] Yuval Peres. Iterating Von Neumann’s procedure for extracting random bits. *The Annals of Statistics*, 20(1):590–597, 1992.
- [Pós76] László Pósa. Hamiltonian circuits in random graphs. *Discrete Mathematics*, 14(4):359–364, 1976.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
- [PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007.
- [Rag88] Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *J. Comput. Syst. Sci.*, 37(2):130–143, 1988.
- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4), 1980.
- [Sha48a] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.
- [Sha48b] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(4):623–656, 1948.

- [Sha56] Claude E. Shannon. The zero error capacity of a noisy channel. *IRE Transactions on Information Theory*, IT-2(3):8–19, 1956.
- [Spe28] Emanuel Sperner. Ein satz über untermengen einer endlichen menge. *Mathematische Zeitschrift*, 27:544–548, 1928.
- [Spe77] Joel Spencer. Asymptotic lower bounds for ramsey functions. *Discrete Mathematics*, 20(1):69–76, 1977.
- [SS87] Eli Shamir and Joel Spencer. Sharp concentration of the chromatic number on random graphs $G_{n,p}$. *Combinatorica*, 7:121–129, 1987.
- [Top07] Flemming Topsøe. Some bounds for the logarithmic function. In Yeol Je Cho, Jong Kyu Kim, and Sever S. Dragomir, editors, *Inequality Theory and Applications*, volume 4, pages 137–151. Nova Science Publishers, New York, 2007.
- [Tse63] Michael Lvovitch Tsetlin. Finite automata and models of simple forms of behaviour. *Russian Mathematical Surveys*, 18(4):1–27, 1963.
- [Vad12] Salil P. Vadhan. *Pseudorandomness*. Now Publishers Inc., Hanover, MA, 2012.
- [VB81] Leslie G. Valiant and Gordon J. Brebner. Universal schemes for parallel communication. In *STOC*, pages 263–277, 1981.
- [von51] John von Neumann. Various techniques used in connection with random digits. In A. S. Householder, G. E. Forsythe, and H. H. Germond, editors, *Monte Carlo Method*, volume 12 of *National Bureau of Standards Applied Mathematics Series*, pages 36–38. U.S. Government Printing Office, Washington, DC, 1951. Summary written by George E. Forsythe.
- [Wal47] Abraham Wald. *Sequential Analysis*. John Wiley & Sons, 1947.
- [YAZ⁺19] Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In *CRYPTO*, pages 147–175, 2019.
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *EUROSAM*, 1979.

