Side Information in Bandit Problems & Low-Density Parity-Check Codes for Non-Symmetric Channels

Chih-Chun Wang

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance

by the Department of

Electrical Engineering

September, 2005

Abstract

Two major topics are discussed in this thesis: side information in the context of learning vs. control and capacity-approaching error correcting codes for non-standard channels.

Side Information in Bandit Problems:

The two-armed bandit problem is a classical model characterizing the conflict between learning and control. The decision maker samples one arm at each time instant in order to maximize the total reward (the control task), while the reward distribution for each individual arm must be learned through the same sampling process as well (the learning task). In this thesis, an extension of this traditional setting is studied, under which the decision maker has access to some side information before deciding which arm to pull. Various degrees of improvement are identified and proven when considering general evenly distributed side information random processes, including independent and identically distributed random processes, Markov chains, and deterministic periodic sequences as special cases. In particular, it is shown that the well-known logarithmic lower bound for the regret in traditional bandit problems can be surpassed by explicit algorithms, achieving new, more favorable performance bounds. These results can be applied to a broad range of problems including clinical trials, early exercise strategies for the American put options, and beam selection in multiple antenna communication systems.

Low-Density Parity-Check Codes on Non-Symmetric Channels:

The near-Shannon-capacity performance of Low-Density Parity-Check (LDPC) and turbo codes has stimulated considerable research in coding theory in recent years, and significant successes have been established for standard memoryless binary-input/symmetric-output channels. In this work, the density evolution method is generalized for non-symmetric memoryless channels. This is achieved by complementing the asymptotic performance concentration theorem with a perfect projection convergence theorem. Some implications of these results include stability conditions for non-symmetric channels, the local optimality of belief propagation decoding, and the typicality of linear LDPC codes, which further justifies the application of other existing tools, e.g. EXtrinsic InformaTion (EXIT) chart analysis, on non-symmetric channels. This new powerful density evolution method successfully bridges the gap between symmetric channels and general memoryless channels, and is demonstrated on a simple non-symmetric model for optical storage channels. Various finite-dimensional bounds on the decodable thresholds are derived, including the best bound for fading channels to date.

Acknowledgements

I would like to express my heartfelt gratitude to my advisors, Prof. Sanjeev R. Kulkarni and Prof. H. Vincent Poor, for the guidance and many insightful discussions during my Ph.D. study. I was lucky to work closely with them and their infectious passions in complementing areas inspired this work in many different ways. Their scholarship also influences me deeply, and I am especially grateful to their unparalleled support during this self-development journey of my Ph.D. study.

I would also like to thank Prof. Bradley W. Dickinson, Prof. Sanjeev R. Kulkarni, Prof. Sun-Yuan Kung, Prof. Bede Liu, Prof. H. Vincent Poor, Prof. Peter J. Ramadge, and Prof. Sergio Verdú for serving in my candidacy and defense committee and for their advices on this thesis.

Thanks to Dongning for his many wise advices as a colleague and as a friend.

This thesis is dedicated to Hsi-Wen, Jr-Yan, my mother and my father. This work would not happen without your love.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

One of the most important issues in telecommunication is to construct reliable point-to-point channels with efficiency. Achieving this goal generally involves selecting a modulation scheme and designing proper error control mechanisms. The aim of modulation is to convert the physical electromagnetic/optical/acoustic signal-carriers to a logical channel sending bit strings, a mixture of 0's and 1's, from the transmitter to the receiver. For the same underlying carriers, different modulation schemes will result in different logical channels, and the system designer is hoping to obtain the best channel by careful selection among available options (or by devising new schemes). Nevertheless, after deciding the best modulation scheme, the resulting channel is usually plagued by noise, corrupting the received symbols from their original values. To counter this noise effect, error correcting codes are often exploited to further reduce the error probability to a negligible value at the expense of lowering the amount of information sent per bit usage.

This thesis focuses on two major topics: bandit problems and error correcting codes on non-symmetric channels. The former is a classical optimal selection/control problem. When applied to selection of different modulation schemes, it deals with situations when the behavior of underlying carriers is unknown, namely, it is about how to select the modulation scheme best suited for the unknown carrier. In the second topic, we will focus on analysis and design of error correcting codes when the logical channel of interest is not symmetric.

1.1 Bandit Problems: An Optimal Control Problem

Since the publication of [94], bandit problems have attracted much attention in various areas of statistics, control, learning, and economics (e.g., see [1, 19, 20, 28, 47, 48, 49, 60, 66, 67, 68]). In the classical two-armed bandit problem, at each time a player selects one of two arms of a jackpot machine and receives a reward drawn from a distribution associated with the arm selected. The objective of the player is to maximize the (expected) total reward. When the underlying distributions are known to the player, this problem becomes trivial and an optimal strategy is to always pull the arm with higher expected reward. Nevertheless, in general scenarios, the underlying reward distributions are unknown and can only be learned through pulling both arms sufficiently often (the learning task), which contradicts the objective of sampling solely the more rewarding arm (the control task). There is a fundamental trade-off between gathering information about the unknown reward distributions and choosing the arm we currently think is more rewarding, which forms the essence of the bandit problem. A rich set of problems arises in trying to find

an optimal/reasonable balance between these two conflicting objectives (also referred to as learning versus control, or exploration versus exploitation). One example of bandit problems is when we consider a modulation selection problem in which the interaction among the environment, the physical carriers, and two candidates of modulation schemes, M1 and M2, is unknown. In this setting, the reward from pulling arm 1 is the success of one error-free transmission packet using scheme M1. The reward from arm 2 is similarly defined. Since the knowledge about which scheme has higher success probability is unknown and can be obtained only by trying both schemes sufficiently often, this decision problem can be characterized as an instance of the bandit problems.

In Chapter 2, we consider an extension of the classical two-armed bandit where we have access to side information before making our decision about which arm to pull. At any time instant, in addition to the history of previous decisions and outcomes, we have access to an auxiliary side information random variable (r.v.) to help us make our current decision. For example, in the previous case of selecting the best modulation scheme, the side information may contain the geographical information about the transmitter-receiver pair or may include environmental variables. We will show that the extent to which this side information can help depends on the relationship among the auxiliary side information r.v. and the reward distributions of both arms. These results will then be generalized for a broader class of evenly distributed random processes, including independently and identically distributed (i.i.d.) random processes, Markov chains of any finite order, and deterministic periodic sequences.

Previous work on bandit problems with side information includes [32, 62, 96, 116, 119]. Woodroofe [116] considered a one-armed bandit in a Bayesian setting, in which an i.i.d.[1] $\{X_\tau\}$ was considered. Contrary to the traditional bandit problems (without side information), Woodroofe proved that with the help of side information, a *myopic* approach becomes asymptotically optimal, assuming the governing conditional distributions $F_{\theta_i}(\cdot|X_t)$ are Gaussian with means $\theta_i + X_t$ and variances 1. Sarkar [96] extended the simple relationship in [116] to exponential families. In [62], Kulkarni considered classes of reward distributions and their effects on performance using results from learning theory. Most of the previous work with side information is on one-armed bandit problems, which can be viewed as a special case of the two-armed setting by letting arm 2 always return zero. Other approaches regarding side information can be found in [32, 62, 119].

In contrast with this previous work, we consider various general settings of side information for a two-armed bandit problem. Our focus is on deriving bounds on achievable rewards and providing bound-achieving algorithms for the various settings. The results and proofs are very much along the lines of [66] and subsequent work in [2, 3, 4, 9, 10] and are published in [113, 114].

1.2 Error Correcting Codes

Error-reduction in communication systems can be achieved by sending automatic repeat request (ARQ) through the feedback channel when errors occur, or by using (forward) error correcting codes (ECCs). The former of these is an efficient and reliable error control scheme based on the two assumptions that a reliable, cost-effective feedback channel exists

[1] In the literature of bandit problems, the commonly used term "i.i.d. side information $\{X_\tau\}$" refers to a marginally i.i.d. $\{X_\tau\}$. Namely, after averaging over $\{Y_\tau^1\}$ and $\{Y_\tau^2\}$, $\{X_\tau\}$ becomes an i.i.d. random process. It does not mean that the side information $\{X_\tau\}$ is independent of the reward sequences $\{Y_\tau^i\}$.

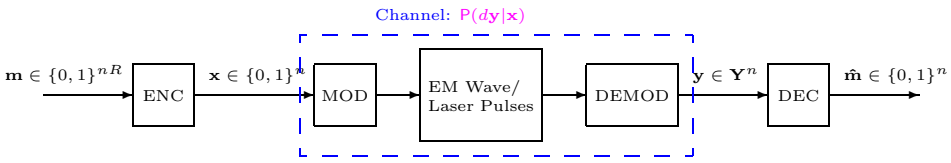

Figure 1.1: A point-to-point logical channel governed by $\mathsf{P}(d\mathbf{y}|\mathbf{x})$.

and the error events during retransmission are independent of those in the initial transmission. On the other hand, a simple point-to-point communication system with ECCs requires no feedback channel, as illustrated in Figure 1.1. ECCs achieve their error control purpose by limiting the legitimate transmission to a subset of the entire signal space. Hence, with very high probability, the corrupted received signal can be mapped inversely back to the transmitted signal, namely, the original signal can be successfully decoded. For modern wireless communication systems, error control is usually achieved by a sophisticated combination of ARQ and ECCs. Nevertheless, for deep-space communications, in which the cost of using a feedback channel is prohibitive, and for storage devices, in which the error events during retransmission are not independent, the system designer has to rely solely on ECCs to ensure reliable communications. This thesis will focus only on ECCs, and the results can benefit systems with or without ARQ.

1.2.1 Capacity-Approaching Codes

In Shannon's ground-breaking paper [98], he showed that for stationary memoryless channels governed by the conditional distribution $\mathsf{P}(d\mathbf{y}|\mathbf{x}) = \prod \mathsf{P}(dy_i|x_i)$, reliable communication requires the code rate R being strictly smaller than the Shannon capacity C, which is defined as $\sup_{\mathsf{P}_X} I(X;Y)$, the supremum of the mutual information over all possible *a priori* distributions P_X. Thereafter, the quest for a pair of *efficient* encoder and decoder achieving the Shannon capacity has become the holy grail for the ECC community.

The first achievability result was demonstrated also by Shannon [98], in which a random code ensemble was considered. It was shown that with unlimited computing power, namely, using the maximum *a posteriori* probability (MAP) decoder, the Shannon capacity of *arbitrary stationary memoryless channels* can be saturated by a signal subset (a codebook) consisting of randomly chosen signals. Elias further proved that a random *linear* code ensemble, a strict subset of the random code ensemble, is able to achieve the Shannon capacity of *binary symmetric channels* (BSCs) with unlimited computing power [39]. The construction of an efficient encoder/decoder pair with near-capacity performance was not achieved until the advent of turbo codes [13, 18]. Recent advancements in the development of low-density parity-check (LDPC) and LDPC-like codes, which can be dated back to the early 1960's [43], demonstrates an efficient, practical encoder-decoder pair with extremely close-to-capacity performance on many channel models, including BSCs, binary erasure channels (BECs) [24], binary additive white Gaussian noise channels (BiAWGNCs) [30, 31, 52], Rayleigh channels [50], high order input alphabet channels [15], bit-interleaved coded modulation (BICM) [51, 83], and inter-symbol interference channels [57]. For BECs, stronger results have also been proved, showing that the performance can be pushed *arbitrarily* close to the capacity [77, 85]. Despite the verified close-to-capacity performance, whether we can saturate the Shannon capacity on channels other than BECs remains an

open question.[2]

1.2.2 Low-Density Parity Check Codes on Non-Symmetric Channels

Since the advent of turbo codes[3] [18] and the rediscovery of low-density parity-check (LDPC) codes [43, 78] in the mid 1990's, graph-based codes[4] [100] have attracted significant attention because of their capacity-approaching error correcting capability and the inherent low-complexity ($\mathcal{O}(n)$ or $\mathcal{O}(n\log(n))$ where n is the codeword length) of message passing decoding algorithms [78]. The near-capacity performance of graph codes is generally based on pseudo-random interconnections and Pearl's belief propagation (BP) algorithm [86], which is a distributed message-passing algorithm for efficiently computing *a posteriori* probabilities in cycle-free inference networks [7, 61].

Because of their simple arithmetic structure, completely parallel decoding algorithms, excellent error correcting capability [30], and linear encoding complexity[5] [93, 101], LDPC codes have been widely and successfully applied to different channels, including BECs [24, 76, 77], BSCs, BiAWGNCs [78, 92], Rayleigh fading channels [50], Markov channels [45], partial response channels/inter-symbol interference channels [57, 65, 72, 106], dirty paper coding [25], bit-interleaved coded modulation (BICM), and multi-level coding [51, 83]. Except for the finite-length analysis of LDPC codes over BECs [36], the analysis of iterative message-passing algorithms is asymptotic (when the block length tends to infinity) [90, 92]. For the MAP decoder, performance analysis is tractable and relies solely on the weight distribution of LDPC codes and turbo-like code ensembles (see e.g. [73, 53], and [70]).

One of the most important analytical tools for LDPC codes is the density evolution (DE) method proposed by Richardson and Urbanke in [92]. In essence, DE is an asymptotic analytical tool for LDPC codes. As the codeword length tends to infinity, the random codebook will be more and more likely to be cycle-free until a certain depth, under which condition the input messages of each node are independently distributed. Therefore the probability density of messages passed can be computed iteratively. A performance concentration theorem and a cycle-free convergence theorem provide the theoretical foundation of DE in [92]. The behavior of codes with block length $n \geq 10^4$ is efficiently and accurately predicted by this technique, and thus DE can be used as a performance metric during code optimization. LDPC codes with near-capacity performance have thus been found in [30] and [90] with the help of DE. In [57], Kavčić *et al.* generalized DE to inter-symbol interference channels by introducing the ensemble of LDPC *coset codes*. Namely, instead of considering only even parity check equations (corresponding to linear codes), a random mixture of odd and even parity check equations is considered. Kavčić *et al.* also proved the corresponding fundamental theorems for this new coset code ensemble.

Due to the symmetry of the BP algorithm and the symmetry of parity check equations in LDPC codes, the decoding error probability will be independent of the transmitted

[2] When combined with unlimited computing power (with ML decoders), LDPC codes are able to saturate the capacity of BSCs [78]. Only when limiting ourselves to practical encoder-decoder pairs, whether we can have arbitrarily close-to-capacity performance becomes uncertain.

[3] Turbo codes can also be viewed as a variation of LDPC codes, as discussed in [78] and [81].

[4] An incomplete list of graph-based codes includes turbo codes, turbo product codes [88], LDPC codes, multi-edge type LDPC codes [91], low-density generating-matrix (LDGM) codes [46], irregular repeat-accumulate (RA) codes [52], concatenated tree codes [87], and the Luby transform (LT) codes [75]

[5] A naive encoder based on the generating matrix requires $\mathcal{O}(n^2)$ computations. Nevertheless, with careful construction of the codebook and regarding the encoding process as decoding on BECs, linear encoding complexity can be easily achieved for practical applications [93].

codeword in the symmetric channel setting. Therefore, in [92], an all-zero transmitted codeword is assumed. However, for symbol-dependent non-symmetric channels, which are the subject of Chapter 4, the noise distribution is codeword-dependent, and thus some codewords are more noise-resistant than others. As a result, the all-zero codeword cannot be assumed. Examples of non-symmetric channels arise in many real world applications. Some of them result from the physical properties of the communication medium, e.g. the z-channel for optical storage devices, and some are due to pre-processing before transmission, which converts symmetric physical channels into non-symmetric logical channels, e.g. the bit-level sub-channels in BICM systems.

In Chapter 4, instead of using a larger coset code ensemble as in [57], we circumvent this difficulty of codeword-dependent performance by averaging over all valid codewords, which is straightforward and has a practical interpretation as the averaged error probability (over the entire codebook). Our results apply to arbitrary binary-input, memoryless, symbol-dependent channels (e.g., z-channels, binary non-symmetric channels (BNSCs), non-symmetric Gaussian channels, etc.) and can be generalized to LDPC codes over $\mathsf{GF}(q)$ or $\mathbb{Z}_m$ [14, 15, 112]. The theorem of convergence to *perfect projection* is provided to justify this codeword-averaged approach in conjunction with the existing theorems. New results on monotonicity, symmetry, stability, and convergence rate analysis of the codeword-averaged density evolution method are also provided. Our approach based on the linear LDPC code ensemble will be linked to that of the coset code ensemble [57] by proving the typicality of linear LDPC codes when the minimum check node degree is sufficiently large, which was first conjectured in [51]. All of the above generalizations are based on the convergence to perfect projection, which will serve also as an essential theoretical foundation for the BP algorithms even when only symmetric channels are considered. To be more precise, the perfect projection condition guarantees the local optimality of BP decoders for either symmetric or non-symmetric channels, which partially explains the high performance of BP when applied to inference networks with cycles. Most of our results in this area are published in [111].

1.2.3 Finite-Dimensional Bounds on LDPC Codes with Belief Propagation Decoding

The BP decoding algorithm [86] (or equivalently the sum-product algorithm [61]) is one of the major components in modern capacity-approaching ECCs. The BP algorithm, which was originally designed to efficiently compute the *a posteriori* probability in cycle-free inference networks, uses distributed local computation to approximate the *a posteriori* probability when the inference network of interest has cycles [118]. As one of the most powerful analysis tools of BP, DE [92] traces the density of the log likelihood ratio (LLR) on an iteration-by-iteration basis. The underlying quantity of interest, the density of LLR, is of infinite dimension and is a sufficient statistic completely describing arbitrary binary-input memoryless channels. As a result, DE admits more accurate performance prediction at the cost of higher computational complexity.

Even after the efficient implementation of DE by moving into the LLR domain and by using sophisticated representation methods [54], a one-dimensional iterative formula (or at most finite-dimensional formulae) to "approximate/substitute" the density evolution is very appealing since it reduces significantly the computational complexity of code degree optimization [30]. Several approximation formulae have been proposed including Gaussian approximations [31, 71], BEC approximations, reciprocal channel approximations [29], and

the EXtrinsic Information Transfer (EXIT) chart analysis [105]. The finite dimensionality also helps the analysis of the message passing decoder [12, 22].

Contrary to the approximations, rigorous iterative upper and lower bounds generally sacrifice the threshold predictability for specific channel models in exchange for guaranteed universal performance for arbitrary channel models. Many results have been found for binary-input/symmetric-output (BI-SO) channels, including Burshtein *et al.* [22] on the expected soft bit value (ESB) for the MAP decoder, Khandekar *et al.* [59] on the Bhattacharyya noise parameter (BNP), and Land *et al.* [69] and Sutskover *et al.* [103] on the mutual information. For binary-input/non-symmetric-output (BI-NSO) channels, a loose one-dimensional iterative upper bound on the BNP was provided in [111], which was used to derive the stability condition of general BI-NSO channels and to upper bound the asymptotic convergence rate of the bit error probability. Bennatan *et al.* [14] used an iterative upper bound to derive the stability conditions for $\mathsf{GF}(q)$-based LDPC codes when q is a power of a prime number.

Our results in Chapter 5 focus either on bounds applicable to a broader class of channels, including BI-NSO channels and $\mathbb{Z}_m$ channels, or on bounds that outperform existing ones, including the best one-dimensional bound to date that is tight for BSCs. All our bounds are based on ESB or BNP of the channel of interest. We also propose a unified framework, which casts the iterative bounding problem into an infinite-dimensional linear programming problem and thus facilitates systematic search for new bounds by solving the corresponding linear optimization problem analytically. One of the most important implications of these finite dimensional bounds is a sufficient stability condition, and we will complement by a matched necessary stability condition for both binary and $\mathbb{Z}_m$ channels. This new stability condition for $\mathbb{Z}_m$ channels suggests a look back to the high-order-alphabet-coded modulation, and a close-to-capacity scheme based on $\mathbb{Z}_m$ LDPC codes is proposed in Chapter 6. This new scheme has higher performance than BICM with similar complexity.

Chapter 2

Side Information in Bandit Problems

2.1 General Formulation

The classical two-armed bandit problem can be described in the context of finding the optimal choice between two slot machines, in which the reward distribution of each machine is unknown. Let Y_t^1 and Y_t^2 denote the respective i.i.d. rewards at time t from machines 1 and 2. The reward function is then defined as follows,

$$W_\phi(t) = \sum_{\tau=1}^{t} \alpha_\tau \left(1_{\{\phi_\tau = 1\}} Y_\tau^1 + 1_{\{\phi_\tau = 2\}} Y_\tau^2 \right),$$

where $1_{\{\cdot\}}$ is the indicator function, ϕ_t, taking values in $\{1, 2\}$, is the player's strategy at time t, and $\{\alpha_\tau\}$ is a predefined discount sequence. It is worth emphasizing that the strategy (decision rule) ϕ_t depends only on the history until time $t-1$, namely, $\{\phi_\tau\}$ is a predictable random process constructed by the decision maker.

With the assumption that the distributions of $\{Y_\tau^1\}$ and $\{Y_\tau^2\}$ are unknown to the player, the knowledge about which arm yields higher reward can only be gathered from sampling both arms often enough. However, this task of learning both arms inevitably limits the opportunity of pulling the more rewarding arm. To be more explicit, the unknown reward distribution is often parameterized as F_θ, where the i.i.d. rewards $\{Y_\tau^1\}$ and $\{Y_\tau^2\}$ are governed by F_{θ_1} and F_{θ_2}. The decision maker has complete knowledge of the entire family $\{F_\theta\}_{\theta \in \mathbf{\Theta}}$, where $\mathbf{\Theta}$ is the set of all possible values of θ. The underlying parameter pair $C_0 = (\theta_1, \theta_2)$, taking values in $\mathbf{\Theta}^2$, is unknown, and must be learned through even sampling. Our goal is to maximize $W_\phi(t)$ under various conditions and discount sequences.

There are three typical choices of the discount sequences, including the finite horizon setting: $\alpha_\tau = 1_{\{\tau \le t_0\}}$, the infinite geometric discount sequence: $\alpha_\tau = r^\tau$, $r < 1$, or the situations with uniform discounting: $\alpha_\tau = 1, \forall \tau \in \mathbb{N}$. One common optimality condition associated with the first two types of discount sequences is to maximize the total expected reward $\lim_{t\to\infty} \mathsf{E}\{W_\phi(t)\}$. Under the finite horizon setting, the bandit problem can be solved optimally by the backward induction method [20]. Gittins showed that under the setting of geometric discount sequences, this optimal decision problem can be converted into an index-computing problem, and the optimal decision rule is to select/pull the arm with the highest index [42, 48, 49]. In this chapter, we will focus on the third choice of discount

sequences: $\alpha_\tau = 1, \forall \tau$, for which the optimality condition becomes maximizing the growth rate of expected reward $\mathsf{E}\{W_\phi(t)\}$.

Let μ_1 and μ_2 denote the expected returns of arms 1 and 2 under distributions F_{θ_1} and F_{θ_2}. By Wald's lemma, $\mathsf{E}\{W_\phi(t)\}$ can be rewritten as:

$$\mathsf{E}\{W_\phi(t)\} = t \cdot \max(\mu_1, \mu_2) - |\mu_1 - \mu_2| \cdot \mathsf{E}\{T_{inf}(t)\}, \tag{2.1}$$

where $T_{inf}(t)$ is the total number of samples taken on the inferior arm up to time t: $T_{inf}(t) = \sum_{\tau=1}^{t} 1\{\phi_\tau \neq \arg\max(\mu_1, \mu_2)\}$. From (2.1), it can be shown that maximizing the growth rate of $\mathsf{E}\{W_\phi(t)\}$ is equivalent to minimizing the growth rate of $\mathsf{E}\{T_{inf}(t)\}$. Since the term $|\mu_1 - \mu_2| \cdot \mathsf{E}\{T_{inf}(t)\}$ represents the expected cost of not knowing the preference between μ_1 and μ_2, it is often called the "regret", and is commonly considered in the literature of bandit problems. For notational simplicity, in this thesis, we will study the expected inferior sampling time $\mathsf{E}\{T_{inf}(t)\}$ instead of regret. It is worth noting that all expectations used in Chapters 2 and 3 depend on the unknown F_{θ_1} and F_{θ_2} and thus are functions of the parameter pair $C_0 = (\theta_1, \theta_2)$. Hence the terms $\mathsf{E}\{T_{inf}(t)\}$ and $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ are used interchangeably.

2.1.1 Uniformly Good Rules

Following the conservative setting of uniform discounting sequences, Lai and Robbins [67] considered bandit problems with no *a priori* distribution on the parameter set $\mathbf{\Theta}^2$. Taking a min-max approach, this formalism leads to the following definition of uniformly good rules.

Definition 2.1 (Uniformly Good Rules [67]) *An allocation rule is uniformly good if for every possible configuration pair $C_0 = (\theta_1, \theta_2)$, $\mathsf{E}_{C_0}\{T_{inf}(t)\} = o(t^\alpha)$, $\forall \alpha > 0$.*

Namely, an adaptive decision rule $\{\phi_\tau\}$ is uniformly good if it is able to simultaneously suppress the growth rates for all possible configurations. There exists a $\log(t)$ lower bound for this simultaneous suppression of $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ (or of the regret), which was proved under various settings [9, 66, 67]. A simple version of these existing $\log(t)$ lower bound theorems is stated as follows.

Theorem 2.1 ($\log(t)$ Lower Bound) *For any uniformly good rule $\{\phi_\tau\}$, $T_{inf}(t)$ satisfies*

$$\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_{inf}(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C_0}}\right) = 1, \ \forall \epsilon > 0,$$

and

$$\liminf_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}},$$

where K_{C_0} is a constant depending on C_0. If $\arg\max(\mu_1, \mu_2) = 2$, then K_{C_0} is defined[1] *as:*

$$K_{C_0} = \inf\{I(\theta_1, \theta) : \forall \theta, \mu_\theta > \mu_{\theta_2}\}, \tag{2.2}$$

where $I(\theta_1, \theta) = \mathsf{E}_{\theta_1}\left\{\log\left(\frac{dF_{\theta_1}}{dF_\theta}\right)\right\}$ is the Kullback-Leibler (K-L) information number between F_{θ_1} and F_θ, and μ_θ is the expected reward under F_θ. The expression for K_{C_0} for the case in which $\arg\max(\mu_1, \mu_2) = 1$ can be obtained by symmetry.

[1] Throughout Chapters 2 and 3, we will adopt the conventions that the infimum of the null set is ∞, and $\frac{1}{\infty} = 0$.

The asymptotic sharpness of the above lower bound is also proved in the above papers:

Theorem 2.2 (Asymptotic Sharpness) *Under certain regularity conditions,*[2] *the above lower bound is asymptotically sharp. That is, given the family of possible distributions* $\{F_\theta\}$*, there exists a single decision rule* $\{\phi_\tau\}$ *such that for all possible configuration pairs* $C_0 = (\theta_1, \theta_2)$*,*

$$\limsup_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

where K_{C_0} *is the same as defined in* Theorem 2.1*.*

Remark: A simple decision rule may sample the inferior arm a *finite* number of times (depending on the sample path), and thereafter stick to the seemingly superior arm indefinitely. For such type of rules, we may have $\lim_{t\to\infty} T_{inf}(t) < \infty$ almost surely for certain values of C_0. However, since no forced sampling is performed after a finite amount of time, one can prove that for some other C_0', $\mathsf{E}_{C_0'}\{T_{inf}(t)\}$ grows linearly, and these rules are thus not *uniformly* good. A uniformly good rule, on the other hand, must always be skeptical, and keeps sampling the other arm infinitely often. Theorem 2.1 guarantees that the probability of the inferior sampling time $T_{inf}(t)$ being no less than $\frac{\log(t)}{K_{C_0}}$ converges to one as t tends to infinity. In other words, the forced sampling times must grow at least on the order of $\log(t)$ with the minimum constant $1/K_{C_0}$.

Henceforth we consider only uniformly good rules and regard other rules as uninteresting. As discussed, by limiting our focus to uniformly good rules, the possibility of almost sure finiteness of $T_{inf}(t)$ is sacrificed,[3] but acceptable performance is guaranteed for all possible C_0. Further results on uniformly good rules within slightly different settings can be found in [2, 3, 4, 9, 10, 56, 63, 67].

The intuition behind the $\log(t)$ lower bound is as follows. Suppose $\mu_1 < \mu_2$ and consider another configuration $C' = (\theta, \theta_2)$ such that $\mu_{\theta'} > \mu_2$. It can be shown that if under configuration $C_0 = (\theta_1, \theta_2)$, $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ is less than the specified $\log(t)$ lower bound, $\mathsf{E}_{C'}\{T_{inf}(t)\}$ must be greater than $\mathcal{O}(t^\alpha)$ for some $\alpha > 0$, which contradicts the assumption that $\{\phi_\tau\}$ is uniformly good. The detailed proof, presented in Appendix B.1 for a more general setting, is similar to the proof of Stein's lemma in hypothesis testing between H_0 and H_1, the change-of-measure argument between P_1(accept H_0) and P_0(accept H_0).

2.1.2 Estimation-Based Uniformly Good Rules (EBUG Rules)

In this subsection, we further characterize a specific subset of uniformly good rules for traditional bandit problems: Estimation-Based Uniformly Good Rules (EBUG Rules). EBUG rules are designed for the cases of a finite parameter space $\mathbf{\Theta}$, and this type of decision rules possesses the following three properties when applied to traditional bandit problems.

1. After time t, an estimate $\hat{C}_t = (\hat{\theta}_1, \hat{\theta}_2)$ is constructed and is used to make the decision ϕ_{t+1} for time $t+1$. To be more explicit, $\hat{C}_t$ is generated by the results for $\tau \in [1, t]$, and ϕ_{t+1} is a function of $\hat{C}_t$. (ϕ_{t+1} does not necessarily to take the myopic approach, and can perform forced sampling as well.)

[2] If the parameter space is finite, Theorem 2.2 always holds. If $\mathbf{\Theta}$ is continuous, the required regularity conditions concern the unboundedness and the continuity of μ_θ w.r.t. θ and the continuity of $I(\theta_1, \theta)$ w.r.t. μ_θ.

[3] It will be shown in Chapters 2 and 3 that under certain scenarios, the almost sure finiteness of $T_{inf}(t)$ can be recovered by exploiting the side information.

2. The expected duration over which $\hat{C}_t \neq C_0$ is finite, namely,

$$\lim_{t\to\infty} \mathsf{E}_{C_0}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau \neq C_0\}\right\} < \infty.$$

3. The expected duration over which $\hat{C}_t = C_0$ and $\phi_{t+1} \neq \arg\max(\mu_1, \mu_2)$ is upper bounded by $\frac{\log(t)}{K_{C_0}}$, namely,

$$\limsup_{t\to\infty} \frac{\mathsf{E}_{C_0}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \phi_{\tau+1} \neq \arg\max(\mu_1,\mu_2)\}\right\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

where K_{C_0} is defined the same as in (2.2).

Definition 2.2 (Estimation-Based Uniformly Good Rules (EBUG Rules)) *For a traditional bandit problem with finite $\mathbf{\Theta}$, a decision rule $\{\phi_\tau\}$ is estimation-based uniformly good (EBUG) if it possesses the above three properties.*

Obviously, an EBUG rule $\{\phi_\tau\}$ is uniformly good and meets the $\log(t)$ lower bound in Theorem 2.1. One detailed construction of a EBUG rule $\{\phi_\tau\}$ can be found in [3]. In Chapters 2 and 3, EBUG rules (for traditional bandit problems) will serve as constituent components in designing new composite decision rules dealing with the side-information-aided bandit problems.

2.1.3 I.i.d. Side Information

A common scenario in practice is that before making the decision ϕ_t, another random variable X_t, which is i.i.d. and takes values in $\mathbf{X}$, can be observed. Suppose at time instant t, $X_t = x$. The rewards (Y_t^1, Y_t^2) are then governed by the conditional distributions[4] $F_{\theta_1}(\cdot|X_t = x)$ and $F_{\theta_2}(\cdot|X_t = x)$, and have conditional expected return $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$. For each configuration pair $C_0 = (\theta_1, \theta_2)$ and each i, the sequence of vectors $\{(X_\tau, Y_\tau^i)\}_\tau$ is i.i.d. and the marginal distribution of each vector is governed by $G_{C_0}(dx)F_{\theta_i}(dy|x)$, where the families $\{G_C\}_{C\in\mathbf{\Theta}^2}$ and $\{F_\theta(\cdot|x)\}_{\theta\in\mathbf{\Theta}}$ are known to the decision maker, but the true value of the corresponding index pair C_0 must be learned through experiments.

Note 1: The concept of the i.i.d. bandit is now extended to the assumption that the sequence of vectors $\{(X_\tau, Y_\tau^i)\}_\tau$ is i.i.d. The unconditioned marginal sequence $\{Y_\tau^i\}$ is still i.i.d. However, rather than facing unconditional marginals, the decision maker is now facing the conditional distribution of Y_t^i, which is a function of the observed side information X_t.

Note 2: Under this new framework, the best arm under configuration C_0 given $X_t = x$ can be defined as $M_{C_0}(x) := \arg\max(\mu_1(x), \mu_2(x))$. The inferior sampling time $T_{inf}(t)$ is defined slightly differently from its traditional counter part (without side information) as

$$T_{inf}(t) = \sum_{\tau=1}^{t} 1\left\{\phi_\tau \neq M_{C_0}(X_\tau)\right\}.$$

[4]The term "side information" implies that the distribution of X_t depends on the upcoming rewards Y_t^1 and Y_t^2. Nevertheless, since X_t is observed before deciding which arm to pull, it is more convenient to reverse the conditional probability using Bayes' formula and view X_t as the basic quantity while letting the distributions of Y_t^1 and Y_t^2 depend on the value of X_t.

For the remaining part of this chapter, we will focus on minimizing the growth rate of $\mathsf{E}\{T_{inf}(t)\}$ under this new framework.

Note 3: The traditional two-armed bandit without side information X_t can be viewed as a degenerate case in which the range of X_t contains only one element: $\mathbf{X} = \{x_0\}$. All of our results collapse to the classical Theorems 2.1 and 2.2 in this degenerate case.

2.1.4 Four Different Types of Interaction

The interactions between the side information and the rewards from each arm can be characterized into the following four cases.

1. **Direct Information:** In this case, X_t provides information directly about the underlying configuration $C_0 = (\theta_1, \theta_2)$, which allows a type of separation between learning and control. This has a dramatic effect on the achievable inferior sampling time. Specifically, estimating $C_0 = (\theta_1, \theta_2)$ by observing $\{X_\tau\}$ and using the estimate $\hat{C} = (\hat{\theta}_1, \hat{\theta}_2)$ to make the decision, results in bounded expected inferior sampling time.

Sometimes, relying on X_t to reveal the information about $C_0 = (\theta_1, \theta_2)$ is overoptimistic. In many cases, the distribution of $\{X_\tau\}$ is not a function of C_0 and we are not able to learn C_0 through $\{X_\tau\}$. However, different values of the side information X_t will result in different conditional distributions of the rewards Y_t^i. By exploiting this new structure (observing X_t in advance), we can hope to do better than the case without any side information.

A physical meaning about this scenario (constant distribution on $\{X_\tau\}$) is that a two-armed bandit with the side information ranging from x_1 to x_n can be viewed as a set of n different two-armed *sub-bandit machines* indexed from x_1 to x_n. The player does not have control over which sub-bandit machine he is going to play, which is determined by tossing a dice with n faces. However, by observing X_t, the player knows which machine (out of the n different ones) he is facing now before selecting which arm to play. The connection between these sub-machines is that they share the same common configuration pair (θ_1, θ_2), so that the rewards observed from one sub-machine provide information on the common (θ_1, θ_2), which can then be applied to *all* of the others (different values of X_t). This is the key aspect that makes this setup distinct from simply having many independent bandit problems with random access opportunity.

It is worth noting that within this setting, the most rewarding arm at time t, $M_{C_0}(X_t)$, is in general a function of both the underlying configuration pair C_0 and the side information X_t. We consider the following three cases of different relationships among the most rewarding arm, C_0, and X_t.

2. **For all C_0, the best arm is a function of X_t:** In this case, for *all* configurations (θ_1, θ_2), arm 1 is preferred for some values of X_t while arm 2 is preferred for other values of X_t. Surprisingly, we exhibit an algorithm that achieves *bounded* expected inferior sampling time in this case. Woodroofe's result [116] can then be viewed as a special case of this scenario.

3. **For all C_0, the best arm is not a function of X_t:** In this case, for *all* configurations (θ_1, θ_2), one of the arms is always preferred regardless of the value of X_t. However, we note that all sub-bandit machines can be used to *learn* $C_0 = (\theta_1, \theta_2)$, and we can then postpone our learning until it is the most advantageous to us. We show that,

asymptotically, our performance will be governed by the most "informative" bandit (among the different values of X_t).

4. **Mixed Case:** This is a general case that combines the previous two, and contains the main contribution of this chapter. For some possible configurations, one arm may always be preferred (for any X_t), while for other possible configurations, the preferred arm is a function of X_t. We exhibit a single algorithm that achieves the best possible in either case. That is, if the best arm is a function of X_t, it achieves bounded expected inferior sampling time as in Case 2, while if the underlying configuration is such that one arm is always preferred, then we achieve the results in Case 3.

In the next four sections, we are going to explore these four cases respectively. For the sake of readability, formal statements of our results are provided in each section, while details of the proofs are included in the appendix.

2.1.5 Notation

Some useful notation is defined as follows and will be used throughout Chapters 2 and 3.

- For any configuration pair $C_0 = (\theta_1, \theta_2)$, we may use $1(C_0) := \theta_1$ and $2(C_0) := \theta_2$ to denote the first and second coordinates of the configuration pair C_0. This notation may look peculiar at the first sight but will be very useful during our development. An illustration of the advantage of this notation is as follows: we can write $\mu_{2(C_0)}(x) = \mu_{\theta_2}(x)$ and $F_{1(C_0)}(\cdot|x) = F_{\theta_1}(\cdot|x)$.
- $\mathbf{\Theta}$, the set of possible values of θ, is a subset of real numbers.
- $T_i(t)$ denotes the number of times arm i is pulled until time t.
- $I(\theta_1, \theta_2|x) := \mathsf{E}_{F_{\theta_1}(\cdot|x)}\left\{\log\left(\frac{F_{\theta_1}(\cdot|x)}{F_{\theta_2}(\cdot|x)}\right)\right\}$ denotes the Kullback-Leibler information number between the conditional distributions $F_{\theta_1}(\cdot|x)$ and $F_{\theta_2}(\cdot|x)$.

Necessary notation and several quantities of interest are summarized in Table 2.1. We assume that all the given expectations exist and are finite.

2.2 Case 1: Direct Information

In this setting, the side information X_t directly reveals information about the underlying configuration pair $C_0 = (\theta_1, \theta_2)$ in the following way.

Direct information: $G_{C_1} = G_{C_2}$ iff $C_1 = C_2$.

As a result, observing the empirical distribution of X_t gives us useful information about the underlying parameter pair C_0. Thus this is a type of identifiability condition.

Examples:

- $\mathbf{\Theta} = (0, 0.5)$ and $\mathbf{X} = \{x_1, x_2, x_3\}$.

$$\mathsf{P}_{(\theta_1,\theta_2)}(X_t = x_k) = \begin{cases} \theta_k & \text{if } k = 1, 2 \\ 1 - \theta_1 - \theta_2 & \text{otherwise} \end{cases}.$$

- $\mathbf{\Theta} = (0, \infty)$ and $\mathbf{X} = [0, 1]$. $X_t \sim \beta(\theta_1, \theta_2)$ is beta distributed with parameters (θ_1, θ_2).

Table 2.1: Glossary of the bandit problem

Not'n	Description
$\boldsymbol{\Theta}, \boldsymbol{\Theta}^2$	$\boldsymbol{\Theta} \subseteq \mathbb{R}$ is the set of all possible θ; $\boldsymbol{\Theta}^2$ is the set of parameter pairs.
$\mathbf{X}$	$\mathbf{X}$ is the set of all possible values of the side information X_t.
$G_C(x)$	The marginal distribution of the i.i.d. $\{X_\tau\}$ under configuration C.
$F_{\theta_i}(y\|x)$	The conditional distribution of Y_t^i, under parameter θ_i.
$\mu_\theta(x)$	The conditional expectation of the reward, $\mu_\theta(x) = \mathsf{E}_\theta\{Y\|x\} = \int y F_\theta(dy\|x)$.
$1(C_0)$, $2(C_0)$	The first and the second coordinates of the configuration pair C_0, i.e. $1(C_0) = \theta_1$, $2(C_0) = \theta_2$. For example: $\mu_{1(C_0)}(x) = \mu_{\theta_1}(x)$.
$M_C(x)$	The index of the preferred arm, i.e. $M_C(x) := \arg\max_{i=1,2}\{\mu_{i(C)}(x)\}$.
ϕ_t	The decision rule taking values in $\{1, 2\}$ and depending only on the past outcomes and the current side information X_t.
$T_{inf}(t)$	The total number of samples taken on the inferior arm up to time t, $T_{inf}(t) = \sum_{\tau=1}^{t} 1_{\{\phi_\tau \neq M_{C_0}(X_\tau)\}}$.
$T_i(t)$	The total number of samples taken on arm i up to time t, $T_i(t) = \sum_{\tau=1}^{t} 1_{\{\phi_\tau = i\}}$.
$I(P, Q)$	The Kullback-Leibler (K-L) information number between distributions P and Q, $I(P, Q) = \mathsf{E}_P\left\{\log\left(\frac{dP}{dQ}\right)\right\}$.
$I(\theta_1, \theta_2\|x)$	The conditional K-L information number, $I(\theta_1, \theta_2\|x) = I(F_{\theta_1}(\cdot\|x), F_{\theta_2}(\cdot\|x))$.

2.2.1 Scheme with Bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$

Consider the following condition.

Condition 2.1 *For any fixed C_0 and any (estimate) sequence $\{\hat{C}_\tau\}$ with limit $\lim_{\tau\to\infty} \hat{C}_\tau = C_0$, there exists t_0 such that for all $x \in \mathbf{X}$ and $t > t_0$, $M_{\hat{C}_t}(x) = M_{C_0}(x)$. Or equivalently,*

$$\inf\left\{\rho(G_{C_0}, G_{C_e}) : C_e \in \boldsymbol{\Theta}^2, \exists x, M_{C_e}(x) \neq M_{C_0}(x)\right\} > 0,$$

for the Prohorov metric[5] *ρ on the space of distributions.*

Here we provide two examples satisfying Condition 2.1:

- *Example 1:* If (1) $\mathbf{X}$ is finite, and (2) $\forall x \in \mathbf{X}$, $\mu_\theta(x)$ is continuous with respect to (w.r.t.) θ, Condition 2.1 is satisfied.
- *Example 2:* If $F_\theta(\cdot|x) \sim \mathcal{N}(\theta x, 1)$ is Gaussian distributed with mean θx and variance 1, then Condition 2.1 is satisfied.

Under this condition, we obtain the following result.

Theorem 2.3 (Bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$) *If* Condition 2.1 *is satisfied, then there exists an allocation rule $\{\phi_\tau\}$, such that $\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$ and $\lim_{t\to\infty} T_{inf}(t) < \infty$ almost surely (a.s.).*

The intuition behind is as follows. Condition 2.1 implies that whenever our estimate $\hat{C}_t$ is close enough to $C_0 = (\theta_1, \theta_2)$, then $\forall x \in \mathbf{X}$, the preference order between $\mu_{1(\hat{C}_t)}(x)$ and $\mu_{2(\hat{C}_t)}(x)$ is the same as that of $\mu_{1(C_0)}(x)$ and $\mu_{2(C_0)}(x)$. Therefore, all myopic decisions based on $\hat{C}_t$ become optimal. We then need only worry about the expected duration for which the

[5]The formal definition and some notes about the Prohorov metric are stated in Appendix A.

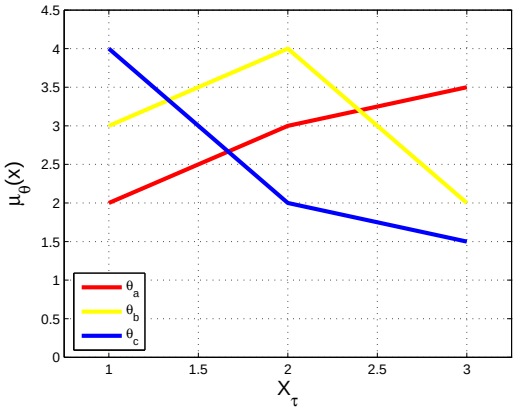

Figure 2.1: The best arm at time t *always* depends on the side information X_t. That is, for any possible pair (θ_1, θ_2) the two curves, $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$, (w.r.t. x) always intersect each other.

estimate $\hat{C}_t$ is far from C_0. By Sanov's theorem,[6] which characterizes the large deviation principle satisfied by empirical measures on $\mathbb{R}$, it is easy to construct an estimate $\hat{C}_t$ such that the probability that $\hat{C}_t$ stays away from C_0 decreases exponentially. This implies that the expected duration of $\left\{|\hat{C}_t - C_0| \geq \epsilon\right\}$ is bounded and $\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$.

Instead of proving Theorem 2.3 directly, a more general version will be stated and proved in Section 3.2 and Appendix C.1.

Note: the information directly revealed by X_t helps the sequential control scheme surpass the $\log(t)$ lower bound stated in Theorem 2.1. This significant improvement (bounded expected inferior sampling time) is due to the fact that the dilemma between learning and control no longer exists in this case.

2.3 Case 2: Best Arm As A Function Of X_t

For all of the following sections (Sections 2.3 through 2.5), we consider only the cases in which observing X_t does not reveal any information about C_0, but reveals only information about the upcoming reward Y_t^i. That is,

No direct information: G_{C_0} is a constant distribution whatever the value of C_0; we use $G := G_{C_0}$ as shorthand notation.

Three further refinements regarding the relationship between $M_C(x)$ and x will be discussed separately (each in one section).

In this section, we assume further that for all possible C, the side information X_t is *always* able to change the preference order as shown in Figure 2.1. Namely,

- For all $C \in \mathbf{\Theta}^2$, there exist x_1 and x_2 such that $M_C(x_1) = 1$ and $M_C(x_2) = 2$.

For future references, if such x_1 and x_2 exist, we say the underlying configuration is *implicitly revealing*. We use Figure 2.1 as an illustrative example and consider the underlying parameter being $C_0 = (\theta_1, \theta_2) = (\theta_a, \theta_b)$. When $X_t = 1$ or 2, arm 1 is more favorable than

[6] The required version of Sanov's theorem is stated in Appendix A

arm 2. When $X_t = 3$, arm 2 is more favorable. The preference order changes when different values of X_t occur, and we say $C_0 = (\theta_a, \theta_b)$ is implicitly revealing. It is easy to check that the other configurations (θ_a, θ_c) and (θ_b, θ_c) are implicitly revealing as well.

The necessary regularity conditions are as follows.

1. $\mathbf{X}$ is a finite set and $\mathsf{P}_G(X_t = x) > 0$ for all $x \in \mathbf{X}$.
2. $\forall \theta_1, \theta_2, x$, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.
3. $\forall x$, $\mu_\theta(x)$ is continuous w.r.t. θ.

The first condition embodies the idea of viewing each possible value of X_t as the index of several sub-bandit machines, which also simplifies our proof. The second condition guarantees that all these sub-bandit problems are non-trivial, with *non-identical* pairs of arms.

Example:

- $\mathbf{\Theta} = (0, \infty)$, $\mathbf{X} = \{-1, 1\}$, and the conditional reward distribution $F_\theta(\cdot|x) \sim \mathcal{N}(\theta x, 1)$.

2.3.1 Scheme with Bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$

Theorem 2.4 (Bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$) *If the above conditions are satisfied, there exists an allocation rule $\{\phi_\tau\}$ such that*

$$\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty, \quad \forall C_0 \in \mathbf{\Theta}^2.$$

Such a rule is obviously uniformly good and surpasses the $\log(t)$ *lower bound for traditional bandit problems.*

Remark: Although the side information X_t does not reveal any information about C_0 in this setting, the alternation of the best arm as the i.i.d. X_t takes on different values x makes it possible to always perform the control part, $\phi_t = M_{\hat{C}_{t-1}}(X_t)$, and simultaneously sample both arms often enough. Since the information about *both* arms will be implicitly revealed (through the alternation of $M_{C_0}(X_t)$), the dilemma of learning versus control no longer exists, and a significant improvement ($\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$) is obtained over the $\log(t)$ lower bound in Theorem 2.1.

We construct an allocation rule with bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ as in Algorithm 1. The intuition as to why the proposed scheme has bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ is as follows. The forced sampling at Line 4 may induce many inferior samplings but on the other hand ensures there are enough samples (at least $\mathcal{O}(\sqrt{t})$) on both arms, which can be used to obtain sufficiently good estimates of C_0. Based on these good estimates, the myopic action, $\phi_{t+1} = M_{\hat{C}_t}(X_{t+1})$ (sampling the seemingly better arm at Line 6), will result in very few inferior samplings. So the remaining question is whether the inferior sampling induced during the forced sampling (Line 4) can be justified by the benefit of having very good estimates.

For traditional bandit problems, in which $|\mathbf{X}| = 1$, the answer is no. The forced sampling (Line 4) makes $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ be of the order of $\mathcal{O}(\sqrt{t})$ which is too often for a uniformly good rule. Contrary to the traditional two-armed bandits, when the underlying configurations are implicitly revealing, the best arm $M_{C_0}(x)$ varies from one outcome of X_t to the other. The myopic action and the even appearances of the i.i.d. $\{X_\tau\}$ will eventually make both $T_1(t)$ and $T_2(t)$ grow linearly with the elapsed time t. As a result, the forced sampling will rarely be triggered and the induced loss will be limited. $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ is bounded in

Algorithm 1 ϕ_{t+1}, the decision at time$t+1$

Variables: Denote $T_i^x(t)$ as the total number of time instants until time t when arm i has been pulled and $X_\tau = x$, i.e.

$$T_i^x(t) := \sum_{\tau=1}^{t} 1_{\{X_\tau = x, \phi_\tau = i\}},$$

and define $x_i^\star := \arg\max_x \{T_i^x(t)\}$ and $T_i^{x^\star}(t) := \max_x \{T_i^x(t)\}$.
Construct

$$\mathbf{C}_t \quad := \quad \left\{ C = (\theta_1, \theta_2) \in \mathbf{\Theta}^2 \,\middle|\, \sigma(C,t) \le \inf\{\sigma(C,t) : C \in \mathbf{\Theta}^2\} + \frac{1}{t} \right\}$$

with

$$\sigma(C,t) \quad := \quad \rho\left(F_{1(C)}(\cdot|x_1^\star), L_1^{x^\star}(t)\right) + \rho\left(F_{2(C)}(\cdot|x_2^\star), L_2^{x^\star}(t)\right),$$

where $L_i^x(t)$ is the empirical measure of rewards sampled from arm i at those time instants $\tau \le t$ when $X_\tau = x$. (As before $\rho(P,Q)$ is the Prohorov metric.) Arbitrarily choose $\hat{C}_t \in \mathbf{C}_t$.

Algorithm:

1: **if** $t+1 \le 6$ **then**
2: $\quad \phi(t+1) = (t \bmod 2) + 1$.
3: **else if** $\exists i$ such that $T_i(t) < \sqrt{t+1}$ **then**
4: $\quad \phi_{t+1} = i$.
5: **else**
6: $\quad \phi_{t+1} = M_{\hat{C}_t}(X_{t+1})$.
7: **end if**

(Note that Line 1 guarantees that there is only one i such that $T_i(t) < \sqrt{t+1}$.)

Algorithm 1, which demonstrates the benefit of exploiting the side information even when the side information reveals no information about C_0.

A detailed analysis of Algorithm 1 is provided in Appendix C.2, showing the boundedness of $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ for this scheme.

2.4 Case 3: Best Arm Is Not A Function Of X_t

Besides the assumption of the constant G, in this section, we consider the case in which for all $C \in \mathbf{\Theta}^2$, $M_C(x)$ is not a function of x. We thus can use $M_C := M_C(x)$ as shorthand notation. Figure 2.2 illustrates this setting.

The necessary regularity conditions are similar to those in Section 2.3:

1. $\mathbf{X}$ is a finite set and $\mathsf{P}_G(X_t = x) > 0$ for all $x \in \mathbf{X}$.
2. $\forall \theta_1, \theta_2, x$, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.

In this case, one arm is always better than the other no matter what value of X_t occurs and the information about $C_0 = (\theta_1, \theta_2)$ is not implicitly revealed as in Section 2.3. Therefore, the conflict between learning and control still exists. As expected, the growth rate of the expected inferior sampling time is again lower bounded by $\log(t)$, but with the additional help of X_t we can still see improvements over the traditional bandit problems.

To greatly simplify the notation, we also assume that

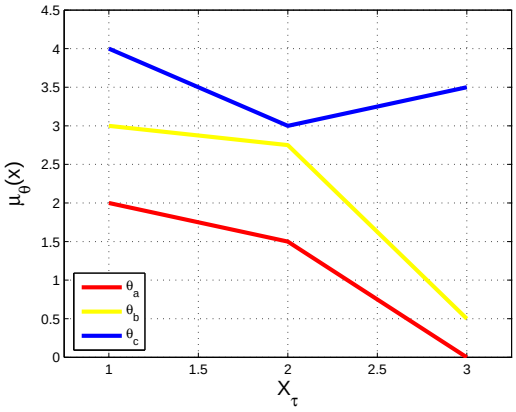

Figure 2.2: The best arm at time t *never* depends on the side information X_t. That is, for any possible pair, (θ_1, θ_2), the two curves, $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$, do not intersect each other. In this case, the optimal strategy is to postpone the forced sampling to the most informative time instants (sub-bandit machines).

3. For all x, the conditional expected reward $\mu_\theta(x)$ is strictly increasing w.r.t. θ.

This condition gives us the notational convenience that the order of $(\mu_{\theta_1}(x), \mu_{\theta_2}(x))$ is simply the same as the order of (θ_1, θ_2).

Example:

- $\mathbf{\Theta} = (1, \infty)$, $\mathbf{X} = \{1, 2, 3\}$, and the conditional reward distribution $F_\theta(\cdot|x) \sim \mathcal{N}(\theta x, 1)$.

2.4.1 A New $\log(t)$ Lower Bound

Theorem 2.5 (New $\log(t)$ Lower Bound) *Under the above assumptions, for any uniformly good rule $\{\phi_\tau\}$, $T_{inf}(t)$ satisfies*

$$\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_{inf}(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C_0}}\right) = 1, \ \forall \epsilon > 0,$$

$$\textit{and} \qquad \liminf_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}, \tag{2.3}$$

where K_{C_0} is a constant depending on C_0. If $M_{C_0} = 2$, then $T_{inf}(t) = T_1(t)$. The constant K_{C_0} is different than that of the traditional bandit problem in Theorem 2.1 *and can be expressed as follows.*

$$K_{C_0} = \inf_{\theta:\theta>\theta_2} \sup_{x\in\mathbf{X}} \{I(\theta_1, \theta|x)\}. \tag{2.4}$$

The expression for K_{C_0} for the case in which $M_{C_0} = 1$ can be obtained by symmetry.

Note 1: If the decision maker is not able to access the side information X_t, the player will then face the *unconditional* reward distribution $\int_x F_{\theta_i}(dy|x)G(dx)$ rather than $F_{\theta_i}(dy|x)$. Let $I(\theta_1, \theta_2)$ denote the Kullback-Leibler information between the *unconditional* reward distributions. By the convexity of the Kullback-Leibler information, we have

$$\sup_x I(\theta_1, \theta|x) \geq \int_x I(\theta_1, \theta|x)G(dx) \geq I(\theta_1, \theta).$$

This shows that the new constant in front of $\log(t)$, in (2.4), is no larger than the corresponding constant in (2.2). Thus, the additional side information X_t generally improves the decision made in the bandit problem.

Note 2: This situation is like having several related bandit machines, whose reward distributions are all determined by the common configuration pair (θ_1, θ_2). The information obtained from one machine is also applicable to the other machines. If arm 2 is *always* better than arm 1 for all possible values of X_t, we can not rely on the random appearance of each value of X_t to direct us sample both arms evenly. The conflict between learning and control still exists. Therefore, we have to sample arm 2 most of the time (the control part), and force sample arm 1 once in a while (the learning part). With the help of the side information X_t, we can pull the seemingly better arm most of the time, and postpone our forced sampling (learning) to the most informative machine $X_t = \arg\max_x I(\theta_1, \theta|x)$. The constant in the $\log(t)$ lower bound in Theorem 2.1 can be further reduced to this new $\frac{1}{K_{C_0}}$ in (2.4).

Note 3: As we would expect, Theorem 2.5 collapses to Theorem 2.1 when $|\mathbf{X}| = 1$.

A detailed proof of Theorem 2.5 is provided in Appendix B.1.

2.4.2 Scheme Achieving the Lower Bound

Consider the following regularity conditions in addition to the first three for the $\log(t)$ lower bound in Theorem 2.5.

4. $\mathbf{\Theta}$ is finite.

5. The existence[7] of the saddle point in the expression of K_{C_0}, (2.4), is assumed. Namely, for all $\theta_1 < \theta_2$,

$$\inf_{\theta:\theta>\theta_2} \sup_{x\in\mathbf{X}} I(\theta_1, \theta|x) = \sup_{x\in\mathbf{X}} \inf_{\theta:\theta>\theta_2} I(\theta_1, \theta|x).$$

   In the game theoretical perspective, the bandit problem can be viewed as a zero-sum two-player game, in which the "nature" tries to select a θ to maximize the constant $\frac{1}{K_{C_0}}$ in front of the $\log(t)$ lower bound. On the other hand, the "decision maker" would like to minimize $\frac{1}{K_{C_0}}$ by selecting a proper x. The existence of the saddle point is then equivalent to the existence of the value of the game.

With the above conditions, we are able to construct a scheme attaining the new $\log(t)$ lower bound in Theorem 2.5. We therefore obtain the asymptotic sharpness result for Theorem 2.5.

Theorem 2.6 (Asymptotic Sharpness) *An adaptive decision rule can be explicitly constructed so that the* $\log(t)$ *lower bound in (2.3) is achieved for all possible* $C_0 \in \mathbf{\Theta}^2$. *This decision rule is thus uniformly good and asymptotically optimal.*

Theorem 2.6 will be proved by explicit construction of such a decision rule. It is worth noting that although the scheme proposed here is asymptotically optimal, such construction is not unique. Different constructions may lead to different initial convergence speed, which

[7] A sufficient condition for the existence of the value of the game is that θ is the dominant factor (compared to x) in determining the conditional distributions $F_\theta(\cdot|x)$. In many cases of interest, the parameter plays a more critical role in determining the distribution than the side information x. Therefore this condition on the value of the game is a reasonable assumption and is generally satisfied.

Algorithm 2 Φ_{t+1}, the composite decision rule at time $t+1$

1: **if** not all $\hat{C}_{x,t}$ are identical, **then**
2: $\quad\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$.
3: **else**
4: $\quad$Denote $\hat{C}_t = (\hat{\theta}_1, \hat{\theta}_2)$ as the common estimate for all $\mathsf{B}_x, \forall x \in \mathbf{X}$. Without loss of generality, we may assume $M_{\hat{C}_t} = 2$. The case that $M_{\hat{C}_t} = 1$ can be obtained by symmetry.
5: $\quad$**if** $X_{t+1} \neq x^* := \arg\max_x \inf_{\{\theta:\theta>\hat{\theta}_2\}} I(\hat{\theta}_1, \theta|x)$, **then**
6: $\quad\quad\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$.
7: $\quad$**else**
8: $\quad\quad\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$.
9: $\quad$**end if**
10: **end if**

‡ A tie-breaking mechanism is necessary while evaluating "arg max" in Line 5, and a natural choice of a randomized tie-breaking mechanism is sufficient for rigorous analysis. However, to minimize the distraction of this minor point, we assume here that no tie exists during the execution of this algorithm.

is of practical value when considering real world applications. We will use the EBUG rules $\{\phi_\tau\}$, designed for traditional bandit problems and discussed in Section 2.1.2, as building blocks to construct a composite $\{\Phi_\tau\}$ achieving this new $\log(t)$ lower bound.

Suppose $|\mathbf{X}| = k < \infty$. Using the values of X_t, we can partition the observed rewards Y_t^1 (or Y_t^2) into k sub-sequences, corresponding to different x's. Consider the sub-sequence obtained when X_t equals some x_0. At those time instants, the decision maker is facing $F_{\theta_1}(\cdot|x_0)$ and $F_{\theta_2}(\cdot|x_0)$, and thus this sub-sequence can be viewed as resulting from a traditional bandit problem with the family of possible distributions being $\{F_\theta(\cdot|x_0)\}_\theta$. For each x_0, we use B_{x_0} to denote the corresponding sub-bandit problem.

For example, if $X_1X_2X_3X_4\cdots = x_a x_b x_a x_c \cdots$, then after time $t = 4$, we have 2 samples in B_{x_a}, 1 sample in B_{x_b}, and 1 sample in B_{x_c}. One straightforward composite decision rule $\{\Phi_\tau\}$ is to apply an EBUG $\{\phi_{x,\tau}\}$ on each sub-bandit B_x. The resulting composite rule is uniformly good but does not yield sharp results matching the new $\log(t)$ lower bound in Theorem 2.5. A more sophisticated composite rule $\{\Phi_\tau\}$ attaining the new $\log(t)$ lower bound is constructed as in Algorithm 2, in which $\hat{C}_{x,t}$ denotes the corresponding estimates from the x-th constituent EBUG rule $\{\phi_{x,\tau}\}$ after time t.

Note: To perform rigorous analysis, the constituent $\{\phi_{x,\tau}\}$ must be fully encapsulated in Algorithm 2. Namely, only those samples obtained from performing $\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$ (Lines 2 and 8) can be counted as valid samples for $\{\phi_{x,\tau}\}$. In other words, the time instants when we let $\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$ (Line 6) must be excluded from the computation of $\hat{C}_{x,t}$ and $\phi_{x,t+1}$. Otherwise it may spoil the tightness of the original $\{\phi_{x,\tau}\}$. For example, suppose $X_1X_2X_3X_4\cdots = x_a x_b x_a x_c \cdots$. At time instants 1 and 2, $\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$, either Line 2 or Line 8, is executed. At time instants 3 and 4, $\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$, Line 6, is executed. From the sub-bandit problem point of view, we have only one sample in B_{x_a}, one sample in B_{x_b}, and no samples in B_{x_c}, and only those samples can be used to generate the corresponding value of $\hat{C}_{x,t}$ and $\phi_{x,t+1}$. Samples made at time instants 3 and 4 will be discarded.

This composite decision rule $\{\Phi_\tau\}$ is designed around the central concept of postponing the forced sampling to the most informative sub-bandit machine when $X_t = x^* := \arg\max_x I(\theta_1, \theta|x)$. A detailed analysis of Algorithm 2 is provided in Appendix C.3.

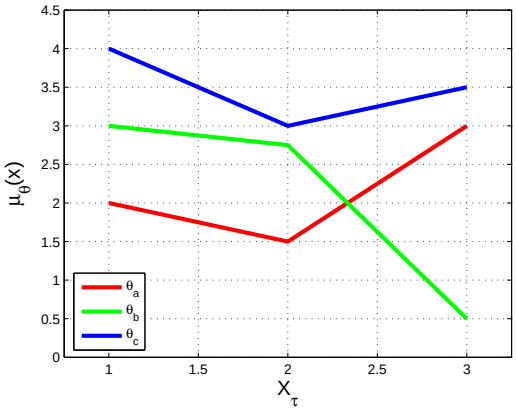

Figure 2.3: If $(\theta_1, \theta_2) = (\theta_a, \theta_b)$, the best arm depends on x, i.e. $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$ intersect each other as in Section 2.3. If $(\theta_1, \theta_2) = (\theta_b, \theta_c)$, the best arm does not depend on x, i.e. $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$ do not intersect each other as in Section 2.4.

2.5 Case 4: Mixed Case

In Sections 2.3 and 2.4, we dealt with the cases in which the distribution of X_t is constant. The main difference between Sections 2.3 and 2.4 is that in one case, for all possible C_0, X_t *always* changes the preference order. Or equivalently, all $C_0 \in \mathbf{\Theta}^2$ are implicitly revealing. While in the other, for all possible C_0, X_t *never* changes the order, namely, no $C_0 \in \mathbf{\Theta}^2$ is implicitly revealing. A more general case is a mixture of these two in which some C_0's are implicitly revealing while some are not. This section will focus on this mixed case and contain one of our major contributions for bandit problems with side information.

This situation is illustrated in Figure 2.3, in which if $C_0 = (\theta_a, \theta_b)$, the best arm depends on x, and the configuration (θ_a, θ_b) is implicitly revealing. If $C_0 = (\theta_b, \theta_c)$, the arm 2 is always more favorable regardless of the value of X_t, and the configuration (θ_b, θ_c) is not implicitly revealing. Without knowledge about the authentic underlying configuration C_0, we do not know whether C_0 is implicitly revealing or not. In view of the results of Sections 2.3 and 2.4, we would like to find a single scheme that is able to achieve bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ when applied to an implicitly revealing C_0, and to have the growth rate being $\mathcal{O}(\log(t))$ when applied to those C_0 which are not implicitly revealing. For the remaining parts of this section, we first show that we cannot expect better improvement by providing a conditional $\log(t)$ lower bound. We then prove this new conditional $\log(t)$ lower bound is attainable.

2.5.1 A New $\log(t)$ Lower Bound

The necessary regularity conditions are the same as those in Sections 2.3 and 2.4:

1. $\mathbf{X}$ is a finite set and $\mathsf{P}_G(X_t = x) > 0$ for all $x \in \mathbf{X}$.
2. $\forall \theta_1, \theta_2, x$, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.

Example:

- $\mathbf{\Theta} = (0, \infty)$, $\mathbf{X} = \{-1, 1\}$ and the conditional reward distribution $F_\theta(\cdot|x) \sim \mathcal{N}(\theta^2 - \theta x, 1)$. Then $C_0 = (\theta_1, \theta_2) = (0.1, 0.2)$ is implicitly revealing, but $C_0 = (0, 10)$ is not.

Theorem 2.7 (New $\log(t)$ Lower Bound) *Under the above assumptions, for any uniformly good rule $\{\phi_\tau\}$, if C_0 is not implicitly revealing, $T_{inf}(t)$ satisfies*

$$\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_{inf}(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C_0}}\right) = 1, \ \forall \epsilon > 0,$$

$$\textit{and} \qquad \liminf_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}, \tag{2.5}$$

where K_{C_0} is a constant depending on C_0. If $M_{C_0} = 2$, $T_{inf}(t) = T_1(t)$. The constant K_{C_0} is different than that of Theorem 2.5 *and can be expressed as follows.*

$$K_{C_0} = \inf_{\{\theta:\exists x_0,\ s.t.\ \mu_\theta(x_0)>\mu_{\theta_2}(x_0)\}} \sup_x \{I(\theta_1,\theta|x)\}.$$

The expression for K_{C_0} for the case in which $M_{C_0} = 1$ can be obtained by symmetry.

Note 1: Theorem 2.7 states only the $\log(t)$ lower bound when C_0 is not implicitly revealing. In the next subsection, we are going to show that when C_0 is implicitly revealing, we can achieve bounded $\mathsf{E}\{T_{inf}(t)\}$.

Note 2: The only difference between the lower bounds (2.3) and (2.5) is that, in (2.5), K_{C_0} has been changed from taking the infimum over $\{\theta \in \mathbf{\Theta} : \forall x, \mu_\theta(x) > \mu_{\theta_2}(x)\}$ to a larger set, $\{\theta \in \mathbf{\Theta} : \exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}$. The reason for this is as follows. Consider a θ for which there exists x_0 such that $\mu_\theta(x_0) > \mu_{\theta_2}(x_0)$. When the true configuration is $C' = (\theta, \theta_2)$ but the decision maker mistakes it as $C_0 = (\theta_1, \theta_2)$, a linear order of incorrect sampling will be incurred. To avoid this type of mistakes, a uniformly good rule has to sample often enough to distinguish C' from C_0, which is the reason why a broader class of competing distributions $C' = (\theta, \theta_2)$ must be considered.

A detailed proof is contained in Appendix B.2.

2.5.2 Scheme Achieving the Lower Bound

Consider the following two additional conditions[8] in addition to the previous two for the $\log(t)$ lower bound in Theorem 2.7.

3. $\mathbf{\Theta}$ is finite.

4. The existence of the saddle point is assumed, that is, for all (θ_1, θ_2),

$$\inf_{\{\theta:\exists x_0,\mu_\theta(x_0)>\mu_{\theta_2}(x_0)\}} \sup_x I(\theta_1,\theta|x) = \sup_x \inf_{\{\theta:\exists x_0,\mu_\theta(x_0)>\mu_{\theta_2}(x_0)\}} I(\theta_1,\theta|x).$$

With the above conditions, we are able to construct a single scheme such that it has bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ if C_0 is implicitly revealing, and saturates the lower bound (2.5) when being applied to such C_0 that is not implicitly revealing. We therefore obtain the asymptotic sharpness result for Theorem 2.7.

Theorem 2.8 (Asymptotic Sharpness) *An adaptive decision rule can be explicitly constructed so that the new conditional* $\log(t)$ *lower bound in (2.5) is achieved if the underlying configuration C_0 is not implicitly revealing. When C_0 is implicitly revealing, the same rule attains bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$.*

[8] These two additional conditions are identical to those in Section 2.4.

Algorithm 3 Φ_{t+1}, the decision at time $t+1$

1: **if** not all $\hat{C}_{x,t}$ are identical, **then**
2: $\quad\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$.
3: **else**
4: $\quad$Denote $\hat{C}_t = (\hat{\theta}_1, \hat{\theta}_2)$ as the common estimate for all B_x.
5: $\quad$**if** $\hat{C}_t$ is implicitly revealing, **then**
6: $\quad\quad$**if** $\ddot{C}_t \neq \hat{C}_t$, **then**
7: $\quad\quad\quad$**if** $\mathrm{ctr}(X_{t+1}, \hat{C}_t, \ddot{C}_t)$ is even, **then**
8: $\quad\quad\quad\quad\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$.
9: $\quad\quad\quad$**else**
10: $\quad\quad\quad\quad\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$.
11: $\quad\quad\quad$**end if**
12: $\quad\quad\quad\mathrm{ctr}(X_{t+1}, \hat{C}_t, \ddot{C}_t) \leftarrow \mathrm{ctr}(X_{t+1}, \hat{C}_t, \ddot{C}_t) + 1$.
13: $\quad\quad$**else**
14: $\quad\quad\quad\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$.
15: $\quad\quad$**end if**
16: $\quad$**else**
17: $\quad\quad$Without loss of generality, we may assume $M_{\hat{C}_t} = 2$. The case in which $M_{\hat{C}_t} = 1$ can be obtained by symmetry.
18: $\quad\quad$**if** $X_{t+1} \neq x^* := \arg\max_x \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\hat{\theta}_2}(x_0)\}} I(\hat{\theta}_1, \theta|x)$, **then**
19: $\quad\quad\quad\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$.
20: $\quad\quad$**else**
21: $\quad\quad\quad\Phi_{t+1} \leftarrow \phi_{X_{t+1},t+1}$.
22: $\quad\quad$**end if**
23: $\quad$**end if**
24: **end if**

One instance of such asymptotically optimal rules is a composite control scheme $\{\Phi_\tau\}$ described in Algorithm 3, the details of which are described in the following paragraphs.

The sub-bandit machines B_x, the corresponding EBUG decision rules $\{\phi_{x,\tau}\}$, and the estimate $\{\hat{C}_{x,\tau}\}$ are as defined in Sections 2.1.2 and 2.4.2, along with a finite number of newly-introduced counters (to be more precisely, $|X| \cdot |\mathbf{\Theta}^2|^2$ counters). These new counters are named $\mathrm{ctr}(x, C', C'')$ and are initially set to zero. The $\ddot{C}_t$ used in Algorithm 3 is an estimate of C_0 generated from the sampling when $\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$ is active, namely, when Line 10, 14, or 19 is executed. On the other hand, those samples when $\Phi_{t+1} \leftarrow \phi_{x,t+1}$ is active, namely, when Line 2, 8, or 21 being executed, are used to generate $\hat{C}_{x,t}$ and $\phi_{x,t+1}$.

For example, suppose $X_1X_2X_3X_4\cdots = x_a x_b x_a x_c \cdots$ and at time instants 1 and 2, $\Phi_{t+1} \leftarrow \phi_{x,t+1}$ (Line 2, 8, or 21), while at time instants 3 and 4, $\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1})$ (Line 10, 14, or 19). As a result, after four pulls of the bandit machine, we have one sample in B_{x_a} to generate $\hat{C}_{x_a,4}$, one sample in B_{x_b} for $\hat{C}_{x_b,4}$, and no samples in B_{x_c} for $\hat{C}_{x_c,4}$. At the same time, we have a total of one sample in B_{x_a}, no samples in B_{x_b} and one sample in B_{x_c} being used to generate $\ddot{C}_4$.

A detailed analysis of Algorithm 3 is given in Appendix C.4, which shows that with any "good" $\ddot{C}_t$, the composite rule $\{\Phi_\tau\}$ described in Algorithm 3 satisfies Theorem 2.8. The definition of a "good" $\ddot{C}_t$ is as follows.

Definition 2.3 (Good Estimate $\ddot{C}_t$) *An estimate $\ddot{\theta}$ is good if there exist $a, b > 0$ such*

that the mis-detection probability $\mathsf{P}_\theta(\ddot{\theta} \neq \theta) \leq a\exp(-bN)$*, where* N *is the number of samples that* $\ddot{\theta}$ *is based upon. An estimate pair* $\ddot{C}_t = (\ddot{\theta}_1, \ddot{\theta}_2)$ *is good if* $\ddot{\theta}_1$ *and* $\ddot{\theta}_2$ *are good estimates for* θ_1 *and* θ_2 *respectively.*

By the large deviation principle and the regularity conditions in Sections 2.5.1 and 2.5.2, a good estimate $\ddot{C}_t$ generally exists.

The performance achieved in both Sections 2.3 and 2.4 has been attained here, and the intuition behind Theorem 2.8 is exactly the mixture of our previous discussions on the pure cases. When the unknown C_0 is implicitly revealing, the side information X_t will direct the player to sample both arms often enough, which leads to bounded $\mathsf{E}\{T_{inf}(t)\}$. If the underlying C_0 is not implicitly revealing, then postponing the forced sampling will reduce the constant $1/K_{C_0}$ in front of the $\log(t)$ lower bound.

2.6 Summary

We have shown that observing additional i.i.d. side information can significantly improve sequential decisions in bandit problems. If the side information itself directly provides information about the underlying configuration, then it resolves the dilemma of forced sampling and optimal control. The expected inferior sampling time will be bounded as t tends to infinity, as shown in Section 2.2. If the side information does not provide information about the underlying configuration (θ_1, θ_2), but *always* affects the preference order (implicitly revealing), then the myopic approach of sampling the seemingly better arm will automatically sample both arms often enough. The expected inferior sampling time is bounded due to the fact that the conflict between learning vs. control is implicitly resolved, as shown in Section 2.3. If the side information *does not* affect the preference order at all, the dilemma still exists. However, by postponing our forced sampling to the most informative time instants, we can reduce the constant in the $\log(t)$ lower bound, as shown in Section 2.4. In Section 2.5, we have combined the settings of Sections 2.3 and 2.4, and obtained a general result. When the underlying configuration C_0 is implicitly revealing, an adaptive decision rule can be constructed to achieve bounded expected inferior sampling time as in Section 2.3. Even if C_0 is not implicitly revealing (in that X_t does not change the preference order), the same decision rule attains the new $\log(t)$ lower as stated in Section 2.4. Our results are summarized in Table 2.2

Table 2.2: Summary of results for bandit problems with i.i.d. side information.

Characterization	**Regularity Conditions**	**Results**
• $G_{C_1} \neq G_{C_2}$ iff $C_1 \neq C_2$.	As $\hat{C}_t \to C_0$, $\forall x$, $M_{\hat{C}_t}(x) = M_{C_0}(x)$.	$\exists\{\phi_\tau\}$ s.t. $\forall C_0$, $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$.
• Constant G_C, i.e., $G_C := G$, • $\forall C, \exists x_1, x_2$, s.t. $M_C(x_1) = 1$, $M_C(x_2) = 2$. Namely, all C are implicitly revealing (i.r.).	(i) $\mathbf{X}$ is finite. (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \vert x) < \infty$. (iii) $\forall x$, $\mu_\theta(x)$ is continuous w.r.t. θ.	$\exists\{\phi_\tau\}$ s.t. $\forall C_0$, $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$.
• Constant G_C, i.e., $G_C := G$, • $\forall C$, $M_C(x)$ only depends on C, not on x. Namely, all C are *not* implicitly revealing (i.r.).	(i) $\mathbf{X}$ is finite. (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \vert x) < \infty$.	For any uniformly good $\{\phi_\tau\}$, we have $\forall C_0$, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}$, $K_{C_0} := \inf_\theta \sup_x I(\theta_1, \theta \vert x)$.
	(i), (ii), and (iii) $\mathbf{\Theta}$ is finite, (iv) The existence of the saddle point: $\inf_\theta \sup_x I(\theta_1, \theta \vert x) = \sup_x \inf_\theta I(\theta_1, \theta \vert x)$.	$\exists\{\phi_\tau\}$, s.t. $\forall C_0$, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \leq \frac{1}{K_{C_0}}$.
• Constant G_C, i.e., $G_C := G$, • The underlying C_0 may be implicitly revealing (i.r.) or not.	(i) $\mathbf{X}$ is finite. (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \vert x) < \infty$.	For any uniformly good $\{\phi_\tau\}$, if C_0 is not i.r., we have $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}$, $K_{C_0} := \inf_\theta \sup_x I(\theta_1, \theta \vert x)$.
	(i), (ii), and (iii) $\mathbf{\Theta}$ is finite, (iv) The existence of the saddle point: $\inf_\theta \sup_x I(\theta_1, \theta \vert x) = \sup_x \inf_\theta I(\theta_1, \theta \vert x)$.	$\exists\{\phi_t\}$ s.t. (1) if C_0 is i.r., $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$, (2) if C_0 is not *i.r.*, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \leq \frac{1}{K_{C_0}}$.

Chapter 3

Arbitrary Side Information in Bandit Problems

Chapter 2 focused solely on i.i.d. $\{X_\tau\}$, and various levels of asymptotic efficiency were proved after characterizing the relationships between $\{X_\tau\}$ and $(\{Y_\tau^1\}, \{Y_\tau^2\})$ into four separate categories, discussed in Sections 2.2 through 2.5 respectively. Compared to those of traditional bandit problems, these results show that side information is able to benefit the decision maker when carefully exploited, and the improvement depends on the relationship between the side information and the bandit problem.

As mentioned in Section 1.1, one of the applications of side-information-aided bandit problems is when considering two different modulation schemes, M1 vs. M2, and the total reward is the number of error-free transmitted packets. In this example, the side information may contain geographical information about the transmitter-receiver pair, such as GPS signals, or some environmental variables. Although the results in Chapter 2 identify the benefit of exploiting i.i.d. side information, it is overoptimistic to assume the side information is i.i.d. For example, the location (obtained from GPS) of a receiving mobile is unlikely to be i.i.d., and a more realistic model would be a Markov chain of a large but finite order. A remaining question is whether such side information is as beneficial as an i.i.d. side information sequence. Furthermore, we would like to extract the essential properties of good side information, based on which new performance bounds and achievability results will be rederived. Chapter 3 is dedicated to this mission. With the generalized results for a wide range of non-i.i.d. side information, one can easily check whether a new type of side information is beneficial and how much improvement can be expected.

This chapter is organized as follows. In Section 3.1, we provide a rigorous formulation of side-observation-aided bandit problems and give formal definitions of several "even distribution" properties, examples of each such property, and relationships among them. In Sections 3.2 through 3.5, we provide results for various relationships among $\{X_\tau\}$, $\{Y_\tau^1\}$, and $\{Y_\tau^2\}$ with the satisfaction of the "even distribution" properties. All results in Chapter 2, obtained under the assumption of i.i.d. $\{X_\tau\}$, hold as special cases under this new framework, which includes many other side observation processes (e.g. Markov chains of any finite order and periodic sequences) as well. Section 3.6 provides several examples, a summary table, and a simple necessary condition concerning the extent of the benefit obtained from observing $\{X_\tau\}$. Section 3.7 concludes this chapter.

3.1 Refined Formulation

3.1.1 Arbitrary Side Information

To characterize explicitly the dependence among $C_0 = (\theta_1, \theta_2)$, arbitrary side information $\{X_\tau\}$, and the rewards $\{Y_\tau^1\}$ and $\{Y_\tau^2\}$, the probability distribution of the two-armed bandit with side information is modelled as follows. At time instants $t_1, \cdots, t_k$, the joint probability distribution of $(X_{t_i}, Y_{t_i}^1, Y_{t_i}^2)_{i=1,\cdots,k}$ is

$$G_{t_1,\cdots,t_k|C_0}(dx_{t_1}, \cdots, dx_{t_k}) \prod_{i=1}^{k} F_{\theta_1}(dy_{t_i}^1|x_{t_i}) F_{\theta_2}(dy_{t_i}^2|x_{t_i}),$$

where $G_{t_1,\cdots,t_k|C_0}(dx_{t_1}, \cdots, dx_{t_k})$ is the finite cylinder distribution of the side information $\{X_\tau\}$, which may or may not depend on C_0. Both families of distributions, $\{G_{\cdots|C}\}_{C\in\mathbf{\Theta}^2}$ and $\{F_\theta(\cdot|x)\}_{\theta\in\mathbf{\Theta}}$, are known to the decision maker, but the true value of C_0 is unknown. There is few restriction on $\mathbf{X}$ and $\mathbf{\Theta}$, the ranges of X_t and θ. Following Chapter 2, both $\mathbf{X}$ and $\mathbf{\Theta}$ are assumed to be subsets of $\mathbb{R}$.

We will reuse most of the definitions for bandit problems with i.i.d. side information. For example, the conditional expected return given $X_t = x$ is still defined as $\mu_{\theta_1}(x)$ and $\mu_{\theta_2}(x)$ when the underlying configuration is $C_0 = (\theta_1, \theta_2)$.

Remark: The i.i.d. side information discussed in Chapter 2 is a special case of this general setting, in which

$$G_{t_1,\cdots,t_k|C_0}(dx_{t_1}, \cdots, dx_{t_k}) = \prod_{i=1}^{k} G_{t_i|C_0}(dx_{t_i}) = \prod_{i=1}^{k} G_{C_0}(dx_{t_i}).$$

3.1.2 Even Distribution Properties

The intuition behind the results in Chapter 2 shows that the decision maker should wait and let the side information $\{X_\tau\}$ direct the myopic sampling of both arms evenly, or postpone the forced sampling to the most informative sub-bandit machine. This optimal strategy therefore suggests that the benefits of side information in bandit problems are not due to the *random* appearance of all values x of the i.i.d. $\{X_\tau\}$, but rather are due to the *evenly* distributed appearance of all possible x.

Our goal is to extract the essential "even distribution" properties of a side information process that are helpful to uniformly good rules. Suppose X_t takes values in a finite state set $\mathbf{X}$, and the relative frequency of x up to time t is denoted by

$$f_r(x, t) = \frac{\sum_{\tau=1}^{t} 1\{X_\tau = x\}}{t}.$$

Four levels of even distribution properties are formally defined as follows.

Definition 3.1 (Evenly Distributed in L^1) *$\{X_\tau\}$ is evenly distributed in L^1 if*

$$\forall x \in \mathbf{X}, \quad \pi(x) := \liminf_{\tau\to\infty} \mathsf{E}\{f_r(x, \tau)\} > 0.$$

Definition 3.2 (Evenly Distributed in Probability) *$\{X_\tau\}$ is evenly distributed in probability if there exists a strictly positive mapping $\pi(\cdot) > 0$ such that*

$$\forall x \in \mathbf{X}, \quad \lim_{\tau \to \infty} \mathsf{P}\left(f_r(x,\tau) < \pi(x)\right) = 0.$$

Definition 3.3 (Evenly Distributed in Probability Series) *$\{X_\tau\}$ is evenly distributed in probability series if there exists a strictly positive mapping $\pi(\cdot) > 0$, such that the duration of the event $\{f_r(x,t) < \pi(x)\}$ has a finite expectation. Namely,*

$$\begin{aligned} \forall x \in \mathbf{X}, \qquad & \mathsf{E}\left\{\sum_{\tau=1}^{\infty} 1\{f_r(x,\tau) < \pi(x)\}\right\} \\ = & \sum_{\tau=1}^{\infty} \mathsf{P}\left(f_r(x,\tau) < \pi(x)\right) < \infty. \end{aligned}$$

By the first Borel-Cantelli lemma , this definition automatically implies that

$$\forall x, \quad \liminf_{t \to \infty} f_r(x,t) \geq \pi(x) \quad a.s.$$

Definition 3.4 (Uniformly Strongly Evenly (u.s.e.) Distributed in L^1) *$\{X_\tau\}$ is u.s.e. distributed in L^1, if for any stopping time T, the conditional expectation of the first hitting time of x after T has a global upper bound. That is, $\exists B < \infty$ such that*

$$\forall T, \forall x \in \mathbf{X}, \ \mathsf{E}\{H_T(x)|T\} \leq B,$$

where $H_T(x) \stackrel{\Delta}{=} \inf\{l > 0 | X_{T+l} = x\}$.

It is easy to verify that these four properties hold for non-degenerate i.i.d. sequences, non-degenerate Markov chains of any finite order, and deterministic periodic sequences, which demonstrates the generality of these classes of random processes.

Remark: These definitions are listed in order from the weakest: Definition 3.1, to the strongest: Definition 3.4. Detailed analysis and further discussion can be found in Appendix D.

The following four sections are devoted to determining even distribution properties that are sufficient for different levels of improvement.

3.2 Case 1: Direct Information

Similar to Section 2.2, we consider the situation in which the side information random process $\{X_\tau\}$ directly reveals information about $C_0 = (\theta_1, \theta_2)$ in the following form.

Direct information If $C_0 \neq C_0'$, then $\exists t_1, \cdots, t_k$, such that $G_{t_1,\cdots,t_k|C_0} \neq G_{t_1,\cdots,t_k|C_0'}$.

Algorithm 4 ϕ_t, the decision at time t

1: Obtain an estimate $\hat{C}_t$ based on the side information $X_1, \cdots, X_t$.
2: Set $\phi_t = M_{\hat{C}_t}(X_t)$.

3.2.1 Scheme of Separating Learning and Control

Since we are able to obtain information about C_0 from $\{X_\tau\}$ in this setting, it is natural to sample only the seemingly better arm while leaving the learning task to $\{X_\tau\}$. A corresponding control scheme $\{\phi_\tau\}$ can be described as Algorithm 4.

With the assumption of Condition 2.1 as well, we can bound the performance of Algorithm 4 by the convergence speed of the estimate $\hat{C}_t$ as in the following theorem.

Theorem 3.1 *Suppose* Condition 2.1 *and the criterion of direct information are valid. For all* C_0 *and any sequence of estimates* $\{\hat{C}_\tau\}$*, there exists* $\epsilon > 0$ *such that* Algorithm 4 *satisfies*

$$\lim_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\sum_{\tau=1}^{t} \mathsf{P}_{C_0}(|\hat{C}_\tau - C_0| > \epsilon)} \leq 1.$$

A detailed proof is given in Appendix C.1.

The above theorem provides an upper bound on the best achievable expected inferior sampling time, and is illustrated in the following examples.

- *Example 1:* Suppose $\{X_\tau\}$ is an i.i.d. sequence with marginal distribution G_{C_0} on $\mathbb{R}$, and the mapping from C_0 to G_{C_0} is one-to-one. Then by Sanov's theorem on $\mathbb{R}$, there exists $\{\hat{C}_\tau\}$ such that $\forall C_0, \epsilon > 0$, $\lim_{t\to\infty} \sum_{\tau=1}^{t} \mathsf{P}_{C_0}(|\hat{C}_\tau - C_0| > \epsilon) < \infty$. By Theorem 3.1, $\forall C_0$, we have $\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$, and the proposed Algorithm 4 is thus uniformly good. Theorem 2.3 becomes a special case of Theorem 3.1.

- *Example 2:* Suppose $\{X_\tau\}$ is a first order Markov chain with transition matrix A_{C_0}, and the mapping from C_0 to A_{C_0} is one-to-one. Then by similar reasoning as in the i.i.d. case, Algorithm 4 becomes a uniformly good rule such that $\forall C_0, \lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$.

- *Example 3:* Consider the case in which $\{X_\tau\}$ is a deterministic sequence denoted by $\{x_{\tau|C_0}\}$. If the mapping from C_0 to $\{x_{\tau|C_0}\}$ is one-to-one, and $\mathbf{\Theta}$ is finite, we can easily find a family of estimates $\{\hat{C}_\tau\}$ such that $\forall C_0, \epsilon > 0$, $\lim_{t\to\infty} \sum_{\tau=1}^{t} \mathsf{P}_{C_0}(|\hat{C}_\tau - C_0| > \epsilon) < \infty$. Hence $\forall C_0, \lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$, and Algorithm 4 is uniformly good.

3.3 Case 2: Best Arm As A Function Of X_t

In Sections 3.3 through 3.5, we consider the situations in which the distribution of $\{X_t\}$ is not a function of C_0. Namely,

No direct information $G_{t_1,\cdots,t_k|C_0} := G_{t_1,\cdots,t_k}$.

The corresponding results for i.i.d. side information random processes can be found in Sections 2.3 through 2.5.

In this section, we consider one refinement of this no-direct-information setting, which is identical to that in Section 2.3. The characterization condition and necessary regularity conditions are re-stated as follows.

Characterization condition:

- For all $C \in \mathbf{\Theta}^2$, there exist x_1 and x_2 such that $M_C(x_1) = 1$ and $M_C(x_2) = 2$.

Regularity conditions:

1. $\mathbf{X}$ is a finite set.
2. $\forall \theta_1, \theta_2, x$, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.
3. $\forall x$, $\mu_\theta(x)$ is continuous w.r.t. θ.

It has been shown in Section 2.3 that for i.i.d. side information, although no information about C_0 is revealed through observing X_t, significant improvement, i.e., bounded $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\}$, can be obtained when the best arm is a function of X_t. In the following theorem, the above result will be generalized for arbitrary evenly distributed side information random processes.

Theorem 3.2 (Bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$) *Suppose the aforementioned conditions are satisfied. If the side information $\{X_\tau\}$ is evenly distributed in probability series, then there exists an allocation rule $\{\phi_\tau\}$ such that*

$$\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty, \quad \forall C_0 \in \mathbf{\Theta}^2.$$

Such a rule is obviously uniformly good and surpasses the $\log(t)$ *lower bound for traditional bandit problems.*

A detailed proof of this theorem will be given in Appendix C.2, in which we demonstrate that the same decision rule in Algorithm 1, designed for i.i.d. side information, also achieves bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ when the side information $\{X_\tau\}$ is evenly distributed in probability series. Our result re-confirms that when the underlying configuration is implicitly revealing, we do not need to perform much forced sampling and any *evenly distributed* side information will help the myopic decision sample both arms often enough. The conflict between learning and control can thus be implicitly resolved.

3.4 Case 3: Best Arm Is Not A Function Of X_t

Following Section 3.3, we assume that no information about C_0 is revealed through the arbitrary side information $\{X_\tau\}$, i.e., $G_{t_1,\cdots,t_k|C_0} = G_{t_1,\cdots,t_k}$. In this section, we consider the case in which $\forall C_0 \in \mathbf{\Theta}^2$, X_t *never* changes the preference order, or equivalently, $M_{C_0}(x)$ is not a function of x and we can use M_{C_0} as shorthand.

In Section 2.4, it has been shown that a new $\log(t)$ lower bound exists, which is less restrictive compared to that for traditional bandit problems. Furthermore, for i.i.d. side information $\{X_\tau\}$, this new bound is attainable and one instance of bound-achieving decision rules is given in Algorithm 2. We will generalize those results in the following subsections.

3.4.1 A New $\log(t)$ Lower Bound

In this subsection, identical settings to those in Section 2.4.1 are considered, except that we are focusing on *non-i.i.d.* side information $\{X_\tau\}$. The necessary regularity conditions are as follows.

1. $\mathbf{X}$ is a finite set.

2. $\forall \theta_1, \theta_2, x$, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.

3. For all x, the conditional expected reward $\mu_\theta(x)$ is strictly increasing w.r.t. θ.

With these conditions, Theorem 2.5 can be generalized as follows.

Theorem 3.3 (New $\log(t)$ Lower Bound) *Consider general non-i.i.d. side information random processes $\{X_\tau\}$. Under the aforementioned conditions, for any uniformly good rule $\{\phi_\tau\}$, $T_{inf}(t)$ satisfies*

$$\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_{inf}(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C_0}}\right) = 1, \ \forall \epsilon > 0,$$

$$\textit{and} \qquad \liminf_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}, \tag{3.1}$$

where K_{C_0} is a constant depending on C_0. If $M_{C_0} = 2$, then $T_{inf}(t) = T_1(t)$, and the constant K_{C_0} can be expressed as follows.

$$K_{C_0} = \inf_{\theta:\theta>\theta_2} \sup_{x\in\mathbf{X}} \{I(\theta_1, \theta|x)\}. \tag{3.2}$$

The expression for K_{C_0} for the case in which $M_{C_0} = 1$ can be obtained by symmetry.

3.4.2 Scheme Achieving the Lower Bound

In addition to the three regularity conditions for the $\log(t)$ lower bound, we need the following conditions to ensure the existence of a $\log(t)$ lower bound achieving scheme.

4. $\mathbf{\Theta}$ is finite, and

5. The existence of the saddle point in the expression of K_{C_0}, (3.2), is assumed. Namely, for all $\theta_1 < \theta_2$,

$$\inf_{\theta:\theta>\theta_2} \sup_{x\in\mathbf{X}} I(\theta_1, \theta|x) = \sup_{x\in\mathbf{X}} \inf_{\theta:\theta>\theta_2} I(\theta_1, \theta|x),$$

With these conditions, the new $\log(t)$ lower bound in Section 3.4.1 is attainable for evenly distributed side information random processes $\{X_\tau\}$.

Theorem 3.4 (Asymptotic Sharpness) *If the side information $\{X_\tau\}$ is u.s.e. distributed in L^1, then an decision rule can be explicitly constructed so that the $\log(t)$ lower bound in (3.1) is achieved for all possible $C_0 \in \mathbf{\Theta}^2$. This decision rule is thus uniformly good and asymptotically optimal.*

This theorem will be proved by showing that the composite decision rule $\{\Phi_\tau\}$ described in Algorithm 2 attains the specified new $\log(t)$ lower bound when the side information $\{X_\tau\}$ is u.s.e. distributed in L^1. In general, it is hard to check whether a random process is u.s.e. distributed in L^1, since by definition, we have to exhaustively verify the global upper boundedness of the conditional expected hitting time for *all stopping times*. However, among Markov chains of any finite order, one can easily identify a u.s.e. distributed $\{X_\tau\}$ due to the inherited strong Markov property.[1] The result in Theorem 3.4 shows that any

[1] Markov chains of any finite order are a broad class of random processes by themselves, including i.i.d. sequences and deterministic periodic sequences as special cases. Nonetheless, the class of u.s.e. distributed in L^1 contains more than Markov chains. One can easily construct a deterministic aperiodic sequence that is not a Markov chain of any finite order but is u.s.e. distributed in L^1.

general evenly distributed side information $\{X_\tau\}$ is still beneficial even when the underlying configuration is not implicitly revealing.

A detailed analysis of $\{\Phi_\tau\}$ in Algorithm 2 is given in Appendix C.3.

3.5 Case 4: Mixed Case

Similar to Section 2.5, we consider the case in which some configurations C_0 are implicitly revealing and some are not while the side information $\{X_\tau\}$ reveals no information about C_0. In this section, we focus on evenly distributed side information $\{X_\tau\}$ and will generalize previous results for i.i.d. $\{X_\tau\}$, including a new conditional $\log(t)$ lower bound when C_0 is not implicitly revealing, and an adaptive rule achieving the best possible for both implicitly-revealing and not-implicitly revealing C_0's.

A formal definition of this mixed case is as follows.

Definition 3.5 (Mixed Condition) *For some $C \in \mathbf{\Theta}^2$, $M_C(x)$ is not a function of x, i.e., $M_C(x) := M_C$. For the remaining C, there exist x_1 and x_2 such that $M_C(x_1) = 1$ and $M_C(x_2) = 2$.*

3.5.1 A New $\log(t)$ Lower Bound

The necessary regularity conditions are as follows, which are identical to those in Section 2.5.1.

1. $\mathbf{X}$ is finite.
2. $\forall\theta_1$, θ_2, and x, $I(\theta_1, \theta_2|x)$ is strictly positive and finite.

We then have the generalized version of Theorem 2.7 as follows.

Theorem 3.5 (New $\log(t)$ Lower Bound) *Suppose the side information $\{X_\tau\}$ is evenly distributed in probability. Under the specified regularity conditions, for any uniformly good rule $\{\phi_\tau\}$, if C_0 is not implicitly revealing, $T_{inf}(t)$ satisfies*

$$\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_{inf}(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C_0}}\right) = 1, \ \forall\epsilon > 0,$$

$$\textit{and} \qquad \liminf_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}},$$

where K_{C_0} is a constant depending on C_0. If $M_{C_0} = 2$, $T_{inf}(t) = T_1(t)$, and the constant K_{C_0} can be expressed as follows.

$$K_{C_0} = \inf_{\{\theta:\exists x_0,\ s.t.\ \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \sup_x \{I(\theta_1, \theta|x)\}.$$

The expression for K_{C_0} for the case in which $M_{C_0} = 1$ can be obtained by symmetry

A detailed proof of this generalized theorem can be found in Appendix B.2.

3.5.2 Scheme Achieving the Lower Bound

In addition to the two regularity conditions for the conditional $\log(t)$ lower bound, the following two conditions are required while devising a bound-achieving decision rule:

3. $\mathbf{\Theta}$ is finite.

4. The existence of a saddle point is assumed, that is, for all (θ_1, θ_2),

$$\inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \sup_{x\in\mathbf{X}} \{I(\theta_1, \theta|x)\} = \sup_{x\in\mathbf{X}} \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \{I(\theta_1, \theta|x)\}.$$

With the aforementioned regularity conditions, it can be shown that by taking the advantage of evenly distributed side information, the new conditional $\log(t)$ lower bound is attainable.

Theorem 3.6 (Asymptotic Sharpness) *Suppose the side information $\{X_\tau\}$ is u.s.e. distributed in L^1. An adaptive decision rule can then be explicitly constructed so that the new conditional* $\log(t)$ *lower bound in (2.5) is achieved if the underlying configuration C_0 is not implicitly revealing. When C_0 is implicitly revealing, the same rule attains bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$.*

We prove this theorem by showing that the adaptive decision rule described in Algorithm 3 satisfies the asymptotic sharpness theorem. A detailed analysis is given in Appendix C.4.

3.6 Examples & Degenerate Situations

3.6.1 Examples

Here we provide several examples illustrating the benefits of side information, in which the range of θ and x are simplified as $\{1, 2, 3, 4\}$ or $\{1, 2, 3\}$, and the governing conditional distributions $F_\theta(\cdot|X_t = x)$ are Bernoulli with success probability $p_{\theta,x}$. The entire family of conditional distributions can then be specified by a matrix $(p_{\theta,x})$, and we will discuss the following three examples.

- *Example 1:*

$$(p_{\theta,x}) = \begin{pmatrix} 0.4 & 0.3 & 0.6 \\ 0.5 & 0.5 & 0.5 \\ 0.6 & 0.7 & 0.4 \end{pmatrix},$$

- *Example 2:*

$$(p_{\theta,x}) = \begin{pmatrix} 0.4 & 0.3 & 0.2 \\ 0.5 & 0.5 & 0.5 \\ 0.6 & 0.7 & 0.8 \end{pmatrix},$$

- *Example 3:*

$$(p_{\theta,x}) = \begin{pmatrix} 0.4 & 0.4 & 0.5 \\ 0.5 & 0.5 & 0.4 \\ 0.6 & 0.6 & 0.6 \\ 0.7 & 0.8 & 0.9 \end{pmatrix}.$$

These three examples possess different internal structures and were characterized as Cases 2, 3, and 4 respectively during our previous discussion.

Although the achievability results require only evenly distributed $\{X_\tau\}$, to be able to compare the improvement with traditional bandit problems, we assume $\{X_\tau\}$ is an i.i.d. sequence with its marginal uniformly distributed among $\{1,2,3\}$. For any parameter θ, if we ignore the side information X_t, the player is then facing a Bernoulli distribution with parameter $p_{\theta-} := \frac{p_{\theta,1}+p_{\theta,2}+p_{\theta,3}}{3}$. Suppose the true parameter pair $C_0 = (\theta_1, \theta_2)$ equals $(1,2)$ (unknown to the player). By Theorem 2.1, $\lim_{t\to\infty} \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}$, where K_{C_0} is $0.0358 = I(p_{1-}, p_{3-})$ for Example 1, $0.3389 = I(p_{1-}, p_{3-})$ for Example 2, and $0.0564 = I(p_{1-}, p_{3-})$ for Example 3.

Our results in Chapters 2 and 3 show that by exploiting X_t, these $\log(t)$ lower bounds can be surpassed. For Example 1, there exists a uniformly good rule $\{\phi_\tau\}$ achieving bounded expected rewards: $\lim_{t\to\infty} \mathsf{E}\{T_{inf}(t)\} < \infty$. For Example 2, the performance is still $\log(t)$ lower bounded, but a smaller constant $\frac{1}{K'_{C_0}}$ can be achieved: $\lim \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K'_{C_0}}$ with $K'_{C_0} = I(p_{1,3}, p_{3,3}) = 0.8318$. The new constant $\frac{1}{K'_{C_0}}$ is only 41% of the original $\frac{1}{K_{C_0}}$. For Example 3, there exists a uniformly good rule admitting bounded $\lim_{t\to\infty} \mathsf{E}\{T_{inf}(t)\} < \infty$, since $(\theta_1, \theta_2) = (1,2)$ is implicitly revealing.

Within the same setting of Example 3, if the unknown (θ_1, θ_2) equals $(2,3)$ instead of $(1,2)$ (the former is not implicitly revealing), it can be proved that no rule can achieve bounded $\mathsf{E}\{T_{inf}(t)\}$ and the minimum regret is still $\log(t)$ lower bounded. The best achievable constant in front of $\log(t)$ becomes $\frac{1}{K'_{C_0}}$ with $K'_{C_0} = I(p_{2,3}, p_{4,3}) = 0.7507$. For comparison, the traditional $\log(t)$ lower bound (ignoring side information) is $\frac{\log(t)}{K_{C_0}}$, $K_{C_0} = I(p_{2-}, p_{4-}) = 0.2716$. The new constant $\frac{1}{K'_{C_0}}$ is only 36% of the original $\frac{1}{K_{C_0}}$.

3.6.2 Degenerate Situations

In Sections 3.3 through 3.5, we have discussed the benefits of having side information under various situations. The main results are summarized in Table 3.1. A question naturally arises as to whether these evenly distributed properties are necessary for the various levels of improvement.

From Theorem 3.1 of Section 3.2, having estimates of C_0 from $\{X_\tau\}$ with appropriate convergence speed provides an upper bound on the attainable expected inferior sampling time, which can help surpassing the $\log(t)$ lower bound in [67]. However, even without a good estimate, if all possible C_0 are implicitly revealing, we are still able to obtain $\lim \mathsf{E}\{T_{inf}(t)\} < \infty$ as described in Section 3.3. Hence, having estimates with appropriate convergence speed is not a necessary condition to surpass the traditional $\log(t)$ lower bound.

Suppose $\{X_\tau\}$ reveals no information about C_0, as in Sections 3.3 to 3.5. We need some minimal amount of even distribution to guarantee the benefit of observing side information can be fully utilized, which is stated as the following result.

Theorem 3.7 (Common Necessary Condition) *For the achievability results in* Theorems 3.2*,* 3.4*, and* 3.6 *to hold for all distribution families* $\{F_\theta(\cdot|x)\}$ *(satisfying the characterization and regularity conditions), we must have*

$$\forall x, \mathsf{P}(\exists \tau,\ \textit{s.t. } X_\tau = x) > 0.$$

Table 3.1: Summary of results for bandit problems with arbitrary side information.

Characterization	**Regularity Cond.**	**Even Distr. Cond.**	**Results**
• $G_{C_1} \neq G_{C_2}$ iff $C_1 \neq C_2$	As $\hat{C}_t \to C_0$, $\forall x$, $M_{\hat{C}_t}(x) = M_{C_0}(x)$.		$\exists\{\phi_\tau\}$ s.t. $\forall C_0$, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\sum_{\tau=1}^{t} \mathsf{P}(\lvert\hat{C}_\tau - C_0\rvert > \epsilon)} \leq 1$.
• Constant G_C, i.e., $G_C := G$, • All $C_0 \in \mathbf{\Theta}^2$ are implicitly revealing (i.r.).	(i) $\mathbf{X}$ is finite, (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \lvert x) < \infty$, (iii) $\forall x$, $\mu_\theta(x)$ is continuous w.r.t. θ.	$\{X_\tau\}$ is evenly distr. in prob. series.	$\exists\{\phi_\tau\}$ s.t. $\forall C_0$, $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$.
• Constant G_C, i.e., $G_C := G$, • All $C_0 \in \mathbf{\Theta}^2$ are not implicitly revealing (i.r.).	(i) $\mathbf{X}$ is finite, (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \lvert x) < \infty$.		For any uniformly good $\{\phi_\tau\}$, we have $\forall C_0$, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}$, $K_{C_0} := \inf_\theta \sup_x I(\theta_1, \theta \lvert x)$.
	(i), (ii), and (iii) $\mathbf{\Theta}$ is finite, (iv) The existence of the saddle point: $\inf_\theta \sup_x I(\theta_1, \theta \lvert x) = \sup_x \inf_\theta I(\theta_1, \theta \lvert x)$.	$\{X_\tau\}$ is u.s.e. distr. in L^1.	$\exists\{\phi_\tau\}$, s.t. $\forall C_0$, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \leq \frac{1}{K_{C_0}}$.
• Constant G_C, i.e., $G_C := G$, • The underlying C_0 may be implicitly revealing (i.r.) or not.	(i) $\mathbf{X}$ is finite, (ii) $\forall \theta_1 \neq \theta_2, x$, $0 < I(\theta_1, \theta_2 \lvert x) < \infty$.	$\{X_\tau\}$ is evenly distri. in prob.	For any uniformly good $\{\phi_\tau\}$, if C_0 is not i.r., we have $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \geq \frac{1}{K_{C_0}}$, $K_{C_0} := \inf_\theta \sup_x I(\theta_1, \theta \lvert x)$.
	(i), (ii), and (iii) $\mathbf{\Theta}$ is finite. (iv) The existence of the saddle point: $\inf_\theta \sup_x I(\theta_1, \theta \lvert x) = \sup_x \inf_\theta I(\theta_1, \theta \lvert x)$.	$\{X_\tau\}$ is u.s.e. distr. in L^1.	$\exists\{\phi_t\}$ s.t. (1) if C_0 is i.r., $\lim_t \mathsf{E}_{C_0}\{T_{inf}(t)\} < \infty$, (2) if C_0 is not *i.r.*, $\lim_t \frac{\mathsf{E}_{C_0}\{T_{inf}(t)\}}{\log(t)} \leq \frac{1}{K_{C_0}}$.

Note that the condition $\forall x$, $\mathsf{P}(\exists\tau, \text{ s.t. } X_\tau = x) > 0$ is the weakest even distribution property we have introduced. In words, the achievability results rely heavily on the evenly distributed $\{X_\tau\}$, and do not hold for degenerate random processes.

To be more explicit, if there exists x_0 such that $\mathsf{P}(\exists\tau, \text{ s.t. } X_\tau = x_0) = 0$, then the range of the side information can be reduced to the positive support of X_t. The benefit of the characterization properties (helpful structure between X_t, Y_t^i) may degenerate to another case with new support $\mathbf{X}' = \mathbf{X} \backslash \{x_0\}$, which severely affects the attainable results. Using Example 1 in Section 3.6.1 as illustration, the implicitly revealing $C_0 = (\theta_1, \theta_2) = (1, 2)$ is no longer implicitly revealing if the support $\mathbf{X} = \{1, 2, 3\}$ is reduced to $\{1, 2\}$. The achievable $\mathsf{E}\{T_{inf}(t)\}$ is then $\mathcal{O}(\log(t))$ lower bounded instead of being upper bounded away from infinity. Theorem 3.7 shows that the benefit of side information indeed comes from the even distribution properties. Similar arguments can be easily demonstrated on Examples 2 and 3 showing that diminishing the support of $\{X_\tau\}$ significantly degrades the achievable performance.

3.7 Summary

In Chapter 2, it was shown that observing additional i.i.d. side information can improve sequential decisions in bandit problems. To further explore the origins of this improvement, in this chapter we have extracted basic properties of the side information processes and proved their efficacy for bandit problems. When the arbitrary side information $\{X_\tau\}$ reveals information about C_0, with a scheme separating the learning and control tasks by observing $\{X_\tau\}$ for learning, and playing arm $M_{\hat{C}_t}(X_t)$ for control, we have proven that $\lim_{t\to\infty} \mathsf{E}\{T_{inf}(t)\} < \infty$ for many types of $\{X_\tau\}$.

If the side information does not provide information about the configuration C_0, three cases have been considered: (1) the best arm is a function of X_t, as in Section 3.3, (2) the best arm is not a function of X_t, as in Section 3.4, and (3) the mixed case as in Section 3.5. For any $\{X_\tau\}$, *regular/even* appearances of all $x \in \mathbf{X}$ guarantee that we can take the full advantage of the beneficial structure/relationship between the side information $\{X_\tau\}$ and the reward process $\{Y_\tau^i\}$. With different levels of "regular/even appearance" properties, the improvements for different cases are as follows. Case (1) leads to bounded expected inferior sampling time, Case (2) leads to asymptotically sharp $\log(t)$ lower bound, and Case (3) leads to $\log(t)$ lower bound for some C_0, and bounded expected inferior sampling time for other C_0. Consequently, a much more general class of side information sequences, including Markov chains of any finite order and all deterministic periodic sequences, has the same impact on bandit problems as that of i.i.d. sequences. All the achievability results are proved by constructing composite decision rules and assuming the existence of the saddle point (the value of a game) on the Kullback-Leibler information.

Finally, we have provided a simple necessary condition, namely $\forall x, \mathsf{P}(\exists \tau, \text{ s.t. } X_\tau = x) > 0$, which is necessary for a side information sequence to fully exploit the inherent structure between X_t and Y_t^i.

Chapter 4

Low-Density Parity Check Codes on Non-Symmetric Channels

In this chapter, we will focus on capacity-approaching low-density parity-check (LDPC) codes with applications on *non-symmetric* memoryless channels. One of the most powerful analytical tools for LDPC codes and graph codes with message passing decoding is the density evolution (DE) method, which was originally devised for applications on symmetric memoryless channels and relied heavily on the assumption of channel symmetry. Due to the success of DE, capacity-approaching LDPC codes for symmetric channels have been constructed [30]. Nonetheless, constructing similarly ultra-powerful codes for non-symmetric channels remains an open problem, and a proper analytical tool like DE will benefit the code development.

We will discuss the difficulty of performing traditional DE on non-symmetric channels, and provide a new DE formula with comparable complexity, which is able to broaden the applications to general memoryless non-symmetric channels, e.g. z-channels, binary non-symmetric channels, etc. The central theorem underpinning this generalization is the convergence to perfect projection of any support tree of fixed size. Several properties of this new DE (on non-symmetric channels) will be discussed, including monotonicity, symmetry, and stability of LDPC codes on non-symmetric channels. Simulations, code optimizations, and possible new applications suggested by this new density evolution method are also provided. This new DE will then be used to prove the typicality of linear LDPC codes among the coset code ensemble when the minimum check node degree is sufficiently large. It will be shown that the convergence to perfect projection is essential to the belief propagation algorithm even when only symmetric channels are considered. Hence the proof of the convergence to perfect projection serves also as a completion of the theory of classical density evolution for symmetric memoryless channels.

This chapter is organized as follows. The formulation of and background on channel models, LDPC code ensembles, the belief propagation algorithm, and density evolution, are provided in Section 4.1. In Section 4.2, an iterative formula is developed for computing the evolution of the codeword-averaged probability density. In Section 4.3, we state and prove the theorem of convergence to perfect projection, which justifies the iterative formula. Monotonicity, symmetry, and stability theorems are stated and proved in Section 4.4. Section 4.5 consists of simulations and discussion of possible applications of our new density evolution method. Section 4.6 proves the typicality of linear LDPC codes and revisits belief propagation for symmetric channels. Section 4.7 concludes this chapter.

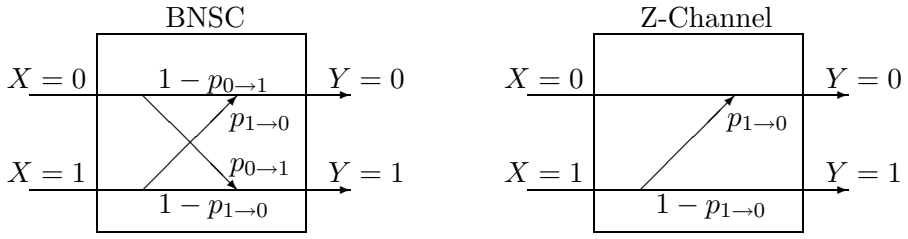

(a) BNSCs (b) Z-Channels

Figure 4.1: Some examples of non-symmetric memoryless channels

4.1 Formulation

4.1.1 Non-Symmetric Memoryless Channels

The memoryless, symbol-dependent channels we consider here are modelled as follows. Let $\mathbf{x}$ and $\mathbf{y}$ denote a transmitted codeword vector and a received signal vector of codeword length n, where x_i and y_i are the i-th transmitted symbol and received signal, respectively, taking values in $\mathsf{GF}(2)$ and the reals, respectively. The channel is memoryless and is governed by the conditional distribution $F_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{n} F(dy_i|x_i)$. Two examples are as follows.

- *Example 1:* [Binary Non-Symmetric Channels (BNSCs)]

$$F(dy|x) = \begin{cases} (1-p_{0\to1})\delta(dy) + p_{0\to1}\delta(d(y-1)) & \text{if } x=0 \\ p_{1\to0}\delta(dy) + (1-p_{1\to0})\delta(d(y-1)) & \text{if } x=1 \end{cases},$$

  where $p_{0\to1}$ and $p_{1\to0}$ are the crossover probabilities from 0 to 1 and from 1 to 0 respectively. $\delta(dy)$ is the Dirac delta probability measure. This example is further illustrated in Figure 4.1(a). When $p_{0\to1}$ is zero, a BNSC collapses to a z-channel, which is named after the shape of the corresponding diagram in Figure 4.1(b). Since the z-channel is the most non-symmetric binary-input/binary-output channel, it will be used extensively in this chapter as an illustrative example.

- *Example 2:* [Composite Binary Additive White Gaussian Channels (Composite Bi-AWGNCs)]

$$F(dy|x) = \begin{cases} \frac{1}{2}\mathcal{N}_{0,\sigma^2}\left(d\left(y-\frac{3}{\sqrt{5}}\right)\right) + \frac{1}{2}\mathcal{N}_{0,\sigma^2}\left(d\left(y+\frac{3}{\sqrt{5}}\right)\right) & \text{if } x=0 \\ \frac{1}{2}\mathcal{N}_{0,\sigma^2}\left(d\left(y-\frac{1}{\sqrt{5}}\right)\right) + \frac{1}{2}\mathcal{N}_{0,\sigma^2}\left(d\left(y+\frac{1}{\sqrt{5}}\right)\right) & \text{if } x=1 \end{cases},$$

  where $\mathcal{N}_{0,\sigma^2}$ is a Gaussian distribution with mean 0 and variance σ^2.

Although these two examples may seem artificial, they do find plenty of applications in real world communication problems. For example, in CD-ROM, an optical storage device, we use a reflective surface to represent bit 1 and a non-reflective surface to represent bit 0. Then any contamination/scratches can only corrupt a bit 1 to be received as a bit 0 but not vice versa, which is illustrated in Figure 4.2 and can be perfectly modelled by z-channels in Example 1.

Example 1 is caused by the non-symmetry within the physical channel of interest. Sometimes, even when the underlying physical channels are symmetric, the pre-processing or the

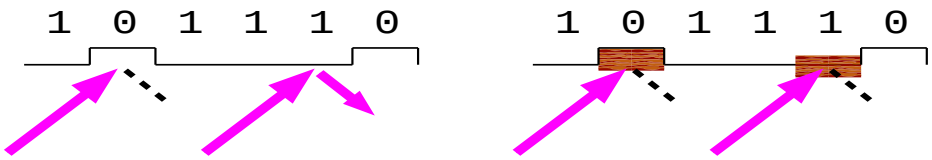

(a) Perfect reception (b) Corrupted reception

Figure 4.2: CD-ROM as an example of z-channels.

modulation before transmission may transform the symmetric physical channels to non-symmetric logical channels instead. For example, consider a 4 Pulse-Amplitude Modulation (4PAM) with Gray mapping, namely, the constellation set is $\left\{-\frac{3}{\sqrt{5}}, -\frac{1}{\sqrt{5}}, \frac{1}{\sqrt{5}}, \frac{3}{\sqrt{5}}\right\}$ and it corresponds to information bit pairs $\{00, 01, 11, 10\}$. We further assume the underlying physical channel is additive white Gaussian with variance σ^2, which is perfectly symmetric. The logical channel for the second bit (the right bit) then becomes non-symmetric, which corresponds to the composite BiAWGNCs in Example 2.

Further discussion on channel symmetry from theoretical perspective can be found in Section 5.1.1.

4.1.2 Achievable Rates of Linear Codes on Non-Symmetric Discrete Memoryless Channels

In this brief subsection, we will digress to discuss some existing results on the achievable rates of linear codes. Although some notation used herein will not be formally defined until later sections, it is essential to have an idea what the performance limit is for general linear codes before discussing ultra powerful, bound-approaching (linear) LDPC codes on non-symmetric channels. We will focus on discrete memoryless channels (DMCs) of the following form: $\mathsf{GF}(p^k) \mapsto \mathsf{GF}(p^k)$, in which p is a prime number.

Let R^* and $R^*_{lin.}$ denote the highest achievable rates of general codes and of linear codes for the DMC of interest. The mutual information of the DMC under the *a priori* distribution P_X is $I(X;Y)$ and the Shannon capacity is defined as $C := \sup_{\mathsf{P}_X} I(X;Y)$. The symmetric mutual information rate smir is defined as $\mathsf{smir} := I_u(X;Y)$ for which P_X is *uniformly* distributed on the input set $\mathsf{GF}(p^k)$. The existing results on achievable rates are summarized as follows.

- Shannon [98] showed that for general DMCs,
$$R^* = C.$$
- Elias [39] proved that for BSCs, a typical linear code is able to achieve the Shannon capacity, namely,
$$R^*_{lin.} = R^* = C = \mathsf{smir}.$$
- Dobrushin [38] further extended Elias' results, showing that for m-ary symmetric channels (MSCs), the m-ary version of BSCs formally defined in Section 5.1.2,
$$R^*_{lin.} = R^* = C = \mathsf{smir}.$$

- Ahlswede [5] and Gemma [6] proved that for binary non-symmetric channels (BNSCs),

$$R^*_{lin.} = \mathsf{smir} < R^* = C. \tag{4.1}$$

  Furthermore, for high order $\mathsf{GF}(p^k)$-based non-symmetric DMCs, they showed that it is possible to have $R^*_{lin.} < \mathsf{smir}$. Fortunately, if we focus on a slightly broader class of codes, the coset codes formally defined in Section 4.6.1, the achievable rate R^*_{coset} satisfies

$$R^*_{lin.} \leq R^*_{coset} = \mathsf{smir} < R^* = C. \tag{4.2}$$

The above results show that using linear codes, one can only hope to achieve smir. For symmetric DMCs, smir equals the Shannon capacity, and is indeed achievable by linear codes. For non-symmetric DMCs, $\mathsf{smir} \neq C$. Linear codes are able to achieve smir when facing BNSCs. For higher order DMCs, there is sometimes a gap between $R^*_{lin.}$ and smir, which, however, can be recovered by focusing on coset codes, or, equivalently, by using the channel symmetrizing argument in Figure 4.11. It is worth emphasizing that the achievable rate of coset codes is still upper bounded by smir and thus is bounded away from the Shannon capacity.

For binary-input/non-symmetric-output channels, the difference between smir and C is generally indistinguishable from the practical point of view. In [79], it was shown that $\frac{\mathsf{smir}}{C} \geq \frac{e\ln(2)}{2} \approx 0.942$. [99] further proved that $|C - \mathsf{smir}| \leq 0.011$ bit/sym.

4.1.3 Linear LDPC Code Ensemble

The linear LDPC codes of length n are actually a special family of parity check codes, such that all codewords can be specified by the following parity check equation in $\mathsf{GF}(2)$:

$$\mathbf{Hx} = \mathbf{0},$$

where $\mathbf{H}$ is an $m \times n$ sparse matrix in $\mathsf{GF}(2)$ with the number of non-zero elements linearly proportional to n. Since the total number of entries in $\mathbf{H}$ is $\mathcal{O}\left(n^2\right)$, when n is sufficiently large, $\mathbf{H}$ becomes sparse or low-density, which is the source of the name of low-density parity-check (LDPC) codes.

A parity check code can always be written in an equivalent graphical representation. As illustrated in Figure 4.3, a parity check code satisfying $\mathbf{Hx} = \mathbf{0}$ is converted to its equivalent bipartite graph. The bipartite graph model consists of a bottom row of variable nodes (corresponding to codeword bits) and a top row of check nodes (corresponding to parity check equations). Let $H_{j,i}$ denote the entry of $\mathbf{H}$ located at the intersection of the j-th row and the i-th column. We use the convention that $H_{j,i} = 1$ if and only if there is an odd number of edges connecting variable node i and check node j. Throughout this chapter, we will interchangeably use the algebraic and the graphical representations.

To facilitate our analysis, we consider a code ensemble, i.e., a set of codes with probabilistic weights, rather than a single code. The linear LDPC code ensemble of interest is generated by equiprobable edge permutation in a regular bipartite graph. Detailed construction is as follows. Suppose we have n variable nodes on the bottom and each of them has d_v sockets. There are $m := \frac{nd_v}{d_c}$ check nodes on the top and each of them has d_c sockets. With these fixed $(n + m)$ nodes, there are a total of $(nd_v)!$ possible configurations obtained by connecting these $nd_v = md_c$ sockets on each side, assuming all sockets are

$$\mathbf{H} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}$$

$$\boxed{\mathbf{Hx} = \mathbf{0}}$$

(a) Algebraic representation

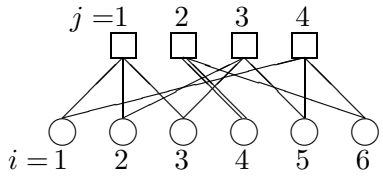

(b) Graphical representation

Figure 4.3: Two methods of representation of parity-check codes.

distinguishable.[1] The resulting graphs[2] will be regular and bipartite with degrees (d_v, d_c), and can be converted back to parity check codes using the aforementioned convention. We construct a regular code ensemble $\mathcal{C}^n(d_v, d_c)$ by putting equal probability on each of the possible configurations of the regular bipartite graphs described above. Since the graphical presentation in Figure 4.3(b) is a regular (2,3) graph with 6 variable nodes, it is an instance of the regular code ensemble $\mathcal{C}^6(2, 3)$. For practical interest, we assume $d_c > 2$.

For future use, we let i and j denote the indices of the i-th variable node and the j-th check node. $\{j_{i_0,c}\}_{c\in[1,d_v]}$ denotes all check nodes connecting to variable node i_0, and similarly does $\{i_{j_0,v}\}_{v\in[1,d_c]}$.

We can further expand the code ensemble by considering irregular bipartite graphs. Let λ and ρ denote the finite order *edge degree distribution* polynomials such that

$$\begin{aligned} \lambda(x) &= \sum_k \lambda_k x^{k-1} \\ \rho(x) &= \sum_k \rho_k x^{k-1}, \end{aligned}$$

where λ_k or ρ_k is the fraction of edges connecting to a degree k variable or check node, respectively. By assigning equal probability to each possible edge permutation of irregular bipartite graphs with fixed degree profiles λ and ρ (similarly to the regular case), we obtain the equiprobable, irregular, bipartite graph ensemble $\mathcal{C}^n(\lambda, \rho)$. For example: $\mathcal{C}^n(3, 6) = \mathcal{C}^n(x^2, x^5)$.

4.1.4 Message Passing Algorithms & Belief Propagation Decoders

The message passing decoding algorithm is a distributed algorithm such that each variable/check node has a processor, which takes all incoming messages from its neighbors as inputs, and outputs new messages back to all its neighbors. The decoding continues in such a distributed fashion, and hopefully, after many rounds of "message passing," the decision on the most probable transmitted codeword can be made based on the final messages. A message passing algorithm can be completely specified by the variable and check node message maps, Ψ_v and Ψ_c, which may or may not be stationary (i.e., the maps remain the same as time evolves) or uniform (i.e., node-independent). The message passing algorithm can be executed sequentially or in parallel depending on the order of the activations of different

[1] When assuming all variable/check node sockets are indistinguishable, the number of configurations can be upper bounded by $\frac{(nd_v)!}{(d_c!)^m}$.

[2] Sometimes the resulting graph is a multigraph, namely, a graph with multiple edges connecting the same pair of nodes.

node processors. Henceforth, we consider only parallel message passing algorithms with stationary uniform message maps. Furthermore, we focus on the message passing algorithms complying with the *extrinsic* principle (adopted from turbo codes), i.e. the new message sending to node i (or j) does not depend on the received message from the same node i (or j) but depends only on other received messages.

A belief propagation (BP) algorithm is one instance of message passing algorithms whose variable and check node message maps are derived from Pearl's inference network [86]. BP has recently been cast into many different forms [7, 8], and is generally considered as an implication of the factor graph technique [61], which is a more general framework for deriving inference algorithms. When the underlying inference network is cycle-free belief propagation calculates the exact marginal *a posteriori* probabilities, and thus we obtain the optimal maximum *a posteriori* probability (MAP) decisions. Let m_0 denote the initial message from the variable nodes, and $\{m_k\}$ denote the messages from its neighbors excluding that from the destination node. The entire belief propagation algorithm with messages representing the corresponding log likelihood ratio (LLR) can be described as follows:

$$\begin{aligned} m_0 &:= \ln \frac{\mathsf{P}(y_i|x_i=0)}{\mathsf{P}(y_i|x_i=1)} \\ \Psi_v(m_0, m_1, \cdots, m_{d_v-1}) &:= \sum_{j=0}^{d_v-1} m_j \qquad (4.3) \\ \Psi_c(m_1, \cdots, m_{d_c-1}) &:= \ln\left(\frac{1+\prod_{i=1}^{d_c-1}\tanh\frac{m_i}{2}}{1-\prod_{i=1}^{d_c-1}\tanh\frac{m_i}{2}}\right). \qquad (4.4) \end{aligned}$$

After many rounds of message passing, the decision is made by the following formula:

$$\hat{x}_i := 1_{\{(\sum_{j=0}^{d_v} m_j)<0\}},$$

where $1_{\{\cdot\}}$ is the indicator function, and $m_1, \cdots, m_{d_V}$ denote all incoming messages of variable node i.

It is worth noting that BP was originally devised as an optimal inference algorithm for cycle-free network. Although BP is still applicable to graphs with cycles, such as the LDPC codes in Figure 4.3(b), it is no longer optimal [41] and some tweaking of the algorithm may result in better performance [41, 104]. Furthermore, only the construction of m_0, the single-bit LLR, depends on the channel model, and can be easily calculated for non-symmetric channels as well. So the entire belief propagation algorithm remains applicable to memoryless, symbol-dependent channels.

- *Example:* For BNSCs,

$$m_0 = \begin{cases} \ln\frac{1-p_{0\to1}}{p_{1\to0}} & \text{if } y_i = 0 \\ \ln\frac{p_{0\to1}}{1-p_{1\to0}} & \text{if } y_i = 1 \end{cases}.$$

We assume that the belief propagation is executed in parallel and each *iteration* is a "round" in which all variable nodes send messages to all check nodes and then the check nodes send messages back. We use l to denote the number of iterations that have been performed.

4.1.5 Density Evolution

For a symmetric channel and any message-passing algorithm, the probability density of the transmitted messages in each iteration can be calculated iteratively with a concrete theoretical foundation [92]. The iterative formula and related theorems are termed "density evolution." Since the belief propagation algorithm performs extremely well under most circumstances and is of great importance, sometimes the term "density evolution" is reserved for the corresponding analytical method for belief propagation algorithms. The derivation of density evolution (DE) for symmetric memoryless channels will be explained in details while considering non-symmetric memoryless channels in Section 4.2.2.

4.2 New Density Evolution: An Iterative Formula

In what follows, we use the belief propagation algorithm as the illustrative example for our new iterative density evolution formula, which is applicable to general message passing algorithms as well.

With the assumption of channel symmetry and the inherent symmetry of the parity check equations in LDPC codes, the probability density of the messages in any symmetric message passing algorithm will be codeword independent. For different codewords, the densities of the messages passed differ only in *parities*, but all of them are of the same *shape* [Lemma 1, [92]]. Therefore, in [92], it was sufficient to assume that the all-zero codeword is transmitted.

In the symbol-dependent setting, symmetry of the channel in general does not hold. Even though the belief propagation mappings remain the same for non-symmetric channels, the densities of the messages for different transmitted codewords are of different shapes and parities. Hence the density obtained from the all-zero codeword assumption cannot represent the behavior when other codewords are transmitted. To circumvent this problem, we *average* the density of the messages over all valid codewords. However, directly averaging over all codewords takes 2^{n-m} times more computations, which ruins the efficiency of the iterative formula for density evolution. Henceforth, we provide a new iterative formula for the codeword-averaged density evolution which increases the number of computations only by a constant factor; the corresponding theoretical foundations are provided in this section and in Section 4.3.

4.2.1 Preliminaries & the Perfect Projection Condition

We consider the density of the message passed from variable node i to check node j. The probability density of this message is denoted by $P^{(l)}_{(i,j)}(\mathbf{x})$ where the superscript l denotes the l-th iteration and the appended argument $\mathbf{x}$ denotes the actual transmitted codeword. For example, $P^{(1)}_{(i,j)}(\mathbf{0})$ is the density of the initial message m_0 from variable node i to check node j assuming the all-zero codeword is transmitted. $P^{(2)}_{(i,j)}(\mathbf{0})$ is the density of the message from i to j during the second iteration, and so on. We also denote by $Q^{(l)}_{(j,i)}(\mathbf{x})$ the density of the message from check node j to variable node i in the l-th iteration.

With the assumption that the corresponding graph is tree-like until depth $2(l-1)$, we define the following quantities. Figure 4.4 illustrates these quantities for the code in Figure 4.3 with $i = j = 1$ and $l = 2$.

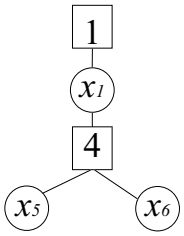

$$\mathbf{X}^l_{(1,1)} := \{x_1x_5x_6 : x_1x_5x_6 = 000, 011, 101, 110\}$$

Figure 4.4: Illustration of $\mathbf{X}^l_{(1,1)}$ and $\mathcal{N}^{2l}_{(1,1)}$ with $l = 2$.

- $\mathcal{N}^{2l}_{(i,j)}$ denotes the tree-like subset of the graph[3] $G = (\mathcal{V}, \mathcal{E})$ with root edge (i, j) and depth $2(l-1)$, named as the supporting tree. A formal definition is: $\mathcal{N}^{2l}_{(i,j)}$ is the subgraph induced by $\mathcal{V}^{2l}_{(i,j)}$, where

$$\mathcal{V}^{2l}_{(i,j)} := \{v \in \mathcal{V} : d(v,i) = d(v,j) - 1 \in [0, 2(l-1)]\}, \tag{4.5}$$

  where $d(v, i)$ is the shortest distance between node v and variable node i. In other words, $\mathcal{N}^{2l}_{(i,j)}$ is the depth $2(l-1)$ tree spanned from edge (i, j). Let $\left|\mathcal{N}^{2l}_{(i,j)}\right|_V$ denote the number of variable nodes in $\mathcal{N}^{2l}_{(i,j)}$ (including variable node i). $\left|\mathcal{N}^{2l}_{(i,j)}\right|_C$ denotes the number of check nodes in $\mathcal{N}^{2l}_{(i,j)}$ (check node j is excluded by definition).

- $\mathbf{X} = \{\mathbf{x} \in \{0,1\}^n : \mathbf{Hx} = \mathbf{0}\}$ denotes the set of all valid codewords, and the information source selects each codeword equiprobably from $\mathbf{X}$.

- $\mathbf{x}|_i$ and $\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}}$ are the projections of codeword $\mathbf{x} \in \mathbf{X}$ on bit i and on the variable nodes in the supporting tree $\mathcal{N}^{2l}_{(i,j)}$, respectively.

- $\mathbf{X}^l_{(i,j)}$ denotes the set of all strings of length $\left|\mathcal{N}^{2l}_{(i,j)}\right|_V$ satisfying the $\left|\mathcal{N}^{2l}_{(i,j)}\right|_C$ check node constraints in $\mathcal{N}^{2l}_{(i,j)}$. $\mathbf{x}^l$ denotes any element of $\mathbf{X}^l_{(i,j)}$ (the subscript (i, j) is omitted if there is no ambiguity). The connection between $\mathbf{X}$, the valid codewords, and $\mathbf{X}^l_{(i,j)}$, the tree-satisfying strings, will be clear in the following remark and in Definition 4.1.

- For any set of codewords (or strings) $\mathbf{W}$, the average operator $\langle\cdot\rangle_{\mathbf{W}}$ is defined as:

$$\langle g(\mathbf{x})\rangle_{\mathbf{W}} = \frac{1}{|\mathbf{W}|}\sum_{\mathbf{x}\in\mathbf{W}} g(\mathbf{x}).$$

- With a slight abuse of notation $P^{(l)}_{(i,j)}(\mathbf{x})$, we define

$$\begin{aligned} P^{(l)}_{(i,j)}(x) &:= \left\langle P^{(l)}_{(i,j)}(\mathbf{x})\right\rangle_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_i = x\}} \\ P^{(l)}_{(i,j)}(\mathbf{x}^l) &:= \left\langle P^{(l)}_{(i,j)}(\mathbf{x})\right\rangle_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}} = \mathbf{x}^l\}}. \end{aligned}$$

[3]The calligraphic $\mathcal{V}$ in $G = (\mathcal{V}, \mathcal{E})$ denotes the set of all vertices, including both variable nodes and check nodes. Namely, a node $v \in \mathcal{V}$ can be a variable or check node.

Namely, $P^{(l)}_{(i,j)}(x)$ and $P^{(l)}_{(i,j)}(\mathbf{x}^l)$ denote the density averaged over all compatible codewords having projections x and $\mathbf{x}^l$, respectively.

Remark: For any tree-satisfying string $\mathbf{x}^l \in \mathbf{X}^l_{(i,j)}$, there may or may not be a codeword $\mathbf{x}$ with projection $\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}} = \mathbf{x}^l$, since the codeword $\mathbf{x}$ must satisfy *all* check nodes, but the string $\mathbf{x}^l$ needs to satisfy only $\left|\mathcal{N}^{2l}_{(i,j)}\right|_C$ constraints. Those check nodes outside $\mathcal{N}^{2l}_{(i,j)}$ may limit the projected space $\mathbf{X}|_{\mathcal{N}^{2l}_{(i,j)}}$ to a strict subset of $\mathbf{X}^l_{(i,j)}$. For example, the second row of $\mathbf{Hx} = \mathbf{0}$ in Figure 4.3 implies $x_6 = 0$. Therefore two of the four elements of $\mathbf{X}^l_{(1,1)}$ in Figure 4.4 are invalid/impossible projections of valid codewords $\mathbf{x} \in \mathbf{X}$. Thus $\mathbf{X}|_{\mathcal{N}^{2l}_{(1,1)}}$ is a proper subset of $\mathbf{X}^l_{(1,1)}$.

To capture this phenomenon, we introduce the notion of a *perfectly projected* $\mathcal{N}^{2l}_{(i,j)}$.

Definition 4.1 (Perfectly Projected $\mathcal{N}^{2l}_{(i,j)}$) *The supporting tree $\mathcal{N}^{2l}_{(i,j)}$ is perfectly projected, if for any $\mathbf{x}^l \in \mathbf{X}^l_{(i,j)}$,*

$$\frac{\left|\{\mathbf{x} \in \mathbf{X} : \mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}} = \mathbf{x}^l\}\right|}{|\mathbf{X}|} = \frac{1}{\left|\mathbf{X}^l_{(i,j)}\right|}. \tag{4.6}$$

That is, if we choose $\mathbf{x} \in \mathbf{X}$ equiprobably, $\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}}$ will appear uniformly among all elements in $\mathbf{x}^l \in \mathbf{X}^l_{(i,j)}$. Thus by looking only at the projections on $\mathcal{N}^{2l}_{(i,j)}$, it is as if we are choosing $\mathbf{x}^l$ from $\mathbf{X}^l_{(i,j)}$ equiprobably and there are only $\left|\mathcal{N}^{2l}_{(i,j)}\right|_C$ check node constraints and no others.

The example in Figures 4.3 and 4.4 is obviously not perfectly projected. By the linearity of LDPC codes, it is straightforward to see that all projections $\mathbf{x}^l$ with at least one compatible codeword $\mathbf{x} \in \mathbf{X}$ will result in the same value of the left-hand side of (4.6). Based on this reasoning, to show an $\mathcal{N}^{2l}_{(i,j)}$ is perfectly projected, we need only to show that all $\mathbf{x}^l \in \mathbf{X}^l_{(i,j)}$ are valid projections with at least one compatible codeword respectively.

Since the message emitted from node i to j in the l-th iteration depends only on the received signals of the supporting tree, $\mathbf{y}|_{\mathcal{N}^{2l}_{(i,j)}}$, the codeword-dependent $P^{(l)}_{(i,j)}(\mathbf{x})$ depends only on the projection $\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}}$, not on the entire codeword $\mathbf{x}$. Namely,

$$P^{(l)}_{(i,j)}(\mathbf{x}) = P^{(l)}_{(i,j)}(\mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}}), \quad \forall \mathbf{x} \in \mathbf{X}. \tag{4.7}$$

An immediate implication of (4.7) and $\mathcal{N}^{2l}_{(i,j)}$ being perfectly projected is

$$\begin{aligned}
P^{(l)}_{(i,j)}(x) &:= \left\langle P^{(l)}_{(i,j)}(\mathbf{x}) \right\rangle_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_i = x\}} \\
&= \frac{1}{|\{\mathbf{x} \in \mathbf{X} : \mathbf{x}|_i = x\}|} \sum_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_i=x\}} P^{(l)}_{(i,j)}(\mathbf{x}) \\
&= \frac{1}{|\{\mathbf{x} \in \mathbf{X} : \mathbf{x}|_i = x\}|} \cdot \left|\{\mathbf{x} \in \mathbf{X} : \mathbf{x}|_{\mathcal{N}^{2l}_{(i,j)}} = \mathbf{x}^l, \mathbf{x}^l|_i = x\}\right| \cdot \sum_{\{\mathbf{x}^l\in\mathbf{X}^l_{(i,j)}:\mathbf{x}^l|_i=x\}} P^{(l)}_{(i,j)}(\mathbf{x}^l) \\
&= \left\langle P^{(l)}_{(i,j)}(\mathbf{x}^l) \right\rangle_{\{\mathbf{x}^l\in\mathbf{X}^l_{(i,j)}:\mathbf{x}^l|_i=x\}}.
\end{aligned} \tag{4.8}$$

Because of these two useful properties, (4.7) and (4.8), throughout this subsection we assume that $\mathcal{N}_{(i,j)}^{2l}$ is perfectly projected. The convergence of $\mathcal{N}_{(i,j)}^{2l}$ to a perfect projection in probability is dealt with in Section 4.3. We will have all the preliminaries necessary for deriving the new density evolution after introducing the following self-explanatory lemma.

Lemma 4.1 (Linearity of Density Transformation) *For any random variable A with distribution P_A, if $g : A \mapsto g(A)$ is measurable, then $B = g(A)$ is a random variable with distribution $P_B = T_g(P_A) := P_A \circ g^{-1}$. Furthermore, the density transformation T_g is linear. I.e. if $P_B = T_g(P_A)$ and $Q_B = T_g(Q_A)$, then $\alpha P_B + (1-\alpha)Q_B = T_g(\alpha P_A + (1-\alpha)Q_A)$, $\forall \alpha \in [0,1]$.*

4.2.2 A New Iterative Formula

In the l-th iteration, the probability of sending an incorrect message (averaged over all possible codewords) from variable node i_0 to check node j_0 is

$$\begin{aligned} p_e^{(l)}(i_0,j_0) &= \frac{1}{|\mathbf{X}|}\left(\sum_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_{i_0}=0\}}\int_{m=-\infty}^{0} P_{(i_0,j_0)}^{(l)}(\mathbf{x})(dm) + \sum_{\{\mathbf{x}\in\mathbf{X}:\mathbf{x}|_{i_0}=1\}}\int_{m=0}^{\infty} P_{(i_0,j_0)}^{(l)}(\mathbf{x})(dm)\right) \\ &= \frac{1}{2}\left(\int_{m=-\infty}^{0} P_{(i_0,j_0)}^{(l)}(0)(dm) + \int_{m=0}^{\infty} P_{(i_0,j_0)}^{(l)}(1)(dm)\right). \end{aligned} \tag{4.9}$$

Motivated by (4.9), we concentrate on finding an iterative formula for the density pair $P_{(i_0,j_0)}^{(l)}(0)$ and $P_{(i_0,j_0)}^{(l)}(1)$.

The cycle free assumption guarantees that no variable node i appears twice in $\mathcal{N}_{(i_0,j_0)}^{2l}$, and therefore all incoming messages depend on disjoint subsets of received signals $\{y_i\}$. Since we are considering memoryless channels, given the transmitted codeword $\mathbf{x}$, all incoming messages are independently distributed. Using an auxiliary function $\gamma(m)$:

$$\gamma(m) := \left(1_{\{m\le 0\}}, \ln\coth\left|\frac{m}{2}\right|\right) \in \mathsf{GF}(2)\times\mathbb{R}^+,$$

the definition of Ψ_c, (4.4), can be rewritten as

$$\Psi_c(m_1,\cdots,m_{d_c-1}) = \gamma^{-1}\left(\sum_{v=1}^{d_c-1}\gamma(m_v)\right). \tag{4.10}$$

By (4.3), (4.10), and the independence among the input messages, there exists an iterative formula for the densities:

$$P_{(i_0,j_0)}^{(l)}(\mathbf{x}) = P_{(i_0,j_0)}^{(0)}(\mathbf{x}) \otimes \left(\bigotimes_{c=1}^{d_v-1} Q_{(j_{i_0,c},i_0)}^{(l-1)}(\mathbf{x})\right) \tag{4.11}$$

$$Q_{(j_{i_0,c},i_0)}^{(l-1)}(\mathbf{x}) = \Gamma^{-1}\left(\bigotimes_{v=1}^{d_c-1}\Gamma\left(P_{(i_{j,v},j_{i_0,c})}^{(l-1)}(\mathbf{x})\right)\right), \tag{4.12}$$

where $\otimes$ denotes the convolution operator on probability density functions, which can be implemented efficiently using the Fourier transform. $\Gamma := T_\gamma$ is the density transformation functional based on γ, defined in Lemma 4.1. When considering memoryless *symmetric*

channels, the all-zero codeword can be assumed ($\mathbf{x} = \mathbf{0}$), and all edges (i, j) become indistinguishable. Based on the above observation, we obtain the classical density evolution (Eq. (9) in [90]) as follows by dropping the codeword arguments and the edge subscripts.

$$\begin{aligned} P^{(l)} &= P^{(0)} \otimes \left(Q^{(l-1)}\right)^{\otimes d_v - 1} \\ Q^{(l-1)} &= \Gamma^{-1}\left(\left(\Gamma\left(P^{(l-1)}\right)\right)^{\otimes d_c - 1}\right). \end{aligned} \tag{4.13}$$

For non-symmetric channels, by (4.7), (4.11), and the perfect projection assumption, we have

$$P^{(l)}_{(i_0,j_0)}(\mathbf{x}^l) = P^{(0)}_{(i_0,j_0)}(\mathbf{x}|_{i_0}) \otimes \left(\bigotimes_{c=1}^{d_v-1} Q^{(l-1)}_{(j_{i_0,c},i_0)}(\mathbf{x}^l)\right). \tag{4.14}$$

Further simplification can be made such that

$$\begin{aligned} P^{(l)}_{(i_0,j_0)}(x) &\overset{(a)}{=} \left\langle P^{(l)}_{(i_0,j_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\ &\overset{(b)}{=} \left\langle P^{(0)}_{(i_0,j_0)}(x) \otimes \left(\bigotimes_{c=1}^{d_v-1} Q^{(l-1)}_{(j_{i_0,c},i_0)}(\mathbf{x}^l)\right)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\ &\overset{(c)}{=} P^{(0)}_{(i_0,j_0)}(x) \otimes \left\langle \bigotimes_{c=1}^{d_v-1} Q^{(l-1)}_{(j_{i_0,c},i_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\ &\overset{(d)}{=} P^{(0)}_{(i_0,j_0)}(x) \otimes \left(\bigotimes_{c=1}^{d_v-1} \left\langle Q^{(l-1)}_{(j_{i_0,c},i_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}}\right) \\ &\overset{(e)}{=} P^{(0)}_{(i_0,j_0)}(x) \otimes \left(\left\langle Q^{(l-1)}_{(j_{i_0,1},i_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}}\right)^{\otimes(d_v-1)}, \end{aligned} \tag{4.15}$$

where (a) follows from (4.8), (b) follows from (4.14), and (c) follows from the linearity of convolutions. The fact that the sub-trees generated by edges $(j_{i_0,c}, i_0)$ are completely disjoint implies that, by the perfect projection assumption on $\mathcal{N}^{2l}_{(i_0,j_0)}$, the distributions of "strings $\mathbf{x}^l$" on different sub-trees are independent. As a result, the average of the convolutional products (over these strings) equals the convolution of the averaged distributions, yielding (d). Finally (e) follows from the fact that the distributions of messages from different subtrees are indistinguishable according to the perfect projection assumption.

To simplify $\left\langle Q^{(l-1)}_{(j_{i_0,1},i_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}}$, we need to define some new notation. We use j_1 to represent $j_{i_0,1}$ for simplicity. Denote by $\left\{\mathcal{N}^{2(l-1)}_{(i_{j_1,v},j_1)}\right\}_{v\in[1,d_c-1]}$ the collection of all $d_c - 1$ subtrees rooted at $(i_{j_1,v}, j_1)$, $v \in [1, d_c - 1]$, and by $\mathbf{X}^{l-1}_{(i_{j_1,v},j_1)}$ the strings compatible to $\mathcal{N}^{2(l-1)}_{(i_{j_1,v},j_1)}$. We then define

$$\mathbf{X}^1(x) := \left\{(x_1, \cdots, x_{d_c-1}) : \left(\sum_{v=1}^{d_c-1} x_v\right) + x = 0\right\}$$

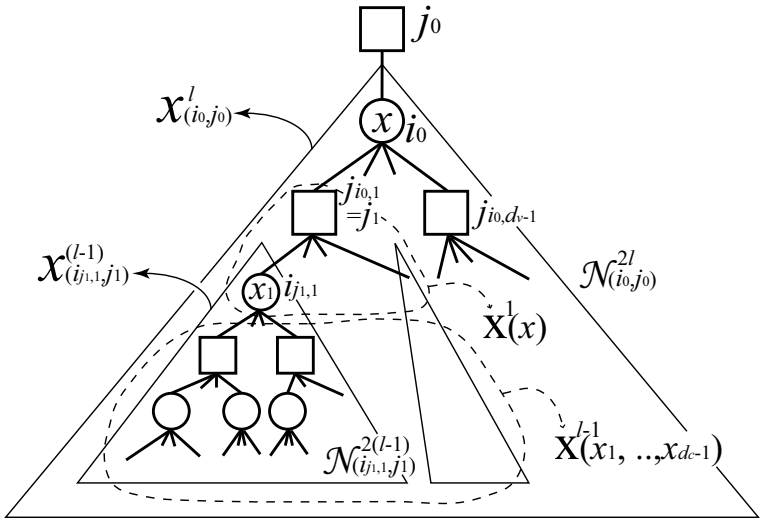

Figure 4.5: Illustration of various quantities used in Section 4.2.

containing the strings satisfying parity check constraint j_1 given $x_{i_0} = x$, and

$$\mathbf{X}^{l-1}(x_1, \cdots, x_{d_c-1}) := \left\{ \left(\mathbf{x}^{l-1}_{(i_{j_1,1},j_1)}, \cdots, \mathbf{x}^{l-1}_{(i_{j_1,d_c-1},j_1)} \right) : \mathbf{x}^{l-1}_{(i_{j_1,1},j_1)}|_{i_{j_1,1}} = x_1, \cdots, \mathbf{x}^{l-1}_{(i_{j_1,d_c-1},j_1)}|_{i_{j_1,d_c-1}} = x_{d_c-1} \right\}$$

is the collection of the concatenations of substrings, in which the leading symbols of the substrings are $(x_1, \cdots, x_{d_c-1})$. All these quantities are illustrated in Figure 4.5.

Note the following two properties: (i) For any v, the message m_v from variable $i_{j_1,v}$ to check node j_1 depends only on $\mathbf{x}^{l-1}_{(i_{j_1,v},j_1)}$; and (ii) With the leading symbols $\{x_v\}_{v\in[1,d_c-1]}$ fixed and the perfect projection assumption, the projection on the strings $\left\{\mathbf{x}^{l-1}_{(i_{j_1,v},j_1)}\right\}_{v\in[1,d_c-1]}$ are independent. Thus the averaged convolution of densities is equal to the convolution of the averaged densities. By repeatedly applying Lemma 4.1 and the above two properties,

we have

$$\begin{aligned}
&\left\langle Q^{(l-1)}_{(j_{i_0,1},i_0)}(\mathbf{x}^l)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\
&= \left\langle \Gamma^{-1}\left(\bigotimes_{v=1}^{d_c-1}\Gamma\left(P^{(l-1)}_{(i_{j,v},j_{i_0,c})}(\mathbf{x}^l)\right)\right)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\
&= \left\langle \Gamma^{-1}\left(\bigotimes_{v=1}^{d_c-1}\Gamma\left(P^{(l-1)}_{(i_{j,v},j_{i_0,c})}(\mathbf{x}^{l-1}_{(i_{j_1,v},j_1)})\right)\right)\right\rangle_{\{\mathbf{x}^l:\mathbf{x}^l|_{i_0}=x\}} \\
&= \Gamma^{-1}\left(\frac{1}{2^{d_c-2}}\sum_{\{\mathbf{x}^1:\mathbf{x}^1\in\mathbf{X}^1(x)\}}\left\langle\bigotimes_{v=1}^{d_c-1}\Gamma\left(P^{(l-1)}_{(i_{j,v},j_{i_0,c})}(\mathbf{x}^{l-1}_{(i_{j_1,v},j_1)})\right)\right\rangle_{\mathbf{X}^{l-1}(\mathbf{x}^1)}\right) \\
&= \Gamma^{-1}\left(\frac{1}{2^{d_c-2}}\sum_{\{\mathbf{x}^1:\mathbf{x}^1\in\mathbf{X}^1(x)\}}\bigotimes_{v=1}^{d_c-1}\Gamma\left(\left\langle P^{(l-1)}_{(i_{j,v},j_{i_0,c})}(\mathbf{x}^{l-1}_{(i_{j_1,v},j_1)})\right\rangle_{\mathbf{X}^{l-1}(\mathbf{x}^1)}\right)\right) \\
&= \Gamma^{-1}\left(\frac{1}{2^{d_c-2}}\sum_{\{\mathbf{x}^1:\mathbf{x}^1\in\mathbf{X}^1(x)\}}\bigotimes_{v=1}^{d_c-1}\Gamma\left(P^{(l-1)}_{(i_{j,v},j_{i_0,c})}(x_v)\right)\right) \qquad (4.16)
\end{aligned}$$

By (4.15), (4.16), and dropping the subscripts during the density evolution, we obtain the desired iterative formulae for $P^{(l)}(0)$ and $P^{(l)}(1)$ as follows.

$$\begin{aligned}
P^{(l)}(x) &= P^{(0)}(x)\otimes\left(Q^{(l-1)}(x)\right)^{\otimes(d_v-1)} \\
Q^{(l-1)}(x) &= \Gamma^{-1}\left(\frac{1}{2^{d_c-2}}\sum_{\mathbf{x}^1\in\mathbf{X}^1(x)}\bigotimes_{v=1}^{d_c-1}\Gamma\left(P^{(l-1)}(x_v)\right)\right) \\
&= \Gamma^{-1}\left(\frac{1}{2^{d_c-2}}\sum_{\{v\in[0,d_c-1]:(-1)^{v+x}=1\}}\binom{d_c-1}{v}\Gamma\left(P^{(l-1)}(0)\right)^{\otimes(d_c-1-v)}\otimes\Gamma\left(P^{(l-1)}(1)\right)^{\otimes v}\right) \\
&\overset{(a)}{=} \Gamma^{-1}\left(\left(\Gamma\left(\frac{P^{(l-1)}(0)+P^{(l-1)}(1)}{2}\right)\right)^{\otimes(d_c-1)}\right. \\
&\qquad \left.+(-1)^x\left(\Gamma\left(\frac{P^{(l-1)}(0)-P^{(l-1)}(1)}{2}\right)\right)^{\otimes(d_c-1)}\right),
\end{aligned}$$

where (a) follows from the linearity of distribution transformations and convolutions. The above formula can be easily generalized to the irregular code ensemble $\mathcal{C}^n(\lambda,\rho)$:

$$\begin{aligned}
P^{(l)}(x) &= P^{(0)}(x)\otimes\lambda\left(Q^{(l-1)}(x)\right) \\
Q^{(l-1)}(x) &= \Gamma^{-1}\left(\rho\left(\Gamma\left(\frac{P^{(l-1)}(0)+P^{(l-1)}(1)}{2}\right)\right)\right. \\
&\qquad \left.+(-1)^x\rho\left(\Gamma\left(\frac{P^{(l-1)}(0)-P^{(l-1)}(1)}{2}\right)\right)\right), \qquad (4.17)
\end{aligned}$$

where all scalar products in λ and ρ are replaced by convolutions. It can be easily seen that our new formula (4.17) has the same complexity as the classical density evolution for symmetric channels (4.13).

Remark: The above derivation relies heavily on the perfect projection assumption, which guarantees that averaging over all codewords is equivalent to averaging over the tree-satisfying strings. Since the tree-satisfying strings are well-structured and symmetric, we are on solid ground to move the average inside the classical density evolution formula.

4.3 New Density Evolution: Fundamental Theorems

As stated in Section 4.2, the cycle-free and the perfect projection assumptions are critical in our analysis. Furthermore, the use of codeword ensembles rather than fixed codes facilitates the analysis but its relationship to fixed codes still needs to be explored. We restate two necessary theorems from [92], and give a novel perfect projection convergence theorem, which is essential to our new density evolution method. With these theorems, a concrete theoretical foundation will be established.

Theorem 4.1 (Convergence to the Cycle-Free Case, [92]) *Fix l, i_0, and j_0. For any (d_v, d_c), there exists a constant $\alpha > 0$, such that for all $n \in \mathbb{N}$, the code ensemble $\mathcal{C}^n(d_v, d_c)$ satisfies*

$$\mathsf{P}\left(\mathcal{N}^{2l}_{(i_0,j_0)} \text{ is cycle-free}\right) \geq 1 - \alpha\left(\frac{\{(d_v-1)(d_c-1)\}^{2l}}{n}\right),$$

where $\mathcal{N}^{2l}_{(i_0,j_0)}$ is the support tree as defined by (4.5).

Theorem 4.2 (Convergence to Perfect Projection in Probability) *Fix l, i_0, and j_0. For any regular, bipartite, equiprobable graph ensemble $\mathcal{C}^n(d_v, d_c)$, we have*

$$\mathsf{P}\left(\mathcal{N}^{2l}_{(i_0,j_0)} \text{ is perfectly projected}\right) = 1 - \mathcal{O}(n^{-0.1}).$$

Note: The above two theorems focus only on the properties of equiprobable regular bipartite graph ensembles, and are independent of the channel model of interest.

Theorem 4.3 (Concentration to the Expectation, [92]) *With fixed transmitted codeword $\mathbf{x}$, let Z denote the number of wrong messages (those m's such that $m(-1)^x < 0$). There exists a constant $\beta > 0$ such that for any $\epsilon > 0$, over the code ensemble $\mathcal{C}^n(d_v, d_c)$ and the channel realizations $\mathbf{y}$, we have*

$$\mathsf{P}\left(\left|\frac{Z - \mathsf{E}\{Z\}}{nd_v}\right| > \frac{\epsilon}{2}\right) \leq 2e^{-\beta\epsilon^2 n}. \tag{4.18}$$

Furthermore, β is independent of $F_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x})$, and thus is independent of $\mathbf{x}$.

Theorem 4.3 can easily be generalized to symbol-dependent channels as in the following corollary.

Corollary 4.1 *Over the equiprobable codebook $\mathbf{X}$, the code ensemble $\mathcal{C}^n(d_v, d_c)$, and channel realizations $\mathbf{y}$, (4.18) still holds.*

Proof: Since the constant β in Theorem 4.3 is independent of the transmitted codeword $\mathbf{x}$, after averaging over the equiprobable codebook $\mathbf{X}$, the inequality still holds. That is,

$$\mathsf{P}\left(\left|\frac{Z-\mathsf{E}\{Z\}}{nd_v}\right|>\frac{\epsilon}{2}\right)=\mathsf{E}_{\mathbf{x}}\left\{\mathsf{P}\left(\left|\frac{Z-\mathsf{E}\{Z\}}{nd_v}\right|>\frac{\epsilon}{2}\Big|\,\mathbf{x}\right)\right\}\leq\mathsf{E}_{\mathbf{x}}\left\{2e^{-\beta\epsilon^2 n}\right\}=2e^{-\beta\epsilon^2 n}.$$

■

Now we have all the prerequisite of proving the theoretical foundation of our codeword-averaged density evolution.

Theorem 4.4 (Validity of Codeword-Averaged Density Evolution) *Consider any regular, bipartite, equiprobable graph ensemble $\mathcal{C}^n(d_v,d_c)$ with fixed l, i_0, and j_0. $p_e^{(l)}(i_0,j_0)$ is derived from (4.9) and the codeword-averaged density evolution. The probability over equiprobable codebook $\mathbf{X}$, the code ensemble $\mathcal{C}^n(d_v,d_c)$, and the channel realizations $\mathbf{y}$, satisfies*

$$\mathsf{P}\left(\left|\frac{Z}{nd_v}-p_e^{(l)}(i_0,j_0)\right|>\epsilon\right)=e^{-\epsilon^2\mathcal{O}(n)},\forall\epsilon>0.$$

Proof: We note that $\frac{Z}{nd_v}$ is bounded between 0 and 1. By observing that

$$\begin{aligned}&\left(\frac{Z}{nd_v}\right)1\{\mathcal{N}^{2l}_{(i_0,j_0)}\text{ is cycle-free and perfectly projected}\}\\&\quad\leq\left(\frac{Z}{nd_v}\right)\\&\quad\leq\left(\frac{Z}{nd_v}-1\right)1\{\mathcal{N}^{2l}_{(i_0,j_0)}\text{ is cycle-free and perfectly projected}\}+1,\end{aligned}$$

and using Theorems 4.1 and 4.2, we have $\lim_{n\to\infty}\mathsf{E}\left\{\frac{Z}{nd_v}\right\}=p_e^{(l)}(i_0,j_0)$. Then by Corollary 4.1, the proof is complete. ■

The proof of Theorem 4.2 will be based on a novel constraint propagation argument, and will be included in Appendix E.1

Remark: Theorem 4.1 focuses on a graphical property of the code, namely, whether the corresponding bipartite graph is cycle-free until a certain depth. On the other hand, Theorem 4.2 provides the missing link between the cycle-free structure, a graphical property, and the averaging over the entire codebook, an algebraic property. This new connection between the graph and the code structures turns out to be critical while analyzing initial performance of BP on LDPC codes, and will be further discussed in Section 4.6.2.

4.4 Monotonicity, Symmetry, & Stability

In this section, we prove the monotonicity, symmetry, and stability of our codeword-averaged density evolution method on belief propagation algorithms. Since the codeword-averaged density evolution collapses to the traditional one when the channel of interest is symmetric, as expected, the theorems introduced herein also collapse to those for traditional DE with symmetric channels; see [92] for reference.

4.4.1 Monotonicity

Proposition 4.1 (Monotonicity with Respect to l) *Let $p_e^{(l)}$ denote the bit error probability of the codeword-averaged density evolution defined in (4.9). Then $p_e^{(l+1)} \leq p_e^{(l)}$, for all $l \in \mathbb{N}$.*

Proof: We first note that the codeword-averaged approach can be viewed as concatenating a bit-to-sequence random mapper with the observation channels, and the larger the tree-structure is, the more observation/information the decision maker has. Since the BP decoder is the optimal MAP decoder for the tree structure of interest, the larger the tree is, the smaller the error probability will be. The proof is thus complete. ■

Proposition 4.2 (Monotonicity with Respect to Physically Degraded Channels) *Let $F(dy|x)$ and $G(dy|x)$ denote two different channel models, such that $G(dy|x)$ is physically degraded with respect to $F(dy|x)$, namely, there exists a conditional distribution $K(dy|z)$ such that $F(dy|x) = \int_z K(dy|z)G(dz|x)$ (see [33] for further references). The corresponding decoding error probabilities, $p_{e,F}^{(l)}$ and $p_{e,G}^{(l)}$, for channels F and G are defined in (4.9). Then for any fixed l, we have $p_{e,F}^{(l)} \leq p_{e,G}^{(l)}$.*

Proof: Since the codeword-averaged approach is a concatenation of a bit-to-sequence random mapper with independent observation channels, this theorem can be easily proved by the channel degradation argument. ■

- *Example:* For BNSCs, suppose the crossover probabilities of channels F and G are $(p_{0\to1}, p_{1\to0})$ and $(p'_{0\to1}, p'_{1\to0})$, and we also assume $(p_{0\to1}, p_{1\to0}) \leq (p'_{0\to1}, p'_{1\to0})$, $p'_{0\to1}+p'_{1\to0} \leq 1$. The channel G is then physically degraded with respect to channel F.

4.4.2 Symmetry

Even though the evolved density is derived from non-symmetric channels, there are still some symmetry properties inherent in the symmetric structure of belief propagation algorithms. We first define the symmetric distribution pair as follows.

Definition 4.2 (Symmetric Distribution Pairs) *Two probability measures P and Q are a symmetric pair if for any integrable function h, we have*

$$\int h(m)d\mathsf{P}(m) \quad = \int e^{-m}h(-m)d\mathsf{Q}(m).$$

A distribution P_s is self-symmetric if $(\mathsf{P}_s, \mathsf{P}_s)$ is a symmetric pair.

Proposition 4.3 *Let $I(m) := -m$ be a parity reversing function, and let $P^{(l)}(0)$ and $P^{(l)}(1)$ denote the resulting density functions from the codeword-averaged density evolution. Then $P^{(l)}(0)$ and $P^{(l)}(1) \circ I^{-1}$ are a symmetric pair for all $l \in \mathbb{N}$.*

Remark: For symmetric channels, $P^{(l)}(0)$ and $P^{(l)}(1)$ differ only in parity (Lemma 1, [92]). Thus, $P^{(l)}(0) = P^{(l)}(1) \circ I^{-1}$ is self-symmetric [Theorem 3 in [90]].

Proof: We note that by the equiprobable codeword distribution and the perfect projection assumption, $P^{(l)}(0)$ and $P^{(l)}(1)$ act on the random variable m, given by

$$m := \ln \frac{\mathsf{P}(x=0|\mathbf{y}^l)}{\mathsf{P}(x=1|\mathbf{y}^l)} = \ln \frac{\mathsf{P}(\mathbf{y}^l|x=0)}{\mathsf{P}(\mathbf{y}^l|x=1)},$$

where $\mathbf{y}^l$ is the received signal on the subset $\mathcal{N}^{2l}$ and P is the distribution over channel realizations and equiprobable codewords. Then by a change of measure,

$$\begin{aligned}\int h(m)P^{(l)}(0)(dm) &= \mathsf{E}_{x=0}\left\{h\left(\ln\frac{\mathsf{P}(\mathbf{y}^l|x=0)}{\mathsf{P}(\mathbf{y}^l|x=1)}\right)\right\}\\ &= \mathsf{E}_{x=1}\left\{\frac{\mathsf{P}(\mathbf{y}^l|x=0)}{\mathsf{P}(\mathbf{y}^l|x=1)}h\left(\ln\frac{\mathsf{P}(\mathbf{y}^l|x=0)}{\mathsf{P}(\mathbf{y}^l|x=1)}\right)\right\}\\ &= \int e^m h(m)P^{(l)}(1)(dm). \qquad (4.19)\end{aligned}$$

This completes the proof. ■

A straightforward corollary of Proposition 4.3 is stated as follows.

Corollary 4.2

$$\langle P^{(l)}\rangle := \frac{P^{(l)}(0)+P^{(l)}(1)\circ I^{-1}}{2}$$

is self-symmetric for all $l\in\mathbb{N}$, i.e. $(\langle P^{(l)}\rangle, \langle P^{(l)}\rangle)$ is a symmetric pair for all $l\in\mathbb{N}$.

4.4.3 Stability

Rather than looking only at the error probability $p_e^{(l)}$ of the evolved densities $P^{(l)}(0)$ and $P^{(l)}(1)$, we also focus on its Bhattacharyya noise parameter (BNP):

$$\mathrm{BNP}^{(l)}(x) := \int_m e^{-\frac{(-1)^x m}{2}} P^{(l)}(x)(dm).$$

For symmetric channels, e.g. BSCs, the physical meaning of the Bhattacharyya noise parameter is as follows. Consider two repetition strings of length n, one of which is the all-zero string $\mathbf{0}$ and the other is the all-one string $\mathbf{1}$. Suppose the *a priori* distribution on these two strings is uniform: $(1/2, 1/2)$. Then for any $\epsilon>0$, the error probability p_e of the MAP detector satisfies $(\mathrm{BNP}-\epsilon)^n \le p_e \le \mathrm{BNP}^n$ for sufficiently large n. From a mathematical perspective, the BNP corresponds to the Chernoff bound value of distinguishing two equally probable hypotheses $x=0$ and $x=1$.

By letting $h(m)=e^{-\frac{m}{2}}$ and by (4.19), we have $\mathrm{BNP}^{(l)}(0)=\mathrm{BNP}^{(l)}(1)$. The averaged $\langle\mathrm{BNP}^{(l)}\rangle$ then becomes

$$\langle\mathrm{BNP}^{(l)}\rangle := \frac{\mathrm{BNP}^{(l)}(0)+\mathrm{BNP}^{(l)}(1)}{2} = \mathrm{BNP}^{(l)}(0) = \mathrm{BNP}^{(l)}(1) = \int e^{\frac{-m}{2}}\langle P^{(l)}\rangle(dm). \qquad (4.20)$$

We state three properties which can easily be derived from the self-symmetry of $\langle P^{(l)}\rangle$. Proofs can be found in [58, 90], and in Section 5.1.4.

- $\langle\mathrm{BNP}^{(l)}\rangle = \min_s \int e^{-s\cdot m}\langle P^{(l)}\rangle(dm)$.
- The density of $e^{-m/2}\langle P^{(l)}\rangle(dm)$ is symmetric with respect to $m=0$.
- $2p_e^{(l)} \le \langle\mathrm{BNP}^{(l)}\rangle \le 2\sqrt{p_e^{(l)}(1-p_e^{(l)})}$. This justifies the use of $\langle\mathrm{BNP}^{(l)}\rangle$ as an alternative performance measure in addition to $p_e^{(l)}$.

Hereafter, we consider $\langle \mathrm{BNP}^{(l)} \rangle$ of the evolved densities $P^{(l)}(0)$ and $P^{(l)}(1)$. With the regularity assumption that $\int_{\mathbf{R}} e^{sm} \langle P^{(0)} \rangle (dm) < \infty$ for all s in some neighborhood of zero, we state the necessary and sufficient stability conditions for the irregular code ensemble $\mathcal{C}(\lambda, \rho)$ as follows.

Theorem 4.5 (Sufficient Stability Condition) *Let* $r := \langle \mathrm{BNP}^{(0)} \rangle = \int_{\mathbf{R}} e^{-m/2} \langle P^{(0)} \rangle (dm)$. *Suppose* $\lambda_2 \rho'(1) r < 1$, *and let* ϵ^* *be the smallest strictly positive root of the following equation.*

$$\lambda(1 - \rho(1 - \epsilon))r = \epsilon.$$

If for some l_0, $\langle \mathrm{BNP}^{(l_0)} \rangle < \epsilon^*$, *then*

$$\langle \mathrm{BNP}^{(l)} \rangle = \begin{cases} \mathcal{O}\left((\lambda_2 \rho'(1) r)^l \right) & \text{if } \lambda_2 > 0 \\ \mathcal{O}\left(e^{-\mathcal{O}\left((k_\lambda - 1)^l \right)} \right) & \text{if } \lambda_2 = 0, \text{ where } k_\lambda = \min\{k : \lambda_k > 0\} \end{cases},$$

and $\lim_{l \to \infty} \langle \mathrm{BNP}^{(l)} \rangle = 0$.

Corollary 4.3 *For any noise distribution* $F(dy|x)$ *with Bhattacharyya noise parameter* $r := \langle \mathrm{BNP}^{(0)} \rangle$, *if there is no* $\epsilon \in (0, r)$ *such that*

$$\lambda(1 - \rho(1 - \epsilon))r = \epsilon,$$

then $\mathcal{C}(\lambda, \rho)$ *will have arbitrarily small bit error rate as* n *tends to infinity.*

The corresponding r in Corollary 4.3 can serve as an inner bound of the achievable region for general non-symmetric memoryless channels. Further discussion of finite dimensional bounds on the achievable region will be included in Chapter 5.

Theorem 4.6 (Necessary Stability Condition) *Let* $r := \langle \mathrm{BNP}^{(0)} \rangle$. *If* $\lambda_2 \rho'(1) r > 1$, *then* $\lim_{l \to \infty} p_e^{(l)} > 0$.

- *Remark 1:* $\langle \mathrm{BNP}^{(0)} \rangle$ is the Bhattacharyya noise parameter and is related to the cutoff rate R_0 by $R_0 = 1 - \log_2(1 + \langle \mathrm{BNP}^{(0)} \rangle)$. Further discussion of $\langle \mathrm{BNP}^{(0)} \rangle$ for turbo-like and LDPC codes can be found in [53, 58] and will be carefully addressed in Section 5.1.4.

- *Remark 2:* The stability results are first stated in [90] without the convergence rate statement and the stability region ϵ^*. Since we focus on general non-symmetric channels (with symmetric channels as a special case), our convergence rate and stability region ϵ^* results also apply to the symmetric channel case. Benefitting from considering the alternative Bhattacharyya noise parameter, we will provide a simple proof, which did not appear in [90].

- *Remark 3:* ϵ^* can be used as a stopping criterion for the iterations of the density evolution. Namely, while performing the density evolution, we periodically check whether $\mathrm{BNP}^{(l)}$ is smaller than ϵ^*. If the answer is positive, we can stop the density evolution, since Theorem 4.5 guarantees the limit of $p_e^{(l)}$ goes to zero after further iterations.

Proof of Theorem 4.5*:* In Section 5.4.1, we will prove the following iterative bound for non-symmetric memoryless channels:

$$\langle \mathrm{BNP}^{(l+1)} \rangle \leq \langle \mathrm{BNP}^{(0)} \rangle \lambda \left(1 - \rho \left(1 - \langle \mathrm{BNP}^{(l)} \rangle\right)\right), \ \forall l \in \mathbb{N}. \tag{4.21}$$

The sufficient stability theorem then follows immediately from (4.21) by taking the infinitesimal analysis. ■

Proof of Theorem 4.6*:* We prove this result by the erasure decomposition technique used in [90].

The erasure decomposition lemma in [90] states that, for any $l_0 > 0$, and any symmetric channel F with log likelihood ratio distribution $P^{(l_0)}$, there exists a BEC with erasure probability ϵ, denoted by $F_{BEC,\epsilon}$, such that F is physically degraded with respect to $F_{BEC,\epsilon}$. Furthermore, F is physically degraded w.r.t. *all* $F_{BEC,\epsilon}$ with $\epsilon \leq 2p_e^{(l_0)}$. It will be shown in Appendix F.2 that this erasure decomposition lemma holds even when F corresponds to a non-symmetric channel with LLR distributions $\{P^{(l_0)}(x)\}_{x=0,1}$ and $p_e^{(l_0)}$ computed from (4.9).

We then use $B^{(l_0)}$ to denote the distribution of the log likelihood ratio when facing $F_{BEC,\epsilon}$. Namely,

$$B^{(l_0)} = \epsilon\delta_0 + (1-\epsilon)\delta_\infty, \tag{4.22}$$

where δ_x denotes a point mass at x. We then assign $B^{(l_0)}(0) := B^{(l_0)}$ and $B^{(l_0)}(1) := B^{(l_0)} \circ I^{-1}$ to distinguish the distributions for different transmitted symbols x, since their parities differ when considering different transmitted x.

Suppose $r\lambda_2\rho'(1) > 1$ and $\lim_{l\to\infty} p_e^{(l)} = 0$. Then for any $\epsilon > 0$, $\exists l_0 > 0$, such that $2p_e^{(l_0)} \leq \epsilon$. For simplicity, we assume $2p_e^{(l_0)} = \epsilon$. If during the iteration procedure (4.17), we replace the density $P^{(l_0)}(x)$ with $B^{(l_0)}(x)$ defined in (4.22), then the resulting density will be

$$\begin{aligned}
P_B^{(l_0+\Delta l)}(0) &= \epsilon \left(\lambda_2\rho'(1)\right)^{\Delta l} P^{(0)}(0) \otimes \left(\langle P^{(0)} \rangle\right)^{\otimes(\Delta l-1)} \\
&\quad + \left(1 - \epsilon \left(\lambda_2\rho'(1)\right)^{\Delta l}\right) \delta_\infty + \mathcal{O}(\epsilon^2) \\
P_B^{(l_0+\Delta l)}(1) &= \epsilon \left(\lambda_2\rho'(1)\right)^{\Delta l} P^{(0)}(1) \otimes \left(\langle P^{(0)} \rangle \circ I^{-1}\right)^{\otimes(\Delta l-1)} \\
&\quad + \left(1 - \epsilon \left(\lambda_2\rho'(1)\right)^{\Delta l}\right) \delta_{-\infty} + \mathcal{O}(\epsilon^2),
\end{aligned}$$

and the averaged error probability $p_{e,B}^{(l_0+\Delta l)}$ is

$$\begin{aligned}
p_{e,B}^{(l_0+\Delta l)} &:= \int_{-\infty}^{0} \frac{P_B^{(l_0+\Delta l)}(0) + P_B^{(l_0+\Delta l)}(1) \circ I^{-1}}{2}(dm) \\
&= \mathcal{O}(\epsilon^2) + \epsilon(\lambda_2\rho'(1))^{\Delta l} \int_{-\infty}^{0} d\left(\langle P^{(0)} \rangle\right)^{\otimes \Delta l}.
\end{aligned} \tag{4.23}$$

By the fact that $r = \langle \mathrm{BNP}^{(0)} \rangle$ is the Chernoff bound on $\int_{-\infty}^{0} d\langle P^{(0)} \rangle$, the regularity condition and the Chernoff theorem, for any $\epsilon' > 0$, there exists a large enough Δl such that

$$\int_{-\infty}^{0} d\left(\langle P^{(0)} \rangle\right)^{\otimes \Delta l} \geq (r - \epsilon')^{\Delta l}.$$

With a sufficiently small ϵ', we have $\lambda_2\rho'(1)(r-\epsilon') > 1$. Thus with large enough Δl, we have

$$p_{e,B}^{(l_0+\Delta l)} \quad > \quad \mathcal{O}(\epsilon^2) + \epsilon.$$

With small enough ϵ or equivalently large enough l_0, we have

$$p_{e,B}^{(l_0+\Delta l)} \quad > \quad \mathcal{O}(\epsilon^2) + \epsilon > \frac{\epsilon}{2} = p_e^{(l_0)}.$$

By the monotonicity with respect to physically degraded channels, we have $p_e^{(l_0+\Delta l)} \geq p_{e,B}^{(l_0+\Delta l)} > p_e^{(l_0)}$, which contradicts the monotonicity of $p_e^{(l)}$ with respect to l. Therefore, the assumption that $\lim_{l\to\infty} p_e^{(l)} = 0$ is incorrect and the proof is complete. ■

Remark: From the sufficient stability condition, for those codes with $\lambda_2 > 0$, the convergence rate of the bit error rate (BER) is exponential in l, i.e. $BER = \mathcal{O}\left((r\lambda_2\rho'(1))^l\right)$. However the number of bits involved in the $\mathcal{N}^{2l}$ tree is $\mathcal{O}\left(((d_v-1)(d_c-1))^l\right)$, which is usually much faster than the reciprocal of the decrease rate of $BER = O\left((r\lambda_2\rho'(1))^l\right)$. As a result, we conjecture that the average performance of the code ensemble with $\lambda_2 > 0$ will have bad *block* error probabilities. This is confirmed in Figure 4.7(b) and theoretically proved for the BEC in [84]. The converse is stated and proved in the following corollary.

Corollary 4.4 *Let* $\mathsf{E}\left\{Z_B^{(l)}\right\}$ *denote the block error probability of codeword length* n *after* l *iterations of the belief propagation algorithm, which is averaged over equiprobable codewords, channel realizations, and the code ensemble* $\mathcal{C}^n(\lambda,\rho)$*. If* $\lambda_2 = 0$ *and* l_n *satisfies* $\ln(\ln(n)) = o(l_n)$ *and* $l_n = o(\ln(n))$*, then*

$$\lim_{n\to\infty} \mathsf{E}\left\{Z_B^{(l_n)}\right\} = 0.$$

Proof: This result can be proven directly by the cycle-free convergence theorem, the super-exponential *bit* convergence rate with respect to l, and the union bound. ■

A similar observation is also made and proved in [53], in which it is shown that the interleaving gain exponent of the block error rate is $-J+2$, where J is the number of parallel constituent codes. On the other hand, for LDPC codes, the variable node degree d_v is the number of parity check equations (parity check sub-codes) in which a variable bit participates. In a sense, an LDPC code is similar to d_v parity check codes interleaved together. With $d_v = 2$, good interleaving gain for the block error probability is not expected.

4.5 Simulation Results

It is worth noting that for non-symmetric channels, different codewords will have different error-resisting capabilities. In this section, we consider only the codeword-averaged performance.

4.5.1 Settings

With the help of the sufficient condition of the stability theorem (Theorem 4.5), we can use ϵ^* to set a stopping criterion for the iterations of the density evolution. We use the 8-bit equally-spaced quantization of the density evolution method with $(-15, 15)$ being the

domain of the LLR messages. We will determine the largest thresholds such that the evolved Bhattacharyya noise parameter $\langle \mathrm{BNP}^{(l)} \rangle$ hits ϵ^* within 100 iterations, i.e. $\langle \mathrm{BNP}^{(100)} \rangle < \epsilon^*$. Better performance can be achieved by allowing more iterations, which, however, is of less practical interest. For example, the 500-iteration threshold of our best code for z-channels, 12B (described below), is 0.2785, which is better than the 100-iteration threshold 0.2731. Five different code ensembles with rate 1/2 are extensively simulated, including regular $(3, 6)$ codes, regular $(4, 8)$ codes, 12A codes, 12B codes, and 12C codes, where

- 12A: 12A is a rate-1/2 code ensemble found by Richardson, *et al.* in [90], which is the best known degree distribution optimized for the symmetric BiAWGNC, having maximum degree constraints $\max d_v \leq 12$ and $\max d_c \leq 9$. Its degree distributions are

$$\begin{aligned} \lambda(x) &= 0.24426x + 0.25907x^2 + 0.01054x^3 + 0.05510x^4 + 0.01455x^7 \\ &\quad + 0.01275x^9 + 0.40373x^{11}, \\ \rho(x) &= 0.25475x^6 + 0.73438x^7 + 0.01087x^8. \end{aligned}$$

- 12B: 12B is a rate-1/2 code ensemble obtained by minimizing the hitting time of ϵ^* in z-channels using our new DE formula. The search for good degree distributions is carried out by a combination of hill-climbing and differential evolution [102]. Differential evolution is a genetic algorithm presumably finding the global maximum of any given objective function ξ. In our application, ξ is the decodable noise threshold of the code ensemble with degree distributions (λ, ρ). The goal is to find the maximum achieving (λ^*, ρ^*) complying with the maximum degree constraints: $\max d_v \leq 12$ and $\max d_c \leq 9$. The differences between 12A and 12B are (1) 12B is optimized for the z-channels based on our codeword-averaged density evolution, and 12A is optimized for the symmetric BiAWGNC. (2) the noise thresholds for 12B are computed by considering the hitting time of ϵ^* while thresholds for 12A are computed by the hitting time of a fixed very small threshold. The degree distributions of 12B are

$$\begin{aligned} \lambda(x) &= 0.236809x + 0.309590x^2 + 0.032789x^3 + 0.007116x^4 + 0.000001x^5 \\ &\quad + 0.413695x^{11}, \\ \rho(x) &= 0.000015x^5 + 0.464854x^6 + 0.502485x^7 + 0.032647x^8. \end{aligned}$$

- 12C: 12C a rate-1/2 code ensemble similar to 12B, but with λ_2 being hard-wired to 0 in order to obtain good block error rate performance, which is suggested by Corollary 4.4. The degree distributions of 12C are

$$\begin{aligned} \lambda(x) &= 0.861939x^2 + 0.000818x^3 + 0.000818x^4 + 0.000818x^5 + 0.000818x^6 \\ &\quad + 0.000818x^7 + 0.000218x^8 + 0.077898x^9 + 0.055843x^{10} + 0.000013x^{11}, \\ \rho(x) &= 0.000814x^4 + 0.560594x^5 + 0.192771x^6 + 0.145207x^7 + 0.100613x^8. \end{aligned}$$

Four different types of channels are considered, including the BEC, BSC, z-channel, and BiAWGNC. Z-channels are simulated by binary non-symmetric channels with very small $p_{0\to1}$ ($p_{0\to1} = 0.00001$) and different values of $p_{1\to0}$. Table 4.1 summarizes the thresholds with precision 10^{-4}. Thresholds are not only presented by their conventional channel parameters, but also by their Bhattacharyya noise parameters (Chernoff bounds). The column "stability" lists the maximum $r := \langle \mathrm{BNP}^{(0)} \rangle$ such that $r\lambda_2\rho'(1) < 1$, which is an upper bound on the $\langle \mathrm{BNP}^{(0)} \rangle$ values of decodable channels.

Codes	BEC		BSC		Z-channels		BiAWGNC		Stability
	ϵ	$\langle$BNP$\rangle$	p	$\langle$BNP$\rangle$	$p_{1\to0}$	$\langle$BNP$\rangle$	σ	$\langle$BNP$\rangle$	$\langle$BNP$\rangle$
(3,6)	0.4294	0.4294	0.0837	0.5539	0.2305	0.4828	0.8790	0.5235	–
(4,8)	0.3834	0.3834	0.0764	0.5313	0.1997	0.4497	0.8360	0.4890	–
12A	0.4682	0.4682	0.0937	0.5828	0.2710	0.5233	0.9384	0.5668	0.6060
12B	0.4753	0.4753	0.0939	0.5834	0.2731	0.5253	0.9362	0.5653	0.6247
12C	0.4354	0.4354	0.0862	0.5613	0.2356	0.4881	0.8878	0.5303	–
Sym. Info. Rate	0.5000	0.5000	0.1100	0.6258	0.2932	0.5415	0.9787	0.5933	–
Capacity	0.5000	0.5000	0.1100	0.6258	0.3035	0.5509	0.9787	0.5933	–

Table 4.1: Thresholds of different codes on symmetric and non-symmetric channels, with precision 10^{-4}.

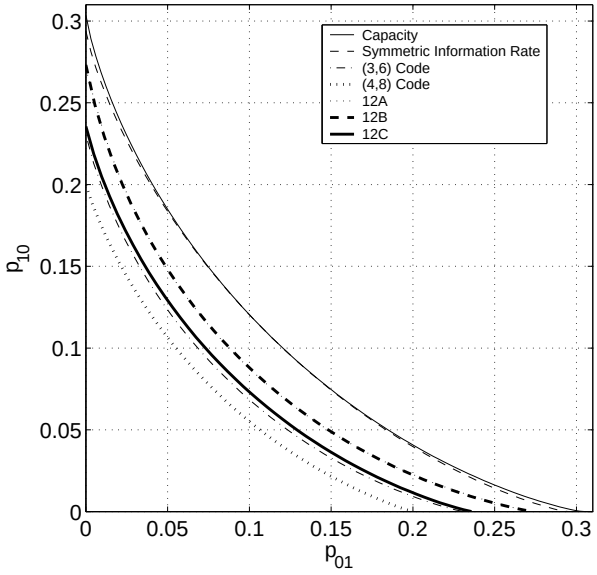

Figure 4.6: Asymptotic thresholds and the achievable regions of different codes on various binary non-symmetric channels.

4.5.2 Discussion

From Table 4.1, we observe that 12A outperforms 12B in Gaussian channels (for which 12A is optimized), but 12B is superior in z-channels for which it is optimized. The above behavior promises room for improvement with codes optimized for different channels, as was also shown in [50].

Figure 4.6 demonstrates the asymptotic thresholds of these codes in binary non-symmetric channels (BNSCs) with the curves of 12A and 12B being very close together. It is seen that 12B is slightly better when either $p_{0\to1}$ or $p_{1\to0}$ is sufficiently small, or when $p_{0\to1} \approx p_{1\to0}$. We notice that all the achievable regions of these codes are bounded by the symmetric mutual information rate (assuming uniform $(1/2, 1/2)$ *a priori* distributions), which was also suggested in [57]. The difference between the symmetric mutual information rate and the capacity for non-symmetric channels is generally indistinguishable from the practical point of view. In [79], it was shown that the ratio between the symmetric mutual information rate and the capacity is lower bounded by $\frac{e \ln 2}{2} \approx 0.942$. [99] further proved that the absolute difference is upper bounded by 0.011 bit/sym. Further discussion of capacity achieving codes with non-uniform *a priori* distributions can be found in [15] and [80].

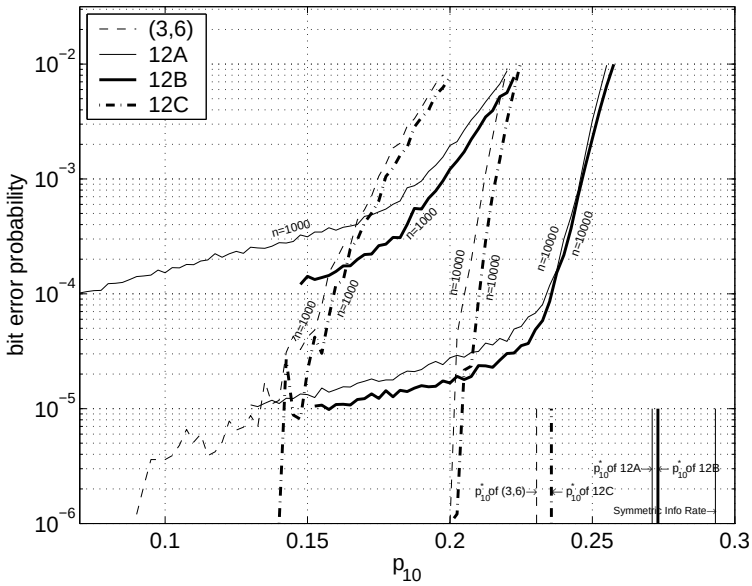

Figure 4.7: Bit error rates versus $p_{1\to0}$ with fixed $p_{0\to1} = 0.00001$. The asymptotic thresholds for symmetric mutual information rate, (3,6), 12A, 12B, and 12C codes are 0.2932, 0.2305, 0.2710, 0.2730, and 0.2356, respectively. 40 iterations of belief propagation decoding were performed. 10,000 codewords were used for the simulations.

Figures 4.7 and 4.8 consider several fixed finite codes in z-channels. We arbitrarily select graphs from the code ensemble with codeword lengths n =1,000 and n =10,000. Then, with these graphs (codes) fixed, we find the corresponding parity matrix $\mathbf{H}$, use Gaussian elimination to find the generator matrix $\mathbf{G}$, and transmit different codewords by encoding equiprobably selected information messages. Belief propagation decoding is used with 40 iterations for each codeword. 10,000 codewords are transmitted, and the overall bit/block error rates versus different $p_{1\to0}$ are plotted for different code ensembles and codeword lengths. Our new density evolution predicts the waterfall region quite accurately when the bit error rates are of primary interest. Though there are still gaps between the performance of finite codes and our asymptotic thresholds, the performance gaps between different codes are very well predicted by the differences between their asymptotic thresholds. From the above observations and the underpinning theorems in Section 4.3, we see that our new density evolution is a successful generalization of the traditional one from both practical and theoretical points of view.

Figure 4.8 exhibits the block error rate of the same 10,000-codeword simulation. The conjecture of bad block error probabilities for $\lambda_2 > 0$ codes is confirmed. Besides the conjectured bad block error probabilities, Figures 4.7 and 4.8 also suggest that codes with $\lambda_2 = 0$ will have a better error floor compared to those with $\lambda_2 > 0$, which can be partly explained by the comparatively slow convergence speed stated in the sufficient stability condition for $\lambda_2 > 0$ codes. 12C is so far the best code we have for $\lambda_2 = 0$. However, its threshold is not as good as those of 12A and 12B. If good block error rate and low error floor are of our major interest, 12C-like codes (with $\lambda_2 = 0$) serve as competitive options. Recent results in [117] shows that the error floor for codes with $\lambda_2 > 0$ can be lowered by carefully arranging the degree two variable nodes in the corresponding graph while keeping a similar waterfall threshold.

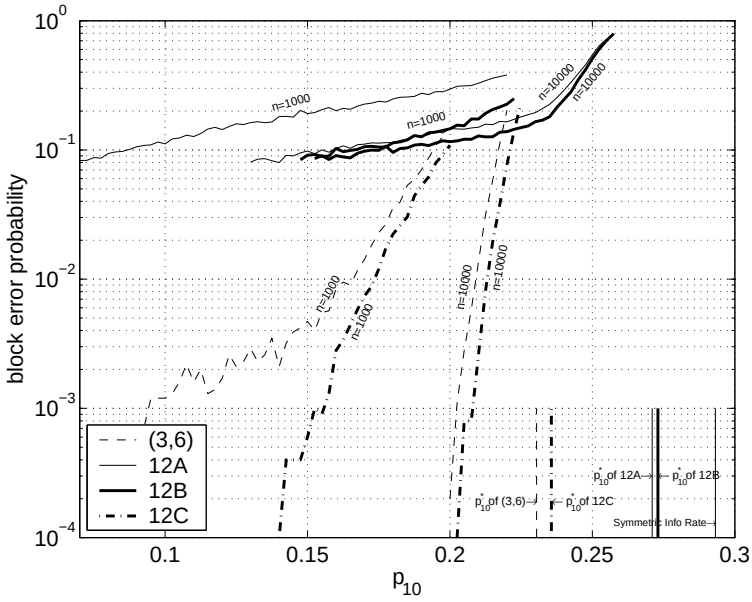

Figure 4.8: Block error rates versus $p_{1\to 0}$ with fixed $p_{0\to 1} = 0.00001$. The asymptotic thresholds for symmetric mutual information rate, (3,6), 12A, 12B, and 12C codes are 0.2932, 0.2305, 0.2710, 0.2730, and 0.2356, respectively. 40 iterations of belief propagation decoding were performed. 10,000 codewords were used for the simulations.

Figures 4.9 and 4.10 illustrate the bit error rates versus different BNSC settings with 2,000 transmitted codewords. Our computed density evolution threshold is again highly correlated with the performance of finite length codes for different non-symmetric channel settings. And it is worth noting that for codes with $\lambda_2 = 0$, namely, 12C and (3,6) codes, no error floor is visible until bit error rate 10^{-6}.

We close this section by highlighting two applications of our results.

1. Error floor analysis: "The error floor" is a characteristic of iterative decoding algorithms, which is of practical importance and may not be able to be determined solely by simulations. More analytical tools are needed to find error floors for corresponding codes. Our convergence rate statements in the sufficient stability condition may shed some light on finding codes with inherent low error floors.

2. Capacity-approaching codes for general non-standard channels: Various *very good* codes (capacity-approaching) are known for standard channels, but very good codes for non-standard channels are not yet known. It is well known that one can construct capacity-approaching codes by incorporating symmetric-information-rate-approaching linear codes with the symbol mapper and demapper as an inner code [15, 80, 89]. In general, after the inner symbol mapper/demapper, the equivalent channel becomes non-symmetric. Understanding density evolution for non-symmetric memoryless channels allows us to construct such symmetric-information-rate-approaching codes (for non-symmetric memoryless channels), and thus to find capacity-approaching codes after concatenating the inner symbol mapper and demapper. Kavčić *et al.* in [57] used the approach of the coset code ensemble to deal with one instance of non-standard channels: the intersymbol interference channel. The coset-code-based approach will be further discussed in Section 4.6.1.

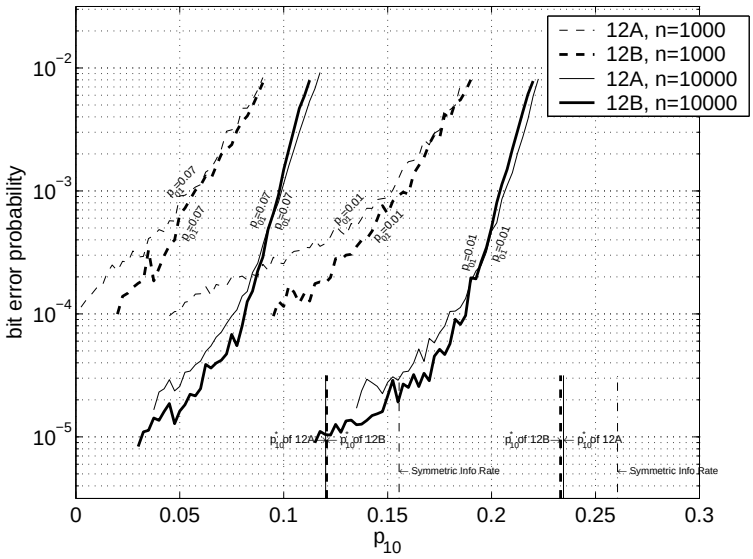

Figure 4.9: Bit error rates versus $p_{1\to0}$ with $p_{0\to1} = 0.01$ and $p_{0\to1} = 0.7$ respectively. The DE thresholds of (12A, 12B) are $(0.2346, 0.2332)$ for $p_{0\to1} = 0.01$ and $(0.1202, 0.1206)$ for $p_{0\to1} = 0.07$. 40 iterations of belief propagation decoding were performed. 2,000 codewords were used for the simulations.

4.6 Further Implications of the Generalized Density Evolution

4.6.1 Typicality of Linear LDPC Codes

One reason that non-symmetric channels are often overlooked is that we can always transform a non-symmetric channel into a symmetric channel. Depending on different points of view, this channel-symmetrizing technique is termed the coset code argument [57] or dithering/the i.i.d. channel adapter [51], as illustrated in Figures 4.11(b) and 4.11(c). Further explanations of this channel symmetrizing technique will be given in Section 5.1.1.

To be more explicit, an LDPC coset code ensemble contains a set of codes satisfying $\mathbf{Hx} = \mathbf{s}$, where $\mathbf{H}$ is obtained from the same equiprobable bipartite graph ensemble as for linear LDPC codes. The only difference is, for linear code ensembles, the coset-defining syndrome $\mathbf{s}$ is hardwired to $\mathbf{0}$, while for coset code ensembles, $\mathbf{s}$ is uniformly randomly chosen from $\{0,1\}^{n(1-R)}$. Our generalized density evolution provides a simple way to directly analyze the linear LDPC code ensemble on non-symmetric channels, as in Figure 4.11(a), instead of using a larger code ensemble to symmetrize the channel.

As shown in Theorems 4.5 and 4.6, the necessary and sufficient stability conditions of linear LDPC codes for non-symmetric channels, Figure 4.11(a), are identical to those of the coset code ensemble, Figure 4.11(c). Monte Carlo simulations based on finite-length codes ($n = 10^4$) [51] further demonstrate that the codeword-averaged performance in Figure 4.11(a) is nearly identical[4] to the performance of Figure 4.11(c) when the same linear encoder/decoder pair is used. The above two facts suggest a close relationship between linear codes and the coset code ensemble, and it was conjectured in [51] that the scheme in Figure 4.11(a) should always have the same/similar performance as those illustrated by

[4]That is, it is within the precision of the Monte Carlo simulation.

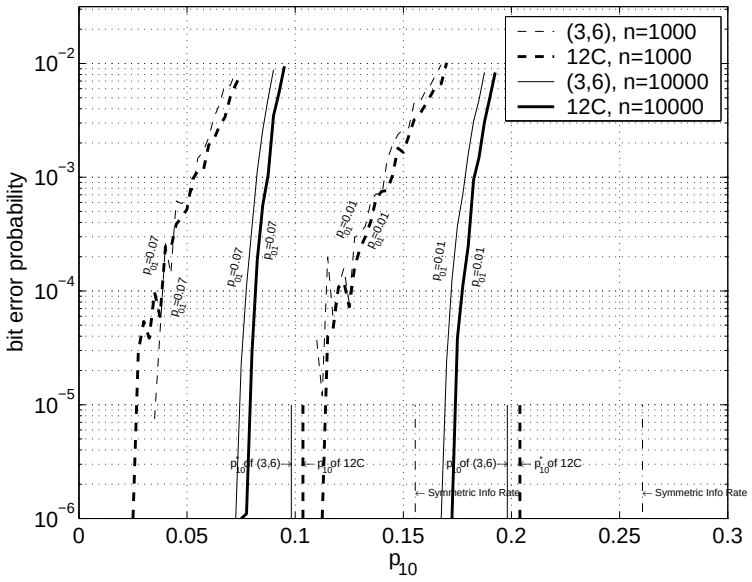

Figure 4.10: Bit error rates versus $p_{1\to 0}$ with $p_{0\to 1} = 0.01$ and $p_{0\to 1} = 0.7$ respectively. The DE thresholds of (12C, (3,6)) are $(0.2039, 0.1981)$ for $p_{0\to 1} = 0.01$ and $(0.1036, 0.0982)$ for $p_{0\to 1} = 0.07$. 40 iterations of belief propagation decoding were performed. 2,000 codewords were used for the simulations.

Figure 4.11(c). This short subsection is devoted to this conjecture. In sum, the performance of the linear code ensemble is very unlikely to be identical to that of the coset code ensemble. However, when the minimum $d_{c,min} := \{k \in \mathbb{N} : \rho_k > 0\}$ is sufficiently large, we can prove that their performance discrepancy is theoretically indistinguishable. In practice, the discrepancy for $d_{c,min} \geq 6$ is extremely small: $< 0.05\%$.

Let $P^{(l)}_{a.p.}(0) := P^{(l)}(0)$ and $P^{(l)}_{a.p.}(1) := P^{(l)}(1) \circ I^{-1}$ denote the two evolved densities with *aligned parity*, and similarly define $Q^{(l)}_{a.p.}(0) := Q^{(l)}(0)$ and $Q^{(l)}_{a.p.}(1) := Q^{(l)}(1) \circ I^{-1}$. Our main result in (4.17) can be rewritten in the following form:

$$\begin{aligned} P^{(l)}_{a.p.}(x) &= P^{(0)}_{a.p.}(x) \otimes \lambda\left(Q^{(l-1)}_{a.p.}(x)\right) \\ Q^{(l-1)}_{a.p.}(x) &= \Gamma^{-1}\left(\rho\left(\Gamma\left(\frac{P^{(l-1)}_{a.p.}(0) + P^{(l-1)}_{a.p.}(1)}{2}\right)\right)\right. \\ &\qquad \left. +(-1)^x \rho\left(\Gamma\left(\frac{P^{(l-1)}_{a.p.}(0) - P^{(l-1)}_{a.p.}(1)}{2}\right)\right)\right). \end{aligned} \tag{4.24}$$

Let $p^{(l)}_{e,linear}$ denote the corresponding bit error probability of the linear codes after l iterations. For comparison, the traditional formula of density evolution for the symmetrized channel (the coset code ensemble) is as follows:

$$\begin{aligned} P^{(l)}_{coset} &= P^{(0)}_{coset} \otimes \lambda\left(Q^{(l-1)}_{coset}\right) \\ Q^{(l-1)}_{coset} &= \Gamma^{-1}\left(\rho\left(\Gamma\left(P^{(l-1)}_{coset}\right)\right)\right), \end{aligned} \tag{4.25}$$

where $P^{(0)}_{coset} = \frac{\sum_{x=0,1} P^{(0)}_{a.p.}(x)}{2}$. Similarly, let $p^{(l)}_{e,coset}$ denote the corresponding bit error probability.

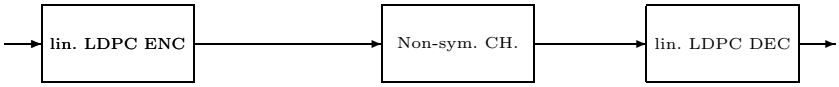

(a) Linear Code Ensemble versus Non-symmetric Channels

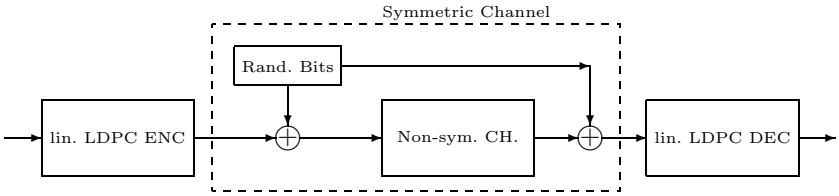

(b) Linear Code Ensemble versus Symmetrized Channels

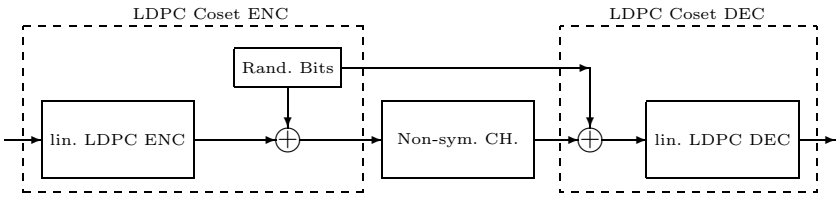

(c) Coset Code Ensemble versus Non-symmetric Channels

Figure 4.11: Comparison of the approaches based on codeword averaging and the coset code ensemble.

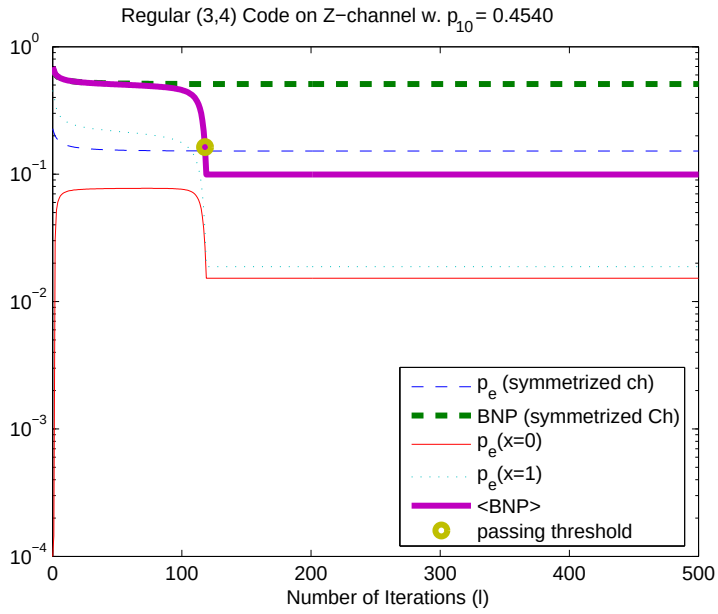

Figure 4.12: Density evolution for z-channels with the linear and the coset code ensembles.

Table 4.2: Threshold ($p^*_{1\to 0}$) comparison between linear and coset LDPC codes on Z-channels

(λ,ρ)	(x^2,x^3)	(x^2,x^5)	$(x^2,0.5x^2+0.5x^3)$	$(x^2,0.5x^4+0.5x^5)$
Linear	0.4540	0.2305	0.5888	0.2689
	↕ 0.29%	↕ 0.043%	↕ 0.17%	↕ 0.037%
Coset	0.4527	0.2304	0.5888	0.2690

It is clear from the above formulae that when the channel of interest is symmetric, namely $P^{(0)}_{a.p.}(0) = P^{(0)}_{a.p.}(1)$, then $P^{(l)}_{coset} = P^{(l)}_{a.p.}(0) = P^{(l)}_{a.p.}(1)$ for all $l \in \mathbb{N}$. However, for non-symmetric channels, since the variable node iteration involves convolution of several densities given the same x value, the difference between $Q^{(l-1)}_{a.p.}(0)$ and $Q^{(l-1)}_{a.p.}(1)$ will be amplified after each variable node iteration. Hence it is very unlikely that the decodable thresholds of linear codes and coset codes will be analytically identical, namely

$$\lim_{l\to\infty} p^{(l)}_{e,linear} = 0 \overset{?}{\Longleftrightarrow} \lim_{l\to\infty} p^{(l)}_{e,coset} = 0.$$

Figure 4.12 demonstrates the traces of the evolved densities for the regular (3,4) code on z-channels. With the one-way crossover probability being 0.4540, the generalized density evolution for linear codes is able to converge within 179 iterations, while the coset code ensemble shows no convergence within 500 iterations. This demonstrates the possible performance discrepancy, though we do not have analytical results proving that the latter will not converge after further iterations. Table 4.2 compares the decodable thresholds such that the density evolution enters the stability region within 100 iterations. Using our codeword-averaged density evolution, we are able to pinpoint the asymptotic thresholds for both linear and coset code ensembles, which is especially important since their performance discrepancy is within the precision of Monte Carlo simulations. Furthermore, we notice a new phenomenon that the larger $d_{c,min}$ is, the smaller the discrepancy is. This phenomenon can be characterized by the following theorem.

Theorem 4.7 *Consider non-symmetric memoryless channels and a fixed pair of finite-degree polynomials λ and ρ. The shifted version of the check node polynomial is denoted as $\rho_\Delta = x^\Delta \cdot \rho$ where $\Delta \in \mathbb{N}$. Let $P^{(l)}_{coset}$ denote the evolved density from the coset code ensemble with degrees (λ,ρ_Δ), and $\langle P^{(l)}\rangle = \frac{1}{2}\sum_{x=0,1} P^{(l)}_{a.p.}(x)$ denote the averaged density from the linear code ensemble with degrees (λ,ρ_Δ). For any $l_0 \in \mathbb{N}$, $\lim_{\Delta\to\infty}\langle P^{(l)}\rangle \overset{\mathcal{D}}{=} P^{(l)}_{coset}$ in distribution for all $l \le l_0$, with the convergence rate for each iteration being $\mathcal{O}\left(\mathsf{const}^\Delta\right)$ for some $\mathsf{const} < 1$.*

To further interpret these results regarding to convergence in distribution, a corollary in terms of decodable thresholds is stated as follows.

Corollary 4.5 (The Typicality Results for Z-Channels) *For any $\epsilon > 0$, there exists a $\Delta \in \mathbb{N}$ such that*

$$\left|\sup\left\{p_{1\to 0} : \lim_{l\to\infty} p^{(l)}_{e,linear} = 0\right\} - \sup\left\{p_{1\to 0} : \lim_{l\to\infty} p^{(l)}_{e,coset} = 0\right\}\right| < \epsilon.$$

Namely, the asymptotic decodable thresholds of the linear and the coset code ensembles are arbitrarily close when the minimum check node degree $d_{c,min}$ is sufficiently large.

Similar corollaries can be constructed for all other non-symmetric channels with different types of noise parameters, e.g., the composite BiAWGNC in Section 4.1.1. A proof of Corollary 4.5 is found in Appendix E.3.

Proof of Theorem 4.7*:* Since the functionals in (4.24) and (4.25) are continuous with respect to convergence in distribution, we need only to show that $\forall l \in \mathbb{N}$,

$$\begin{aligned}\lim_{\Delta\to\infty} Q_{a.p.}^{(l-1)}(0) &\overset{\mathcal{D}}{=} \lim_{\Delta\to\infty} Q_{a.p.}^{(l-1)}(1) \\ &\overset{\mathcal{D}}{=} \Gamma^{-1}\left(\rho\left(\Gamma\left(\frac{P_{a.p.}^{(l-1)}(0)+P_{a.p.}^{(l-1)}(1)}{2}\right)\right)\right) \\ &= \frac{Q_{a.p.}^{(l-1)}(0)+Q_{a.p.}^{(l-1)}(1)}{2}, \end{aligned} \tag{4.26}$$

where $\overset{\mathcal{D}}{=}$ denotes convergence in distribution. Then by inductively applying this weak convergence argument, for any bounded l_0, $\lim_{\Delta\to\infty}\langle P^{(l)}\rangle \overset{\mathcal{D}}{=} P_{coset}^{(l)}$ in distribution for all $l \le l_0$. Without loss of generality,[5] we may assume $\rho_\Delta = x^\Delta$ and prove the weak convergence of distributions on the domain

$$\gamma(m) := \left(1_{\{m\le 0\}}, \ln\coth\left|\frac{m}{2}\right|\right) = (\gamma_1,\gamma_2) \in \mathsf{GF}(2)\times\mathbb{R}^+,$$

on which the check node iteration becomes

$$\gamma_{out,\Delta} = \gamma_{in,1}+\gamma_{in,2}+\cdots+\gamma_{in,\Delta},$$

where $\gamma_{in,i}$, $i = 1,\cdots,\Delta$, represent the Δ incoming messages of the check node of interest. Let P_0' denote the density of $\gamma_{in}(m)$ given that the distribution of m is $P_{a.p.}^{(l-1)}(0)$ and let P_1' similarly correspond to $P_{a.p.}^{(l-1)}(1)$. Similarly let $Q_{0,\Delta}'$ and $Q_{1,\Delta}'$ denote the output distributions on $\gamma_{out,\Delta}$ when the check node degree is $\Delta+1$. It is worth noting that any pair of $Q_{0,\Delta}'$ and $Q_{1,\Delta}'$ can be mapped bijectively to the LLR distributions $Q_{a.p.}^{(l-1)}(0)$ and $Q_{a.p.}^{(l-1)}(1)$.

Let $\Phi_{P'}(k,r) := \mathsf{E}_{P'}\left\{(-1)^{k\gamma_1}e^{ir\gamma_2}\right\}, \forall k\in\mathbb{N}, r\in\mathbb{R}$, denote the Fourier transform of the density P'. Proving (4.26) is equivalent to showing that

$$\forall k\in\mathbb{N}, r\in\mathbb{R},\ \lim_{\Delta\to\infty}\Phi_{Q_{0,\Delta}'}(k,r) = \lim_{\Delta\to\infty}\Phi_{Q_{1,\Delta}'}(k,r).$$

However, to deal with the strictly growing average of the "limit distribution", we concentrate on the distribution of the normalized output $\frac{\gamma_{out,\Delta}}{\Delta}$ instead. We then need to prove that

$$\forall k\in\mathbb{N}, r\in\mathbb{R},\ \lim_{\Delta\to\infty}\Phi_{Q_{0,\Delta}'}(k,\frac{r}{\Delta}) = \lim_{\Delta\to\infty}\Phi_{Q_{1,\Delta}'}(k,\frac{r}{\Delta}).$$

We first note that for all $x = 0,1$, $Q_{x,\Delta}'$ is the averaged distribution of $\gamma_{out,\Delta}$ when the inputs $\gamma_{in,i}$ are governed by $P_{a.p.}^{(l)}(x_i)$ satisfying $\sum_{i=1}^{\Delta} x_i = x$. From this observation, we can derive the following iterative equations: $\forall\Delta\in\mathbb{N}$,

$$\begin{aligned}\Phi_{Q_{0,\Delta}'}(k,\frac{r}{\Delta}) &= \frac{\Phi_{Q_{0,\Delta-1}'}(k,\frac{r}{\Delta})\Phi_{P_0'}(k,\frac{r}{\Delta})+\Phi_{Q_{1,\Delta-1}'}(k,\frac{r}{\Delta})\Phi_{P_1'}(k,\frac{r}{\Delta})}{2} \\ \Phi_{Q_{1,\Delta}'}(k,\frac{r}{\Delta}) &= \frac{\Phi_{Q_{0,\Delta-1}'}(k,\frac{r}{\Delta})\Phi_{P_1'}(k,\frac{r}{\Delta})+\Phi_{Q_{1,\Delta-1}'}(k,\frac{r}{\Delta})\Phi_{P_0'}(k,\frac{r}{\Delta})}{2}.\end{aligned}$$

[5] We also need to assume that $\forall x, P_{a.p.}^{(l-1)}(x)(m=0) = 0$ so that $\ln\coth\left|\frac{m}{2}\right| \in \mathbb{R}^+$ almost surely. This assumption can be relaxed by separately considering the event that one of the input message is zero.

By induction, the difference thus becomes

$$\begin{aligned}
&\Phi_{Q'_{0,\Delta}}(k,\frac{r}{\Delta})-\Phi_{Q'_{1,\Delta}}(k,\frac{r}{\Delta})\\
&= \left(\Phi_{Q'_{0,\Delta-1}}(k,\frac{r}{\Delta})-\Phi_{Q'_{1,\Delta-1}}(k,\frac{r}{\Delta})\right)\left(\frac{\Phi_{P'_0}(k,\frac{r}{\Delta})-\Phi_{P'_1}(k,\frac{r}{\Delta})}{2}\right)\\
&= 2\left(\frac{\Phi_{P'_0}(k,\frac{r}{\Delta})-\Phi_{P'_1}(k,\frac{r}{\Delta})}{2}\right)^{\Delta}. \qquad (4.27)
\end{aligned}$$

By Taylor's expansion and the BNSC decomposition argument, which will be introduced in Section 5.1.2, we can show that for all $k \in \mathbb{N}$, $r \in \mathbb{R}$, and for all possible P'_0 and P'_1, the quantity in (4.27) converges to zero with convergence rate $\mathcal{O}\left(\mathsf{const}^{\Delta}\right)$ for some $\mathsf{const} < 1$. A detailed derivation of the convergence rate is given in Appendix E.4. Since the point-wise limit of the right-hand side of (4.27) is zero, the proof of weak convergence is complete. The exponentially fast convergence rate $\mathcal{O}\left(\mathsf{const}^{\Delta}\right)$ also justifies the fact that even for moderate $d_{c,min} \geq 6$, the performances of linear and coset LDPC codes are very close ($< 0.05\%$). ■

Remark 1: Consider any non-perfect message distribution, namely, $\exists x_0$ such that

$$P_{a.p.}^{(l-1)}(x_0) \neq \delta_{\infty}.$$

A persistent reader may notice that $\forall x, \lim_{\Delta\to\infty} Q_{a.p.}^{(l-1)}(x) \overset{\mathcal{D}}{=} \delta_0$, namely, as Δ becomes large, all information is erased after passing a check node of large degree. If this convergence (erasure effect) occurs earlier than the convergence of $Q_{a.p.}^{(l-1)}(0)$ and $Q_{a.p.}^{(l-1)}(1)$, the performances of linear and coset LDPC codes are "close" only when the code is "useless."[6] To quantify the convergence rate, we consider again the distributions on γ and their Fourier transforms. For the average of the output distributions $Q_{a.p.}^{(l-1)}(x)$, we have

$$\begin{aligned}
&\frac{\Phi_{Q'_{0,\Delta}}(k,\frac{r}{\Delta})+\Phi_{Q'_{1,\Delta}}(k,\frac{r}{\Delta})}{2}\\
&= \left(\frac{\Phi_{Q'_{0,\Delta-1}}(k,\frac{r}{\Delta})+\Phi_{Q'_{1,\Delta-1}}(k,\frac{r}{\Delta})}{2}\right)\left(\frac{\Phi_{P'_0}(k,\frac{r}{\Delta})+\Phi_{P'_1}(k,\frac{r}{\Delta})}{2}\right)\\
&= \left(\frac{\Phi_{P'_0}(k,\frac{r}{\Delta})+\Phi_{P'_1}(k,\frac{r}{\Delta})}{2}\right)^{\Delta}. \qquad (4.28)
\end{aligned}$$

By Taylor's expansion and the BNSC decomposition argument, one can show that the limit of (4.28) exists and the convergence rate is $\mathcal{O}(\Delta^{-1})$. (A detailed derivation is included in Appendix E.4.) This convergence rate is much slower than the exponential rate $\mathcal{O}\left(\mathsf{const}^{\Delta}\right)$ of (4.27). Therefore, we do not need to worry about the case in which the required Δ for the convergence of $Q_{a.p.}^{(l-1)}(0)$ and $Q_{a.p.}^{(l-1)}(1)$ is excessively large so that $\forall x \in \mathsf{GF}(2), Q_{a.p.}^{(l-1)}(x) \overset{\mathcal{D}}{\approx} \delta_0$.

Remark 2: The intuition behind Theorem 4.7 is that when the minimum d_c is sufficiently large, the parity check constraint (on valid bit strings) becomes less stringent. Thus we can approximate the density of the outgoing messages for linear codes by assuming all bits involved in that particular parity check equation are "independently" distributed among

[6] To be more precise, it corresponds to an extremely high-rate code and the information is erased after every check node iteration.

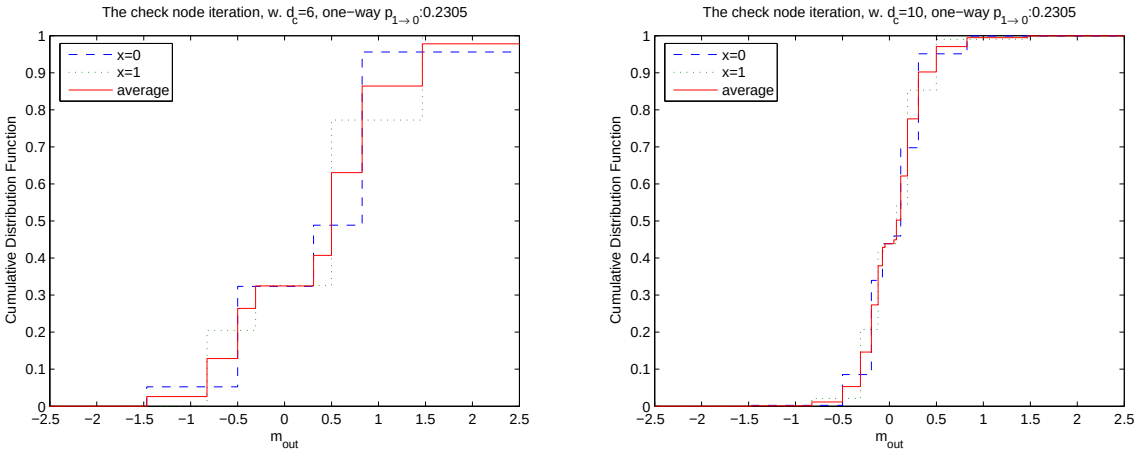

Figure 4.13: Illustration of the weak convergence of $Q_{a.p.}^{(l-1)}(0)$ and $Q_{a.p.}^{(l-1)}(1)$. One can see that the convergence of $Q_{a.p.}^{(l-1)}(0)$ and $Q_{a.p.}^{(l-1)}(1)$ is faster than the convergence of $\frac{Q_{a.p.}^{(l-1)}(0)+Q_{a.p.}^{(l-1)}(1)}{2}$ and δ_0.

$\{0,1\}$, which leads to the formula for the coset code ensemble. On the other hand, extremely large d_c is required for a check node iteration to *completely* destroy all information coming from the previous iteration. This explains the difference between their convergence rates: $\mathcal{O}\left(\mathsf{const}^{\Delta}\right)$ versus $\mathcal{O}(\Delta^{-1})$.

Figure 4.13 illustrates the weak convergence predicted by Theorem 4.7 and depicts the convergence rates of $Q_{a.p.}^{(l-1)}(0) \longrightarrow Q_{a.p.}^{(l-1)}(1)$ and $\frac{Q_{a.p.}^{(l-1)}(0)+Q_{a.p.}^{(l-1)}(1)}{2} \longrightarrow \delta_0$.

Our typicality result can be viewed as a complementing theorem of the concentration theorem in [Corollary 2.2 of [57]], where a constructive method of finding a typical coset-defining syndrome was not specified. Besides the theoretical importance, we are now on a solid basis to interchangeably use the linear LDPC codes and the LDPC coset codes when the check node degree is of moderate size. For instance, from the implementation point of view, the hardware uniformity of linear codes makes them a superior choice compared to any other coset code. We can then use the fast density evolution [54] plus the coset code ensemble to optimize the degree distribution for the linear LDPC codes. Or instead of simulating the codeword-averaged performance of linear LDPC codes, we can simulate the error probability of the all-zero codeword in the coset code ensemble, for which the efficient LDPC encoder [93] is not necessary. Our results also show that there is not much room for improvement by selecting the "best" coset-defining syndrome $\mathbf{s}$ since all $\mathbf{s}$ will lead to similar performance when $d_{c,min}$ is of moderate size.

4.6.2 Local Optimality of the Belief Propagation Decoder

Two known facts about the BP algorithm and the density evolution method are as follows. First, the BP algorithm is optimal for any cycle-free network, since it exploits the independence of the incoming LLR message. Second, by the cycle-free convergence theorem, the traditional density evolution is able to predict the behavior of the BP algorithm (designed for the tree structure) for l_0 iterations, even when we are focusing on a Tanner graph of a finite-length LDPC code, which inevitably has many cycles. The performance of BP, predicted by density evolution, is outstanding so that we "implicitly assume" that the BP

(designed for the tree structure) is optimal for the first l_0 iterations in terms of minimizing the *codeword-averaged* BER. Theoretically, to be able to minimize the codeword-averaged BER, the optimal decision rule inevitably must exploit the global knowledge about all possible codewords, which is, however, not available to the BP decoder. A question of interest is whether BP is indeed optimal for the first l_0 iterations. Namely, with only local knowledge about possible codewords, does BP have the same performance as the optimal detector with the global information about the entire codebook and unlimited computational power when we are only interested in the first l_0 iterations? The answer is a straightforward corollary to Theorem 4.2, the convergence to perfect projection, which provides the missing link regarding the optimality of BP when only local observations (on the $\mathcal{N}^{2l}$) are available.

Theorem 4.8 (Local Optimality of the BP Decoder) *Fix $i, l_0 \in \mathbb{N}$. For sufficiently large codeword length n, almost all instances in the random code ensemble have the property that the BP decoder for x_i after l_0 iterations, $\hat{X}_{BP}(\mathbf{Y}^{l_0})$, coincides with the optimal MAP bit detector $\hat{X}_{MAP,l_0}(\mathbf{Y}^{l_0})$, where l_0 is a fixed integer. The MAP bit detector $\hat{X}_{MAP,l_0}(\cdot)$ uses the same number of observations as in $\hat{X}_{BP}(\cdot)$ but is able to exploit the global knowledge about the entire codebook.*

Proof: When the support tree $\mathcal{N}^{2l_0}_{(i,j)}$ is perfectly projected, the local information about the tree-satisfying strings is equivalent to the global information about the entire codebook. Therefore, the extra information about the entire codebook does not benefit the decision maker, and $\hat{X}_{BP}(\cdot) = \hat{X}_{MAP,l_0}(\cdot)$. Theorem 4.2 shows that $\mathcal{N}^{2l_0}_{(i,j)}$ converges to perfect projection in probability, which in turn implies that for sufficiently large n, BP decoder is locally optimal for almost all instances of the code ensemble. ■

Note: Even when limiting ourselves to symmetric memoryless channels, this local optimality of BP can be proved[7] only by the convergence to perfect projection. Theorem 4.8 can thus be viewed as a completion of the classical density evolution for symmetric memoryless channels.

4.7 Summary

In this chapter, we have developed a codeword-averaged density evolution, which allows analysis of general *non-symmetric* memoryless channels. An essential perfect projection convergence theorem has been proved by a constraint propagation argument and by analyzing the behavior of random matrices. With this perfect projection convergence theorem, the theoretical foundation of the codeword-averaged density evolution is well established. Most of the properties of symmetric density evolution have been generalized and proved for the codeword-averaged density evolution on non-symmetric channels, including monotonicity, distribution symmetry, and stability. Besides a necessary stability condition, a sufficient stability condition has been stated with convergence rate arguments and a simple proof.

The typicality of linear codes among the LDPC coset code ensemble has been proved by the weak convergence (w.r.t. d_c) of the evolved densities in our codeword-averaged density evolution. Namely, when the check node degree is sufficiently large (e.g. $d_c \geq 6$), the performance of the linear LDPC code ensemble is very close to (e.g. within 0.05%) the performance of the LDPC coset code ensemble. In addition to a new iterative DE formula,

[7] The existing cycle-free convergence theorem along does not guarantee the local optimality of BP. One can easily construct a cycle-free $\mathcal{N}^{2l}$ that is not perfectly projected and on which the BP is not optimal.

another important corollary to the perfect projection convergence theorem is the optimality of the belief propagation algorithms when the global information about the entire codebook is accessible. This can be viewed as a completion of the theory of classical density evolution for symmetric memoryless channels.

Extensive simulations have been presented, the degree distribution has been optimized for z-channels, and possible applications of our results have been discussed as well. From both practical and theoretical points of view, this codeword-averaged density evolution offers a straightforward and successful generalization of the traditional symmetric density evolution for general non-symmetric memoryless channels.

Chapter 5

Finite-Dimensional Bounds on LDPC Codes with Belief Propagation Decoding

The benefit of a finite dimensional upper/lower bound on the decodable threshold is two-fold. It can be used as an alternative development tool with superior computational efficiency, while it can also serve as theoretical basis for closed form analysis. This chapter focuses on finite-dimensional bounds for $\mathbb{Z}_m$ and binary low-density parity-check (LDPC) codes, assuming belief propagation (BP) decoding on memoryless channels. Two noise measures will be considered: the Bhattacharyya noise parameter (BNP) and the expected soft bit (ESB) value, the latter of which is obtained from a maximum *a posteriori* probability (MAP) decoder on the uncoded channel. For $\mathbb{Z}_m$ LDPC codes, an iterative m-dimensional bound is derived for m-ary-input/symmetric-output (MI-SO) channels, which gives a sufficient stability condition for $\mathbb{Z}_m$ LDPC codes and will be complemented by a matched necessary stability condition introduced herein. Applications to coded modulation and to codes with non-equiprobable distributed codewords will also be discussed.

For binary codes, two new lower bounds are provided for binary-input/symmetric-output (BI-SO) channels, including a two-dimensional iterative bound and a one-dimensional non-iterative bound, the latter of which is the best known bound that is tight for binary symmetric channels (BSCs). By adopting the reverse channel perspective, upper and lower bounds based on BNP are derived for binary-input/non-symmetric-output (BI-NSO) channels, which coincide with the existing bounds for BI-SO channels.

This chapter is organized as follows. The necessary definitions and background knowledge will be provided in Section 5.1, including the definitions of the symmetric channels and the noise measures of interest. Section 5.2 will provide the framework for the iterative bounding problem and review some existing results. A BNP-based bound and a pair of stability conditions will be provided for $\mathbb{Z}_m$ LDPC codes in Section 5.3, which will lead to the application-oriented discussion in Chapter 6 on $\mathbb{Z}_m$ LDPC coded modulation. For binary LDPC codes, Sections 5.4 and 5.5 are devoted to the iterative and non-iterative bounds respectively, the former of which include a one-dimensional bound for BI-NSO channels and a two-dimensional bound for BI-SO channels, while the latter of which provides the best (tightest) known bound for binary symmetric channels (BSCs). Performance comparisons are provided in Section 5.6. Section 5.7 summarizes this chapter.

5.1 Formulation

As in Chapter 4, we consider only memoryless channels with discrete input alphabets.

5.1.1 Symmetric Channels

A BI-SO $\mathbf{X} \mapsto \mathbf{Y}$ channel is conventionally defined as a channel with binary[1] input set $\mathbf{X} = \{0, 1\}$ and real output set $\mathbf{Y} = \mathbb{R}$, such that $\mathsf{P}(Y \in dy|X = 0) = \mathsf{P}(Y \in -dy|X = 1)$ where X and Y denote the channel input and output, respectively. In the literature of LDPC codes (e.g., [92]), an equivalent definition is that the BI-SO channel satisfies $P(dm) = e^m P(-dm)$, where $P(\cdot)$ is the density of the log likelihood ration (LLR) messages, $m := \log \frac{\mathsf{P}(Y|X=0)}{\mathsf{P}(Y|X=1)}$, given $X = 0$.

Let $\mathbb{Z}_m := \{0, 1, \cdots, m-1\}$ denote the integer ring modulo m. A more general, theoretical definition for m-ary-input/symmetric-output (MI-SO) channels is given as follows.

Definition 5.1 (MI-SO Channels) *An m-ary-input channel $\mathbb{Z}_m \mapsto \mathbf{Y}$ is (circularly) symmetric if there exists a bijective transform $\mathcal{T} : \mathbf{Y} \mapsto \mathbf{Y}$ such that $\forall y \in \mathbf{Y}, \mathcal{T}^m(y) = (y)$ and*

$$\forall x \in \mathbb{Z}_m, F(dy|0) = F(\mathcal{T}^x(dy)|x),$$

where $F(dy|x)$ is the conditional distribution of $Y \in dy$ given $X = x$. When $m = 2$, this definition coincides with that of the conventional BI-SO channel.

Remark 1: In this definition, m can be any strictly positive integer. Since when $m = 1$ the channel carries no information, we assume $m \geq 2$. When $m = 2$, Definition 5.1 becomes a formal definition of BI-SO channels.

Remark 2: This definition characterizes the essence of channel symmetry and is much broader than the conventional definition of BI-SO channels, since there is no constraint on $\mathbf{Y}$, the range of the channel output. For example, in phase-shift keying (PSK) or quadrature amplitude modulation (QAM) scenarios, $\mathbf{Y} = \mathbb{R}^2$.

Remark 3: There are many different levels of "channel symmetry." An incomplete list contains strongly and weakly symmetric channels [33], symmetric channels in [44], and the circularly symmetric channels introduced herein. The former three definitions were derived from an information theoretic point of view, while the circularly symmetric channel is derived from a coding perspective. This definition of "circularly symmetric channels" coincides with the definition of the "matched signal set" introduced in [74].

Since the BP decoding algorithm on LDPC codes is also circularly symmetric, one immediate benefit of considering a symmetric channel with LDPC codes is that all codewords have the same error resiliency. Therefore we can use the all-zero codeword as a representative, which facilitates the simulation of codes with finite length. Further discussion can be found in [74, 92] and in Section 4.2.

Another advantage of Definition 5.1 is that we can immediately prove the channel symmetrizing argument as first mentioned in Section 4.6.1. That is, the channels in Figures 5.1(a) and 5.1(b) have the same channel capacity while the new channel $X \mapsto (W, Y)$ in Figure 5.1(b) is circularly symmetric by Definition 5.1. Using this channel symmetrizing

[1] Another common setting is to consider $\mathbf{X} = \{+1, -1\}$, which reflects coherent binary phase shift keying (BPSK) modulation. However, to be compatible with the algebra on which the parity check equations are defined, we assume $\mathbf{X} = \{0, 1\}$ instead of $\{+1, -1\}$.

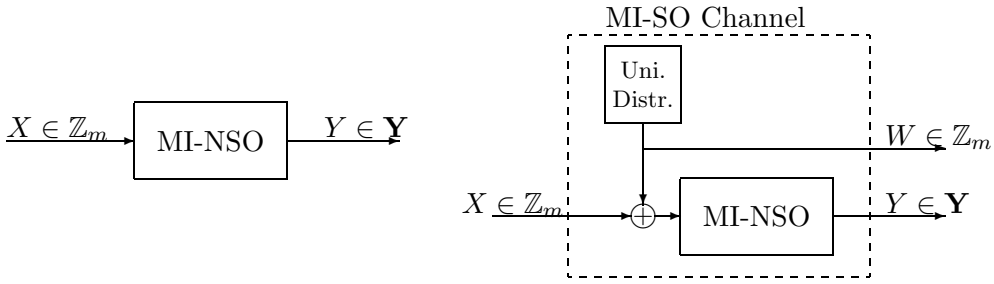

(a) An MI-NSO channel (b) A symmetrized MI-NSO channel.

Figure 5.1: Channel symmetrization.

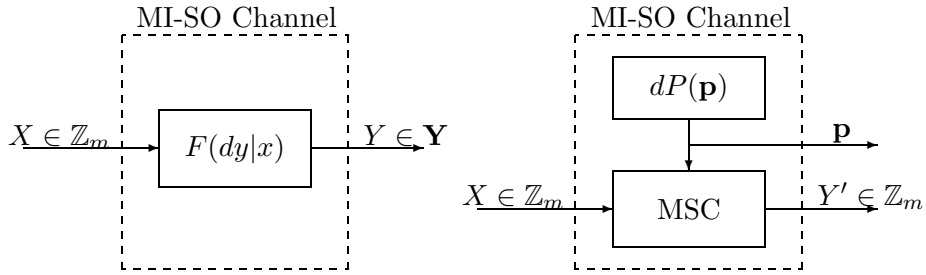

(a) By conditional distributions (b) By probabilistic combinations

Figure 5.2: Different methods of representation for the m-ary-input/symmetric-output channels.

argument, we can assume all channels are circularly symmetric as long as the additional complexity of channel symmetrizing[2] is bearable.

5.1.2 MSC Decomposition

One of the simplest MI-SO channels is the m-ary symmetric channel (MSC), which is a $\mathbb{Z}_m \mapsto \mathbb{Z}_m$ channel and can be fully specified by a parameter vector $\mathbf{p} = (p_0, p_1, \cdots, p_{m-1})$ such that the conditional probability $\mathsf{P}(Y = x + i | X = x) = p_i$, $\forall x, i \in \mathbb{Z}_m$. By partitioning the set of channel outputs, $\mathbf{Y}$, into disjoint subsets of m elements[3] according to the bijective transform $\mathcal{T}$ in Definition 5.1, it can be shown that any MI-SO channel can be *uniquely* expressed as a probabilistic combination of different MSCs, while the parameter $\mathbf{p}$ of each MSC is observed by the receiver as side information. This decomposition is illustrated in Figures 5.2(a) and 5.2(b), in which the probabilistic weight of different vectors $\mathbf{p}$ is denoted by $dP(\mathbf{p})$. This method of representation will be used extensively throughout this chapter.

When $m = 2$, an MSC collapses to a BSC and the channel-specifying vector $\mathbf{p}$ equals $(1 - p, p)$, where p is the crossover probability. For simplicity, we sometimes use a scalar parameter p rather than a two-dimensional vector $\mathbf{p}$ to specify a BSC.

Remark 1: The mapping from the conditional distribution $F(dy|x)$ to the probabilistic weight $dP(\mathbf{p})$ is surjective. Furthermore, the equivalence classes of MI-SO channels with the same probabilistic weight $dP(\mathbf{p})$ are identical to the equivalence classes obtained from the detection point of view.

[2]This channel symmetrizing technique is equivalent to considering the LDPC coset code ensemble [57].

[3]Each subset contains $y, \mathcal{T}(y), \cdots, \mathcal{T}^{m-1}(y)$. In some cases, there are repeated elements in the subset of interest.

Remark 2: The probabilistic weight $dP(\mathbf{p})$ does not depend on the *a priori* input distribution on $\mathbf{X}$, but only depends on the channel model $F(dy|x)$. This observation will be used in the proof of the non-iterative bound in Section 5.5.

A Note on MNSC Decomposition

Similar to the MSC case, the simplest MI-NSO channel is the m-ary non-symmetric channel (MNSC), which is a $\mathbb{Z}_m \mapsto \mathbb{Z}_m$ channel and can be specified by a matrix $\mathbf{P} = (p_{i,j})$ such that the conditional distribution $\mathsf{P}(y = j|x = i) = p_{i,j}$. It can be shown that any general MI-NSO channel is equivalent to a probabilistic combination $dP(\mathbf{P})$ of many MNSCs with parameter $\mathbf{P}$ from a detection point of view. This representation method is very helpful when considering MI-NSO channels.

Nevertheless, unlike MI-SO channels, the MNSC decomposition for general MI-NSO channels is not unique. A channel, specified by $F(dy|x)$, can be converted to different $dP(\mathbf{P})$ and $dP'(\mathbf{P})$. Furthermore the probabilistic weight $dP(\mathbf{P})$ depends on the *a priori* distribution on X, which is another major difference between the MSC and MNSC decomposition.

5.1.3 Noise Measures

Binary Input Channels

For a binary channel $\mathbf{X} = \{0, 1\} \mapsto \mathbf{Y}$, we use $p(x|y)$ to denote the *a posteriori* probability $\mathsf{P}(X = x|Y = y)$, and consider the following two noise measures:

- [The Bhattacharyya Noise Parameter (BNP)]

$$\begin{aligned} \text{BNP} &:= \mathsf{E}_{X,Y}\left\{\sqrt{\frac{p(\bar{X}|Y)}{p(X|Y)}}\right\} \\ &= \sum_{x\in\{0,1\}} \mathsf{P}(X = x)\mathsf{E}\left\{\sqrt{\frac{p(\bar{x}|Y)}{p(x|Y)}}\middle| X = x\right\}, \end{aligned} \tag{5.1}$$

  A discussion of the BNP parameter in turbo-like codes can be found in [37]. With uniform distribution on X, BNP can be related to the cutoff rate R_0 by $R_0 = 1 - \log_2(\text{BNP} + 1)$.

- [The Expected Soft Bit Value (ESB)]

$$\begin{aligned} \text{ESB} &:= 2\mathsf{E}_{X,Y}\left\{p(\bar{X}|Y)\right\} \\ &= 2\sum_{x\in\{0,1\}} \mathsf{P}(X = x)\mathsf{E}\left\{p(\bar{x}|Y)|X = x\right\}, \end{aligned} \tag{5.2}$$

  which was used in the bounds of [22].

Each of the above noise measures has the property that the condition BNP = 0 (or ESB = 0) represents the noise-free channel, while BNP = 1 (or ESB = 1), implies a completely random channel. It is worth noting that both BNP and ESB are well defined even for BI-NSO channels with non-uniform input distributions. Throughout this chapter, unless further mentioned, we assume the *a priori* distribution is uniform.

For BSCs, BNP $= 2\sqrt{p(1-p)}$ and ESB $= 4p(1-p)$ where p is the crossover probability. By the BSC decomposition argument in Section 5.1.2, the value of BNP or ESB for any BI-SO channel is simply the probabilistic average of the corresponding values of the constituent BSCs, that is,

$$\begin{aligned} \text{BNP} &= \int 2\sqrt{p(1-p)}dP(p) \\ \text{ESB} &= \int 4p(1-p)dP(p). \end{aligned}$$

The above formulae will be extensively used in our derivation of finite dimensional bounds. In the context of density evolution [92], BNP and ESB can be expressed as

$$\begin{aligned} \text{BNP} &= \int_{m=-\infty}^{\infty} e^{-\frac{m}{2}} P(dm) \\ \text{ESB} &= \int_{m=-\infty}^{\infty} \frac{2}{1+e^m} P(dm), \end{aligned} \tag{5.3}$$

where $m := \log\left(\frac{\mathsf{P}(\mathbf{Y}|X=0)}{\mathsf{P}(\mathbf{Y}|X=1)}\right)$ is the passing LLR message and $P(dm) := \mathsf{P}(dm|X=0)$ is the density of m given $X=0$.
Examples (listed in order from the most BSC-like to the most BEC-like):

1. The BSC with crossover probability p:

$$\begin{aligned} \text{BNP} &= 2\sqrt{p(1-p)} \stackrel{\Delta}{=} \text{BNP}(p) \\ \text{ESB} &= 4p(1-p) \stackrel{\Delta}{=} \text{ESB}(p). \end{aligned}$$

2. The binary-input Laplace channel (BiLC) with variance $2\lambda^2$, i.e. $p_L(y) = \frac{1}{2\lambda}\exp\left(-\frac{|y|}{\lambda}\right)$:

$$\begin{aligned} \text{BNP} &= \frac{1+\lambda}{\lambda}\exp\left(-\frac{1}{\lambda}\right) \\ \text{ESB} &= \frac{\exp\left(-\frac{1}{\lambda}\right)}{\cosh\left(\frac{1}{\lambda}\right)} + 2\exp\left(-\frac{1}{\lambda}\right)\arctan\left(\tanh\left(\frac{1}{2\lambda}\right)\right). \end{aligned}$$

3. The binary-input additive white Gaussian noise channel (BiAWGNC) with noise variance σ^2:

$$\begin{aligned} \text{BNP} &= \exp\left(-\frac{1}{2\sigma^2}\right) \\ \text{ESB} &= \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{1}{2\sigma^2}\right)\int_{-\infty}^{\infty}\frac{\exp\left(-\frac{x^2}{2\sigma^2}\right)}{\cosh\left(\frac{x}{\sigma^2}\right)}dx. \end{aligned}$$

4. The binary-input Rayleigh fading channel with unit input energy and noise variance σ^2, i.e. the density function of the signal amplitude is $p_A(a) = 2a\exp(-a^2)$ and the additive noise distribution is Gaussian $\mathcal{N}(0,\sigma^2)$:

$$\begin{aligned} \text{BNP} &= \frac{1}{1+\frac{1}{2\sigma^2}} \\ \text{ESB} &= \frac{2}{\sqrt{2\pi\sigma^2}}\int_{a=0}^{\infty}\int_{x=-\infty}^{\infty} a\exp\left(-a^2\right)\frac{\exp\left(-\frac{x^2+a^2}{2\sigma^2}\right)}{\cosh\left(\frac{xa}{\sigma^2}\right)}dxda. \end{aligned}$$

5. The BEC with erasure probability ϵ:

$$\mathrm{BNP} = \mathrm{ESB} = \epsilon.$$

One example of a BI-NSO channel is the binary non-symmetric channel (BNSC), for which we have

- The BNSC is a $\{0,1\} \mapsto \{0,1\}$ channel with crossover probabilities $(p_{0\to1}, p_{1\to0})$:

$$\begin{aligned} \mathrm{BNP} &= \sqrt{(1-p_{0\to1})p_{1\to0}} + \sqrt{p_{0\to1}(1-p_{1\to0})} \\ \mathrm{ESB} &= 2\frac{(1-p_{0\to1})p_{1\to0}}{1-p_{0\to1}+p_{1\to0}} + 2\frac{p_{0\to1}(1-p_{1\to0})}{1+p_{0\to1}-p_{1\to0}}. \end{aligned}$$

m-ary Input Channels

For m-ary-input channels, we define the pairwise Bhattacharyya noise parameter from x to x' as follows:

$$\mathrm{BNP}(x \to x') := \mathsf{E}\left\{ \sqrt{\frac{\mathsf{P}(x'|Y)}{\mathsf{P}(x|Y)}} \middle| X = x \right\}. \tag{5.4}$$

Considering any MI-SO channel, we immediately have

$$\begin{aligned} \text{Symmetry:} \quad & \mathrm{BNP}(x \to x') = \mathrm{BNP}(x' \to x) \\ \text{Stationarity:} \quad & \mathrm{BNP}(x \to x') = \mathrm{BNP}(0 \to x' - x). \end{aligned} \tag{5.5}$$

By stationarity, we can use $\mathsf{BNP} := \{\mathrm{BNP}(0 \to x')\}_{x' \in \mathbb{Z}_m}$ as the representing vector for all $\mathrm{BNP}(x \to x')$. With the uniform distribution on X, the cutoff rate R_0 and BNP are related as follows [27]:

$$R_0 = \log_2 m - \log_2 \left(\sum_{x' \in \mathbb{Z}_m} \mathrm{BNP}(0 \to x') \right). \tag{5.6}$$

Example:

- For an MSC with parameter $\mathbf{p}$, we have

$$\mathrm{BNP}(0 \to x) = \sum_{y \in \mathbb{Z}_m} \sqrt{p_y p_{y+x}}. \tag{5.7}$$

  When $m = 2$, the representing vector becomes $\mathsf{BNP} = (1, \mathrm{BNP})$, where $\mathrm{BNP} = 2\sqrt{p(1-p)}$ is the traditional Bhattacharyya noise parameter for BSCs.

- Consider an 8PSK system with natural mapping, namely, $\mathbf{Y} = \mathbb{R}^2$ and the output Y satisfies

$$Y = \left(\cos\left(\frac{\pi X}{4}\right), \sin\left(\frac{\pi X}{4}\right) \right) + (N_i, N_q),$$

  where $X \in \mathbb{Z}_8$, and N_i and N_q are independent additive Gaussian noises with variance $\sigma^2/2$. Then using the MSC-decomposition method, the BNP values of this 8PSK system can be computed by Monte-Carlo simulations as in Algorithm 5.

Algorithm 5 Compute BNP values for an 8PSK system with variance $\sigma^2\mathbf{I}$.

1: $n_s \leftarrow 0, \mathsf{BNP}_{total} \leftarrow \mathbf{0}$.
2: **repeat**
3: Sample one (n_r, n_i) from independent zero-mean Gaussian noises with variances σ^2.
4: $y \leftarrow (1,0) + (n_r, n_i)$.
5: Construct the likelihood values

$$l_x = \frac{1}{\pi\sigma^2} e^{-\frac{||y-(\cos(\frac{\pi x}{4}),\sin(\frac{\pi x}{4}))||^2}{\sigma^2}}, \quad \forall x \in \mathbb{Z}_8.$$

6: Construct $\mathbf{p}$ such that $\mathbf{p} \propto \{l_x\}$ and $\sum_{x\in\mathbb{Z}_8} p_x = 1$.
7: Using $\mathbf{p}$ and (5.7) to construct BNP_{sample}.
8: $\mathsf{BNP}_{total} \leftarrow \mathsf{BNP}_{total} + \mathsf{BNP}_{sample}$
9: **until** $n_s \geq N_s$
10: Output $\mathsf{BNP} \leftarrow \frac{\mathsf{BNP}_{total}}{N_s}$.

5.1.4 Error Probability vs. BNP vs. ESB

Let $p_e = \mathsf{P}\left(X \neq \hat{X}_{\mathrm{MAP}}(Y)\right)$ denote the error probability of the MAP decoder. The relationships between p_e and the above noise measures BNP (or BNP) and ESB are stated by the following lemmas.

Lemma 5.1 *For general BI-NSO channels and arbitrary a priori input distributions, we have*

$$\begin{aligned} 2p_e \leq\ & \mathrm{BNP} \leq 2\sqrt{p_e(1-p_e)} \\ 2p_e \leq\ & \mathrm{ESB} \leq 4p_e(1-p_e) \\ \textit{and} \qquad & \mathrm{ESB} \leq \mathrm{BNP} \leq \sqrt{\mathrm{ESB}}. \end{aligned}$$

Lemma 5.2 *For any MI-SO channel with uniform input distribution, we have*

$$2p_e \leq \sum_{x\in\mathbb{Z}_m\backslash\{0\}} \mathrm{BNP}(0 \to x).$$

If $p_e \leq 1/2$, *then*[4]

$$\max_{x\in\mathbb{Z}_m\backslash\{0\}} \mathrm{BNP}(0 \to x) \leq 2\sqrt{p_e(1-p_e)}.$$

Lemma 5.1 guarantees that the three statements: $p_e \to 0$, BNP $\to 0$, and ESB $\to 0$ are equivalent. Lemma 5.2 guarantees $p_e \to 0$ is equivalent to the statement that $\forall x \in \mathbb{Z}_m\backslash\{0\}$, $\mathrm{BNP}(0 \to x) \to 0$. Detailed proofs of Lemmata 5.1 and 5.2 are provided in Appendix F.1.

5.2 The Support Tree Channel & Existing Bounds

We consider the equiprobable bipartite graph ensemble for LDPC codes as discussed in Section 4.1.3. For the following subsections, we will revisit the cycle-free support tree $\mathcal{N}^{2l}$ by viewing it as a binary-input/vector-output channel. Existing finite-dimensional bounds will be discussed by adopting this new perspective.

[4]When $m = 2$, p_e is guaranteed to be no larger than $1/2$. However, when $m > 2$, there are situations when $p_e > 1/2$.

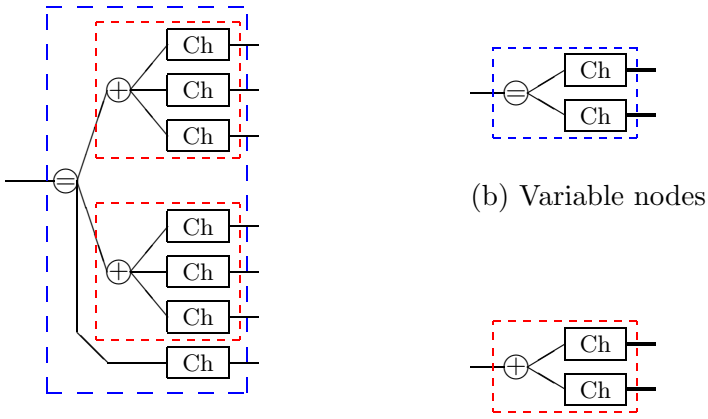

(a) Support tree of a regular (3,4) code (b) Variable nodes (c) Check nodes

Figure 5.3: Viewing the support tree as iterative concatenation of simple variable node and check node channels.

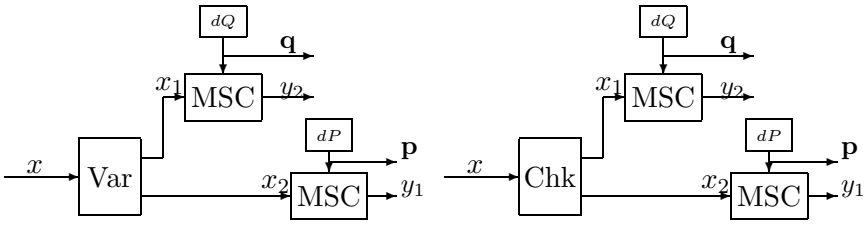

(a) Variable node channels (b) Check node channels

Figure 5.4: Separate consideration of variable and check nodes.

5.2.1 The Support Tree Channel

The belief propagation algorithm on LDPC codes can be broken down into tree structures, and Figure 5.3(a) demonstrates such tree structure for a regular (3,4) code. It can be clearly seen in Figure 5.3(a) that BP on the support tree becomes equivalent to a MAP detector on a $\mathbb{Z}_m \mapsto \mathbf{Y}^7$ vector channel. By the MSC decomposition argument in Figure 5.2(b), each observation channel can be converted from a $\mathbb{Z}_m \mapsto \mathbf{Y}$ channel to a $\mathbb{Z}_m \mapsto ([0,1]^m \times \mathbb{Z}_m)$ channel. Therefore, the overall $\mathbb{Z}_m \mapsto \mathbf{Y}^7$ vector channel can be converted into an equivalent $\mathbb{Z}_m \mapsto ([0,1]^m \times \mathbb{Z}_m)^7$.

Our target problem is to bound the (finite-dimensional) noise measures of the $\mathbb{Z}_m \mapsto ([0,1]^m \times \mathbb{Z}_m)^7$ vector output channel, given constraints on noise measures of the constituent channel distribution $dP_i(\mathbf{p})$, $i = 1, 2, \cdots, 7$. Once upper/lower bounds on the output noise measures are obtained, we can iteratively apply these bounds and test whether the noise measures converge to zero or are bounded away from zero as iteration goes on, which in turn gives us finite-dimensional lower/upper bounds on the decodable threshold. To simplify the problem further, we consider the variable node and the check node channels respectively as in Figure 5.3(b) and 5.3(c), and as in Figure 5.4.

For variable/check nodes with degrees $d > 3$, if we take the marginal approach (focusing on one input constituent channel while leaving other constituent channels fixed), all the effects of the fixed inputs can be grouped into a single input message. Therefore, it is as if

we take the marginal approach on a variable/check node with degree equal three. Since the analysis of nodes of degree one or two is trivial, only nodes of degree $d = 3$ will be discussed in detail.

5.2.2 Existing Results for Binary LDPC Codes

For BI-SO channels, the best way to explain the existing results in [22, 59, 69], and [103] is by the idea of "transfer functions" and the convexity/concavity analysis. In this subsection, we will consider only BI-SO channels and the noise measure BNP for example, which will lead to the iterative upper bound in [59] and a new iterative lower bound. Similar arguments can be used to derive the results in [22] (or in [69, 103]), if we substitute the noise measure BNP with ESB (or with the conditional entropy).

Check Nodes

For a check node as in Figure 5.4(b), the problem of finding an iterative upper/lower bound becomes an optimization problem as follows.

$$\begin{aligned}
\text{maximize/minimize} \quad & \mathrm{BNP}_{out} = \int \sqrt{4p(1-p)+4q(1-q)-4p(1-p)4q(1-q)}\,dP(p)dQ(q) \\
& \triangleq \int \mathrm{BNP}_{chk}(p,q)dP(p)dQ(q), \qquad (5.8) \\
\text{subject to} \quad & \mathrm{BNP}_{in,1} = \int 2\sqrt{p(1-p)}dP(p) \triangleq \int \mathrm{BNP}_1(p)dP(p) \\
& \mathrm{BNP}_{in,2} = \int 2\sqrt{q(1-q)}dQ(q) \triangleq \int \mathrm{BNP}_2(q)dQ(q),
\end{aligned}$$

where $\mathrm{BNP}_{chk}(p,q)$ denotes the value of BNP for the tree-like check node channel if both the constituent channels are BSCs with parameters p and q, respectively. Using some simple algebra and omitting the input parameters p and q, we can rewrite BNP_{chk} as

$$\mathrm{BNP}_{chk} = \sqrt{\mathrm{BNP}_1^2 + \mathrm{BNP}_2^2 - \mathrm{BNP}_1^2\mathrm{BNP}_2^2}, \qquad (5.9)$$

which is the BNP-based "transfer function" of the check node. Since BNP_{chk} is a convex function of BNP_1, this maximization/minimization problem is easy. The maximizing distribution $dP^*(p)$ is obtained by letting all probability weights concentrate on both extreme ends $p = 0$ and $p = 1/2$, that is

$$dP^*(p) = \begin{cases} 1 - \mathrm{BNP}_{in,1} & \text{if } p = 0 \\ \mathrm{BNP}_{in,1} & \text{if } p = 1/2 \\ 0 & \text{otherwise} \end{cases}.$$

Note: dP^* is a probabilistic combination of a noise-free channel ($p = 0$) and a completely random channel ($p = 1/2$), which corresponds to a BEC with erasure probability $\epsilon = \mathrm{BNP}_{in,1}$.

By Jensen's inequality, the minimizing distribution $dP^\dagger(p)$ is obtained by letting all probability weights concentrate on a single point with the same $\mathrm{BNP}_{in,1}$, that is

$$dP^\dagger(p) = \begin{cases} 1 & \text{if } \mathrm{BNP}_1(p) = 2\sqrt{p(1-p)} = \mathrm{BNP}_{in,1} \\ 0 & \text{otherwise} \end{cases}.$$

Note: $dP^\dagger$ corresponds to a BSC.

The same arguments can be applied to find dQ^* and $dQ^\dagger$. By replacing both dP and dQ in (5.8) with the maximizing dP^* and dQ^*, we prove that for all BI-SO constituent channels,

$$\mathrm{BNP}_{out} \le \mathrm{BNP}_{in,1} + \mathrm{BNP}_{in,2} - \mathrm{BNP}_{in,1}\mathrm{BNP}_{in,2}.$$

By replacing both dP and dQ in (5.8) with the minimizing $dP^\dagger$ and $dQ^\dagger$, we also have

$$\mathrm{BNP}_{out} \ge \sqrt{\mathrm{BNP}_{in,1}^2 + \mathrm{BNP}_{in,2}^2 - \mathrm{BNP}_{in,1}^2\mathrm{BNP}_{in,2}^2}.$$

By a straightforward extension to check nodes of higher degree $d_c \ge 3$, an iterative upper bound can be obtained by replacing all constituent channels with BECs having the same values of $\mathrm{BNP}_{in,i}$. That is,

$$\mathrm{BNP}_{out} \le 1 - \prod_{i=1}^{d_c-1}(1 - \mathrm{BNP}_{in,i}). \tag{5.10}$$

An iterative lower bound can be obtained by replacing all constituent channels with BSCs having the same values of $\mathrm{BNP}_{in,i}$. That is,

$$\mathrm{BNP}_{out} \ge \sqrt{1 - \prod_{i=1}^{d_c-1}(1 - \mathrm{BNP}_{in,i}^2)}. \tag{5.11}$$

Variable Nodes

For a variable node as shown in Figure 5.4(a), the problem of finding an iterative upper/lower bound is

$$\begin{aligned}
\text{maximize/minimize} \qquad & \mathrm{BNP}_{out} = \int 4\sqrt{p(1-p)q(1-q)}dP(p)dQ(q) \\
& \triangleq \int \mathrm{BNP}_{var}(p,q)dP(p)dQ(q), \\
\text{subject to} \qquad & \mathrm{BNP}_{in,1} = \int 2\sqrt{p(1-p)}dP(p) \triangleq \int \mathrm{BNP}_1(p)dP(p) \\
& \mathrm{BNP}_{in,2} = \int 2\sqrt{q(1-q)}dQ(q) \triangleq \int \mathrm{BNP}_2(q)dQ(q),
\end{aligned}$$

where $\mathrm{BNP}_{var}(p,q)$ denotes the value of BNP for the tree-like variable node channel if both the constituent channels are BSCs with parameters p and q, respectively. We can rewrite BNP_{var} as

$$\mathrm{BNP}_{var} = \mathrm{BNP}_1\mathrm{BNP}_2, \tag{5.12}$$

which is the BNP-based "transfer function" of the variable node. Since BNP_{var} is a concave[5] function of BNP_1, this maximization/minimization problem is easy. By Jensen's inequality,

[5]Actually BNP_{var} is a linear function of BNP_1. The reason we still view it as a concave function is to keep the argument reusable when we are considering other types of noise measures (e.g., ESB and the conditional entropy).

the maximizing distribution $dP^*(p)$ is obtained by letting all probability weights concentrate on a single point with the same $\mathrm{BNP}_{in,1}$, that is

$$dP^*(p) = \begin{cases} 1 & \text{if } \mathrm{BNP}_1(p) = 2\sqrt{p(1-p)} = \mathrm{BNP}_{in,1} \\ 0 & \text{otherwise} \end{cases},$$

which corresponds to a BSC. The minimizing distribution $dP^\dagger(p)$ is obtained by letting all probability weights concentrate on both extreme ends $p = 0$ and $p = 1/2$, that is

$$dP^\dagger(p) = \begin{cases} 1 - \mathrm{BNP}_{in,1} & \text{if } p = 0 \\ \mathrm{BNP}_{in,1} & \text{if } p = 1/2 \\ 0 & \text{otherwise} \end{cases},$$

which corresponds to a BEC. As a result, by replacing all constituent BI-SO channels with BSCs having the same values of $\mathrm{BNP}_{in,i}$, we obtain an iterative upper bound for the variable node:

$$\mathrm{BNP}_{out} \leq \prod_{i=1}^{d_v - 1} \mathrm{BNP}_{in,i}. \tag{5.13}$$

By replacing all constituent channels with BECs having the same values of $\mathrm{BNP}_{in,i}$, we obtain an iterative lower bound for the variable node:

$$\mathrm{BNP}_{out} \geq \prod_{i=1}^{d_v - 1} \mathrm{BNP}_{in,i}. \tag{5.14}$$

Combined Results

Consider BI-SO channels and the irregular code ensemble with degree polynomials λ and ρ. By combining (5.10) and (5.13) and averaging over the degree distributions, we have

$$\mathrm{BNP}^{(l+1)} \leq \mathrm{BNP}^{(0)} \lambda\left(1 - \rho\left(1 - \mathrm{BNP}^{(l)}\right)\right), \tag{5.15}$$

where $\mathrm{BNP}^{(l)}$ is the value of BNP after l iterations, namely, the value of BNP for the support tree of depth $2l$. This is the result of Khandekar *et al.* in [59].

By combining (5.11) and (5.14) and averaging over (λ, ρ), we have a new iterative lower bound.

Theorem 5.1 *For BI-SO channels,*

$$\mathrm{BNP}^{(l+1)} \geq \mathrm{BNP}^{(0)} \lambda\left(\sum_k \rho_k \sqrt{1 - \left(1 - \left(\mathrm{BNP}^{(l)}\right)^2\right)^{k-1}}\right). \tag{5.16}$$

Given $\mathrm{BNP}^{(0)}$, the initial BNP value over the uncoded channel, we can iteratively compute $\mathrm{BNP}^{(l)}$ by (5.15), which is guaranteed to be an upper bound for the actual BNP value after l iterations. If $\lim_{l\to\infty} \mathrm{BNP}^{(l)} = 0$, then all channels with that initial BNP value $\mathrm{BNP}^{(0)}$ are guaranteed to be decodable. So $\mathrm{BNP}^* := \sup\{\mathrm{BNP}^{(0)} : \lim_{l\to\infty} \mathrm{BNP}^{(l)} = 0\}$

serves as a lower bound of the decodable thresholds. Similarly, if we use (5.16) to compute $\mathrm{BNP}^{(l)}$, then $\mathrm{BNP}^{\dagger} := \inf\{\mathrm{BNP}^{(0)} : \lim_{l\to\infty} \mathrm{BNP}^{(l)} > 0\}$ is an upper bound of the decodable thresholds.

Similar arguments can be applied to other types of noise measures. In each iteration, replacing constituent channels of a check node with BECs/BSCs having the same value of ESB, and replacing variable node constituent channels with BSCs/BECs having the same value of ESB, we can reproduce the iterative upper/lower bound on ESB in [22]. By considering the conditional entropy instead of ESB, we can reproduce the iterative upper/lower bound on the mutual information found in [69, 103].

This chapter will focus on developing new bounds or strengthening existing bounds for the cases in which this simple convexity/concavity analysis of the transfer function does not hold.

5.3 Iterative Bounds for $\mathbb{Z}_m$ LDPC Codes

5.3.1 Code Ensemble

The $\mathbb{Z}_m$-based LDPC code ensemble can be described as follows. We consider only parity check matrices $\mathbf{H}$ with the values of non-zero entries being all ones. The random parity check matrix ensemble is then identical to the ensemble of binary LDPC codes introduced in Section 4.1.3. The only difference is that the parity check equation $\mathbf{Hx} = \mathbf{0}$ is now evaluated in $\mathbb{Z}_m$. A further deviation from the binary code ensemble is the $\mathsf{GF}(q)$-based code ensemble. For the $\mathsf{GF}(q)$ code ensemble, besides evaluating $\mathbf{Hx} = \mathbf{0}$ in $\mathsf{GF}(q)$, the non-zero entries in $\mathbf{H}$ are uniformly distributed between $\{1, 2, \cdots, q-1\}$. Further discussion of the $\mathbb{Z}_m$ and $\mathsf{GF}(q)$ LDPC code ensembles can be found in [14, 15].

5.3.2 Iterative Bounds

Variable Nodes

We focus on a variable node with degree $d_v = 3$ as in Figure 5.4(a). We will first consider $\mathbf{p}$ and $\mathbf{q}$ being fixed (non-random) parameters and then extend our analysis to accommodate the random parameter generators $dP(\mathbf{p})$ and $dQ(\mathbf{q})$.

By grouping the outputs Y_1 and Y_2 into a $2m$-dimensional vector $\mathbf{Y} = (Y_1, Y_2)$, the variable node becomes a $\mathbb{Z}_m \mapsto \mathbb{Z}_m^2$ channel, and it can be verified that it is still symmetric. By definition (5.4), the resulting $\mathrm{BNP}_{var}(0 \to x)$ for the vector output channel is

$$\begin{aligned}
\mathrm{BNP}_{var}(0 \to x) &= \sum_{\mathbf{y}\in\mathbb{Z}_m^2} \sqrt{(p_{y_1} q_{y_2})(p_{y_1 - x} q_{y_2 - x})} \\
&= \left(\sum_{y_1 \in \mathbb{Z}_m} \sqrt{p_{y_1} p_{y_1 - x}} \right) \cdot \left(\sum_{y_2 \in \mathbb{Z}_m} \sqrt{q_{y_2} q_{y_2 - x}} \right) \\
&= \mathrm{BNP}_{in,1}(0 \to x) \cdot \mathrm{BNP}_{in,2}(0 \to x).
\end{aligned}$$

A compact vector representation using the component-wise product "$\bullet$" then becomes

$$\mathsf{BNP}_{var} = \mathsf{BNP}_1 \bullet \mathsf{BNP}_2.$$

By iteratively applying the above inequality for variable nodes with $d_v \geq 3$, we have

$$\mathsf{BNP}_{var} = \prod_{j=1}^{d_v-1} \mathsf{BNP}_{in,j}, \tag{5.17}$$

where the $\prod$ represents the component-wise product. Consider general MI-SO constituent channels with random parameter generators $dP_j(\mathbf{p}^j)$, where $\mathbf{p}^j$ denotes the parameter vector for the j-th constituent channel and its distribution is denoted by $dP_j(\cdot)$. Since the parameter vectors $\mathbf{p}^j$ are independently distributed for different values of j, the probabilistic average of the product in (5.17) is the product of individual averages. This implies that (5.17) holds for general MI-SO channels as well.

Check Nodes

Consider a check node with degree $d_c = 3$, namely, two constituent MSCs with parameters $\mathbf{p}$ and $\mathbf{q}$, as illustrated in Figure 5.4(b). By definition, $\mathrm{BNP}_{chk}(0 \to x)$ for the $\mathbb{Z}_m \mapsto \mathbb{Z}_m^2$ channel is given as follows:

$$\begin{aligned}
\mathrm{BNP}_{chk}(0 \to x) &= \sum_{w=0}^{m-1} \sqrt{\left(\sum_{y_1+y_2=w}^{m-1} p_{y_1} q_{y_2}\right)\left(\sum_{y_1+y_2=w+x}^{m-1} p_{y_1} q_{y_2}\right)} \\
&= \sum_{w=0}^{m-1} \sqrt{\left(\sum_{y_1=0}^{m-1} p_{y_1} q_{w-y_1}\right)\left(\sum_{y_1=0}^{m-1} p_{y_1} q_{x+w-y_1}\right)}.
\end{aligned} \tag{5.18}$$

Each summand in (5.18) can be upper bounded by

$$\begin{aligned}
&\sqrt{\left(\sum_{y_1=0}^{m-1} p_{y_1} q_{w-y_1}\right)\left(\sum_{y_1=0}^{m-1} p_{y_1} q_{x+w-y_1}\right)} \\
&= \sqrt{\left(\sum_{y=0}^{m-1} p_y q_{w-y}\right)\left(\sum_{z=0}^{m-1} p_z q_{x+w-z}\right)} \\
&= \sqrt{\sum_{y=0}^{m-1}\sum_{z=0}^{m-1} p_y q_{w-y} p_z q_{x+w-z}} \\
&\leq \sum_{y=0}^{m-1}\sum_{z=0}^{m-1} \sqrt{p_y p_z}\sqrt{q_{w-y} q_{x+w-z}},
\end{aligned} \tag{5.19}$$

where the inequality follows from the fact that $\sqrt{\sum x_i} \leq \sum \sqrt{x_i}$ if $x_i \geq 0, \forall i$. By combining (5.18) and (5.19), we have

$$\begin{aligned} \mathrm{BNP}_{chk}(0 \to x) &\leq \sum_{w=0}^{m-1} \sum_{y=0}^{m-1} \sum_{z=0}^{m-1} \sqrt{p_y p_z} \sqrt{q_{w-y} q_{x+w-z}} \\ &= \sum_{w'=0}^{m-1} \sum_{y=0}^{m-1} \sum_{z'=0}^{m-1} \sqrt{p_y p_{y+z'}} \sqrt{q_{w'} q_{w'+x-z'}} \qquad (5.20) \\ &= \sum_{z'=0}^{m-1} \mathrm{BNP}_{in,1}(0 \to z') \mathrm{BNP}_{in,2}(0 \to x - z'), \end{aligned}$$

where (5.20) follows from the change of variables: $w' = w - y$ and $z' = z - y$. A compact vector representation using circular convolution "$\otimes$" then becomes

$$\mathsf{BNP}_{chk} \leq \mathsf{BNP}_{in,1} \otimes \mathsf{BNP}_{in,2}. \qquad (5.21)$$

By iteratively applying the above inequality and noting the monotonicity of the convolution operator (given all operands are component-wise non-negative), we have

$$\mathsf{BNP}_{chk} \leq \bigotimes_{j=1}^{d_c - 1} \mathsf{BNP}_{in,j}. \qquad (5.22)$$

Since the circular convolution is a linear operator and the $\mathbf{p}^j$ are independently distributed for different values of j, the probabilistic average of the circular convolution in (5.22) is the circular convolution of individual averages. This implies that (5.22) holds for general MI-SO channels as well.

Note: (5.21) is loose for the binary case ($m = 2$). For $m > 2$, there are many non-trivial cases in which (5.21) is tight. For example, suppose $m = 6$, $\mathbf{p}^1 = (0.5, 0.5, 0, 0, 0, 0)$, and $\mathbf{p}^2 = (0.5, 0, 0.5, 0, 0, 0)$. We have $\mathsf{BNP}_{chk} = (1, 0.75, 0.5, 0.5, 0.5, 0.75)$, $\mathsf{BNP}_{in,1} = (1, 0.5, 0, 0, 0, 0.5)$ and $\mathsf{BNP}_{in,2} = (1, 0, 0.5, 0, 0.5, 0)$, which attains the equality in (5.21).

Combined Results

Consider general MI-SO channels and the irregular code ensemble with degree polynomials λ and ρ. By combining (5.17) and (5.22) and averaging over the degree distributions, we have proven the following theorem.

Theorem 5.2 *Let* $\mathsf{BNP}^{(l)}$ *denote the value of* BNP *for the support tree channel after* l *iterations. Then we have*

$$\mathsf{BNP}^{(l+1)} \leq \mathsf{BNP}^{(0)} \lambda\left(\rho\left(\mathsf{BNP}^{(l)}\right)\right), \qquad (5.23)$$

where the scalar products within $\lambda(\cdot)$ *are replaced by component-wise products, and the scalar products within* $\rho(\cdot)$ *are replaced by circular convolutions.*

For a $\mathbb{Z}_m$ code ensemble with degree polynomials (λ, ρ), we can first fix an arbitrary[6] BNP^*, let $\mathsf{BNP}^{(0)} = \mathsf{BNP}^*$, and iteratively compute the upper bound by (5.23)[7] . Suppose

[6] Any valid BNP must satisfy the symmetry condition in (5.5) and the condition that $\mathrm{BNP}(0 \to x) \in [0, 1]$, $\forall x \in \mathbb{Z}_m$.

[7] During the iterations, we may further strengthen (5.23) by $\mathsf{BNP}^{(l+1)} \leq \min\left\{\mathbf{1}, \mathsf{BNP}^{(0)} \lambda\left(\rho\left(\mathsf{BNP}^{(l)}\right)\right)\right\}$, since any valid BNP value is upper bounded by 1.

for all $x \neq 0$, $\lim_{l\to\infty} \mathrm{BNP}^{(l)}(0 \to x) = 0$. By Lemma 5.2, any MI-SO channel with $\mathsf{BNP} \leq \mathsf{BNP}^*$ is guaranteed to be decodable by the BP algorithm when sufficiently long codes are used.

5.3.3 Stability Conditions

Just like the binary case in Section 4.4.3, the sufficient stability condition for $\mathbb{Z}_m$ LDPC codes can be obtained as a corollary to Theorem 5.2.

Corollary 5.1 (Sufficient Stability Condition) *Consider any MI-SO channel with noise measure* BNP *and any* $\mathbb{Z}_m$ *LDPC code ensemble with degree polynomials* (λ, ρ)*. If*

$$\lambda_2\rho'(1)\mathrm{BNP}(0 \to x) < 1, \quad \forall x \in \mathbb{Z}_m\backslash\{0\},$$

then this $\mathbb{Z}_m$ *code is stable under the BP decoder. Namely, there exists* $\epsilon^* > 0$ *such that if after* l_0 *iterations,* $\max_{x\in\mathbb{Z}_m\backslash\{0\}} \mathrm{BNP}^{(l_0)}(0 \to x) < \epsilon^*$*, then* $\lim_{l\to\infty} \mathrm{BNP}^{(l)}(0 \to x) = 0$ *for all* $x \in \mathbb{Z}_m\backslash\{0\}$*. (Or equivalently* $\lim_{l\to\infty} p_e^{(l)} = 0$*.) Furthermore, the convergence rate of* $\mathrm{BNP}^{(l)}(0 \to x)$ *is exponential or superexponential depending on whether* $\lambda_2 > 0$ *or* $\lambda_2 = 0$*.*

Proof: By noting the fact that $\mathrm{BNP}^{(l)}(0 \to 0) = 1$ for all $l \in \mathbb{N}$ and taking the infinitesimal analysis of (5.23), this corollary can be easily proved. ■

A matching necessary stability condition can be proved as follows.

Theorem 5.3 (Necessary Stability Condition) *Consider any MI-SO channel with noise measure* BNP *and any* $\mathbb{Z}_m$ *LDPC code ensemble with degree polynomials* (λ, ρ)*. If*

$$\exists x_0 \in \mathbb{Z}_m\backslash\{0\}, \text{ such that } \lambda_2\rho'(1)\mathrm{BNP}(0 \to x_0) > 1,$$

then this $\mathbb{Z}_m$ *code is not stable under the BP decoder. Namely, there exists* $x_0 \in \mathbb{Z}_m\backslash\{0\}$ *such that* $\lim_{l\to\infty} \mathrm{BNP}^{(l)}(0 \to x_0) > 0$*, or equivalently,* $\lim_{l\to\infty} p_e^{(l)} > 0$*.*

A detailed proof using a channel degradation argument similar to [14, 92] is provided in Appendix G.1.

We close this subsection by showing the stability results for $\mathsf{GF}(q)$ LDPC codes in [14] can be derived as a corollary to the above stability conditions for $\mathbb{Z}_m$ LDPC codes.

Consider a MI-SO channel with noise measure BNP, and a $\mathsf{GF}(q)$-based LDPC code with degree polynomials (λ, ρ), where q is a prime number. We then have

Corollary 5.2 (Sufficient Stability Condition for $\mathsf{GF}(q)$ LDPC Codes [14]) *If*

$$\lambda_2\rho'(1)\frac{\sum_{x\in\mathsf{GF}(q)\backslash\{0\}} \mathrm{BNP}(0 \to x)}{q-1} < 1,$$

then this $\mathsf{GF}(q)$ *code is stable under the BP decoder.*

Corollary 5.3 (Necessary Stability Condition for $\mathsf{GF}(q)$ LDPC Codes [14]) *If*

$$\lambda_2\rho'(1)\frac{\sum_{x\in\mathsf{GF}(q)\backslash\{0\}} \mathrm{BNP}(0 \to x)}{q-1} > 1,$$

then this $\mathsf{GF}(q)$ *code is not stable under the BP decoder.*

We note the two facts that the stability conditions of the $\mathbb{Z}_m$ LDPC codes rely only on the pairwise error probability, and the multiplication of the uniformly distributed edge weight $w \in \mathsf{GF}(q)\backslash\{0\}$ is equivalent to a uniform permutation of all non-zero entries. Therefore, the stability conditions of a $\mathsf{GF}(q)$ code are equivalent to those of a $\mathbb{Z}_m$ code with the "error pattern" averaged over all non-zero entries. As a result, all results for $\mathbb{Z}_m$ codes involving only $\mathrm{BNP}(0 \to x)$ hold for $\mathsf{GF}(q)$ codes as well with each $\mathrm{BNP}(0 \to x)$ being replaced by the average $\frac{1}{q-1}\sum_{x=1}^{q-1}\mathrm{BNP}(0 \to x)$. The above corollaries then become simply the restatements of Corollary 5.1 and Theorem 5.3, the stability conditions for $\mathbb{Z}_m$ LDPC codes.

5.3.4 Applications

There are many applications of high order LDPC codes, including the following ones.

1. Improving the code performance [34]: By grouping two bits into a $\mathsf{GF}(4)$ symbol, the finite length performance of binary codes can be enhanced by using higher dimensional $\mathsf{GF}(4)$ codes even though the equivalent codeword length (in terms of number of symbols) is reduced. This improvement is obtained at the cost of additional computational cost for each processing unit, which is $\mathcal{O}(m\log(m))$ or $\mathcal{O}(m)$ depending on whether we are looking at the complexity per symbol or per bit. It is worth mentioning that due to the reduction of codeword length, we can use a smaller interleaver, which offsets the increasing complexity of the processing unit.

2. High order coded modulation: Berkmann used high order codes for coded modulation with the natural code-to-symbol mapping [16, 17]. Using high order codes, we can convert the coded modulation into an m-ary-input memoryless channel. With optimized coded-alphabet-to-constellation mapper, we can fully utilize the additional error correcting capability of high order codes, and performance superior to bit-interleaved coded modulation (BICM) can be achieved. Further discussion of this issue will be included in Chapter 6.

3. Constructing mutual-information-achieving codes with non-uniform coded bit distribution: The basic idea is to concatenate a $\mathbb{Z}_m$ (or $\mathsf{GF}(q)$) code with a symbol mapper $U : \mathbb{Z}_m \mapsto \{0, 1, \cdots, k-1\}$ with $k < m$. Since for $\mathbb{Z}_m$ codes, the marginal symbol distribution is uniform on $\mathbb{Z}_m$, by carefully selecting (m, U), we can reproduce/approximate any designated *a priori* symbol distribution on $\{0, 1, \cdots, k-1\}$. If the original $\mathbb{Z}_m$ code achieves the symmetric mutual information rate, then the concatenation of the $\mathbb{Z}_m$ code and the symbol mapper U achieves the mutual information with the designated *a priori* distribution. Among the applications here are the following.

   - Constructing codes for cases in which the capacity achieving *a priori* distribution is not uniform [80, 15].
   - Designing optimal superposition codes for broadcast channels.
   - Designing codes for optical channels with cross talk [89].

4. Lossless data compression: For discrete memoryless channels, a capacity-approaching LDPC error correcting code can serve also as an entropy-approaching LDPC compressor (of the noise vector) due to the duality between source and channel coding. In [26], both high order LDPC codes and binary LDPC coded multi-level schemes

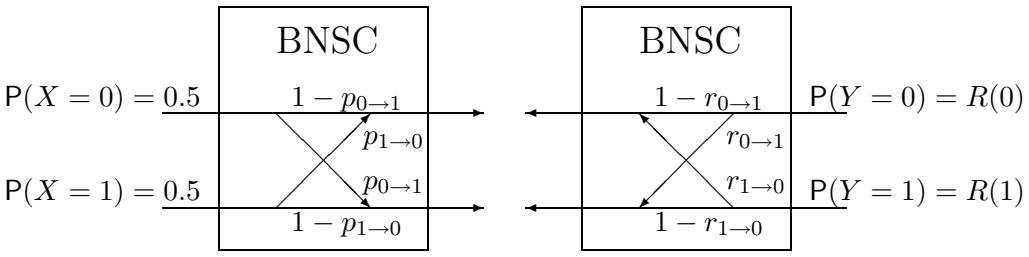

(a) Normal Perspective (b) Reverse Perspective

Figure 5.5: The probabilistic models of the BNSC.

are proposed for lossless, byte-level data compression, and very close-to-entropy compression ratios have been reported. For data with memory, further improvement can be achieved by incorporating the Burrows-Wheeler block sorting transform, and the results outperform the standard gzip and bzip algorithms based on the Lempel-Ziv algorithm and adaptive arithmetic coding.

Other references on high order LDPC codes can be found in [40, 71]

5.4 Iterative Bounds for Binary LDPC Codes

In this section, we will first show that the existing BNP-based iterative bounds for BI-SO channels also hold for BI-NSO channels. Then we will strengthen the existing BNP-based and ESB-based bounds by providing a two-dimensional (BNP, ESB)-based iterative upper bound for BI-SO channels.

5.4.1 A BNP-based Bound on BI-NSO Channels

Assuming the uniform *a priori* distribution on $\mathbf{X} = \{0, 1\}$, BNP is well defined in (5.1) for BI-NSO channels. We will prove that the inequalities (5.15) and (5.16) hold even for BI-NSO channels by adopting the reverse channel perspective.

Theorem 5.4 (BNP**-based Bounds for BI-NSO Channels)** *For the irregular LDPC code with degree polynomials* (λ, ρ)*, the iterative upper and lower bounds (5.15) and (5.16) hold for BI-NSO channels.*

Proof: We first focus on the simplest BNSC as illustrated in Figure 5.5(a). Since any BI-NSO channel can be regarded as the probabilistic combination of many BNSCs, our results for BNSC can then be generalized to arbitrary BI-NSO channels.

Any BNSC can be specified by two scalar parameters $p_{0\to1}$ and $p_{1\to0}$, where $p_{x\to y}$ denotes the conditional probability $\mathsf{P}(Y = y|X = x)$. BNP thus becomes

$$\text{BNP} = \sqrt{p_{0\to1}(1-p_{1\to0})} + \sqrt{p_{1\to0}(1-p_{0\to1})}. \tag{5.24}$$

We can also represent this $X \mapsto Y$ BNSC using the reverse channel perspective $Y \mapsto X$ as in Figure 5.5(b), such that

$$\begin{aligned} r_{0\to 1} &:= \frac{p_{1\to 0}}{1-p_{0\to 1}+p_{1\to 0}} \\ r_{1\to 0} &:= \frac{p_{0\to 1}}{1+p_{0\to 1}-p_{1\to 0}} \\ R(0) &:= \frac{1-p_{0\to 1}+p_{1\to 0}}{2} \\ R(1) &:= \frac{1+p_{0\to 1}-p_{1\to 0}}{2}, \end{aligned}$$

where $r_{y\to x} := \mathsf{P}(X=x|Y=y)$ and $R(y) = \mathsf{P}(Y=y)$. Then by definition, we have

$$\begin{aligned} \text{BNP} &:= \mathsf{E}_{X,Y}\left\{\sqrt{\frac{p(\bar{X}|Y)}{p(X|Y)}}\right\} \\ &= R(0)2\sqrt{r_{0\to 1}(1-r_{0\to 1})} + R(1)2\sqrt{r_{1\to 0}(1-r_{1\to 0})} \\ &= R(0)\text{BNP}(r_{0\to 1}) + R(1)\text{BNP}(r_{1\to 0}), \end{aligned}$$

in which $\text{BNP}(p) \stackrel{\Delta}{=} 2\sqrt{p(1-p)}$ computes the value of BNP for a BSC with crossover probability p. This representation separates the entangled expression of BNP in (5.24) so that it is as if there are two BSCs with parameters $r_{0\to 1}$ and $r_{1\to 0}$ respectively, and these two BSCs are probabilistically combined with coefficients $R(0)$ and $R(1)$. We use $(R(\cdot), \{r_{\cdot\cdot}\})$ to represent this reverse channel perspective.

Consider the variable/check nodes of degree 3 with two constituent BNSCs in reverse form, namely, Ch1:$(R(\cdot), \{r_{\cdot\cdot}\})$ and Ch2:$(S(\cdot), \{s_{\cdot\cdot}\})$, such that

$$\begin{aligned} \text{BNP}_{in,1} &:= \sum_{y_1\in\{0,1\}} R(y_1)\text{BNP}(r_{y_1\to 0}), \\ \text{BNP}_{in,2} &:= \sum_{y_2\in\{0,1\}} S(y_2)\text{BNP}(s_{y_2\to 0}). \end{aligned}$$

For a variable node of degree $d_v = 3$, by definition and after some simple algebra, we have

$$\begin{aligned} \text{BNP}_{var} &= \sum_{y_1\in\{0,1\}}\sum_{y_2\in\{0,1\}} 4R(y_1)S(y_2)\sqrt{r_{y_1\to 0}r_{y_1\to 1}s_{y_2\to 0}s_{y_2\to 1}} \\ &= \left(\sum_{y_1\in\{0,1\}} R(y_1)2\sqrt{r_{y_1\to 0}r_{y_1\to 1}}\right)\left(\sum_{y_2\in\{0,1\}} S(y_2)2\sqrt{s_{y_2\to 0}s_{y_2\to 1}}\right) \\ &= \text{BNP}_{in,1}\cdot\text{BNP}_{in,2}. \end{aligned} \tag{5.25}$$

By noting that (5.25) possesses the same form as in (5.12), all our previous analyses for variable nodes with BI-SO constituent channels hold for BI-NSO channels as well.

Consider a check node of degree $d_c = 3$, which is similar to Figure 5.4(b) except that the constituent channels are now BNSCs. By definition, some simple algebra, and the

observation that $X = X_1 + X_2$, we have

$$\begin{aligned}
&\mathrm{BNP}_{chk} \\
&= \sum_{y_1\in\{0,1\}} \sum_{y_2\in\{0,1\}} \mathsf{P}\left(\mathbf{Y} = (y_1, y_2)\right) \\
&\qquad 2\sqrt{\mathsf{P}\left(X_1 + X_2 = 0 | \mathbf{Y} = (y_1, y_2)\right) \mathsf{P}\left(X_1 + X_2 = 1 | \mathbf{Y} = (y_1, y_2)\right)} \\
&= \sum_{y_1\in\{0,1\}} \sum_{y_2\in\{0,1\}} R(y_1)S(y_2)\sqrt{4r_{y_1\to 0}r_{y_1\to 1} + 4s_{y_2\to 0}s_{y_2\to 1} - 16r_{y_1\to 0}r_{y_1\to 1}s_{y_2\to 0}s_{y_2\to 1}} \\
&= \sum_{y_1\in\{0,1\}} \sum_{y_2\in\{0,1\}} R(y_1)S(y_2)\sqrt{\left(\mathrm{BNP}(r_{y_1\to 0})\right)^2 + \left(\mathrm{BNP}(s_{y_2\to 0})\right)^2 - \left(\mathrm{BNP}(r_{y_1\to 0})\right)^2 \cdot \left(\mathrm{BNP}(s_{y_2\to 0})\right)^2}.
\end{aligned} \tag{5.26}$$

Note that (5.26) possesses the same form as in (5.8) and (5.9). Thus, it is as if we are facing two constituent BI-SO channels and each of them is a probabilistic combination of two BSCs with weights $R(0)$ and $R(1)$ (or $S(0)$ and $S(1)$). Therefore, all our previous analyses for check nodes with BI-SO constituent channels hold for BI-NSO channels as well.

Since our previous analyses for both variable and check nodes (with BI-SO channels) hold for BI-NSO channels as well, we have proven Theorem 5.4. ■

5.4.2 A Two-Dimensional Upper Bound on BI-SO Channels

In this section, we develop a two-dimensional upper bound on the (BNP, ESB) pair of a BI-SO channel, for which convexity/concavity analysis of the transfer function is not sufficient. Similar to the one-dimensional results in Section 5.2.2, we consider variable nodes and check nodes separately.

Check Nodes

Suppose the check node channel has two constituent BSCs with crossover probabilities $p \in [0, 1/2]$ and $q \in [0, 1/2]$ as shown in Figure 5.4(b), where p and q have distributions $dP(p)$ and $dQ(q)$, respectively. Let $\mathrm{BNP}_{in,1}$ and $\mathrm{ESB}_{in,1}$ denote upper bounds on the values of BNP and ESB for the first constituent channel and let $\mathrm{BNP}_{in,2}$ and $\mathrm{ESB}_{in,2}$ denote corresponding upper bounds for the second constituent channel. We would like to develop an upper bound on the pair (BNP, ESB) for the support tree channel. This iterative bounding problem thus becomes:

$$\begin{aligned}
\text{maximize} \qquad & \mathrm{BNP}_{out} = \int \sqrt{4p(1-p) + 4q(1-q) - 4p(1-p)4q(1-q)}dP(p)dQ(q) \\
& \mathrm{ESB}_{out} = \int 4p(1-p) + 4q(1-q) - 4p(1-p)4q(1-q)dP(p)dQ(q) \\
\text{subject to} \qquad & \int 2\sqrt{p(1-p)}dP(p) \le \mathrm{BNP}_{in,1} \\
& \int 4p(1-p)dP(p) \le \mathrm{ESB}_{in,1} \\
& \int 2\sqrt{q(1-q)}dQ(q) \le \mathrm{BNP}_{in,2} \\
& \int 4q(1-q)dQ(q) \le \mathrm{ESB}_{in,2}.
\end{aligned} \tag{5.27}$$

Using some simple algebra, we can show that $\mathrm{ESB}_{out} \leq \mathrm{ESB}_{in,1} + \mathrm{ESB}_{in,2} - \mathrm{ESB}_{in,1}\mathrm{ESB}_{in,2}$ and the equality can be attained by all feasible $dP(p)$ and $dQ(q)$ meeting the constraints. The remaining problem thus reduces to the maximization of BNP_{out} subject to two input constraints on each of dP and dQ. Solving this maximization problem, the maximizing dP^* and dQ^* can be expressed as follows.

$$\begin{aligned} dP^*(p) &= \begin{cases} 1 - \frac{\mathrm{BNP}_{in,1}}{t} & \text{if } p = 0 \\ \frac{\mathrm{BNP}_{in,1}}{t} & \text{if } 2\sqrt{p(1-p)} = t \\ 0 & \text{otherwise} \end{cases}, \\ \text{where} \qquad t &= \frac{\mathrm{ESB}_{in,1}}{\mathrm{BNP}_{in,1}}. \end{aligned}$$

dQ^* can be obtained by replacing $\mathrm{BNP}_{in,1}$, $\mathrm{ESB}_{in,1}$, and p in the above equation with $\mathrm{BNP}_{in,2}$, $\mathrm{ESB}_{in,2}$, and q, respectively. A proof of the optimality of dP^* and dQ^* is given in Appendix G.2.

By substituting all constituent BI-SO channels with channels of the same form as dP^*, we obtain an upper bound on $(\mathrm{BNP}, \mathrm{ESB})$ in check node iterations as follows.

Theorem 5.5 ($UB_{\mathrm{BNP,ESB}}$ in Check Node Iterations) *Suppose the check node degree is d_c and the input $(\mathrm{BNP}, \mathrm{ESB})$ pair is upper bounded by $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$. Then the pair $(\mathrm{BNP}_{out}, \mathrm{ESB}_{out})$ of the check node iteration is bounded by*

$$\begin{aligned} \mathrm{ESB}_{out} &\leq 1 - (1 - DCB_{in})^{d_c - 1} \\ \mathrm{BNP}_{out} &\leq \sum_{i=1}^{d_c-1} \binom{d_c - 1}{i} \sqrt{1 - \left(1 - \left(\frac{\mathrm{ESB}_{in}}{\mathrm{BNP}_{in}}\right)^2\right)^i} \\ &\qquad \left(1 - \frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}}\right)^{d_c - 1 - i} \left(\frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}}\right)^i. \end{aligned} \tag{5.28}$$

Corollary 5.4 *For the check node iteration of any (λ, ρ) irregular LDPC codes, we have*

$$\begin{aligned} \mathrm{ESB}_{out} &\leq 1 - \rho\left(1 - \mathrm{ESB}_{in}\right) \\ \mathrm{BNP}_{out} &\leq \sum_k \rho_k \sum_{i=1}^{k-1} \binom{k-1}{i} \sqrt{1 - \left(1 - \left(\frac{\mathrm{ESB}_{in}}{\mathrm{BNP}_{in}}\right)^2\right)^i} \\ &\qquad \left(1 - \frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}}\right)^{k - 1 - i} \left(\frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}}\right)^i. \end{aligned}$$

Note: By incorporating the ESB_{in} constraint, the BNP bound (5.28) is now tight for both the BEC and BSC cases, which is a strict improvement over the BNP-only bound (5.10). (The bound (5.10) is obtained by linearization and is tight for the BEC case but loose for the BSC case.)

Variable Nodes

We consider a variable node of degree $d_v = 3$. Given that the $(\mathrm{BNP}, \mathrm{ESB})$ values of the constituent channels are upper bounded by $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$ and $(\mathrm{BNP}_{in,2}, \mathrm{ESB}_{in,2})$,

respectively, the iterative upper bounding problem becomes

$$\begin{aligned}
\text{maximize} \quad & \mathrm{BNP}_{out} = \int 2\sqrt{p(1-p)} \cdot 2\sqrt{q(1-q)} dP(p) dQ(q) \\
& \mathrm{ESB}_{out} = \int \frac{4p(1-p)4q(1-q)}{4p(1-p)(1-4q(1-q)) + 4q(1-q)} dP(p) dQ(q) \\
\text{subject to} \quad & \int 2\sqrt{p(1-p)} dP(p) \le \mathrm{BNP}_{in,1} \\
& \int 4p(1-p) dP(p) \le \mathrm{ESB}_{in,1} \\
& \int 2\sqrt{q(1-q)} dQ(q) \le \mathrm{BNP}_{in,2} \\
& \int 4q(1-q) dQ(q) \le \mathrm{ESB}_{in,2}.
\end{aligned} \tag{5.29}$$

By some simple algebra, it can be shown that $\mathrm{BNP}_{out} \le \mathrm{BNP}_{in,1}\mathrm{BNP}_{in,2}$ and the equality can be attained by all feasible solutions $dP(p)$ and $dQ(q)$ meeting the constraints. Unfortunately, for the remaining maximization problem on ESB_{out}, the maximizing distribution dP^* depends on both $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$ and $(\mathrm{BNP}_{in,2}, \mathrm{ESB}_{in,2})$. Simple replacement of each constituent channel with a maximizing counterpart does not work this time. To circumvent this difficulty, we provide an upper bounding distribution dP^{**} depending only on $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$, such that the objective value of any feasible solutions dP and dQ is no larger than the objective value obtained from dP^{**} and dQ. The distinction between the upper bounding distribution dP^{**} and the maximizing distribution dP^* is that dP^{**} may not be feasible and thus may serve merely the bounding purpose.

For simplicity, we express dP^{**} by dropping the subscript 1 in the vector constraint $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$.

$$dP^{**}(p) = \begin{cases} (1 - f_{\mathrm{ESB}}) \frac{t}{t + \mathrm{BNP}_{in}} & \text{if } 2\sqrt{p(1-p)} = \mathrm{BNP}_{in} \\ f_{\mathrm{ESB}} & \text{if } 2\sqrt{p(1-p)} = \sqrt{\mathrm{ESB}_{in}} \\ (1 - f_{\mathrm{ESB}}) \frac{\mathrm{BNP}_{in}}{t + \mathrm{BNP}_{in}} & \text{if } 2\sqrt{p(1-p)} = t \\ 0 & \text{otherwise} \end{cases}, \tag{5.30}$$

where

$$\begin{aligned}
t &= \frac{\mathrm{ESB}_{in}}{\mathrm{BNP}_{in}}, \\
f_{\mathrm{ESB}} &= \begin{cases} 0 & \text{if } 2\sqrt{\mathrm{ESB}_{in}} - t + \sqrt{\mathrm{BNP}_{in}(2t - \mathrm{BNP}_{in})} \ge 0 \\ \frac{\eta(w^*)}{2t(t - \mathrm{BNP}_{in})^2} & \text{otherwise} \end{cases}, \\
\eta(w) &= w^3 - 2tw^2 + (t - \mathrm{BNP}_{in})^2 w, \quad \text{and} \\
w^* &= \begin{cases} 2\sqrt{\mathrm{ESB}_{in}} & \text{if } \eta'(2\sqrt{\mathrm{ESB}_{in}}) \le 0 \\ \frac{2t - \sqrt{4t^2 - 3(t - \mathrm{BNP}_{in})^2}}{3} & \text{otherwise} \end{cases}.
\end{aligned} \tag{5.31}$$

The upper bounding distribution dQ^{**} for the second constituent channel can be obtained by symmetry. A derivation of dP^{**} is included in Appendix G.3. It is worth noting that when there is no constraint on BNP_{in} (namely, when $\mathrm{BNP}_{in} = \sqrt{\mathrm{ESB}_{in}}$ by Lemma 5.1), dP^{**} collapses to a BSC, which coincides with the ESB-based bound in [22]. Hence the upper bound distribution dP^{**} is a strict improvement over the existing ESB-based bound.

Using this upper bounding distribution dP^{**}, an upper bound for $(\mathrm{BNP}, \mathrm{ESB})$ for variable node iterations is given as follows.

Theorem 5.6 ($UB_{\mathrm{BNP,ESB}}$ in Variable Node Iterations) *Suppose the variable node degree is d_v, the input $(\mathrm{BNP}, \mathrm{ESB})$ pair is upper bounded by $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$, and the uncoded channel has noise measures $(\mathrm{BNP}_0, \mathrm{ESB}_0)$. Then, the output of the variable node iteration is upper bounded by*

$$\mathrm{BNP}_{out} \leq \mathrm{BNP}_0 (\mathrm{BNP}_{in})^{d_v - 1},$$

and

$$\mathrm{ESB}_{out} \leq \Phi_{d_v - 1} \left((\mathrm{BNP}_0, \mathrm{ESB}_0), (\mathrm{BNP}_{in}, \mathrm{ESB}_{in}) \right),$$

where Φ_{d_v-1} computes the value of ESB *for a variable node channel with one constituent Ch_0 channel and $(d_v - 1)$ constituent Ch_{in} channels. Here, Ch_{in} and Ch_0 are of the form of dP^{**} and can be uniquely specified by $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$ and $(\mathrm{BNP}_0, \mathrm{ESB}_0)$ respectively.*

Corollary 5.5 *For the variable node iteration of any (λ, ρ) irregular LDPC codes, we have*

$$\begin{aligned} \mathrm{BNP}_{out} &\leq \mathrm{BNP}_0 \lambda (\mathrm{BNP}_{in}) \\ \mathrm{ESB}_{out} &\leq \sum_k \rho_k \Phi_{k-1} \left((\mathrm{BNP}_0, \mathrm{ESB}_0), (\mathrm{BNP}_{in}, \mathrm{ESB}_{in}) \right). \end{aligned}$$

An explicit expression for Φ_{k-1} is given in Appendix G.4, which involves a direct sum of various terms and the complexity of the formula grows at the order of k^3. A more practical, fast implementation is via the fast Fourier transform (FFT), which is similar to that used in density evolution. We first calculate the LLR message distribution $dP(m)$ for Ch_0 and Ch_{in} from the upper bounding distribution $dP^{**}(p)$ in (5.30). Since the output LLR is the summation of input LLRs, the distribution of the output LLR is the convolution of the input LLRs, which can be calculated by FFT. At the end, we can use (5.3) to compute the corresponding output ESB value.

Two-dimensional Iterative Upper Bound $UB_{\mathrm{BNP,ESB}}$

By combining the aforementioned upper bounds for the variable node and the check node iterations, we obtain a two-dimensional iterative upper bound $UB_{\mathrm{BNP,ESB}}$. Since this two-dimensional bound is based on separate analysis of variable nodes and parity check nodes, it can be applied to any LDPC-like codes with graph-based ensembles without modification, including regular/irregular repeat accumulate (RA) codes [52], and joint-edge-distribution LDPC codes [55].

We omit the explicit expression for this two-dimensional bound since it is a direct concatenation of Theorems 5.5 and 5.6. By iteratively computing the upper bound $(\mathrm{BNP}^{(l)}, \mathrm{ESB}^{(l)})$ and testing whether it converges to $(0, 0)$, we can lower bound the decodable threshold for general BI-SO channels. The performance comparison of this procedure to existing results will be discussed in Section 5.6

5.5 A One-Dimensional Non-Iterative Bound on BI-SO Channels

In this section, we construct a non-iterative upper bound for BI-SO channels, which is the best known bound that is tight for BSCs.

First we introduce some new notation. We first recall that $p_e^{(l)}$ denote the bit error probability of the belief propagation after l iterations. To distinguish between the types of BI-SO channels on which we are focusing, we append an argument F_C to the end of $p_e^{(l)}$. That is, $p_e^{(l)}(F_C)$ denotes the bit error probability after l iterations with the conditional distribution of the BI-SO channel being $F_C(dy|x)$. In a similar fashion, we define $\mathrm{BNP}^{(l)}(F_C)$ as the Bhattacharyya noise parameter after l iterations with the BI-SO channel being F_C, and $\mathrm{ESB}^{(l)}(F_C)$ is defined similarly. Following this definition, $\mathrm{BNP}(F_C) := \mathrm{BNP}^{(0)}(F_C)$ denotes the Bhattacharyya noise parameter of the uncoded BI-SO channel F_C. For simplicity, we use $F_{BSC,p}$ to denote the F_C of a BSC with crossover probability p, and similarly we define $F_{BEC,\epsilon}$. Suppose for some $F_{BSC,\tilde{p}}$, the LDPC code ensemble is decodable. By the channel degradation argument, one can show that all BI-SO channels F_C with $p_e^{(0)}(F_C) \leq \tilde{p}$ are guaranteed to be decodable. This result, though being tight for BSCs, generally gives a very loose bound for other channels. We strengthen this result by the following theorem.

Theorem 5.7 *Suppose F_C is a BI-SO channel and $F_{BSC,\tilde{p}}$ is a BSC. If* $\mathrm{ESB}(F_C) = \mathrm{ESB}(F_{BSC,\tilde{p}})$*, then for any $l \in \mathbb{N}$ and any irregular (λ, ρ) LDPC codes,*

$$\mathrm{BNP}^{(l)}(F_{BSC,\tilde{p}}) \geq \mathrm{BNP}^{(l)}(F_C).$$

In Theorem 5.7, it is possible that $p_e^{(l)}(F_{BSC,\tilde{p}}) < p_e^{(l)}(F_C)$. This phenomenon shows that Theorem 5.7 is stronger than the existing result and cannot be derived from the channel degradation argument. One corollary to the above theorem is stated as follows.

Corollary 5.6 *If a (λ, ρ) irregular LDPC code is decodable for a BSC with crossover probability p^*, then any BI-SO channel F_C with* $\mathrm{ESB}(F_C) \leq \mathrm{ESB}(F_{BSC,p^*}) = 4p^*(1-p^*)$ *is decodable under the same (λ, ρ) code.*

Proof: For any symmetric channel F_C with $\mathrm{ESB}(F_C) \leq \mathrm{ESB}(F_{BSC,p^*})$, we consider a $F_{BSC,\tilde{p}}$ such that $\mathrm{ESB}(F_{BSC,\tilde{p}}) = \mathrm{ESB}(F_C)$. Since F_{BSC,p^*} is physically degraded w.r.t. $F_{BSC,\tilde{p}}$, $F_{BSC,\tilde{p}}$ is also decodable, namely, $\lim_{l\to\infty} p_e^{(l)}(F_{BSC,\tilde{p}}) = 0$. By the relationship between p_e and BNP in Lemma 5.1 and by Theorem 5.7, we have

$$2p_e^{(l)}(F_C) \leq \mathrm{BNP}^{(l)}(F_C) \leq \mathrm{BNP}^{(l)}(F_{BSC,\tilde{p}}) \leq 2\sqrt{p_e^{(l)}(F_{BSC,\tilde{p}})} = o(1).$$

This completes the proof. ■

Corollary 5.6 can be used as a tight one-dimensional upper bound, which is denoted by UB^*_{ESB}.

A proof of Theorem 5.7 is given in Appendix G.5.

Comparisons of lower bounds on decodable thresholds.

	DE	UB_{ESB}	UB_{BNP}	UB_{info}	$UB_{\text{BNP,ESB}}$	$UB^*_{p_e}$	UB^*_{ESB}
Decodable Thresholds	—	$\text{ESB}^* \geq 0.2632$	$\text{BNP}^* \geq 0.4294$	$h^* \geq 0.3644$	—	$p_e \geq 0.0837$	$\text{ESB}^* \geq 0.3068$
BEC (ϵ^*)	≈ 0.4294	≥ 0.2632	≥ 0.4294	≥ 0.3644	≥ 0.4294	≥ 0.0837	≥ 0.3068
Rayleigh (σ^*)	≈ 0.644	≥ 0.5191	≥ 0.6134	≥ 0.6088	≥ 0.6148	≥ 0.471	≥ 0.5804
Z-channels (p^*_{10})	≈ 0.2304	—	≥ 0.1844	—	—	≥ 0.1674	—
BiAWGNC (σ^*)	≈ 0.8790	≥ 0.7460	≥ 0.7690	≥ 0.8018	≥ 0.7826	≥ 0.7244	≥ 0.8001
BiLC(λ^*)	≈ 0.65	≥ 0.5610	≥ 0.5221	≥ 0.5864	≥ 0.5670	≥ 0.5595	≥ 0.6146
BSC (p^*)	≈ 0.0837	≥ 0.0708	≥ 0.0484	≥ 0.0696	≥ 0.0710	≥ 0.0837	≥ 0.0837

Table 5.1: Comparison of various lower bounds derived from finite-dimensional iterative upper bounds.

5.6 Performance Comparisons

In this section, we compare the tightness of various lower bounds on the asymptotic decodable thresholds, obtained from the existing results and our results of Sections 5.4.1, 5.4.2, and 5.5.

Three existing results are included in Table 5.1, including one based on the Bhattacharyya noise parameter [59], denoted as UB_{BNP}, one on the expected soft bit value [22], denoted as UB_{ESB}, and one on the conditional entropy $H(X|Y)$ [69, 103], denoted as UB_{info}. $UB_{\text{BNP,ESB}}$ denotes the two-dimensional (BNP, ESB)-based bound provided in Section 5.4.2, and UB^*_{ESB} denotes the non-iterative tight bound given in Section 5.5. The DE column lists the asymptotic decodable thresholds computed from density evolution [92]. In Section 5.4.1, UB_{BNP} has been generalized for arbitrary BI-NSO channels. Therefore, the non-symmetric z-channel is also included for comparison, for which the asymptotic threshold is obtained from the generalized density evolution method for BI-NSO channels in Chapter 4.

As proved in Section 5.4.2 and evidenced in Table 5.1, the two dimensional bound $UB_{\text{BNP,ESB}}$ provides strict improvement over existing UB_{BNP} and UB_{ESB}. For channels that are neither BSC-like nor BEC-like, e.g. BiAWGNC and BiLC, the bound UB_{info} found by Sutskover *et al.*, is tighter than $UB_{\text{BNP,ESB}}$ while the two dimensional $UB_{\text{BNP,ESB}}$ is tighter at both extreme ends. This phenomenon can be explained by the convexity/concavity analysis of the transfer functions. For UB_{BNP} the bounding inequality resides in the check node iteration, in which BECs attain the equality. Therefore, UB_{BNP} is the tightest when channels are BEC-like. For UB_{ESB}, the bounding inequality resides in the variable node iteration, in which BSCs attain the equality, so UB_{ESB} is preferred for BSC-like channels. Since $UB_{\text{BNP,ESB}}$ is strictly tighter than both UB_{BNP} and UB_{ESB}, it has better performance in both extreme cases. On the other hand, UB_{info} invokes bounding inequalities in both the variable node and the check node iterations. Since the absolute values of the curvatures of the transfer function is smaller when changing the underlying variables to the mutual information, better predictability is obtained when the channel of interest is neither BSC nor BEC-like, e.g., the BiAWGN channel.

By Lemma 5.1, the feasible (BNP, ESB) pairs satisfy $\text{BNP} \geq \text{ESB}$ and $(\text{BNP})^2 \leq \text{ESB}$. By plotting general BI-SO channels according to their (BNP, ESB) values, the set of decodable channels forms a "decodable region" and Figure 5.6 demonstrates the decodable region of regular $(3, 6)$ codes. The thin, gray boundary corresponds to the suggested decodable region, which is plotted using the channels considered in Table 5.1. The density evolution method does not guarantee that all channels with (BNP, ESB) inside the region are decod-

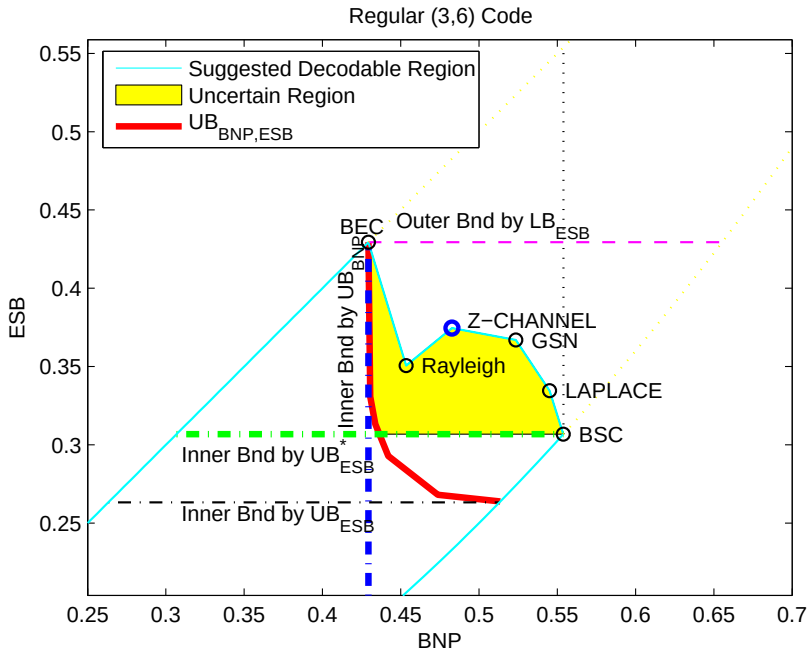

Figure 5.6: The decodable region of the regular $(3, 6)$ code in the $(\mathrm{BNP}, \mathrm{ESB})$ domain and several inner bounds of the decodable region.

able. It is possible that two types of channels have the same $(\mathrm{BNP}, \mathrm{ESB})$ values but one is decodable while the other is not.

The vertical line in Figure 5.6 marked by UB_{BNP} represents the inner bound of the decodable BNP* threshold [59]. The horizontal line marked by UB_{ESB} represents the inner bound of the decodable ESB* threshold. Our results on the two-dimensional bound and the non-iterative tight bound greatly push the inner bounds of the decodable region toward its boundary (the curve marked $UB_{\mathrm{BNP,ESB}}$ and the horizontal line marked by UB^*_{ESB}). These bounds guarantee that all BI-SO channels with $(\mathrm{BNP}, \mathrm{ESB})$ within the inner bounds are decodable under belief propagation decoding. Therefore, the uncertain region has been suppressed to the yellow area. In Section 5.4.1, we have also shown that the vertical line UB_{BNP} holds as an inner bound even for BI-NSO channels.

Burshtein *et al.* [22] proved a tight outer bound of the decodable region, the horizontal line marked by LB_{ESB}. Based on the mathematical symmetry between BNP and ESB in variable node and check node iterations, we conjecture the existence of a tight outer bound in terms of BNP, which remains an open problem.

5.7 Summary

Finite dimensional bounds on the decodable thresholds find applications in both theoretical analysis and practical approximations. In this chapter, we have modelled the iterative bounding problem by considering its probabilistic decomposition, which enables systematic searches for more finite-dimensional bounds as (analytically) solving an infinite dimensional linear programming problem. Based on this new framework, we have developed a new iterative upper bound for $\mathbb{Z}_m$-based LDPC codes on MI-SO channels, which leads to a

sufficient stability condition and provides insight into the analytical structure of $\mathbb{Z}_m$ LDPC codes. Combined with a matching necessary stability condition proved herein, our stability condition pair can be used to reproduce the existing stability conditions for $\mathsf{GF}(q)$-based LDPC codes.

Two new bounds for binary codes on BI-SO channels have also been constructed with focus on two types of noise measures, the Bhattacharyya noise parameter (BNP) and the expected soft bit value (ESB). These bounds push the existing inner bounds of the decodable region toward its boundary. We have further generalized one existing BNP-based bound to BI-NSO channels, which is the first bound for BI-NSO channels and demonstrates the strong connection between the BNP and the iterative BP decoding for general memoryless channels. The performance discrepancy among various bounds can be explained by the tightness of different bounds during the variable node and the check node iterations. Besides the implied uniform good performance over all types of channels, these new finite dimensional bounds and the proposed framework provide a useful tool for studying the behavior of iterative decoding. For instance, until present, all existing results on sufficient stability conditions for iterative BP decoding are derived by the infinitesimal analysis on deliberately constructed finite dimensional bounds.

Chapter 6

$\mathbb{Z}_m$ LDPC Coded Modulation

In wireless communications, transmitting high data rate bit strings through the bandwidth limited environment has become a critical and challenging topic for all designers of protocols for physical layers. Two parallel approaches have been considered in the communications community. One is to limit the interference from other bandwidth users, including but not limited to multi-user detection, orthogonal frequency division multiplexing, and ultra wide band schemes. The other approach is to use high order signal constellations to compress more information bits into one symbol usage. Modern systems are usually based on a mixture of these two methods and both approaches require sophisticated signal processing to fully unlock the hidden bandwidth efficiency. $\mathbb{Z}_m$ LDPC coded modulation belongs to the second method, and we will demonstrate that using high order LDPC codes allows full utilization of the high order constellation, and close-to-capacity performance can be achieved.

6.1 System Design

The basic idea of combining coding and modulation can be dated back to mid 1970's, and Ungerboeck later proposed trellis coded modulation (TCM) based on trellis codes and the set partition mapping [107]. Since then, many other bandwidth efficient schemes have been provided, including multi-level codes (MLCs) with parallel/multi-stage decoding (PID/MSD) [108], and bit-interleaved coded modulation (BICM) [27]. Recently, these schemes have been used with ultra-powerful turbo codes and LDPC codes, including turbo TCM [95], turbo-code-based BICM [11], LDPC-code-based MLC, and LDPC-code-based BICM [51, 83].

The major advantage of BICM is its simplicity, although BICM suffers capacity loss due to breaking down the coded modulation symbol-wise channel into bit-level sub-channels [27]. This drawback was offset by its computational efficiency when used with traditional convolutional codes or Reed-Solomon codes. However, when the goal is to approach the capacity with the help of ultra powerful error correcting codes, this inherent disadvantage becomes more noticeable.

The performance difference between MLC with PID and MLC with MSD is that the former suffers from the same capacity loss as in BICM while the latter is able to recover the capacity loss by staged decoding. Unfortunately, the multi-stage decoding (MSD) inevitably induces longer delay, which cannot be overcome by using parallel computing hardware. The remaining choice, turbo TCM, does not have any capacity loss or delay problem.

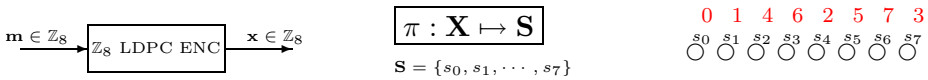

$\mathbb{Z}_m$ LDPC encoder symbol-to-signal mapper constellation points

Figure 6.1: System diagram of $\mathbb{Z}_m$ coded modulation with 8PAM constellation

However, it is generally more challenging to devise turbo codes with extremely close-to-capacity performance, compared to designing ultra powerful LDPC codes.

High order LDPC codes generally have better performance compared to binary codes when considering the same number of information bits (as mentioned in Section 5.3.4). We therefore can directly combine high order LDPC codes with the constellation set as in Figure 6.1, in which a $\pi : \mathbb{Z}_m \mapsto \mathbb{Z}_m$ symbol mapper bridges the physical signal domain with the logical symbols used in the iterative decoding. The advantage of this system is two-fold: the additional error correcting power of $\mathbb{Z}_m$ is exploited and no capacity loss is inherited in this type of systems. The major drawback of this system is the decoding complexity, which is $\mathcal{O}(m)$ times the binary-code-based system. Fortunately, there is no staged structure and parallel computing is admissible.

Decoding of this system can be obtained by first converting the received signals to the *a posteriori* probability. Then the BP decoding is performed by assuming cycle-free structure. Each message is an m-vector, and the message initialization and message maps for one-dimensional BiAWGNCs are as follows.

$$\begin{aligned}
\mathbf{m}_0 &= \propto \left(e^{-\frac{(y-\pi(0))^2}{2\sigma^2}}, e^{-\frac{(y-\pi(1))^2}{2\sigma^2}}, \cdots, e^{-\frac{(y-\pi(m-1))^2}{2\sigma^2}} \right) \\
\Psi_v &= \propto \left(\mathbf{m}_0 \prod_{j=1}^{d_v-1} \mathbf{m}_j \right) \\
\Psi_c &= \mathcal{R} \left(\bigotimes_{i=1}^{d_c-1} \mathbf{m}_i \right),
\end{aligned}$$

where $\pi(0), \cdots, \pi(m-1)$ represent the signal points after the symbol mapper, $\propto$ is a normalization operator so that the sum of all components is one, and $\mathcal{R}$ is a flipping operator so that $(x_0, x_1, x_2, \cdots, x_{m-1})$ becomes $(x_0, x_{m-1}, x_{m-2}, \cdots, x_1)$. $\prod$ and $\bigotimes$ represent the component-wise product and the circular convolution as discussed in Section 5.3.2. $\mathbf{m}_1$ through $\mathbf{m}_{d_v-1}$ (or $\mathbf{m}_{d_c-1}$) represent the incoming $d_v - 1$ (or $d_c - 1$) messages coming to the variable node (or the check node), excluding the one coming from the destination node. After many rounds of message passing, the decision is made by the following formula:

$$\hat{x} := \arg\max_x \left\{ \text{the } x\text{-th component of } \mathbf{m}_0 \prod_{j=1}^{d_v} \mathbf{m}_j \right\}.$$

In general, this system can be viewed as a MI-NSO channel, for which the pairwise Bhattacharyya noise parameters $\mathrm{BNP}(x \to x')$ can be computed. The symbol mapper π shuffles the pairwise error pattern by the following equation:

$$\mathrm{BNP}_\pi(\pi(x) \to \pi(x')) \leftarrow \mathrm{BNP}(x \to x').$$

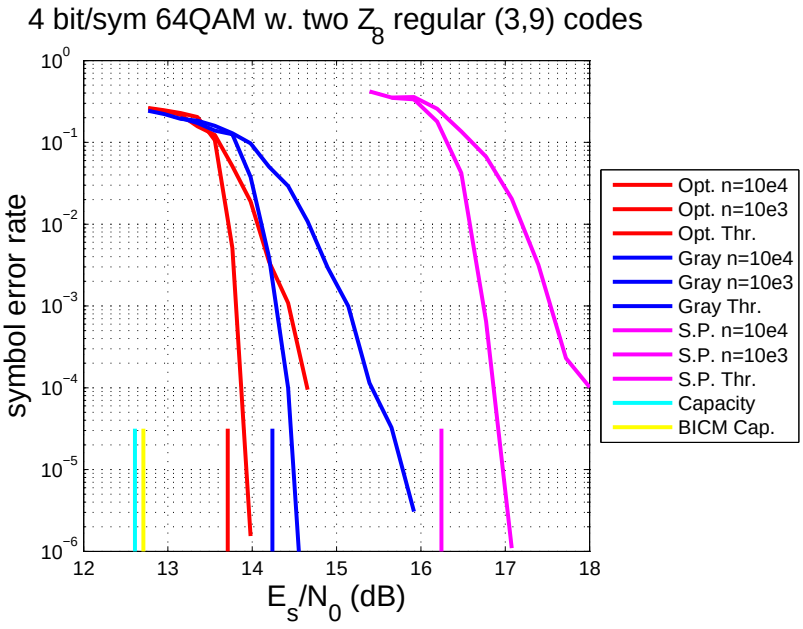

Figure 6.2: Performance comparison of using two regular (3,9) codes on 64QAM with different symbol mappers

Suggested by the typicality results in Section 4.6.1, the performance for this new MI-NSO channel will be extremely close to the performance for its symmetrized version in Figure 5.1. And the pairwise $\mathrm{BNP}_{\pi,sym}(0 \to x')$ for the symmetrized channel becomes

$$\mathrm{BNP}_{\pi,sym}(0 \to x') = \frac{\sum_{x \in \mathbb{Z}_m} \mathrm{BNP}_{\pi}(x \to x + x')}{m}.$$

In Section 5.3.3, it has been shown that the necessary and sufficient stability conditions for $\mathbb{Z}_m$ codes depend on the maximum of the Bhattacharyya noise parameter $\max_x \mathrm{BNP}(0 \to x)$ for all MI-SO channels. Since different choices of $\pi(\cdot)$ will result in different values of $\mathsf{BNP}_{\pi,sym}$, the above stability result suggests selecting a symbol mapper $\pi(\cdot)$ such that $\max_x \mathrm{BNP}_{\pi,sym}(0 \to x)$ is kept minimal. In the next section, we will use simulation to demonstrate that a significant amount of improvement can be obtained by selecting a good symbol mapper. It is worth noting that the symbol mapper $\pi(\cdot)$ is realized by permutation. Therefore selecting the optimal symbol mapper induces no additional computational/hardware cost for the entire system.

6.2 Simulation

The major simulation results are summarized in Figure 6.2, in which we consider applying two $\mathbb{Z}_8$ codes on the in-phase and quadrature components of a 64QAM system. Since regular (3,9) codes have rate 2/3, each coded symbol corresponds to $\log_2(64) \cdot 2/3 = 4$ information bits. The codeword length equals 10^3 or 10^4 symbols, or equivalently, 6×10^3 or 6×10^4 coded bits. We first picked up code instances from the random code ensembles $\mathcal{C}^{1000}(3, 9)$ and $\mathcal{C}^{10000}(3, 9)$ respectively, which were fixed and used throughout the simulation. As suggested in Section 4.6.1, we perform simulation on a symmetrized channel, as in Figure 4.11(c), so that we can assume the all-zero codeword is transmitted and do not need to implement

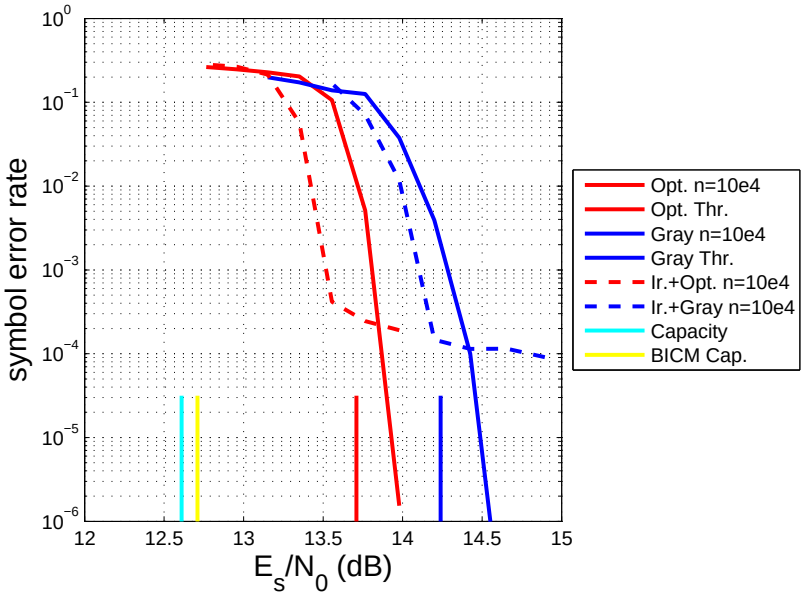

Figure 6.3: Comparison between regular codes and irregular codes optimized for BECs.

the encoding part. 150 iterations of BP are performed, after which we check whether the all-zero codeword is successfully decoded. For each signal to noise ratio E_s/N_0, the simulation proceeds until 80 block errors occur, and the averaged symbol error rate is computed accordingly.

The "Opt." curves correspond to the optimized[1] symbol mapper, $\{0, 1, 4, 6, 2, 5, 7, 3\}$, specified also in Figure 6.1. The Gray mapping is $\{0, 1, 3, 2, 6, 7, 5, 4\}$ and the set partition mapping, S.P., happens to be the natural mapping $\{0, 1, 2, 3, 4, 5, 6, 7\}$. The thresholds are obtained by Monte-Carlo-based density evolution with 40 iterations. The capacity for 64QAM and for 64QAM with BICM plus the Gray mapping are also plotted for comparison.

Our optimized symbol mapper provides a significant 0.5dB gain over the apparent choice: the Gray mapping, and 2.9dB gain over the natural set partition mapping. This result shows that the "optimal" Gray mapping used in BICM is not good when considering high order codes. If we are interested in symbol error rate 10^{-3} with codeword length 10^4 symbols, regular $\mathbb{Z}_8$ codes plus 64QAM achieve thresholds 1.1dB away from capacity, which is roughly 0.3dB better than BICM with regular binary codes and 0.3dB worse than BICM with irregular binary codes [51].

Since we consider only regular codes here, there is still room of improvement using irregular codes. Due to the inefficiency of Monte-Carlo-based density evolution, we use the approximation technique that in the high SNR regime, the behavior of BiAWGNCs is similar to that of BECs. Therefore, we use the off-the-shelf degree distributions optimized for BECs[2] for the proposed $\mathbb{Z}_8$ coded modulation. Detailed information about this optimized

[1]Besides the fact that a near-random symbol mapper would be the most stable, the exact relationship between the symbol mapper and the code performance is still an open problem. The specified symbol mapper is optimized by exhaustively searching over all 8! possible mappers with Monte-Carlo-based density evolution.

[2]This code degree distribution is generated by the online LDPC degree optimization tool: http://lthcwww.epfl.ch/research/ldpcopt/. The additional optimization constraints are rate = 0.67, $\max d_v \leq 12$, and $\rho(x) = x^{10}$. The code reference number is 2311.

code is as follows.

$$\begin{aligned}\lambda(x) &= 0.310801x + 0.113333x^2 + 0.173092x^3 + 0.0468565x^4 + 0.355918x^{11}\\ \rho(x) &= x^{10}\\ \text{rate} &= 0.67\\ \epsilon^* &= 0.3257.\end{aligned}$$

With the optimized mapper/the Gray mapping, the performance of the predetermined irregular code is illustrated in Figure 6.3. Another 0.3dB is obtained when combined with the optimized symbol mapper. Compared to using the same irregular code with the Gray mapping, most performance gain is obtained by the cost-free optimization over symbol mappers, while the degree optimization induces more complexity and worse error floor with less threshold improvement.

6.3 Discussion & Summary

Prompted by the stability conditions, great improvement has been shown by selecting a proper symbol mapper for high order LDPC coded modulation. With zero cost, good symbol mappers contribute to even bigger performance gain than considering irregular codes. With no inherent capacity loss (over BICM) and no extra delay (over MLC/MSD), $\mathbb{Z}_m$ LDPC coded modulation becomes a competitive choice especially with the progress of semi-conductor technologies. With even higher order coded modulation, say 32PAM, we can apply high order LDPC codes only to the most correlated components. For example, we can use a $\mathbb{Z}_{32}$ code for 32PAM, or use one $\mathbb{Z}_8$ and one $\mathbb{Z}_4$ to decoupled the correlated signals. Although the latter combination suffers from capacity loss, this however will be much less then the capacity loss of BICM systems. Using high order codes, the designer is able to flexibly balance decoding complexity and the code performance for general communication systems.

The reason why the high order coded modulation is overlooked can probably be explained by the paper of Caire *et al.* [27], in which it has been show that the cutoff rate $R_{0,\mathbb{Z}_m}$ for $\mathbb{Z}_m$ codes is smaller than the cutoff rate $R_{0,bin}$ for BICM plus the Gray mapping. Taking the $\mathbb{Z}_8$ codes as example, it has been shown that

$$\begin{aligned} & R_{0,\mathbb{Z}_8} = \log_2 8 - \log_2\left(1 + \sum_{x\in\mathbb{Z}_8\setminus\{0\}} \mathrm{BNP}(0\to x)\right)\\ < \;& R_{0,bin} = R_{0,b1} + R_{0,b2} + R_{0,b3} = \sum_{j=1,2,3}\left(\log_2 2 - \log_2\left(1+\mathrm{BNP}(F_{bj})\right)\right),\end{aligned}$$

where $\mathrm{BNP}(F_{bj})$ is the BNP value with respect to the bit-wise sub-channel F_{bj} for the j-th bit resulted from the Gray mapping. Our stability results in Section 5.3.3 show that the performance depends mainly on the maximum of $\mathrm{BNP}(0\to x)$ instead of the summation of $\mathrm{BNP}(0\to x)$. This explains why the cutoff rate $R_{0,\mathbb{Z}_m}$ is no longer a good performance measure of $\mathbb{Z}_m$ LDPC codes when $m \geq 3$. (When $m = 2$, $R_{0,bin}$ can be mapped bijectively to the BNP value so that $R_{0,bin}$ is highly related to the code performance as shown in our discussion of BNP-based bounds in Section 5.4.) Furthermore, the origin of the cutoff rate can be traced back to the Gallager bound [97], which is a union bound based on pairwise symbol error events. When moving to higher order alphabets, this bounding technique

becomes more and more pessimistic and overestimates the error probability. Therefore, the cutoff rate should not be used to compare the performance of codes with different alphabet sizes. Even for codes having the same alphabet, the performance prediction is less accurate for high order LDPC codes, since the performance depends more on $\max_x \mathrm{BNP}(0 \to x)$ than on $\sum_x \mathrm{BNP}(0 \to x)$.

Besides having superior performance compared to binary-code-based systems, another advantage of the high order coded modulation is the compatibility with most equalization systems, since the high order modulation is directly considered as a channel with high order input alphabet. For instance, a matched spectral null code can be concatenated as an inner code to improve the performance when used in inter-symbol interference (ISI) environments [57]. Joint iterative decoding and equalization (also known as turbo equalization) can be performed based on the soft symbol detection.

Chapter 7

Conclusion and Future Work

Bandit Problems

In the first part of this thesis, we have considered the two-armed bandit problem with the help of i.i.d. side information. We take a relatively conservative setting by assuming no *a priori* distribution on the possible configuration parameters (θ_1, θ_2) and by focusing on minimizing the growth rate of the regret with a uniform discount sequence. Our results show that different values of side information x can be viewed as the indices of many sub-bandit machines, where nature decides which sub-bandit machine is accessible at the current time t, revealing the index to the decision maker by the side information random variable $X_t = x$.

Under this framework, the bandit problem can be viewed as a two-player zero-sum game such that various degrees of improvement can be obtained, including achieving bounded regret, or reducing the constant in front of the traditional $\log(t)$ lower bound, depending on whether the underlying configuration is implicitly revealing or not. We then extract the essential even distribution properties of a beneficial side information random process. Using coin-flipping as an example, our results show that fairness of the coin (the side information) is more important to the decision maker than the randomness of the coin. All our previous results of performance improvement apply to a very broad class of side information, including i.i.d. sequences, Markov chains of any finite order and deterministic periodic sequences as special cases. Relationships between different levels of even distribution properties have been discussed therein.

The Gittins index, on the other hand, takes a more aggressive approach, assuming an *a priori* distribution on (θ_1, θ_2) and focusing on the total expected reward with geometric discount sequences. Due to the aggressive nature of this setting, it is more suitable for financial applications, e.g. the early exercise strategy of the American put options, while our conservative setting is more appropriate to clinical trials. Our future research will focus on exploring the interaction between the side information and the Gittins' index. Whether the game theoretic perspective still applies under different settings is a valid question of substantial theoretical and practical value.

LDPC Codes

The second part of this thesis has focused on LDPC codes with applications to non-symmetric channels. The major difficulty in analyzing LDPC codes for non-symmetric channels is due to the codeword dependent error resiliency. We have overcome this difficulty by computing the tree-structure averaged density evolution, and prove that the

tree-structure averaged density evolution is indeed depicting the codeword-averaged performance by the perfect projection convergence theorem. Using our new density evolution formula, we further prove the typicality of linear codes among the LDPC coset code ensemble when the minimum check node degree is sufficiently large. Besides being the foundation of the tree-structure averaged density evolution, the perfect projection condition further guarantees the optimality of the belief propagation decoder when only local observation is available. We also have provided a simple counter example showing that assuming that the girth of the graph is sufficiently large cannot guarantee this local optimality of belief propagation decoders. In other words, using the perfect projection convergence theorem, we have shown that for the first few iterations, no other distributed decoding algorithm, or equivalently, no other message passing algorithm, is able to outperform the belief propagation decoder, which was designed for cycle-free networks rather than LDPC codes with many cycles. Our results serve as both generalization and completion of the classical density evolution theorems.

The analysis of arbitrary (non-symmetric) memoryless channels is well established by our codeword averaged approach. Furthermore, for other non-standard channels such as intersymbol interference channels, broadcast channels, and multiple access channels, the performance is in general codeword-dependent. This codeword dependence hampers the analysis even though the locally tree-like structure still exists due to the use of long random interleavers. Our approach can easily be generalized for those situations once the corresponding perfect projection convergence theorem is proved. The perfect projection condition identifies the necessary and sufficient link between tree-based density evolution and the average of the codeword-dependent performance.

All our analyses are asymptotic, assuming sufficiently large codeword length n. Although the performance when $n \geq 10^4$ can be well predicted by the asymptotic thresholds, certain code properties like error floors cannot be predicted by the asymptotic approach. Until present, the most successful finite length analysis is the stopping/trapping set analysis, which identifies the "bad" channel realizations deteriorating the finite code performance of iterative decoders. Since it is impossible to enumerate all bad configurations, this technique focuses on providing a performance lower bound (on the word-error rate), of which the tightness is confirmed by empirical data. Complementing the existing results, a rigorous upper bound is especially important when we are interested in extremely low bit-error-rate (BER) performance which can not be verified by Monte-Carlo simulations nor by empirical data.

Motivated by the iterative bounding techniques in Chapter 5, we are interested in developing semi-asymptotic results by quantifying the correlation among input messages for finite code analysis. Either a density evolution method or a finite-dimensional bound for finite codes will be very beneficial, which can serve as the performance metric for optimizing the code/graph structure in the low-BER regime. Figure 7.1 consists of an example of a simple finite-dimensional bound for a finite code of length 4 bits on a BEC with erasure probability ϵ. For a code with its Tanner graph as in Figure 7.1(a), a bit-level ML decoder for bit 1 achieves BER $p_{e,1} = 2\epsilon^3 - \epsilon^4$ and a BP decoder has $p_{e,1} = 3\epsilon^3 - 2\epsilon^4$, which demonstrates the performance discrepancy between ML and BP decoders. On the other hand, the asymptotic DE analysis, as in Figure 7.1(b), assumes that the Tanner graph can be "expanded as a tree" and the channel observations of all nodes are independent. The BER

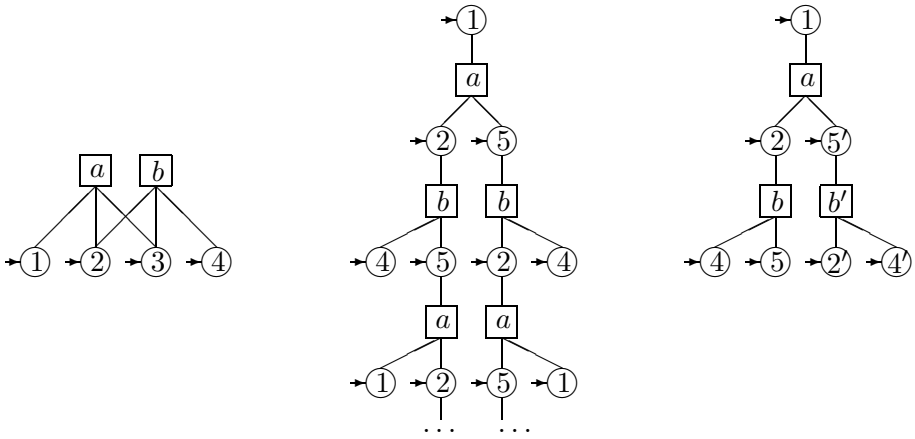

ML: $p_{e,1} = 2\epsilon^3 - \epsilon^4$
BP: $p_{e,1} = 3\epsilon^3 - 2\epsilon^4$

DE: $p_{e,1} = 2\epsilon^3 + o(\epsilon^3)$

Upper bound of BP:
$p_{e,1} \leq 4\epsilon^3 - 2\epsilon^4$

(a) The finite code (b) Asymptotic tree (c) Trimmed tree

Figure 7.1: An example of tree-based upper bounds of BP for finite codes.

predicted by DE[1] is then

$$p_{e,1} = \frac{2\epsilon^3 - 2\epsilon^4 + \epsilon^5}{(1 - \epsilon + \epsilon^2)^2} = 2\epsilon^3 + o(\epsilon^3).$$

For finite codes, the BER predicted by DE is in general much smaller than that of BP and provides little insight to the low-BER performance. To rigorously upper bound the BER of BP, we take the tree-based approach with the following simple trimming rules:

1. A variable node u is an ancestor of a variable node v if u is on the path from v to the root of the tree.

2. If a single message (in the finite Tanner graph) enters both variable nodes u and v, and u is an ancestor of v, then the message enters v can be replaced by an erasure.

3. If two input messages are dependent, we can replace one of them by a completely negatively correlated message. Namely, the probability that both messages are erasure is kept minimal.

4. Once the original infinite tree has been trimmed according to the above rules, we can compute the BER of the resulting finite tree, which is an upper bound of the BER of BP.

In Figure 7.1(a), the initial message of variable node 1 enters three different nodes in Figure 7.1(b). According to the second rule, we can replace the messages entering node 1's in the bottom level by erasures. For each of the two check node a's, there is at least one erasure message. Therefore we can drop both check node a's, resulting the trimmed structure as in Figure 7.1(c). The messages $2 \to a$ and $5 \to a$ in the top level are dependent, since there are overlapped descendent nodes in the corresponding sub-trees. By the third

[1]The DE analysis in this example corresponds to an analysis on a multi-edge-type ensemble [91] instead of a irregular graph ensemble.

rule, we can replace the nodes on the right sub-tree by nodes $5'$, $2'$ and $4'$ so that the messages $2 \to a$ and $5' \to a$ become completely negatively correlated. The BER of the trimmed tree is then $p_{e,1} \leq 4\epsilon^3 - 2\epsilon^4$, which upper bounds the BER of BP.

More sophisticated rules are required for longer, more complicated codes, and the rules have to be efficiently checked since the support tree grows exponentially fast. Some of our preliminary results show that the structure of the trimmed tree is closely related to the subgraph induced by the "stopping set" used in lower bounding the block error probability. The connection between the message correlation argument, a symbol-based approach, and the stopping/trapping set analysis, a block-based analysis, is well worth investigation, and the result would be helpful in designing efficient trimming rules.

This thesis has also discussed high order LDPC codes, with which a close-to-capacity high order coded modulation scheme has been proposed. However, the benefit of high order codes is not fully utilized due to the lack of efficient, non-Monte-Carlo-based, density evolution for degree optimization. A continuing next step of our research on $\mathbb{Z}_m$ coded modulation is to devise a fast density evolution method for $\mathbb{Z}_m$ codes. Once the code degree is optimized according to this new tool, we are expecting extremely close-to-capacity performance as in the binary case.

Besides high order LDPC codes, another very important example of codes with high order alphabets is the algebraic Reed-Solomon code, which is widely used in modern commercial products due to its guaranteed performance and simple/ultra-efficient algebraic decoding method. Recent results show that using soft iterative decoders instead of algebraic decoders, the performance gain can be as high as 2dB for BiAWGNCs. Future research should be conducted either on pushing the performance gain or on enhancing the efficiency of the soft algorithms. Our results on the iterative bounds and stability conditions for $\mathbb{Z}_m$ LDPC codes will be a good starting point toward this ambitious goal.

Appendix A

Sanov's Theorem and the Prohorov Metric

For two distributions P and Q on the reals, the Prohorov metric is defined as follows.

Definition A.1 (The Prohorov metric) *For any closed set $A \subset \mathbb{R}$ and $\epsilon > 0$, define A^ϵ, the ϵ-flattening of A, as*

$$A^\epsilon := \left\{ x \in \mathbb{R} : \inf_{y \in A} |x - y| < \epsilon \right\}.$$

The Prohorov metric ρ is then defined as follows.

$$\rho(P, Q) := \inf \left\{ \epsilon > 0 : P(A) \leq Q(A^\epsilon) + \epsilon, \textit{for all closed } A \subset \mathbb{R}. \right\}.$$

The Prohorov metric generates the topology corresponding to convergence in distribution. Throughout Chapters 2 and 3, the open/closed sets on the space of distributions are thus defined accordingly.

Theorem A.1 (Sanov's theorem) *Let $L_X(n)$ denote the empirical measure of the real-valued i.i.d. random variables $X_1, X_2, \cdots, X_n$. Suppose X_i is of distribution P and consider any open set A and closed set B from the topological space of distributions generated by the Prohorov metric. We have*

$$\begin{aligned} \liminf_{n \to \infty} \frac{1}{n} \log \left(\mathsf{P}_P \left(L_X(n) \in A \right) \right) &\geq -\inf_{Q \in A} I(Q, P) \\ \limsup_{n \to \infty} \frac{1}{n} \log \left(\mathsf{P}_P \left(L_X(n) \in B \right) \right) &\leq -\inf_{Q \in B} I(Q, P), \end{aligned}$$

where $I(Q, P) := \mathsf{E}_Q \left\{ \log \left(\frac{dQ}{dP} \right) \right\}$ is the Kullback-Leibler information number.

The above is not the most general form of Sanov's theorem, but is sufficient for this thesis. Further discussion of the Prohorov metric and Sanov's theorem can be found in [21, 35].

Appendix B

Proofs of the New $\log(t)$ Lower Bounds

In this appendix, we are going to prove all our results of $\log(t)$ lower bounds on $\mathsf{E}\{T_{inf}(t)\}$ for bandit problems with *i.i.d./arbitrary* side information X_t, which are stated in Sections 2.1.1, 2.4.1, 2.5.1, 3.3 and 3.5.1.

B.1 Proof of Theorems 2.5 and 3.3

The proof is inspired by [9], and this proof holds for *arbitrary* side information, including i.i.d. sequences, Markov chains, deterministic periodic sequences, and any general random processes one can conceive. Therefore this single proof is valid for both Theorems 2.5 and 3.3. Rigorous descriptions of the assumptions can be found in Section 3.1.

Proof: Without loss of generality, we assume $M_{C_0} = 2$, which immediately implies $T_{inf}(t) = T_1(t)$. Fix a θ with $\mu_\theta > \mu_{\theta_2}$, and define $C' = (\theta, \theta_2)$ as the competing configuration. This proof is mainly based on a change-of-measure argument as follows. Consider a uniformly good rule $\{\phi_\tau\}$ and suppose under the distribution P_{C_0}, the probability of the event that the inferior sampling time falls below the specified $\log(t)$ lower bound is bounded away from zero. Then under the competing distribution $\mathsf{P}_{C'}$, the same decision rule $\{\phi_\tau\}$ will result in $\mathcal{O}(t^\alpha)$ growth rate of $\mathsf{E}\{T_{inf}(t)\}$ for some $\alpha > 0$, which contradicts the assumption that $\{\phi_\tau\}$ is uniformly good.

Let $\lambda(n)$ denote the log likelihood ratio between θ_1 and θ based on the first n observed rewards of arm 1. That is

$$\lambda(n) := \sum_{m=1}^{n} \log\left(\frac{dF_{\theta_1}(Y^1_{\tau_m}|X_{\tau_m})}{dF_{\theta}(Y^1_{\tau_m}|X_{\tau_m})}\right),$$

where τ_m is a random variable corresponding to the time index of the m-th pull of arm 1.

We first note that by conditioning on the sequence of $\{X_{\tau_m}\}$, $\lambda(n)$ is a sum of independent r.v.'s, i.e., $\left\{\log\left(\frac{dF_{\theta_1}(Y^1_{\tau_m}|X_{\tau_m})}{dF_{\theta}(Y^1_{\tau_m}|X_{\tau_m})}\right)\right\}$. Let $K_{C'} := \sup_{x\in\mathbf{X}}\{I(\theta_1, \theta|x)\}$, and suppose there exists $\delta > 0$ such that

$$\limsup_{n\to\infty} \frac{\lambda(n)}{n} > K_{C'} + \delta,$$

with strictly positive probability. Then with strictly positive probability, there exists an x_0 such that the average of the subsequence for which $X_{\tau_m} = x_0$, will be larger than $K_{C'} + \delta$.

This, however, contradicts the strong law of large numbers since the subsequence is i.i.d. with marginal expectation $I(\theta_1, \theta|x_0)$. We thus obtain

$$\limsup_{n\to\infty} \frac{\lambda(n)}{n} \leq K_{C'}, \quad \mathsf{P}_{C_0} - a.s. \tag{B.1}$$

The inequality (B.1) is equivalent to the statement that with probability one, there are finitely many n such that $\lambda(n) > n(K_{C'} + \delta)$ for some $\delta > 0$. And since $K_{C'} > 0$, this in turn implies there are at most finitely many n such that $\max_{m\leq n} \lambda(m) > n(K_{C'} + \delta)$. As a result, we have

$$\limsup_{n\to\infty} \frac{\max_{m\leq n} \lambda(m)}{n} \leq K_{C'}, \quad \mathsf{P}_{C_0} - a.s.,$$

Since almost sure convergence implies convergence in probability, we have

$$\lim_{n\to\infty} \mathsf{P}_{C_0} \left(\exists m \leq n, \lambda(m) \geq (1+\delta) n K_{C'}\right) = 0. \tag{B.2}$$

Henceforth, we proceed using contradiction. Suppose

$$\limsup_{t\to\infty} \mathsf{P}_{C_0} \left(T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}} \right) > 0. \tag{B.3}$$

By (B.2), we have

$$\limsup_{t\to\infty} \mathsf{P}_{C_0} \left(T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}}, \lambda(T_1(t)) \leq (1+\delta)\frac{\log(t)}{(1+2\delta)K_{C'}} K_{C'} \right) > 0. \tag{B.4}$$

Using A_1 and A_2 as shorthand to denote events $A_1 := \left\{ T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}} \right\}$ and $A_2 := \left\{ \lambda(T_1(t)) \leq \frac{(1+\delta)\log(t)}{(1+2\delta)} \right\}$, we have

$$\begin{aligned}
\mathsf{E}_{C'} \left\{ T_{inf}(t) \right\} &\overset{(a)}{=} \mathsf{E}_{C'} \left\{ T_2(t) \right\} \\
&\overset{(b)}{=} \mathsf{E}_{C'} \left\{ t - T_1(t) \right\} \\
&\overset{(c)}{\geq} \left(t - \frac{\log(t)}{(1+2\delta)K_{C'}} \right) \mathsf{P}_{C'} \left(A_1 \right) \\
&\overset{(d)}{\geq} \left(t - \frac{\log(t)}{(1+2\delta)K_{C'}} \right) \mathsf{P}_{C'} \left(A_1 \cap A_2 \right) \\
&\overset{(e)}{\geq} \left(t - \frac{\log(t)}{(1+2\delta)K_{C'}} \right) e^{-\frac{(1+\delta)\log(t)}{1+2\delta}} \mathsf{P}_{C_0} \left(A_1 \cap A_2 \right) \\
&\overset{(f)}{=} \mathcal{O}\left(t^{\frac{\delta}{1+2\delta}} \right).
\end{aligned} \tag{B.5}$$

The equality marked (a) follows from $M_{C'} = 1$ and (b) follows from the fact that $T_1(t) + T_2(t) = t$. (c) and (d) follow from elementary probability inequalities. (e) follows from the change-of-measure formula and the definition of A_2, within which $\lambda(T_1(t)) \leq \frac{(1+\delta)\log(t)}{(1+2\delta)}$. (f) comes from simple arithmetic and (B.4).

The inequality (B.5) contradicts the assumption that $\{\phi_\tau\}$ is uniformly good. Thus (B.3) is incorrect, and by letting the δ satisfy $\frac{1}{1+2\delta} = 1 - \epsilon$, we have proven

$$\lim_{t\to\infty} \mathsf{P}_{C_0} \left(T_1(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C'}} \right) = 0, \quad \forall \epsilon > 0.$$

By choosing the θ in $C' = (\theta, \theta_2)$ with the minimizing configuration of $\inf_{\theta>\theta_2} \sup_x I(\theta_1, \theta|x)$, we complete the proof of the first statement of Theorem 2.5. The second statement in Theorem 2.5 can be obtained by simply applying Markov's inequality and the first statement. ■

B.2 Proof of Theorems 2.7 and 3.5

We first note that Theorem 2.7 is a special case of Theorem 3.5, and therefore we will focus on the latter, under which the side information is evenly distributed in probability.

Proof: This proof is basically a variation of that for Theorem 3.3 in Appendix B.1, with the major difference being that the competing configuration $C' = (\theta, \theta_2)$ is now from a different set: $\{\theta : \exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}$.

Consider C_0 that is not implicitly revealing. By definition and by symmetry, we may assume $M_{C_0}(x) = 2, \forall x$, which immediately implies $T_{inf}(t) = T_1(t)$. Fix a θ such that $\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)$, and define $C' = (\theta, \theta_2)$ as the competing configuration. Consider a uniformly good rule $\{\phi_\tau\}$, and let $\lambda(n)$ denote the log likelihood ratio between θ_1 and θ based on the first n observed rewards of arm 1. That is

$$\lambda(n) := \sum_{m=1}^{n} \log\left(\frac{dF_{\theta_1}(Y^1_{\tau_m}|X_{\tau_m})}{dF_{\theta}(Y^1_{\tau_m}|X_{\tau_m})}\right),$$

where τ_m is a random variable corresponding to the time index of the m-th pull of arm 1. By the same argument as in Appendix B.1, we have

$$\lim_{n\to\infty} \mathsf{P}_{C_0}\left(\exists m \le n, \lambda(m) \ge (1+\delta) n K_{C'}\right) = 0. \tag{B.6}$$

Hereafter, we proceed using contradiction. Suppose

$$\limsup_{t\to\infty} \mathsf{P}_{C_0}\left(T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}}\right) > 0. \tag{B.7}$$

By (B.6), we have

$$\limsup_{t\to\infty} \mathsf{P}_{C_0}\left(T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}}, \lambda(T_1(t)) \le (1+\delta)\frac{\log(t)}{(1+2\delta)K_{C'}}K_{C'}\right) > 0.$$

Using A_1 and A_2 as shorthand to denote events $A_1 := \left\{T_1(t) < \frac{\log(t)}{(1+2\delta)K_{C'}}\right\}$ and $A_2 := \left\{\lambda(T_1(t)) \le \frac{(1+\delta)\log(t)}{(1+2\delta)}\right\}$, and by the assumption that $\{X_\tau\}$ is evenly distributed in probability, there exists a $\pi > 0$ such that

$$\limsup_{t\to\infty} \mathsf{P}_{C_0}\left(A_1 \cap A_2 \cap \left\{\left(\sum_{\{x:M_{C'}(x)=1\}} f_r(x,t)\right) \ge \pi\right\}\right) > 0. \tag{B.8}$$

Using A_3 as shorthand for the event $\left\{\left(\sum_{\{x:M_{C'}(x)=1\}} f_r(x,t)\right) \geq \pi\right\}$, we have

$$
\begin{aligned}
& \mathsf{E}_{C'}\left\{T_{inf}(t)\right\} \\
\geq\ & \mathsf{E}_{C'}\left\{\sum_{\tau=1}^{t} 1\{\phi_\tau = 2, M_{C'}(X_\tau)=1\}\right\} \\
=\ & \mathsf{E}_{C'}\left\{\sum_{\tau=1}^{t} 1\{M_{C'}(X_\tau)=1\} - \sum_{\tau=1}^{t} 1\{\phi_\tau = 1, M_{C'}(X_\tau)=1\}\right\} \\
\overset{(a)}{\geq}\ & \left(\mathsf{E}_{C'}\left\{\sum_{\tau=1}^{t} 1\{M_{C'}(X_\tau)=1\} - \sum_{\tau=1}^{t} 1\{\phi_\tau = 1, M_{C'}(X_\tau)=1\}\middle| A_1 \cap A_3\right\}\right)\mathsf{P}_{C'}(A_1 \cap A_3) \\
\overset{(b)}{\geq}\ & \left(\pi \cdot t - \mathsf{E}_{C'}\left\{T_1(t)|\, A_1 \cap A_3\right\}\right)\mathsf{P}_{C'}(A_1 \cap A_3) \\
\overset{(c)}{\geq}\ & \left(\pi t - \frac{\log(t)}{(1+2\delta)K_{C'}}\right)\mathsf{P}_{C'}\left(A_1 \cap A_3\right) \\
\overset{(d)}{\geq}\ & \left(\pi t - \frac{\log(t)}{(1+2\delta)K_{C'}}\right)\mathsf{P}_{C'}\left(A_1 \cap A_2 \cap A_3\right) \\
\overset{(e)}{\geq}\ & \left(\pi t - \frac{\log(t)}{(1+2\delta)K_{C'}}\right) e^{-\frac{(1+\delta)\log(t)}{1+2\delta}} \mathsf{P}_{C_0}\left(A_1 \cap A_2 \cap A_3\right) \\
\overset{(f)}{=}\ & \mathcal{O}\left(t^{\frac{\delta}{1+2\delta}}\right),
\end{aligned}
\tag{B.9}
$$

The equality marked (a) follows from an elementary probability inequality and (b) follows from the facts that A_3 is satisfied and $T_1(t) \geq \sum_{\tau=1}^{t} 1\{\phi_\tau = 1, M_{C'}(X_\tau)=1\}$. (c) and (d) follow from elementary probability inequalities. (e) follows from the change-of-measure formula and the definition of A_2, within which $\lambda(T_1(t)) \leq \frac{(1+\delta)\log(t)}{(1+2\delta)}$. (f) follows from simple arithmetic and (B.8).

The inequality (B.9) contradicts the assumption that $\{\phi_\tau\}$ is uniformly good. Thus (B.7) is incorrect, and by letting the δ satisfy $\frac{1}{1+2\delta} = 1-\epsilon$, we have proven

$$
\lim_{t\to\infty} \mathsf{P}_{C_0}\left(T_1(t) \geq \frac{(1-\epsilon)\log(t)}{K_{C'}}\right) = 0, \quad \forall \epsilon > 0.
$$

By choosing the θ in $C' = (\theta, \theta_2)$ with the minimizing configuration of

$$
\inf_{\{\theta:\exists x_0,\ \text{s.t.}\ \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \sup_x I(\theta_1, \theta | x),
$$

the proof of the first statement in Theorem 3.5. The second statement in Theorem 3.5 can be obtained by simply applying Markov's inequality and the first statement. ■

Appendix C

Proofs of the Achievability Results

C.1 Proof of Theorems 2.3 and 3.1

Since Theorem 2.3 is a special case of Theorem 3.1 as discussed in Section 3.2, we will focus solely on the proof of Theorem 3.1.

Proof of Theorem 3.1*:* For each underlying configuration pair $C_0 = (\theta_1, \theta_2)$, define the error set $\mathbf{C}_e$ as follows.

$$\mathbf{C}_e \;:=\; \bigcup_{x \in \mathbf{X}} \{C \in \mathbf{\Theta}^2 : M_C(x) \neq M_{C_0}(x)\}. \tag{C.1}$$

Let $\overline{\mathbf{C}}_e$ denote the closure of $\mathbf{C}_e$. By Condition 2.1, we have $C_0 \notin \overline{\mathbf{C}}_e$ and there exists $\epsilon > 0$ such that $\overline{\mathbf{C}}_e \subseteq \{C : |C - C_0| > \epsilon\}$. For any $\tau \geq 1$,

$$\begin{aligned}
\mathsf{P}_{C_0}\left(\phi_\tau \neq M_{C_0}(X_\tau)\right) &= \mathsf{P}_{C_0}\left(M_{\hat{C}_\tau}(X_\tau) \neq M_{C_0}(X_\tau)\right) \\
&\leq \mathsf{P}_{C_0}\left(\exists x, M_{\hat{C}_\tau}(x) \neq M_{C_0}(x)\right) \\
&= \mathsf{P}_{C_0}\left(\hat{C}_\tau \in \mathbf{C}_e\right) \\
&\leq \mathsf{P}_{C_0}\left(\hat{C}_\tau \in \overline{\mathbf{C}}_e\right) \\
&\leq \mathsf{P}_{C_0}\left(|\hat{C}_\tau - C_0| > \epsilon\right).
\end{aligned}$$

From the above inequality, we have

$$\begin{aligned}
\mathsf{E}_{C_0}\{T_{inf}(t)\} &= \sum_{\tau=1}^{t} \mathsf{P}_{C_0}(\phi_\tau \neq M_{C_0}(X_\tau)) \\
&\leq \sum_{\tau=1}^{t} \mathsf{P}_{C_0}\left(|\hat{C}_\tau - C_0| > \epsilon\right).
\end{aligned}$$

This completes the proof. ■

C.2 Analysis of Algorithm 1

With $\{X_\tau\}$ being evenly distributed in probability series, we will show that the decision rule $\{\phi_\tau\}$ in Algorithm 1 achieves bounded $\mathsf{E}_{C_0}\{T_{inf}(t)\}$ for all $C_0 \in \mathbf{\Theta}^2$. Since non-degenerate

i.i.d. random processes are evenly distributed in probability series, this analysis provides proofs for both Theorems 2.4 and 3.2.

We define $\mathbf{C}_e$ similarly to (C.1). A necessary lemma is stated and proved as follows.

Lemma C.1 *With the regularity conditions specified in Section 2.3 and* $x^\star$ *defined in* Algorithm 1, $\exists a_1, a_2 > 0$ *such that* $\mathsf{P}_{C_0}(\hat{C}_t \in \mathbf{C}_e) \leq a_1 \exp\left(-a_2 \min\left\{T_1^{x^\star}(t), T_2^{x^\star}(t)\right\}\right)$.

Proof of Lemma C.1*:*

By the continuity of $\mu_\theta(x)$ w.r.t. θ and the assumption of finite $\mathbf{X}$, it can be shown that $C_0 \notin \overline{\mathbf{C}}_e$, the closure of $\mathbf{C}_e$. Therefore there exists a neighborhood of C_0, $\mathbf{C}_\delta = (\theta_1 - \delta, \theta_1 + \delta) \times (\theta_2 - \delta, \theta_2 + \delta)$, such that $\mathbf{C}_\delta \subseteq \left(\mathbf{\Theta}^2 \backslash \overline{\mathbf{C}}_e\right)$, or equivalently $\overline{\mathbf{C}}_e \subseteq \left(\mathbf{\Theta}^2 \backslash \mathbf{C}_\delta\right)$. Define

$$\epsilon := \frac{1}{4} \inf\left\{\rho\left(F_{\hat{\theta}_i}(\cdot|x), F_{\theta_i}(\cdot|x)\right) : \forall x \in \mathbf{X}, i \in \{1,2\}, (\hat{\theta}_1, \hat{\theta}_2) \notin \mathbf{C}_\delta\right\}, \tag{C.2}$$

and from the construction of $\mathbf{C}_\delta$, we have $\epsilon > 0$. We would like to prove that for sufficiently large $t > \frac{1}{\epsilon}$, the following relationship holds.

$$\left\{\hat{C}_t \notin \mathbf{C}_\delta\right\} \subseteq \left\{\exists i, \rho\left(L_i^{x^\star}(t), F_{\theta_i}(\cdot|x^\star)\right) > \epsilon\right\}.$$

We consider the case when the empirical measure $L_i^{x^\star}(t)$ is close to the true distribution $F_{\theta_i}(\cdot|x^\star)$ and use proof by contradiction. Suppose $\rho\left(L_i^{x^\star}(t), F_{\theta_i}(\cdot|x^\star)\right) \leq \epsilon$ for both $i = 1, 2$. By the definition of $\sigma(C_0, t)$, we have

$$\sigma(C_0, t) \leq 2\epsilon. \tag{C.3}$$

However, for those $\hat{C}_t \notin \mathbf{C}_\delta$, by the definition of ϵ in (C.2), there exists $i \in \{1, 2\}$ such that

$$\begin{aligned}
\sigma(\hat{C}_t, t) &= \rho(F_{\hat{\theta}_1}(\cdot|x^\star), L_1^{x^\star}(t)) + \rho(F_{\hat{\theta}_2}(\cdot|x^\star), L_2^{x^\star}(t)) \\
&\geq \rho(F_{\hat{\theta}_i}(\cdot|x^\star), L_i^{x^\star}(t)) \\
&\geq \rho(F_{\hat{\theta}_i}(\cdot|x^\star), F_{\theta_i}(\cdot|x^\star)) - \rho(F_{\theta_i}(\cdot|x^\star), L_i^{x^\star}(t)) \\
&\geq 3\epsilon,
\end{aligned} \tag{C.4}$$

which contradicts the definition of $\mathbf{C}_t$ since (C.3) and (C.4) imply $\sigma(\hat{C}_t, t) > \frac{1}{t} + \sigma(C_0, t)$. As a result, for sufficiently large $t > \frac{1}{\epsilon}$, we have

$$\begin{aligned}
\left\{\hat{C}_t \in \mathbf{C}_e\right\} &\subseteq \left\{\hat{C}_t \notin \mathbf{C}_\delta\right\} \\
&\subseteq \left\{\exists i, \rho\left(L_i^{x^\star}(t), F_{\theta_i}(\cdot|x^\star)\right) > \epsilon\right\} \\
&= \bigcup_{i=1,2} \left\{\rho\left(L_i^{x^\star}(t), F_{\theta_i}(\cdot|x^\star)\right) > \epsilon\right\}.
\end{aligned} \tag{C.5}$$

By Sanov's theorem, the probability of each term in the union of the right-hand side of (C.5) is exponentially bounded w.r.t. $T_i^{x^\star}(t)$. As a result, the probability of this finite union is upper bounded by $a_1 \exp\left(-a_2 \min\left\{T_1^{x^\star}(t), T_2^{x^\star}(t)\right\}\right)$ for some a_1, $a_2 > 0$. The proof is complete.

■

Analysis of Algorithm 1*:*

By the definition of $\mathbf{C}_e$, when $\hat{C}_t$ is not in $\mathbf{C}_e$, the estimate is accurate enough that the myopic decision is simply the optimal decision, namely, $\forall x$, $M_{\hat{C}_t}(x) = M_{C_0}(x)$. Hence we have

$$\begin{aligned}\{\phi_{t+1} \neq M_{C_0}(X_{t+1})\} &= \left\{\phi_{t+1} \neq M_{C_0}(X_{t+1}), \hat{C}_t \in \mathbf{C}_e\right\} \\ &\quad \cup \left\{\phi_{t+1} \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\right\} \\ &\subseteq \left\{\hat{C}_t \in \mathbf{C}_e\right\} \cup \left\{\phi_{t+1} \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\right\} \\ &\triangleq A_{t+1} \cup B_{t+1}. \end{aligned} \tag{C.6}$$

We first show that $\forall t \geq 6, T_i(t) \geq \sqrt{t}$, by induction. This statement is true for $t = 6$. Suppose $T_t(t-1) \geq \sqrt{t-1}$. If $T_i(t-1) \geq \sqrt{t}$, by the monotonicity of $T_i(t)$ with respect to t, we have $T_i(t) \geq T_i(t-1) \geq \sqrt{t}$. If $T_i(t-1) < \sqrt{t}$, by the forced sampling mechanism, $T_i(t) = T_i(t-1) + 1 \geq \sqrt{t-1} + 1 \geq \sqrt{t}$. As a result, $T_i(t) \geq \sqrt{t}$, which in turn implies $\min_i T_i(t) \geq \sqrt{t}$ and $\min_i T_i^{x^\star}(t) \geq \frac{\sqrt{t}}{|\mathbf{X}|}$. By Lemma C.1, we have $\mathsf{P}_{C_0}(A_{t+1}) \leq a_1 e^{-a_2 \frac{\sqrt{t}}{|\mathbf{X}|}}$, and hence $\sum_{t+1=7}^{\infty} \mathsf{P}_{C_0}(A_{t+1}) < \infty$.

For B_{t+1}, we have

$$\begin{aligned} B_{t+1} &= \{\phi_{t+1} \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\} \\ &= \{\phi_{t+1} = 1 \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\} \cup \{\phi_{t+1} = 2 \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\} \\ &\triangleq B_{t+1}^1 \cup B_{t+1}^2, \end{aligned}$$

where B_{t+1}^1 and B_{t+1}^2 correspond to $\phi_{t+1} = 1, 2$, separately. We then have

$$\begin{aligned} B_{t+1}^1 &= \left\{\exists s \in [\sqrt{t}, t-1] \text{ s.t. } \hat{C}_s \in \mathbf{C}_e, \phi_{t+1} = 1 \neq M_{C_0}(X_{t+1}), \hat{C}_t \notin \mathbf{C}_e\right\} \\ &\qquad \cup \left\{\forall s \in [\sqrt{t}, t], \hat{C}_s \notin \mathbf{C}_e, \phi_{t+1} = 1 \neq M_{C_0}(X_{t+1})\right\} \\ &\subseteq \left\{\exists s \in [\sqrt{t}, t-1] \text{ s.t. } \hat{C}_s \in \mathbf{C}_e\right\} \cup B^{1.1}. \end{aligned} \tag{C.7}$$

The above subset sign follows from modifying the first term of the union and using $B^{1.1}$ as shorthand for the second term. The probability of the first term in (C.7) will later be upper bounded by the union bound and the large deviation principle. To upper bound $B^{1.1}$, we need some new notation:

$$\begin{aligned} N_1 &:= \sum_{s \in [1,t]} 1\{M_{C_0}(X_s) = 1\} \\ N_{1\to 2} &:= \sum_{s \in [1,t]} 1\{M_{C_0}(X_s) = 1, \phi_s = 2\} \\ \text{and} \quad N_{2\to 1} &:= \sum_{s \in [1,t]} 1\{M_{C_0}(X_s) = 2, \phi_s = 1\}. \end{aligned}$$

By definition, we have $T_1(t) = N_1 - N_{1\to 2} + N_{2\to 1}$. Suppose $\forall s \in [\sqrt{t}, t]$, $\hat{C}_s \notin \mathbf{C}_e$, namely,

the first condition of $B^{1.1}$ is satisfied. We notice the following inequalities,

$$\begin{aligned} N_{1\to 2}+N_{2\to 1} &= \sum_{s\in[1,\sqrt{t}]} 1\{\phi_s \neq M_{C_0}(X_s)\} + \sum_{s\in[\sqrt{t}+1,t]} 1\{\phi_s \neq M_{C_0}(X_s)\} \\ &\leq \sqrt{t} + \sum_{s\in[\sqrt{t}+1,t]} 1\{\phi_s \neq M_{\hat{C}_{s-1}}(X_s)\} \\ &\leq 2\sqrt{t}. \end{aligned} \tag{C.8}$$

The equality is obvious and the first inequality is true since $\forall s \in [\sqrt{t}, t]$, $\hat{C}_s \notin \mathbf{C}_e$ and thus $M_{\hat{C}_s}(\cdot) = M_{C_0}(\cdot)$. The second inequality follows from the fact that the total number of forced samples up to time t cannot be greater than $\sqrt{t}$. The reason is according to Line 3 in Algorithm 1, the forced sampling can only happen at time $\tau \leq t$ when $\sqrt{\tau - 1}$ is an integer. From the above reasoning, the number of times $\phi_s \neq M_{\hat{C}_{s-1}}(X_s)$, $s \leq t$, is smaller than $\sqrt{t}$.

If the second condition of $B^{1.1}$, $\phi_{t+1} = 1 \neq M_{\hat{C}_t}(X_{t+1})$, is satisfied, it implies that the player performs the forced sampling at time $t+1$, or equivalently $T_1(t) < \sqrt{t+1}$. Since $\forall i, T_i(t) \geq \sqrt{t}$, it follows that $T_1(t) = N_1 - N_{1\to 2} + N_{2\to 1} = \sqrt{t}$. Combining the result in (C.8), we conclude that

$$\begin{aligned} B^{1.1} &\subseteq \{N_1 \leq 3\sqrt{t}\} \\ &= \left\{ \left(\sum_{s\in[1,t]} 1\{M_{C_0}(X_s) = 1\} \right) \leq 3\sqrt{t} \right\}. \end{aligned} \tag{C.9}$$

Let $\mathbf{X}^1_{C_0} := \{x \in \mathbf{X} : M_{C_0}(x) = 1\}$ denote the set of possible values (of X_t) such that arm 1 is favorable. From (C.7) we have

$$\begin{aligned} \mathsf{P}(B^1_{t+1}) &\leq \left(\sum_{s\in[\sqrt{t},t-1]} \mathsf{P}(\hat{C}_s \in \mathbf{C}_e) \right) + \mathsf{P}(B^{1.1}) \\ &\leq \left(\sum_{s\in[\sqrt{t},t-1]} a_1 e^{-a_2\sqrt{s}} \right) + \mathsf{P}\left(\frac{\sum_{s\in[1,t]} 1\{X_s \in \mathbf{X}^1_{C_0}\}}{t} \leq \frac{3\sqrt{t}+1}{t} \right), \end{aligned} \tag{C.10}$$

where the second inequality follows from the application of Lemma C.1 to the first term, and the second term follows from (C.9). By simple algebra, we have

$$\sum_{t+1=7}^{\infty} \sum_{s\in[\sqrt{t},t-1]} a_1 e^{-a_2\sqrt{s}} < \infty. \tag{C.11}$$

And by the assumption that $\{X_\tau\}$ is evenly distributed in probability series, we have

$$\sum_{t+1=7}^{\infty} \mathsf{P}\left(\frac{\sum_{s\in[1,t]} 1\{X_s \in \mathbf{X}^1_{C_0}\}}{t} \leq \frac{3\sqrt{t}+1}{t} \right) < \infty. \tag{C.12}$$

From (C.10), (C.11), (C.12), and the symmetry between B^1_{t+1} and B^2_{t+1}, we conclude

$$\sum_{t+1=7}^{\infty} \mathsf{P}(B_{t+1}) \leq \sum_{t+1=7}^{\infty} \left(\mathsf{P}(B^1_{t+1}) + \mathsf{P}(B^2_{t+1}) \right) < \infty,$$

and by (C.6),

$$\lim_{t\to\infty} \mathsf{E}_{C_0}\{T_{inf}(t)\} \leq 6 + \sum_{t+1=7}^{\infty} (\mathsf{P}(A_{t+1}) + \mathsf{P}(B_{t+1})) < \infty.$$

The analysis is complete. ■

C.3 Analysis of Algorithm 2

We will show that Algorithm 2 is able to achieve the new $\log(t)$ lower bound specified in Theorems 2.6 and 3.4 for i.i.d. or u.s.e. distributed side information $\{X_\tau\}$. Since i.i.d. sequences are a special case of u.s.e. distributed in L^1 random processes, we assume $\{X_\tau\}$ is u.s.e. distributed in L^1 throughout this appendix and the following analysis serves as a proof for both Theorems 2.6 and 3.4.

We first state the following lemma, which will be used in the later proof.

Lemma C.2 *Consider a random process $\{X_\tau\}$ and a sequence of stopping time pairs $\{(S_j, T_j)\}$, where for all $j \in \mathbb{N}$, $S_j \leq T_j \leq S_{j+1}$ are stopping times taking values in $\mathbb{N} \cup \{\infty\}$. Denote*

$$\begin{aligned} U &:= \sup\{j \in \mathbb{N} | S_j < \infty\} \\ \textit{and} \quad \mathsf{sum} &:= \sum_{j=1}^{U} (T_j - S_j + 1). \end{aligned}$$

Suppose for some $B < \infty$ and $K < \infty$, we have $\mathsf{E}\{U\} \leq K$, and $\forall j$, $\mathsf{E}\{T_j - S_j + 1 | S_j\} \leq B$. It follows that $\mathsf{E}\{\mathsf{sum}\} \leq K \cdot B < \infty$.

Proof: The proof is similar to that of Wald's Lemma. Using the convention that $0 \cdot \infty = 0$, we can rewrite sum in the following form:

$$\begin{aligned} \mathsf{sum} &= \sum_{j=1}^{\infty} 1\{S_j < \infty\} (T_j - S_j + 1) \\ \Longrightarrow \mathsf{E}\{\mathsf{sum}\} &= \sum_{j=1}^{\infty} \mathsf{E}\left\{1\{S_j < \infty\} \cdot \mathsf{E}\{T_j - S_j + 1 | S_j\}\right\} \\ &\leq \sum_{j=1}^{\infty} \mathsf{E}\{1\{S_j < \infty\}\} \cdot B \\ &= \left(\sum_{j=1}^{\infty} \mathsf{P}(U \geq j)\right) \cdot B = K \cdot B. \end{aligned}$$

■

With the help of Lemma C.2, we prove Theorem 3.4 by making the following arguments regarding to Algorithm 2.

- ARGUMENT 1: The expected duration over which $\hat{C}_\tau$ does not exist is finite, i.e.,

$$\lim_{t\to\infty} \mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau \text{ does not exists}\}\right\} < \infty,$$

where $\hat{C}_\tau$ in Algorithm 2 is defined as the common estimate $\hat{C}_\tau := \hat{C}_{x,\tau}, \forall x$ from the EBUG rules $\{\phi_{x,\tau}\}$ on all sub-bandit machines.

For simplicity, we use $\chi\{\hat{C}_\tau\} = 0$ as shorthand notation for the condition that $\hat{C}_\tau$ does not exist.

- ARGUMENT 2: The expected duration over which $\hat{C}_t \neq C_0$ is finite, i.e.,

$$\lim_{t\to\infty} \mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau \neq C_0\}\right\} < \infty.$$

- ARGUMENT 3: The expected duration over which $\hat{C}_t = C_0$ and $\Phi_{t+1} \neq M_{C_0}(X_{t+1})$ is upper bounded by $\frac{\log(t)}{K_{C_0}}$, i.e.,

$$\lim_{t\to\infty} \frac{\mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\}\right\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

where $K_{C_0} = \inf_{\theta>\theta_2} \sup_x I(\theta_1, \theta|x)$ if $M_{C_0} = 2$.

Once we establish ARGUMENTS 1 through 3, it is straightforward to show that Algorithm 2 attains the specified $\log(t)$ lower bound.

Proof of ARGUMENT 1*:* To discuss stopping times, we first define the filtration $\mathcal{F}_t$ in an explicit way, that is, $\mathcal{F}_t$ is the σ-algebra generated by the past outcomes of the rewards $1\{\Phi_\tau = 1\}Y_\tau^1 + 1\{\Phi_\tau = 2\}Y_\tau^2$ for $\tau \in [1, t]$, and the observations X_τ for $\tau \in [1, t+1]$. For instance, by definition we have $\hat{C}_t \in \mathcal{F}_t$, $X_{t+1} \in \mathcal{F}_t$ and $\phi_{t+1} \in \mathcal{F}_t$.

For any $x \in \mathbf{X}$, we iteratively define the stopping time pairs $S_{x,j}$ and $T_{x,j}$ as follows.

$$\begin{aligned} S_{x,j} \;:=\; \inf\Big\{ & t > S_{x,j-1} : X_t = x, \chi\{\hat{C}_t\} = 0, \\ & \text{and either } \chi\{\hat{C}_{t-1}\} = 1 \text{ or } \mathbf{X} = \bigcup_{\tau\in(S_{x,j-1},t)} \{X_\tau\} \Big\}, \end{aligned}$$

and

$$T_{x,j} \;:=\; \inf\left\{ t > S_{x,j} : \text{ either } \chi\{\hat{C}_t\} = 1 \text{ or } \mathbf{X} = \bigcup_{\tau\in(S_{x,j},t]} \{X_\tau\} \right\},$$

where $S_{x,0} = 0$. Note that $S_{x,j} < T_{x,j}$ are basically dividing the duration over which $\chi\{\hat{C}_t\} = 0$ into disjoint[1] intervals, with x specifying the value of the side observation X_t at the leading time instant $S_{x,j}$. We then have

$$\sum_{\tau=1}^{\infty} 1\{\chi\{\hat{C}_\tau\} = 0\} \leq \sum_x \sum_{j\in\mathbb{N}} (T_{x,j} - S_{x,j} + 1).$$

Since

$$T_{x,j} \leq \inf\left\{ t > S_{x,j} : \mathbf{X} = \bigcup_{\tau\in(S_{x,j},t]} \{X_\tau\} \right\},$$

[1]For any $x \in \mathbf{X}$, we have $S_{x,1} < T_{x,1} < S_{x,2} < \cdots < S_{x,j} < \cdots$, which forms a set of disjoint intervals. Nonetheless, for different x and x', we may have $[S_{x,j}, T_{x,j}] \cap [S_{x',j'}, S_{x',j'}] \neq \emptyset$ for some $j, j' \in \mathbb{N}$.

and by the assumption that $\{X_\tau\}$ is u.s.e. distributed in L^1, there exists $B < \infty$ such that $\forall x, j$, $\mathsf{E}\{T_{x,j} - S_{x,j} + 1 | S_{x,j}\} < B$. If we can show

$$\forall x, \mathsf{E}\{\sup\{j \in \mathbb{N} : S_{x,j} < \infty\}\} < \infty, \tag{C.13}$$

then by Lemma C.2, we have $\mathsf{E}\left\{\sum_{t=1}^{\infty} 1\{\chi\{\hat{C}_t\} = 0\}\right\} < \infty$ and ARGUMENT 1 is established.

We prove (C.13) by case study. For any x_0, j, and after time instant $t := S_{x_0,j}$, since $\chi\{\hat{C}_t\} = 0$ and $X_t = x_0$, we must be in one of the following two cases.

- $\hat{C}_{x_0,t} \neq C_0$:
  - If $\chi\{\hat{C}_{t-1}\} = 0$, then $\Phi_t \leftarrow \phi_{x_0,t}$. By the assumption that the constituent $\{\phi_{x_0,\tau}\}$ is EBUG, the expected duration of the event $\left\{X_t = x_0, \Phi_t \leftarrow \phi_{x_0,t}, \hat{C}_{x_0,t} \neq C_0\right\}$ must be finite. So this case can only contribute finite expectation.
  - If $\chi\{\hat{C}_{t-1}\} = 1$, the only condition resulting in $\chi\{\hat{C}_t\} = 0$ is that $\hat{C}_{x_0,t-1} \neq \hat{C}_{x_0,t}$, which in turn implies $\Phi_t \leftarrow \phi_{x_0,t}$. By the assumption of EBUG $\{\phi_{x_0,\tau}\}$, the expected duration of the event $\left\{X_t = x, \Phi_t \leftarrow \phi_{x_0,t}, \hat{C}_{x_0,t} \neq C_0\right\}$ must be finite. So this case only contributes finite expectation.
- $\hat{C}_{x_0,t} = C_0$:
  By observing $\sup\{j \in \mathbb{N} : S_{x_0,j} < \infty\} \leq \sup\{j \in \mathbb{N} : T_{x_0,j} < \infty\} + 1$, we choose to show the latter has bounded expectation.

  Suppose $T_{x_0,j} < \infty$, and note that $\chi\{\hat{C}_t\} = 0$ implies there exists $x' \neq x_0$ such that $\hat{C}_{x',t} \neq C_0$. There are only two sub-cases as follows.
  - $\exists t' \in (S_{x_0,j}, T_{x_0,j}]$ such that $X_{t'} = x'$ and $\hat{C}_{x',t'-1} \neq C_0$.
  - $X_{T_{x_0,j}} = x_0$ and $\hat{C}_{x_0,T_{x_0,j}} \neq \hat{C}_{x_0,t} = C_0$.

  The reason why there are only two sub-cases follows because if there exists no such t' as in the first case, then $\hat{C}_{x',s}$ remains unchanged within the interval $(S_{x_0,j}, T_{x_0,j}]$. So the only situation leading to $T_{x_0,j} < \infty$ is when $\hat{C}_{x_0,T_{x_0,j}} \neq C_0 = \hat{C}_{x_0,T_{x_0,j}-1}$. Since for all $\tau \in (S_{x_0,j}, T_{x_0,j}]$ the decision rule is $\Phi_\tau \leftarrow \phi_{X_\tau,\tau}$, we then have

$$\begin{aligned}
&\sup\{j \in \mathbb{N} : T_{x_0,j} < \infty\} \\
\leq\ & \sum_{x' : x' \neq x_0} \sum_{\tau=1}^{\infty} 1\{X_\tau = x', \Phi_\tau \leftarrow \phi_{x',\tau}, \hat{C}_{x',\tau-1} \neq C_0\} \\
&+ \sum_{\tau=1}^{\infty} 1\{X_\tau = x_0, \Phi_\tau \leftarrow \phi_{x_0,\tau}, \hat{C}_{x_0,\tau} \neq C_0\}.
\end{aligned}$$

  By the assumption of EBUG constituent $\{\phi_{x,\tau}\}$, the above must have finite expectation.

From the previous discussions, we have proven $\mathsf{E}\{\sup\{j \in \mathbb{N} : S_{x,j} < \infty\}\} < \infty, \forall x \in \mathbf{X}$, and thus established ARGUMENT 1. ■

Proof of ARGUMENT 2*:* Consider a fixed $C' := (\theta_1', \theta_2') \neq C_0$ and let x^* denote the maximizing argument of $\inf_{\theta:\theta>\theta_2'} I(\theta_1', \theta|x)$. We then iteratively define the stopping time pairs $S_{C',j}$ and $T_{C',j}$ as follows.

$$\begin{aligned} S_{C',j} &:= \inf\Big\{t > S_{C',j-1} : \hat{C}_t = C', \\ &\qquad \text{and either } \chi\{\hat{C}_{t-1}\} = 0, \text{ or } \hat{C}_{t-1} \neq C', \text{ or } X_{t-1} = x^*\Big\}, \end{aligned}$$

and

$$T_{C',j} := \inf\Big\{t > S_{C',j} : \text{ either } \chi\{\hat{C}_t\} = 0, \text{ or } \hat{C}_t \neq C', \text{ or } X_t = x^*\Big\},$$

where $S_{C',0} = 0$. Note that $S_{C',j}$ and $T_{C',j}$ are basically dividing the duration of the event $\{\hat{C}_t \neq C_0\}$ into disjoint intervals while C' is specifying the value of the common estimate $\hat{C}_t$ during those intervals. Then we have

$$\sum_{t=1}^{\infty} 1\{\hat{C}_t \neq C_0\} \leq \sum_{C' \neq C_0} \sum_{j \in \mathbb{N}} \left(T_{C',j} - S_{C',j} + 1\right).$$

Since

$$T_{C',j} \leq \inf\left\{t > S_{C',j} : X_t = x^*\right\},$$

and by the assumption that $\{X_\tau\}$ is u.s.e. distributed in L^1, there exists $B < \infty$ such that $\forall x, j$, $\mathsf{E}\left\{T_{C',j} - S_{C',j} + 1 | S_{C',j}\right\} < B$. If we can show

$$\forall x,\ \mathsf{E}\left\{\sup\{j \in \mathbb{N} : S_{C',j} < \infty\}\right\} < \infty,$$

then by Lemma C.2, we have $\mathsf{E}\left\{\sum_{t=1}^{\infty} 1\{\hat{C}_t \neq C_0\}\right\} < \infty$.

By observing $\sup\{j \in \mathbb{N} : S_{C',j} < \infty\} \leq \sup\{j \in \mathbb{N} : T_{C',j} < \infty\} + 1$, we choose to show the latter has bounded expectation. We first note that there is some redundancy in the definition of $T_{C',j}$ since when $\hat{C}_t$ exists, the only possible situation under which $\hat{C}_t$ will change or become non-existential is when $X_t = x^*$. So $T_{C',j}$ can be rewritten as follows.

$$T_{C',j} := \inf\left\{t > S_{C',j} : X_t = x^*\right\}.$$

By this new definition, if $T_{C',j} < \infty$, we have $X_{T_{C',j}} = x^*$, $\hat{C}_{x^*,T_{C',j}-1} = C' \neq C_0$, and $\Phi_{T_{C',j}} \leftarrow \phi_{x^*,T_{C',j}}$. Using these facts, we have

$$\sup\{j \in \mathbb{N} : T_{C',j} < \infty\} \leq \sum_{\tau=1}^{\infty} 1\{X_\tau = x^*, \Phi_\tau \leftarrow \phi_{x^*,\tau}, \hat{C}_{x^*,\tau-1} \neq C_0\}.$$

By the assumption of EBUG constituent $\{\phi_{x,\tau}\}$, the above has finite expectation and we have proven ARGUMENT 2. ■

Proof of ARGUMENT 3*:* Suppose $C_0 = (\theta_1, \theta_2)$. Without loss of generality, we may assume $M_{C_0} = 2$ and let $x^* = \arg\max_x \inf_{\theta:\theta>\theta_2} I(\theta_1, \theta|x)$. We then have

$$\begin{aligned} &\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\} \\ &= \sum_{\tau=1}^{t} 1\{\hat{C}_\tau = \hat{C}_{x^*,\tau} = C_0, X_{\tau+1} = x^*, \Phi_{\tau+1} \leftarrow \phi_{x^*,\tau+1} \neq M_{C_0}(X_{\tau+1})\} \\ &\leq \sum_{\tau=1}^{t} 1\{\hat{C}_{x^*,\tau} = C_0, X_{\tau+1} = x^*, \Phi_{\tau+1} \leftarrow \phi_{x^*,\tau+1} \neq M_{C_0}(X_{\tau+1})\}. \end{aligned}$$

By the assumptions of EBUG constituent $\{\phi_{x,\tau}\}$ and the existence of the value of the game, we have

$$\lim_{t\to\infty} \frac{\mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\}\right\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

where

$$K_{C_0} = \inf_{\theta>\theta_2} I(\theta_1, \theta|x^*) = \inf_{\theta>\theta_2} \sup_x I(\theta_1, \theta|x).$$

The proof of ARGUMENT 3 is complete. ■

C.4 Analysis of Algorithm 3

With the help of Lemma C.2, we prove Theorems 2.8 and 3.6 by showing that for all $\{X_\tau\}$ being u.s.e. distributed in L^1, the following arguments hold in Algorithm 3.

- ARGUMENT 1: The expected duration over which $\hat{C}_t$ does not exist is finite, namely,

$$\lim_{t\to\infty} \mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau \text{ does not exists}\}\right\} < \infty.$$

  Again we use $\chi\{\hat{C}_t\} = 0$ as shorthand for the situation in which $\hat{C}_t$ does not exist.

- ARGUMENT 2: The expected duration over which $\hat{C}_t \neq C_0$ is finite, namely,

$$\lim_{t\to\infty} \mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau \neq C_0\}\right\} < \infty.$$

- ARGUMENT 3: If C_0 is implicitly revealing, the expected duration over which $\hat{C}_t = C_0$ and $\Phi_{t+1} \neq M_{C_0}(X_{t+1})$ is finite, namely,

$$\lim_{t\to\infty} \mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\}\right\} < \infty.$$

- ARGUMENT 4: If C_0 is not implicitly revealing, the expected duration over which $\hat{C}_t = C_0$ and $\Phi_{t+1} \neq M_{C_0}(X_{t+1})$ is upper bounded by $\frac{\log(t)}{K_{C_0}}$, namely,

$$\lim_{t\to\infty} \frac{\mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\}\right\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

  where $K_{C_0} = \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \sup_x I(\theta_1, \theta|x)$ if $M_{C_0} = 2$.

With the above four arguments, it is straightforward to show that the $\{\Phi_\tau\}$ described in Algorithm 3 satisfies the statements in Theorems 2.8 and 3.6.

Proof of ARGUMENT 1*:* This proof follows word by word the proof of ARGUMENT 1 in Appendix C.3. ■

Proof of ARGUMENT 2*:* Since

$$\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau \neq C_0\} = \sum_{C' \neq C_0} \sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0\},$$

we would like to prove that for any $C' \neq C_0$, $\mathsf{E}_{C_0}\left\{\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0\}\right\}$ is bounded. For those C' that are not implicitly revealing, the proof follows word by word the proof of ARGUMENT 2 in Appendix C.3.

So we may assume that C' is implicitly revealing, and by conditioning on whether or not $\hat{C}_\tau = \ddot{C}_\tau$, we have

$$\begin{aligned}\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0\} &= \sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0, \ddot{C}_\tau \neq \hat{C}_\tau\} \\ &\quad + \sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0, \ddot{C}_\tau = \hat{C}_\tau\}.\end{aligned}$$

These two summations will be considered separately.

Let $C'' \neq C'$ denote another implicitly revealing parameter pair, and construct the stopping time pairs $S_{x,C',C'',j}$ and $T_{x,C',C'',j}$ iteratively as follows.

$$\begin{aligned}S_{x,C',C'',j} &:= \inf\Big\{t > S_{x,C',C'',j-1} : X_{t+1} = x, \hat{C}_t = C', \ddot{C}_t = C'', \\ &\qquad \text{and either } \hat{C}_{t-1} \neq C' \text{ or } \ddot{C}_{t-1} \neq C'' \text{ or } X_t \neq x\Big\},\end{aligned}$$

and

$$T_{x,C',C'',j} := \inf\Big\{t > S_{x,C',C'',j} : \text{ either } \hat{C}_t \neq C', \text{ or } \ddot{C}_t \neq C'', \text{ or } X_{t+1} = x\Big\},$$

where $S_{x,C',C'',0} = 0$. Note that $S_{x,C',C'',j}$ and $T_{x,C',C'',j}$ are basically dividing the duration over which $\{\hat{C}_t = C', \ddot{C}_t = C''\}$ into disjoint intervals when x specifies the value of the side observation X_{t+1} at the leading time instants of those intervals. Thus we have

$$\begin{aligned}\sum_{t=1}^{\infty} 1\{\hat{C}_t = C' \neq C_0, \ddot{C}_t \neq \hat{C}_t\} &= \sum_{C''} \sum_{t=1}^{\infty} 1\{\hat{C}_t = C', \ddot{C}_t = C''\} \\ &\leq \sum_{x,C''} \sum_{j \in \mathbb{N}} \left(T_{x,C',C'',j} - S_{x,C',C'',j} + 1\right).\end{aligned}$$

Since

$$T_{x,C',C'',j} \leq \inf\left\{t > S_{x,C',C'',j} : X_{t+1} = x\right\},$$

and by the assumption that $\{X_\tau\}$ is u.s.e. distributed in L^1, there exists a $B < \infty$ such that $\forall x, j$, $\mathsf{E}\left\{T_{x,C',C'',j} - S_{x,C',C'',j} + 1 | S_{x,C',C'',j}\right\} < B$. It we can show that

$$\forall x, C'', \exists K, \text{ s.t. } \mathsf{E}\left\{\sup\{j \in \mathbb{N} : S_{x,C',C'',j} < \infty\}\right\} < K,$$

then by Lemma C.2, we have $\mathsf{E}\left\{\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C' \neq C_0, \ddot{C}_\tau \neq \hat{C}_\tau\}\right\} < \infty$.

When $t = S_{x,C',C'',j}$, by the definition of Algorithm 3, for odd j the decision rule results in $\Phi_{t+1} \leftarrow \phi_{X_{t+1},t}$ (since at time t, $\mathrm{ctr}(x, C', C'') = j - 1$). Thus we have

$$\begin{aligned}\sup\left\{j \in \mathbb{N} : S_{x,C',C'',j} < \infty\right\} &= \sum_{j=1}^{\infty} 1\{S_{x,C',C'',j} < \infty\} \\ &\leq\ 2\sum_{\tau=1}^{\infty} 1\{X_{\tau+1} = x, \hat{C}_\tau = C' \neq C_0, \ddot{C}_\tau = C'', \Phi_{\tau+1} \leftarrow \phi_{x,\tau+1}\}.\end{aligned}$$

By the assumption of EBUG $\{\phi_{x,\tau}\}$, the above right-hand side has finite expectation.

For the case in which $\hat{C}_t = \ddot{C}_t = C' \neq C_0$, we construct the stopping time pairs as follows.

$$\begin{aligned}S_{x,C',j} \ :=\ \inf\Big\{&t > S_{x,C',j-1} : X_{t+1} = x, \hat{C}_t = \ddot{C}_t = C', \\ &\text{and either } \hat{C}_{t-1} \neq C' \text{ or } \ddot{C}_{t-1} \neq C' \text{ or } \{1,2\} = \bigcup_{\tau \in (S_{x,C',j-1},t]} \{M_{C'}(X_\tau)\}\Big\},\end{aligned}$$

and

$$T_{x,C',j} \ :=\ \inf\left\{t > S_{x,C',j} :\ \text{either } \hat{C}_t \neq C',\ \text{or } \ddot{C}_t \neq C',\ \text{or } \{1,2\} = \bigcup_{\tau \in (S_{x,C',j},t]} \{M_{C'}(X_\tau)\}\right\},$$

where $S_{x,C',0} = 0$. We then have

$$\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = \ddot{C}_\tau = C' \neq C_0\} \ \leq\ \sum_{x \in \mathbf{X}} \sum_{j \in \mathbb{N}} \left(T_{x,C',j} - S_{x,C',j} + 1\right).$$

Since

$$T_{x,C',j} \leq \inf\left\{t > S_{x,C',j} : \mathbf{X} = \bigcup_{\tau \in (S_{x,C',j},t]} \{X_\tau\}\right\},$$

and by the assumption that $\{X_\tau\}$ is u.s.e. distributed in L^1, there exists a $B < \infty$ such that $\forall x, C', j$, $\mathsf{E}\left\{T_{x,C',j} - S_{x,C',j} + 1 | S_{x,C',j}\right\} \leq B$. If we can show

$$\forall x \in \mathbf{X}, \mathsf{E}\left\{\sup\{j \in \mathbb{N} : S_{x,C',j} < \infty\}\right\} < \infty, \tag{C.14}$$

then by Lemma C.2, we have $\mathsf{E}\left\{\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = \ddot{C}_\tau = C' \neq C_0\}\right\} < \infty$.

We prove (C.14) by case study. For any fixed pair of (x, C') considered here, without loss of generality, we may assume $M_{C'}(x) = 1$. Recalling that $1(C)$ denotes the first coordinate of the configuration pair C, we consider the cases as follows.

- $1(C') \neq 1(C_0)$: When $t = S_{x,C',j}$, we then have $X_{t+1} = x$, $\Phi_{t+1} \leftarrow M_{\hat{C}_t}(X_{t+1}) = M_{C'}(x) = 1$, and

$$\begin{aligned}\sup\{j \in \mathbb{N} : S_{x,C',j} < \infty\} &= \sum_{j=1}^{\infty} 1\{S_{x,C',j} < \infty\} \\ &\leq \sum_{\tau=1}^{\infty} 1\{X_{\tau+1} = x, 1(\ddot{C}_\tau) = 1(C') \neq 1(C_0), \Phi_{\tau+1} \leftarrow M_{C'}(x) = 1\} \triangleq D_1.\end{aligned}$$

Since every time the event $\{X_{\tau+1} = x, 1(\ddot{C}_\tau) = 1(C') \neq 1(C_0), \Phi_{\tau+1} \leftarrow M_{C'}(x) = 1\}$ occurs, the effective sample size of arm 1 (used to generate $\{\ddot{C}_\tau\}$) increases by one. Because $\{\ddot{C}_\tau\}$ is a *good* estimate, the expectation of D_1 must be bounded. Thus, this case can at most contribute finite expectation.

- $1(C') = 1(C_0)$: This condition implies that $2(C') \neq 2(C_0)$. By noting that $\sup\{j \in \mathbb{N} : S_{x,C',j} < \infty\} \leq \sup\{j \in \mathbb{N} : T_{x,C',j} < \infty\}+1$, we prove that the latter can have at most finite expectation. When $t = T_{x,C',j} < \infty$, it follows that we have either $M_{C'}(X_t) = 2$ or $1(\ddot{C}_t) \neq 1(C_0)$. The reason is if $T_{x,C',j} < \infty$ is due to $\ddot{C}_t \neq C'$, then it is either $1(\ddot{C}_t) \neq 1(C_0)$ or $2(\ddot{C}_t) \neq 2(C')$. The latter implies $\Phi_t \leftarrow M_{\ddot{C}_{t-1}}(X_t) = M_{C'}(X_t) = 2$ since $2(\ddot{C}_{t-1}) = 2(C')$ changes after time t. Another possibility is when $T_{x,C',j} < \infty$ is due to $\{1, 2\} = \bigcup_{\tau\in(S_{x,C',j},t]}\{M_{C'}(X_\tau)\}$. Since $M_{C'}(X_{S_{x,C',j}+1}) = M_{C'}(x) = 1$, it implies also $M_{C'}(X_t) = 2$. From the above discussion, we have

$$\begin{aligned}
&\sup\{j \in \mathbb{N} : T_{x,C',j} < \infty\} = \sum_{j=1}^{\infty} 1\{T_{x,C',j} < \infty\} \\
&\leq \sum_{\tau=1}^{\infty} 1\{1(\ddot{C}_\tau) \neq 1(C_0), \ddot{C}_{\tau-1} = C', \Phi_\tau \leftarrow M_{\ddot{C}_{\tau-1}}(X_\tau) = 1\} \\
&\quad + \sum_{x':M_{C'}(x')=2} \sum_{\tau=1}^{\infty} 1\{X_\tau = x', \ddot{C}_{\tau-1} = C', 2(\ddot{C}_{\tau-1}) \neq 2(C_0), \Phi_\tau \leftarrow M_{\ddot{C}_{\tau-1}}(X_\tau) = 2\}.
\end{aligned}$$

Since the estimate $\{\ddot{C}_\tau\}$ is *good*, each infinite sum in the above equation has finite expectation. Thus we have proven that this case can contribute at most finite expectation.

From our treatment of the three cases: $\hat{C}_t$ is not implicitly revealing, $\hat{C}_t$ is implicitly revealing but $\hat{C}_t \neq \ddot{C}_t$, and $\hat{C}_t = \ddot{C}_t$ is implicitly revealing, the proof of ARGUMENT 2 is complete.

■

Proof of ARGUMENT 3*:* When $\hat{C}_t = C_0$, the only situation of sampling the inferior arm is $\ddot{C}_t \neq \hat{C}_t = C_0$. For any fixed $C' \neq C_0$, construct the stopping time pairs as follows.

$$\begin{aligned}
S_{C',j} \;:=\; \inf\Big\{ & t > S_{C',j-1} : \hat{C}_t = C_0, \ddot{C}_t = C', \\
& \text{and either } \hat{C}_{t-1} \neq C_0 \text{ or } \ddot{C}_{t-1} \neq C', \\
& \text{or } \{1, 2\} = \bigcup_{\tau\in\mathbf{S}_{j-1,t-1}} \{M_{C_0}(X_\tau)\} \Big\},
\end{aligned}$$

where $S_{C',0} = 0$ and

$$\mathbf{S}_{j-1,t-1} := \{\tau \in (S_{C',j-1}, t-1] : \text{the line } \Phi_\tau \leftarrow M_{\hat{C}_{\tau-1}}(X_\tau) = M_{C_0}(X_\tau) \text{ is active}\}.$$

For $T_{C',j}$, we have the following definition:

$$T_{C',j} \;:=\; \inf\left\{ t > S_{C',j} : \text{ either } \hat{C}_t \neq C_0, \text{ or } \ddot{C}_t \neq C', \text{ or } \{1, 2\} = \bigcup_{\tau\in\mathbf{S}_{j,t}} \{M_{C_0}(X_\tau)\} \right\}.$$

Since $S_{C',j}$ and $T_{C',j}$ partition the duration over which $\{\hat{C}_t = C_0, \ddot{C}_t = C'\}$ into disjoint intervals, we then have

$$\begin{aligned}\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C_0 \neq \ddot{C}_\tau\} &\leq \sum_{C' \neq C_0} \sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C_0, \ddot{C}_\tau = C'\} \\ &\leq \sum_{C' \neq C_0} \sum_{j \in \mathbb{N}} \left(T_{C',j} - S_{C',j} + 1\right).\end{aligned}$$

By Line 7 in Algorithm 3, for any $X_{t+1} = x$, $\hat{C}_t = C_0$, $\ddot{C}_t = C'$, the decision rule Φ_{t+1} is alternating between $\phi_{x,t}$ and $M_{C_0}(x)$. As a result, we have

$$T_{C',j} \leq \inf\{t > S_{C',j} : \forall x \in \mathbf{X}, \exists \tau_1 \neq \tau_2 \in (S_{C',j}, t] \text{ s.t. } X_{\tau_1} = X_{\tau_2} = x\}.$$

By the assumption that $\{X_\tau\}$ is u.s.e. distributed in L^1, there exists a $B < \infty$ such that $\forall C', j, \mathsf{E}\left\{T_{C',j} - S_{C',j} + 1 \middle| S_{C',j}\right\} \leq B$. If we can show

$$\forall x \in \mathbf{X}, C', \mathsf{E}\left\{\sup\{j \in \mathbb{N} : S_{C',j} < \infty\}\right\} < \infty,$$

then by Lemma C.2, we have $\mathsf{E}\left\{\sum_{t=1}^{\infty} 1\{\hat{C}_t = C_0 \neq \ddot{C}_t\}\right\} < \infty$

Since $\sup\{j : S_{C',j} < \infty\} \leq \sup\{j : T_{C',j} < \infty\} + 1$, equivalently, we can focus on proving $\mathsf{E}\left\{\sup\{j \in \mathbb{N} : T_{C',j} < \infty\}\right\} < \infty$. For any $j \in \mathbb{N}$, let $t := T_{C',j} < \infty$. Then one of the following situations must be true.

- $\Phi_t \leftarrow \phi_{X_t,t}$: The only situation under which the interval ends right after triggering $\Phi_t \leftarrow \phi_{X_t,t}$ is due to $\hat{C}_t \neq C_0$. Since the constituent $\{\phi_{x,\tau}\}$ is EBUG, this part contributes at most bounded expectation.
- $\Phi_t \leftarrow M_{C_0}(X_t)$: There are two ways in which the interval will end in this situation:
  - $\{1, 2\} = \bigcup_{\tau \in \mathbf{S}_{j,t}} \{M_{C_0}(X_\tau)\}$: In this case, both the samples of arm 1 and arm 2 used by $\{\ddot{C}_\tau\}$ must have increased by 1 while $\ddot{C}_{t-1} = C' \neq C_0$. Since $\{\ddot{C}_\tau\}$ is a good estimate, this portion contributes at most bounded expectation.
  - $\ddot{C}_t \neq C'$: Without loss of generality, we may assume $\Phi_t \leftarrow M_{C_0}(X_t) = 1$ and thus $\{1\} = \bigcup_{\tau \in \mathbf{S}_{j,t}} \{M_{C_0}(X_\tau)\}$. Two sub-cases are as follows:
    - $1(C') \neq 1(C_0)$: Since $\Phi_t \leftarrow M_{C_0}(X_t) = 1$, the number of samples from arm 1 used to generate $\{\ddot{C}_\tau\}$ must increase by 1 during the interval $[S_{C',j}, T_{C',j}]$. By the assumption that $\{\ddot{C}_\tau\}$ is good, that portion contributes at most finite expectation.
    - $1(C') = 1(C_0)$: Since $\Phi_t \leftarrow M_{C_0}(X_t) = 1$, for each j, the number of samples of arm 1 (used by $\{\ddot{C}_\tau\}$) during the interval $[S_{C',j}, T_{C',j}]$ increases by at least one. We also note that $\ddot{C}_t \neq C' = \ddot{C}_{t-1}$. Because at time t, $\Phi_t \leftarrow M_{C_0}(X_t) = 1$, only the first coordinate of $\ddot{C}_{t-1}$ may change, we obtain $1(\ddot{C}_t) \neq 1(\ddot{C}_{t-1}) = 1(C_0)$. Combining the above observations and the assumption that $\{\ddot{C}_\tau\}$ is good, this portion can contribute at most bounded expectation.

From the above discussions, we have

$$\mathsf{E}\left\{\sum_{\tau=1}^{\infty} 1\{\hat{C}_\tau = C_0 \neq \ddot{C}_\tau\}\right\} < \infty.$$

■

Proof of ARGUMENT 4*:* Suppose $C_0 = (\theta_1, \theta_2)$, $M_{C_0} = 2$ and define x^* as the maximizing argument of $\inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} I(\theta_1, \theta|x)$. We then have

$$\begin{aligned}
&\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\} \\
&= \sum_{\tau=1}^{t} 1\{\hat{C}_\tau = \hat{C}_{x^*,\tau} = C_0, X_{\tau+1} = x^*, \Phi_{\tau+1} \leftarrow \phi_{x^*,\tau+1} \neq M_{C_0}(X_{\tau+1})\} \\
&\leq \sum_{\tau=1}^{t} 1\{\hat{C}_{x^*,\tau} = C_0, X_{\tau+1} = x^*, \Phi_{\tau+1} \leftarrow \phi_{x^*,\tau+1} \neq M_{C_0}(X_{\tau+1})\}.
\end{aligned}$$

By the assumptions of EBUG constituent $\{\phi_{x,\tau}\}$ and by the existence of the saddle point, we have

$$\lim_{t\to\infty} \frac{\mathsf{E}\left\{\sum_{\tau=1}^{t} 1\{\hat{C}_\tau = C_0, \Phi_{\tau+1} \neq M_{C_0}(X_{\tau+1})\}\right\}}{\log(t)} \leq \frac{1}{K_{C_0}},$$

where

$$\begin{aligned}
K_{C_0} &= \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} I(\theta_1, \theta|x^*) \\
&= \sup_x \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} I(\theta_1, \theta|x) \\
&= \inf_{\{\theta:\exists x_0, \mu_\theta(x_0) > \mu_{\theta_2}(x_0)\}} \sup_x I(\theta_1, \theta|x).
\end{aligned}$$

The proof of ARGUMENT 4 is then complete. ■

Appendix D

Relationships Among Evenly Distribution Properties

Let $\mathcal{X}_{u.s.e.}$ denote the collection of random processes $\{X_\tau\}$ that are u.s.e. distributed in L^1, and $\mathcal{X}_{p.s.}$, $\mathcal{X}_{p.}$ and $\mathcal{X}_{L^1}$ denote the corresponding collections such that $\{X_\tau\}$ is evenly distributed in probability series, evenly distributed in probability, and evenly distributed in L^1, respectively. We then have the following proposition.

Proposition D.1 *Using the above notation, it can be shown that*

$$\mathcal{X}_{u.s.e.} \subsetneq \mathcal{X}_{p.s.} \subsetneq \mathcal{X}_{p.} \subsetneq \mathcal{X}_{L^1}.$$

Proof of $\mathcal{X}_{p.} \subsetneq \mathcal{X}_{L^1}$: We first prove that $\mathcal{X}_{p.} \subset \mathcal{X}_{L^1}$. Suppose $\{X_\tau\} \in \mathcal{X}_{p.}$. By definition, $\exists \pi(\cdot) > 0$ such that

$$\lim_{\tau\to\infty} \mathsf{P}(f_r(x,\tau) < \pi(x)) = 0 \Longleftrightarrow \lim_{\tau\to\infty} \mathsf{P}(f_r(x,\tau) \geq \pi(x)) = 1.$$

And since

$$\mathsf{E}\{f_r(x,\tau)\} \geq \pi(x) \cdot \mathsf{P}(f_r(x,\tau) \geq \pi(x)),$$

we have

$$\liminf_{\tau\to\infty} \mathsf{E}\{f_r(x,\tau)\} \geq \pi(x).$$

From the above reasoning, we have $\{X_\tau\} \in \mathcal{X}_{L^1}$ and $\mathcal{X}_{p.} \subset \mathcal{X}_{L^1}$.

An $\{X_\tau\} \in \mathcal{X}_{L^1} \backslash \mathcal{X}_{p.}$ can be constructed as follows. Suppose $\mathbf{X} = \{0, 1\}$ and consider the following two deterministic sequences $\mathbf{u}$ and $\mathbf{v}$, such that $\mathbf{u} = 010101\cdots$ is alternating between 0 and 1 while $\mathbf{v} = 000\cdots$ is all zero. Let $\{X_\tau\}$ equal $\mathbf{u}$ and $\mathbf{v}$, each with probability 1/2. It is easy to verify that $\{X_\tau\} \in \mathcal{X}_{L^1}$ but not in $\mathcal{X}_{p.}$.

■

Proof of $\mathcal{X}_{p.s.} \subsetneq \mathcal{X}_{p.}$: By definition, for any $\{X_\tau\} \in \mathcal{X}_{p.s.}$, it follows that there exists a strictly positive mapping $\pi(\cdot)$ such that for all possible $x \in \mathbf{X}$,

$$\sum_{\tau=1}^{\infty} \mathsf{P}(f_r(x,\tau) < \pi(x)) < \infty \implies \lim_{\tau\to\infty} \mathsf{P}(f_r(x,\tau) < \pi(x)) = 0.$$

Thus $\{X_\tau\} \in \mathcal{X}_{p.}$.

An $\{X_\tau\} \in \mathcal{X}_{p.} \backslash \mathcal{X}_{p.s.}$ can be constructed as follows. Suppose $\mathbf{X} = \{0, 1\}$ and define a family of deterministic sequences $\{\mathbf{u}_{(j)}\}$ as follows. $\mathbf{u}_{(0)} = 101010\cdots$ is periodically alternating between 0 and 1. $\mathbf{u}_{(j)}$ is obtained by appending a prefix 0 to $\mathbf{u}_{(j-1)}$, namely, $\mathbf{u}_{(j)} = 0\mathbf{u}_{(j-1)}$. We then have

$$\begin{aligned}
\mathbf{u}_{(1)} &= \{0, 1, 0, 1, 0, 1, 0, 1, 0, 1, \cdots\} \\
\mathbf{u}_{(2)} &= \{0, 0, 1, 0, 1, 0, 1, 0, 1, 0, \cdots\} \\
\mathbf{u}_{(3)} &= \{0, 0, 0, 1, 0, 1, 0, 1, 0, 1, \cdots\}.
\end{aligned}$$

Then a non-ergodic $\{X_\tau\}$ is defined as $\mathsf{P}(\{X_\tau\} = \mathbf{u}_{(j)}) = K \cdot \frac{1}{j^2}$, where K is a normalization constant such that $\sum_{j=1}^{\infty} \mathsf{P}(\{X_\tau\} = \mathbf{u}_{(j)}) = 1$. For such $\{X_\tau\}$, one can easily show that any strictly positive mapping $\pi(\cdot) > 0$ satisfies

$$\lim_{\tau \to \infty} \mathsf{P}(f_r(x, \tau) < \pi(x)) = 0,$$

but

$$\sum_{\tau=1}^{\infty} \mathsf{P}(f_r(x, \tau) < \pi(x)) = \infty.$$

Therefore, $\{X_\tau\} \in \mathcal{X}_{p.} \backslash \mathcal{X}_{p.s.}$, and the proof is complete. ■

Proof of $\mathcal{X}_{u.s.e.} \subsetneq \mathcal{X}_{p.s.}$:

We prove $\mathcal{X}_{u.s.e.} \subset \mathcal{X}_{p.s.}$ by showing that $\forall \{X_\tau\} \in \mathcal{X}_{u.s.e.}$, $\forall x \in \mathbf{X}$, $\exists \pi_0, a > 0$ such that

$$\mathsf{P}\left(f_r(x, \tau) < \pi\right) \leq e^{-a\tau}, \quad \forall 0 < \pi < \pi_0. \tag{D.1}$$

Then, by definition, $\{X_\tau\}$ belongs to $\mathcal{X}_{p.s.}$.

To that end, we first show that there exist an $\epsilon > 0$ and a $B' < \infty$ such that for any stopping time T,

$$\mathsf{E}\left\{ e^{\lambda H_T(x)} \middle| T \right\} \leq B', \quad \forall \lambda \in (0, \epsilon). \tag{D.2}$$

By defining $T_{(j)}$ as the j-th time instant such that $X_\tau = x$, we have

$$\begin{aligned}
\mathsf{P}\left(f_r(x, \tau) < \pi\right) &= \mathsf{P}\left(T_{(\pi\tau)} > \tau\right) \\
&\leq \left(\frac{B'}{e^{\lambda \frac{1}{\pi}}}\right)^{\pi\tau},
\end{aligned}$$

where the last inequality follows from the Chernoff bound. By selecting a sufficiently small π such that $B' < e^{\lambda \frac{1}{\pi}}$, we have proven (D.1) and thus $\{X_\tau\} \subset \mathcal{X}_{p.s.}$

(D.2) is proved as follows. Consider any stopping time T and a corresponding stopping time S, where S is defined as $T + 10B$ iff[1] $\forall \tau \in (T, T + 10B]$, $X_\tau \neq x$, and B is the global upper bound of the conditional hitting time in the definition of $\mathcal{X}_{u.s.e.}$. By definition,

$$\begin{aligned}
\mathsf{E}\left\{ e^{\lambda H_T(x)} \middle| T \right\} &\leq e^{\lambda 10B} + \mathsf{P}\left(H_T(x) > 10B \middle| T\right) \mathsf{E}\left\{ e^{\lambda H_T(x)} \middle| T, H_T(x) > 10B \right\} \\
&\leq e^{\lambda 10B} + \frac{1}{10} \mathsf{E}\left\{ e^{\lambda(10B + H_S(x))} \middle| S \right\} \\
&= e^{\lambda 10B} + \frac{e^{\lambda 10B}}{10} \mathsf{E}\left\{ e^{\lambda H_S(x)} \middle| S \right\}.
\end{aligned} \tag{D.3}$$

[1] If $\exists \tau \in (T, T + 10B]$ such that $X_\tau = x$, then S is not defined.

By recursively applying the same procedure on $\mathsf{E}\left\{ e^{\lambda H_S(x)} \middle| S \right\}$, it can thus be shown[2] that when λ is sufficiently small such that $\frac{e^{\lambda 10B}}{10} < 1$, we have

$$\mathsf{E}\left\{ e^{\lambda H_T(x)} \middle| T \right\} \leq \frac{e^{\lambda 10B}}{1 - \frac{e^{\lambda 10B}}{10}}.$$

The proof of (D.2) is then complete.

An $\{X_\tau\} \in \mathcal{X}_{p.s.} \backslash \mathcal{X}_{u.s.e.}$ can be explicitly constructed. Suppose $\mathbf{X} = \{0, 1\}$. Let $\mathbf{u}$ be a deterministic sequence of the following form: $\mathbf{u} = 0^{s_1}1^{t_1}0^{s_2}1^{t_2}\cdots$, where 0^s (1^s) represents a sequence of all zeros (ones) of length s. For instance, $\mathbf{u} = 0^1 1^2 0^2 1^3 \cdots = 01100111\cdots$. We can then sequentially define s_i and t_i as follows, and thus a specific $\mathbf{u}_0$ is constructed.

$$\begin{aligned} s_i &= \min\left\{ s > 0 : \frac{s + \sum_{j=1}^{i-1} s_j}{s + \sum_{j=1}^{i-1}(s_j + t_j)} > 2/3 \right\} \\ t_i &= \min\left\{ t > 0 : \frac{\sum_{j=1}^{i} s_j}{t + s_i + \sum_{j=1}^{i-1}(s_j + t_j)} < 1/3 \right\}. \end{aligned}$$

Following this construction, $s_1 = 1$, $t_1 = 3$, $s_2 = 6$, $t_2 = 12$, and the first 12 bits of $\mathbf{u}_0$ is $011100000011\cdots$. The resulting $f_r(0,t)$ keeps growing until it hits $2/3$, then starts to decrease until it hits $1/3$, then keeps growing again, and repeats this cycle indefinitely. Based on this construction, one can easily show that $\mathbf{u}_0 \in \mathcal{X}_{p.s.}$ but not in $\mathcal{X}_{u.s.e.}$. ■

It is worth mentioning that if $\{X_\tau\}$ is u.s.e. distributed in L^1 then the hitting time after any fixed time t has bounded expectation. However, the converse statement is not true. This relationship is formally stated as follows.

Definition D.1 (Uniformly Evenly (u.e.) Distributed in L^1) *$\{X_\tau\}$ is u.e. distributed in L^1, if for any deterministic time t_0, the expectation of the first hitting time of x after t_0 has a global upper bound. That is, $\exists B < \infty$ such that*

$$\forall t_0 \in \mathbb{N}, \forall x \in \mathbf{X},\ \mathsf{E}\{H_{t_0}(x)\} \leq B,$$

where $H_{t_0}(x) \triangleq \inf\{l > 0 | X_{t_0+l} = x\}$.

Let $\mathcal{X}_{u.e.}$ denote the collection of random processes being uniformly evenly distributed in L^1. We have the following proposition.

Proposition D.2

$$\mathcal{X}_{u.s.e.} \subsetneq \mathcal{X}_{u.e.} \tag{D.4}$$

Proof: By definition, $\mathcal{X}_{u.s.e.} \subset \mathcal{X}_{u.e.}$. For the following, we will construct an example showing that the $\mathcal{X}_{u.s.e.} \neq \mathcal{X}_{u.e.}$. Let $\mathbf{u}_{(i)}$ denote a collection of deterministic periodic sequences with period equal to $i+1$. In each period, $\mathbf{u}_{(i)}$ contains a length $(i+1)$ interval starting with a 1, and followed by all 0's. This is illustrated for $i = 1, 2, 3$:

$$\begin{aligned} \mathbf{u}_{(1)} &= \{1,0,1,0,1,0,1,0,1,0,\cdots\} \\ \mathbf{u}_{(2)} &= \{1,0,0,1,0,0,1,0,0,1,\cdots\} \\ \mathbf{u}_{(3)} &= \{1,0,0,0,1,0,0,0,1,0,\cdots\}. \end{aligned}$$

[2] A rigorous argument requires using the truncated hitting time $H_{T,n}(x) = \min(n, H_T(x))$, so that the recursive application of (D.3) will terminate. A global upper bound B' is then established for the truncated $H_{T,n}(x)$. By letting $n \to \infty$, by using the monotone convergence theorem, and by the fact that $H_T(x) < \infty$ almost surely, we can established (D.2) for non-truncated $H_T(x)$.

We then construct a $\{X_\tau\}$ such that $\mathsf{P}(\{X_\tau\} = \mathbf{u}_{(i)}) = K \cdot \frac{1}{i^3}$, where K is a normalization factor such that $\sum_{i=1}^{\infty} \mathsf{P}(\{X_\tau\} = \mathbf{u}_{(i)}) = 1$.

According to the above construction, it is easy to check that $\{X_\tau\} \in \mathcal{X}_{u.e.}$. However, for any $B < \infty$, we can choose our stopping time T as the first time instant at which we have found B consecutive 0's, followed by a single 1. We can then easily show that

$$\mathsf{E}\{H_T(1)|T\} \geq B + 1,$$

and thus $\{X_\tau\} \notin \mathcal{X}_{u.s.e.}$ ■

Remarks:

- The example is nontrivial since $\mathsf{P}(T < \infty) > 0$ for all B.
- The above is a good example to illustrate why the definition of the desired even distribution involves stopping times when considering the benefit of observing side information $\{X_\tau\}$. In particular, for k large, the unevenly distributed sequence $\mathbf{u}_{(k)}$ will cause a significant amount of inferior sampling times for the decision scheme.

Appendix E

Proofs of the Perfect Projection Convergence and the Typicality Theorems

E.1 Proof of Theorem 4.2

We first introduce a straightforward corollary of Theorem 4.1:

Corollary E.1 (Cycle-free Convergence for Growing Trees) *For a sequence*

$$l_n = \frac{4}{9}\frac{\ln n}{\ln(d_v-1)+\ln(d_c-1)},$$

we have for any i_0, j_0*,*

$$\mathsf{P}\left(\mathcal{N}^{2l_n}_{(i_0,j_0)} \text{ is cycle-free}\right) = 1 - \mathcal{O}\left(n^{-1/9}\right).$$

With this corollary, the proof of Theorem 4.2 proceeds as follows. Throughout this proof, the subscript (i_0, j_0) will be omitted for notational simplicity.

Proof of Theorem 4.2*:* We first notice that if for any $l_n \geq l$, $\mathcal{N}^{2l_n}$ is perfectly projected, then so is $\mathcal{N}^{2l}$. Choose $l_n = \frac{4}{9}\frac{\ln n}{\ln(d_v-1)+\ln(d_c-1)}$. By Corollary E.1, we have

$$\begin{aligned}
&\mathsf{P}(\mathcal{N}^{2l} \text{ is perfectly projected})\\
\geq\ &\mathsf{P}(\mathcal{N}^{2l_n} \text{ is perfectly projected})\\
\geq\ &\mathsf{P}\left(\mathcal{N}^{2l_n} \text{ is perfectly projected}\middle|\, \mathcal{N}^{2(l_n+1)} \text{ is cycle-free}\right)\mathsf{P}\left(\mathcal{N}^{2(l_n+1)} \text{ is cycle-free}\right)\\
=\ &\mathsf{P}\left(\mathcal{N}^{2l_n} \text{ is perfectly projected}\middle|\, \mathcal{N}^{2(l_n+1)} \text{ is cycle-free}\right)\left(1-\mathcal{O}\left(n^{-1/9}\right)\right).
\end{aligned}$$

We then need only to show that

$$\mathsf{P}(\mathcal{N}^{2l_n} \text{ is perfectly projected}|\mathcal{N}^{2(l_n+1)} \text{ is cycle-free}) = 1 - \mathcal{O}\left(n^{-0.1}\right). \tag{E.1}$$

To prove (E.1), we take a deeper look at the incidence matrix (the parity check matrix) $\mathbf{H}$, and use the $(3, 5)$ regular code as our illustrative example. The proof is nonetheless general for all regular code ensembles. Conditioning on the event that the graph is cycle-free until

depth $2 \cdot 2$, we can transform $\mathbf{H}$ into the following form by row and column swaps, with $\otimes$ denoting the Kronecker product (whether it represents convolution or Kronecker product should be clear from the context).

$\mathbf{H} \quad =$

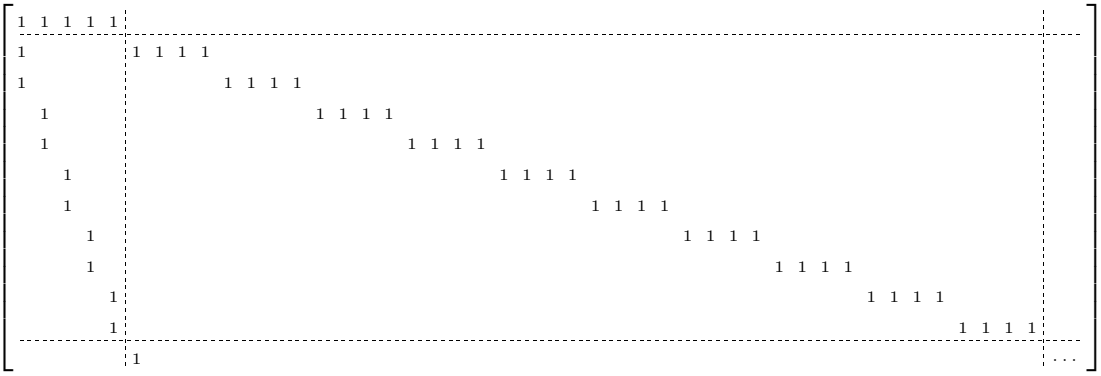

$$= \left(\begin{array}{c|c|c} \mathbf{I}_{1\times 1} \otimes (1,1,1,1,1) & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{I}_{5\times 5} \otimes \binom{1}{1} & \mathbf{I}_{10\times 10} \otimes (1,1,1,1) & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{I}_{40\times 40} \otimes \binom{1}{1} & \mathbf{H}'_{80\times(n-45)} \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{H}''_{(\frac{3n}{5}-91)\times(n-45)} \end{array}\right),$$

where $\mathbf{I}_{a\times a}$ denotes the $a \times a$ identity matrix, $\begin{pmatrix} \mathbf{H}'_{80\times(n-45)} \\ \mathbf{H}''_{(\frac{3n}{5}-91)\times(n-45)} \end{pmatrix}$ is the incidence matrix of the equiprobable, bipartite subgraph, in which all $(n-45)$ variable nodes have degree d_v, 80 check nodes have degree $d_c - 1$, and $(\frac{3n}{5} - 91)$ check nodes have degree d_c. Conditioning on a more general event that the graph is cycle free until depth $2(l_n + 1)$ instead of $2 \cdot 2$, we will have

$$\mathbf{H} = \left(\begin{array}{c|c|c} \mathbf{H}_{l_n} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} \;\big|\; \mathbf{I}_{(5\cdot 8^{l_n-1})\times(5\cdot 8^{l_n-1})} \otimes \binom{1}{1} & \mathbf{I}_{(10\cdot 8^{l_n-1})\times(10\cdot 8^{l_n-1})}(1,1,1) & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{I}_{(5\cdot 8^{l_n})\times(5\cdot 8^{l_n})} \otimes \binom{1}{1} & \mathbf{H}' \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{H}'' \end{array}\right),$$

where $\mathbf{H}_{l_n}$ corresponds to the incidence matrix of the cycle-free graph of depth $2l_n$. $\binom{\mathbf{H}'}{\mathbf{H}''}$ is the incidence matrix with rows (check nodes) in $\mathbf{H}'$ and $\mathbf{H}''$ having degree $(d_c - 1)$ and d_c. For convenience, we denote the blocks in $\mathbf{H}$ as

$$\mathbf{H} = \left(\begin{array}{c|c|c} \mathbf{H}_{l_n} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} \;\big|\; \mathbf{T}_{l_n} & \mathbf{U}_{l_n+1} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{T}_{l_n+1} & \mathbf{H}' \\ \hline \mathbf{0} & \mathbf{0} & \mathbf{H}'' \end{array}\right).$$

Then $\mathcal{N}^{2l_n}$ is *not* perfectly projected if and only if there exists a *non-zero* row vector $(\mathbf{r}|\mathbf{0}|\mathbf{0})$ such that

$$(\mathbf{r}|\mathbf{0}|\mathbf{0}) \in \text{RowSpace}\left(\begin{array}{cc|c|c} \mathbf{0} & \mathbf{T}_{l_n} & \mathbf{U}_{l_n+1} & \mathbf{0} \\ \hline \multicolumn{2}{c|}{\mathbf{0}} & \mathbf{T}_{l_n+1} & \mathbf{H}' \\ \hline \multicolumn{2}{c|}{\mathbf{0}} & \mathbf{0} & \mathbf{H}'' \end{array}\right), \tag{E.2}$$

and

$$\begin{aligned} &\mathbf{r} \text{ is not in the row space of } \mathbf{H}_{l_n}, \\ &\text{or equivalently } (\mathbf{r}|\mathbf{0}|\mathbf{0}) \text{ is not in RowSpace}(\mathbf{H}_{l_n}|\mathbf{0}|\mathbf{0}). \end{aligned} \tag{E.3}$$

(E.3) and (E.2), in this order, say that there exists a constraint $\mathbf{r}$ on the variable nodes of $\mathcal{N}^{2l_n}$, which is not from the linear combination of those check node equations within $\mathcal{N}^{2l_n}$, but rather is imposed by the parity check equations outside $\mathcal{N}^{2l_n}$. It can be easily proved that if the matrix $\binom{\mathbf{H}'}{\mathbf{H}''}$ is of full row rank, then no such $\mathbf{r}$ exists and $\mathcal{N}^{2l_n}$ is perfectly projected.[1] Instead of proving $\binom{\mathbf{H}'}{\mathbf{H}''}$ is of full rank, we take a different approach, which uses a novel constraint propagation argument.

From (E.2), we know that, for $(\mathbf{r}|\mathbf{0}|\mathbf{0})$ to exist, there must exist a *non-zero* row vector $(\mathbf{0}|\mathbf{s}|\mathbf{0})$ such that

$$(\mathbf{0}|\mathbf{s}|\mathbf{0}) \in \text{RowSpace}\left(\begin{array}{c|c|c} \mathbf{0} & \mathbf{T}_{l_n+1} & \mathbf{H}' \\ \mathbf{0} & \mathbf{0} & \mathbf{H}'' \end{array}\right), \tag{E.4}$$

and

$$\mathbf{s} \in \text{RowSpace}(\mathbf{U}_{l_n+1}) = \text{RowSpace}\left(I_{(10\cdot 8^{l_n-1})\times(10\cdot 8^{l_n-1})} \otimes (1,1,1,1)\right). \tag{E.5}$$

From (E.5), the 1's in $\mathbf{s}$ must be aligned such that four neighboring bits should have the same value; for example, a possible choice of $\mathbf{s}$ is $\mathbf{s} = (11110000111100000000111\cdots 00001111)$.

Any non-zero $\mathbf{s}$ satisfying (E.4) is generated by $\mathbf{T}_{l_n+1}$. By applying the row symmetry in $\mathbf{H}'$, we see that the 1's in any $\mathbf{s}$ are uniformly distributed among all these $|\mathbf{s}| = 5\cdot 8^{l_n}$ bits. Therefore, conditioning on the event that there exists a not-all-one $\mathbf{s}$ satisfying (E.4), the probability that $\mathbf{s}$ satisfies (E.5) is

$$\begin{aligned} &\mathsf{P}\left(\mathbf{s} \text{ satisfies (E.5)}|\exists \mathbf{s} \text{ satisfies (E.4) and is not } \mathbf{1}\right) \\ =\ &\mathsf{P}\left(\text{the 1's in } \mathbf{s} \text{ are aligned}|\exists \mathbf{s} \text{ satisfies (E.4) and is not } \mathbf{1}\right) \\ =\ &\sum_{a=1}^{10\cdot 8^{l_n-1}-1} \frac{\binom{10\cdot 8^{l_n-1}}{a}}{\binom{5\cdot 8^{l_n}}{4a}} \cdot \mathsf{P}(\text{there are } 4a \text{ ones in } \mathbf{s}) \\ \le\ &\frac{\binom{10\cdot 8^{l_n-1}}{1}}{\binom{5\cdot 8^{l_n}}{4}} = \mathcal{O}\left(\left(\frac{1}{((d_v-1)(d_c-1))^{l_n}}\right)^{d_c-2}\right) = \mathcal{O}\left(n^{-\frac{4}{9}(d_c-2)}\right). \end{aligned} \tag{E.6}$$

The last inequality follows from the assumption that $\mathbf{s}$ is neither all-zero nor all-one. The reason why we can exclude the case that $\mathbf{s}$ is all-one is that, if d_v is odd, then there is an even number of 1's in each column of $\mathbf{T}_{l_n}$. Since there is only one 1 in each column of $\mathbf{U}_{l_n+1}$, by (E.2), an all-one $\mathbf{s}$ can only generate an all-zero $\mathbf{r}$, which puts no constraints on $\mathcal{N}^{2l_n}_{(i_0,j_0)}$. If d_v is even, by the same reasoning, an all-one $\mathbf{s}$ will generate $\mathbf{r}$ of the form $(00\cdots 0\overbrace{11\cdots 1}^{5\cdot 8^{l_n-1}})$. Nevertheless, when d_v is even, this specific type of $\mathbf{r}$ is in the row space of $\mathbf{H}_{l_n}$, which does not fulfill the requirement in (E.3). From the above reasoning, we can exclude the all-one $\mathbf{s}$.

Let m_r denote the number of rows of $\binom{\mathbf{H}'}{\mathbf{H}''}$ minus $\mathsf{Rank}(\binom{\mathbf{H}'}{\mathbf{H}''})$. The number of vectors $\mathbf{s}$ satisfying (E.4) is upper bounded by 2^{m_r}. By (E.6), Proposition E.1 (which will be formally

[1] Unfortunately, $\binom{\mathbf{H}'}{\mathbf{H}''}$ is *not* of full row rank. We can only show that with sufficiently large n, the row rank of $\binom{\mathbf{H}'}{\mathbf{H}''}$ converges to the number of rows minus one by methods similar to those in [82]. A simple constraint propagation argument is still necessary for this approach.

stated and proved later), and the union bound, we have

$$\begin{aligned}
&\mathsf{P}(\mathcal{N}^{2l_n}_{(i_0,j_0)} \text{ is not perfectly projected}|\mathcal{N}^{2(l_n+1)}_{(i_0,j_0)} \text{ is cycle-free}) \\
&= \mathsf{P}(\exists \mathbf{r} \text{ satisfying (E.2) and (E.3)}) \\
&= \mathsf{P}(\exists \mathbf{s}, \text{ which satisfies (E.4) and (E.5), but is not all-one}) \\
&\leq n^{1.1} \cdot \mathsf{P}\left(\mathbf{s} \text{ satisfies (E.5)}|\exists \mathbf{s} \text{ satisfies (E.4) and is not } \mathbf{1}\right) \\
&\qquad \cdot \mathsf{P}\left(\#\text{ of } \mathbf{s} \text{ is smaller than } n^{1.1}\right) \\
&\qquad + \mathsf{P}\left(\#\text{ of } \mathbf{s} \text{ is larger than } n^{1.1}\right) \\
&= n^{1.1}\mathcal{O}(n^{-\frac{4}{9}(d_c-2)}) + \mathsf{P}(2^{m_r} > n^{1.1}) = \mathcal{O}(n^{-0.1}), \quad \forall d_c \geq 5. \qquad \text{(E.7)}
\end{aligned}$$

To prove the cases in which $3 \leq d_c < 5$, we focus on the probability that the constraints propagate two levels rather than just one level, i.e. instead of (E.1), we focus on proving the following statement:

$$\mathsf{P}(\mathcal{N}^{2l_n}_{(i_0,j_0)} \text{ is perfectly projected}|\mathcal{N}^{2(l_n+2)}_{(i_0,j_0)} \text{ is cycle-free}) = 1 - \mathcal{O}\left(n^{-0.1}\right).$$

Most of the analysis remains the same. The conditional probability in (E.6) will be replaced by

$$\begin{aligned}
&\mathsf{P}((\mathbf{0}|\mathbf{s}|\mathbf{0}) \text{ is able to propagate two levels}|\exists \mathbf{s} \text{ satisfying (E.4)}) \\
&= \mathsf{P}((\mathbf{0}|\mathbf{s}|\mathbf{0}) \text{ propagates the 2nd level}|(\mathbf{0}|\mathbf{s}|\mathbf{0}) \text{ propagates the 1st level}, \exists \mathbf{s} \text{ satisfying (E.4)}) \\
&\qquad \cdot \mathsf{P}((\mathbf{0}|\mathbf{s}|\mathbf{0}) \text{ propagates the 1st level}|\exists \mathbf{s} \text{ satisfying (E.4)}) \\
&= \sum_{a,b} \frac{\binom{10\cdot 8^{l_n-1}}{a}}{\binom{5\cdot 8^{l_n}}{4a}} \frac{\binom{10\cdot 8^{l_n}}{b}}{\binom{5\cdot 8^{l_n+1}}{4b}} \mathsf{P}(\text{there are } 4a \text{ and } 4b \text{ 1's to pass through the 2nd and 1st levels}) \\
&\overset{(a)}{\leq} \frac{\binom{10\cdot 8^{l_n-1}}{1}}{\binom{5\cdot 8^{l_n}}{4}} \frac{\binom{10\cdot 8^{l_n 1}}{4}}{\binom{5\cdot 8^{l_n+1}}{4\cdot 4}} \\
&= \mathcal{O}\left(\left(\frac{1}{((d_v-1)(d_c-1))^{l_n}}\right)^{d_c-2}\left(\frac{1}{((d_v-1)(d_c-1))^{l_n}}\right)^{(d_c-1)(d_c-2)}\right) \\
&= \mathcal{O}\left(n^{-\frac{4}{9}(d_c^2-2d_c)}\right),
\end{aligned}$$

where the inequality marked (a) follows from an analysis of the minimum number of bits required for the constraint propagation, which is similar to that for the single level case. By this stronger inequality and a bounding inequality similar to that in (E.7), we thus complete the proof of the case $d_c \geq 3$ for all regular codes of practical interest.

■

Note: This constraint propagation argument shows that the convergence to a perfectly projected tree is very strong. Even for codes with redundant check node equations (not of full row rank), it is probabilistically hard for the external constraints to propagate inward and impose on the variable nodes within $\mathcal{N}^{2l}$. This property is helpful when we consider belief propagation decoding on the alternative graph representation as in [64].

We close this section by stating the proposition regarding m_r, the number of linearly dependent rows in $\binom{\mathbf{H}'}{\mathbf{H}''}$. The proof is left to Appendix E.2.

Proposition E.1 *Consider the semi-regular code ensemble $\mathcal{C}^n_{m',m''}(d_v, d_c)$ generated by equiprobable edge permutation on a bipartite graph with n variable nodes of degree d_v, and m' and*

m'' check nodes with respective degrees $(d_c - 1)$ and d_c. The corresponding parity check matrix is $\mathbf{H} = \binom{\mathbf{H}'}{\mathbf{H}''}$. With m_r denoting the number of linearly dependent rows in $\mathbf{H}$, i.e. $m_r := m' + m'' - \mathsf{Rank}(\mathbf{H})$, we have

$$\mathsf{E}\{2^{m_r}\} = \mathcal{O}(n),$$

which automatically implies $\mathsf{P}\left(2^{m_r} > n^{1+\alpha}\right) = \mathsf{P}\left(m_r > \frac{(1+\alpha)\ln n}{\ln 2}\right) = \mathcal{O}(n^{-\alpha})$, for any $\alpha > 0$.

Corollary E.2 *Let R denote the rate of a regular LDPC code ensemble $\mathcal{C}^n(d_v, d_c)$, i.e., $R = \frac{n - \mathsf{Rank}(\mathbf{H})}{n}$, where $\mathbf{H}$ is the corresponding parity check matrix. Then R converges to $(n-m)/n$ in L^1, i.e.*

$$\lim_{n\to\infty} \mathsf{E}\left\{\left|R - \frac{n-m}{m}\right|\right\} = 0.$$

Proof: It is obvious that $R \geq \frac{n-m}{n}$. To show that $\limsup_{n\to\infty} \mathsf{E}\{R - \frac{n-m}{n}\} = 0$, we let $m_1 = 0$ and rewrite $R = \frac{n - \mathsf{Rank}(\mathbf{H})}{n} = \frac{n-m}{n} + \frac{m_r}{n}$. By Proposition E.1 and the fact that $\frac{m_r}{n} \leq 1$, we have $\lim_{n\to\infty} \mathsf{E}\{\frac{m_r}{n}\} = 0$. This completes the proof. ■

A stronger version of the convergence of R, proved by the code weight distribution argument, can be found in [82].

E.2 Proof of Proposition E.1

We finish the proof of Proposition E.1 by first stating the following lemma.

Lemma E.1 *For all $0 < k \in \mathbb{N}, 0 < i < n \in \mathbb{N}$, we have*

$$\frac{\binom{n}{i}}{\binom{kn}{ki}} \leq \sqrt{k} e^{\frac{1}{6}} 2^{-(k-1)nH_2(i/n)},$$

where $H_2(\cdot)$ is the binary entropy function: $H_2(p) = -p\log_2(p) - (1-p)\log_2(1-p)$.

Proof: By Stirling's double inequality,

$$\sqrt{2\pi} n^{(n+\frac{1}{2})} e^{(-n+\frac{1}{12n+1})} < n! < \sqrt{2\pi} n^{(n+\frac{1}{2})} e^{(-n+\frac{1}{12n})},$$

we can prove

$$\frac{1}{\sqrt{2\pi}} 2^{nH_2(i/n)} \sqrt{\frac{n}{i(n-i)}} e^{-\frac{1}{6}} < \binom{n}{\theta n} < \frac{1}{\sqrt{2\pi}} 2^{nH_2(i/n)} \sqrt{\frac{n}{i(n-i)}},$$

which immediately leads to the desired inequality. ■

Proof of Proposition E.1*:* By the definition of m_r, we have

$$2^{m_r} = (\text{total } \# \text{ of codewords})/2^{n-m}, \text{ where } m = m' + m''.$$

Then

$$\mathsf{E}\left\{\frac{2^{m_r}}{n}\right\} \leq \frac{2}{n2^{n-m}} + \sum_{i=1}^{n-1} \frac{\mathsf{E}\{\# \text{ of codewords of weight } i\}}{n2^{n-m}}.$$

Using the enumerating function as in [23, 82], the above quantity can be further upper bounded as follows.

$$\begin{aligned}
\mathsf{E}\left\{\frac{2^{m_r}}{n}\right\} &\leq \frac{2}{n2^{n-m}} + \sum_{i=1}^{n-1} \frac{\frac{\binom{n}{i}}{\binom{nd_v}{id_v}} \inf_{x>0} \frac{\left(\frac{(1+x)^{d_c-1}+(1-x)^{d_c-1}}{2}\right)^{m'} \left(\frac{(1+x)^{d_c}+(1-x)^{d_c}}{2}\right)^{m''}}{x^{id_v}}}{n2^{n-m}} \\
&\leq \frac{2}{n2^{n-m}} + \sqrt{d_v}e^{1/6} \sum_{i=1}^{n-1} \frac{1}{n} \frac{\inf_{x>0} \frac{\left(\frac{(1+x)^{d_c-1}+(1-x)^{d_c-1}}{2}\right)^{m'} \left(\frac{(1+x)^{d_c}+(1-x)^{d_c}}{2}\right)^{m''}}{x^{id_v}}}{2^{(d_v-1)nH_2(i/n)+n-m}} \\
&\leq \frac{2}{n2^{n-m}} + \sqrt{d_v}e^{1/6} \sum_{i=1}^{n-1} \frac{1}{n} \frac{\inf_{x>0} \frac{\left(\frac{(1+x)^{d_c-1}+(1-x)^{d_c-1}}{2}\right)^{m'} \left(\frac{(1+x)^{d_c}+(1-x)^{d_c}}{2}\right)^{m''}}{x^{id_v}}}{2^{d_v nH_2(i/n)-m}},
\end{aligned} \tag{E.8}$$

where the second inequality follows from Lemma E.1 and the third inequality follows from the fact that the binary entropy function $H_2(\cdot)$ is upper bounded by 1.

By defining

$$f_n(i,x) := 2^{m-d_v nH_2(i/n)} \frac{\left(\frac{(1+x)^{d_c-1}+(1-x)^{d_c-1}}{2}\right)^{m'} \left(\frac{(1+x)^{d_c}+(1-x)^{d_c}}{2}\right)^{m''}}{x^{id_v}},$$

the summation in (E.8) is upper bounded[2] by

$$\sum_{i=1}^{n-1} \frac{1}{n} \inf_{x>0} f_n(i,x) \leq \max_{i\in[0,n]} \inf_{x>0} f_n(i,x) \leq \inf_{x>0} \max_{i\in[0,n]} f_n(i,x) \leq \max_{i\in[0,n]} f_n(i,1).$$

By simple calculus, $\max_{i\in[0,n]} f_n(i,1)$ is attained when $i = n/2$. Since $f_n(n/2,1) = 1$, the summation in (E.8) is bounded by 1 for all n, and therefore

$$\limsup_{n\to\infty} \mathsf{E}\left\{\frac{2^{m_r}}{n}\right\} \leq \sqrt{d_v}e^{1/6}.$$

The proof is complete.

■

E.3 Proof of Corollary 4.5

We prove one direction that

$$\begin{aligned}
p^*_{1\to 0,linear} &:= \sup\left\{p_{1\to 0} > 0 : \lim_{l\to\infty} p^{(l)}_{e,linear} = 0\right\} \\
&> \sup\left\{p_{1\to 0} > 0 : \lim_{l\to\infty} p^{(l)}_{e,coset} = 0\right\} - \epsilon := p^*_{1\to 0,coset} - \epsilon.
\end{aligned}$$

The other direction, $p^*_{1\to 0,coset} > p^*_{1\to 0,linear} - \epsilon$, can be easily obtained by symmetry.

[2]The range of i is expanded here from a discrete integer set to a continuous interval.

By definition, for any $\epsilon > 0$, we can find a sufficiently large $l_0 < \infty$ such that for a z-channel with one-way crossover probability $p_{1\to 0} := p^*_{1\to 0,coset} - \epsilon$, $P^{(l_0)}_{coset}$ is in the interior of the stability region. We note that the stability region depends only on the Bhattacharyya noise parameter of $P^{(l_0)}_{coset}$, which is a continuous function with respect to convergence in distribution. Therefore, by Theorem 4.7, there exists a $\Delta \in \mathbb{N}$ such that $\left\langle P^{(l_0)} \right\rangle$ is also in the stability region. By the definition of the stability region, we have $\lim_{l\to\infty} p^{(l)}_{e,linear} = 0$, which implies $p^*_{1\to 0,linear} \geq p_{1\to 0}$. The proof is thus complete.

E.4 The Convergence Rates of (4.27) and (4.28)

Convergence rate analysis of (4.27): We will consider the cases that $k = 0$ and $k = 1$ separately. By the BNSC decomposition argument in Section 5.1.2, namely, all non-symmetric channels can be decomposed as the probabilistic combination of many BNSCs, we can limit our attention to simple BNSCs rather than general memoryless BI-NSO channels. Suppose $P^{(l-1)}_{a.p.}(0)$ and $P^{(l-1)}_{a.p.}(1)$ correspond to a BNSC with crossover probabilities $p_{0\to 1}$ and $p_{1\to 0}$. Without loss of generality, we may assume $p_{0\to 1} + p_{1\to 0} < 1$ because of the assumption that $\forall x \in \mathsf{GF}(2), P^{(l-1)}_{a.p.}(x)(m = 0) = 0$. We then have

$$\begin{aligned} \Phi_{P'_0}(k, \frac{r}{\Delta}) &= (1 - p_{0\to 1}) e^{i\frac{r}{\Delta} \ln \frac{1-p_{0\to 1}+p_{1\to 0}}{1-p_{0\to 1}-p_{1\to 0}}} + (-1)^k p_{0\to 1} e^{i\frac{r}{\Delta} \ln \frac{1+p_{0\to 1}-p_{1\to 0}}{1-p_{0\to 1}-p_{1\to 0}}} \\ \text{and } \Phi_{P'_1}(k, \frac{r}{\Delta}) &= (1 - p_{1\to 0}) e^{i\frac{r}{\Delta} \ln \frac{1+p_{0\to 1}-p_{1\to 0}}{1-p_{0\to 1}-p_{1\to 0}}} + (-1)^k p_{1\to 0} e^{i\frac{r}{\Delta} \ln \frac{1-p_{0\to 1}+p_{1\to 0}}{1-p_{0\to 1}-p_{1\to 0}}}. \end{aligned}$$

By Taylor's expansion, for $k = 0$, (4.27) becomes

$$\begin{aligned} & 2\left(\frac{\Phi_{P'_0}(0, \frac{r}{\Delta}) - \Phi_{P'_1}(0, \frac{r}{\Delta})}{2}\right)^{\Delta} \\ &= 2\left(i\left(\frac{1 - p_{0\to 1} - p_{1\to 0}}{2}\right)\left(\frac{r}{\Delta}\right) \ln\left(\frac{1 - p_{0\to 1} + p_{1\to 0}}{1 + p_{0\to 1} - p_{1\to 0}}\right) + \mathcal{O}\left(\left(\frac{r}{\Delta}\right)^2\right)\right)^{\Delta}, \end{aligned}$$

which converges to zero with convergence rate $\mathcal{O}\left(\mathcal{O}(\Delta)^{-\Delta}\right)$. For $k = 1$, we have

$$\begin{aligned} & 2\left(\frac{\Phi_{P'_0}(1, \frac{r}{\Delta}) - \Phi_{P'_1}(1, \frac{r}{\Delta})}{2}\right)^{\Delta} \\ &= 2\left((p_{1\to 0} - p_{0\to 1}) + \frac{i}{2}\left(\frac{r}{\Delta}\right) f(p_{0\to 1}, p_{1\to 0}) + \mathcal{O}\left(\left(\frac{r}{\Delta}\right)^2\right)\right)^{\Delta}, \end{aligned} \tag{E.9}$$

where

$$\begin{aligned} f(p_{0\to 1}, p_{1\to 0}) &:= (1 - p_{0\to 1} + p_{1\to 0}) \ln\left(\frac{1 - p_{0\to 1} + p_{1\to 0}}{1 - p_{0\to 1} - p_{1\to 0}}\right) \\ &\quad - (1 + p_{0\to 1} - p_{1\to 0}) \ln\left(\frac{1 + p_{0\to 1} - p_{1\to 0}}{1 - p_{0\to 1} - p_{1\to 0}}\right). \end{aligned}$$

(E.9) converges to zero with convergence rate $\mathcal{O}(\mathsf{const}^{\Delta})$, where const satisfies $|p_{1\to 0} - p_{0\to 1}| < \mathsf{const} < 1$. Since the convergence rate is determined by the slower of the above two cases ($k = 0, 1$), we have proven that (4.27) converges to zero with rate $\mathcal{O}(\mathsf{const}^{\Delta})$ for some $\mathsf{const} < 1$. ■

Convergence rate analysis of (4.28): Since we assume that the input is not perfect, we have $\max(p_{0\to1}, p_{1\to0}) > 0$. For $k = 0$, by Taylor's expansion, we have

$$\begin{aligned}&\left(\frac{\Phi_{P_0'}(0,\frac{r}{\Delta})+\Phi_{P_1'}(0,\frac{r}{\Delta})}{2}\right)^{\Delta}\\ =\;&\left(1+\frac{i}{2}\left(\frac{r}{\Delta}\right)g(p_{0\to1},p_{1\to0})+\mathcal{O}\left(\left(\frac{r}{\Delta}\right)^2\right)\right)^{\Delta},\end{aligned}\tag{E.10}$$

where

$$\begin{aligned}g(p_{0\to1},p_{1\to0}) \;:=\;&(1-p_{0\to1}+p_{1\to0})\ln\left(\frac{1-p_{0\to1}+p_{1\to0}}{1-p_{0\to1}-p_{1\to0}}\right)\\ &+(1+p_{0\to1}-p_{1\to0})\ln\left(\frac{1+p_{0\to1}-p_{1\to0}}{1-p_{0\to1}-p_{1\to0}}\right).\end{aligned}$$

(E.10) converges to $e^{i\left(\frac{r}{2}\right)g(p_{0\to1},p_{1\to0})}$ with rate $\mathcal{O}\left(\Delta^{-1}\right)$. For $k = 1$, we have

$$\begin{aligned}&\left(\frac{\Phi_{P_0'}(1,\frac{r}{\Delta})+\Phi_{P_1'}(1,\frac{r}{\Delta})}{2}\right)^{\Delta}\\ =\;&\left((1-p_{0\to1}-p_{1\to0})\left(\frac{e^{i\frac{r}{\Delta}\ln\frac{1-p_{0\to1}+p_{1\to0}}{1-p_{0\to1}-p_{1\to0}}}+e^{i\frac{r}{\Delta}\ln\frac{1+p_{0\to1}-p_{1\to0}}{1-p_{0\to1}-p_{1\to0}}}}{2}\right)\right)^{\Delta},\end{aligned}$$

which converges to zero with rate $\mathcal{O}\left((1-p_{0\to1}-p_{1\to0})^{\Delta}\right)$. Since the overall convergence rate is the slower of the above two, we have proven that the convergence rate is $\mathcal{O}\left(\Delta^{-1}\right)$. ■

Appendix F

Implications of the BSC, BNSC and MSC Decomposition Lemmas

F.1 The Relationship among BNP, ESB, and p_e

Without loss of generality, we assume the conditional probability $\mathsf{P}(Y|X)$ is discrete, and all our derivations can be easily generalized to continuous/mixed situations.

Proof of Lemma 5.1*:* We use $q_{x,y} := \mathsf{P}(X = x, Y = y)$ to denote the joint probability of $X = x$ and $Y = y$. By definition, we have

$$\begin{aligned} p_e &= \sum_{y\in\mathbf{Y}} \min(q_{0,y}, q_{1,y}) \\ \text{BNP} &= 2\sum_{y\in\mathbf{Y}} \sqrt{q_{0,y}q_{1,y}} \\ \text{ESB} &= 2\sum_{y\in\mathbf{Y}} \frac{2q_{0,y}q_{1,y}}{q_{0,y}+q_{1,y}}. \end{aligned}$$

Since for any $x, y > 0$, $\min(x, y) \le \frac{1}{\frac{1}{2}\left(\frac{1}{x}+\frac{1}{y}\right)} \le \sqrt{xy}$, we immediately have $2p_e \le \text{ESB} \le \text{BNP}$. By Jensen's inequality and the concavity of the square root function, we can rewrite BNP as

$$\begin{aligned} \text{BNP} &= \sum_{y\in\mathbf{Y}} (q_{0,y}+q_{1,y}) \sqrt{\frac{4q_{0,y}q_{1,y}}{(q_{0,y}+q_{1,y})^2}} \\ &\le \sqrt{\sum_{y\in\mathbf{Y}} (q_{0,y}+q_{1,y}) \frac{4q_{0,y}q_{1,y}}{(q_{0,y}+q_{1,y})^2}} \\ &= \sqrt{\text{ESB}}. \end{aligned} \tag{F.1}$$

Again by Jensen's inequality and the concavity of the polynomial $f(x) = 4x(1-x)$, we have

$$\begin{aligned}
\text{ESB} &= \sum_{y\in\mathbf{Y}} (q_{0,y}+q_{1,y}) 4 \frac{q_{0,y}}{q_{0,y}+q_{1,y}} \frac{q_{1,y}}{q_{0,y}+q_{1,y}} \\
&= \sum_{y\in\mathbf{Y}} (q_{0,y}+q_{1,y}) \left[4 \frac{\min(q_{0,y},q_{1,y})}{q_{0,y}+q_{1,y}} \left(1 - \frac{\min(q_{0,y},q_{1,y})}{q_{0,y}+q_{1,y}} \right) \right] \\
&\le 4x(1-x)\big|_{x=\sum_{y\in\mathbf{Y}}(q_{0,y}+q_{1,y})\frac{\min(q_{0,y},q_{1,y})}{q_{0,y}+q_{1,y}}} \\
&= 4p_e(1-p_e). \qquad \text{(F.2)}
\end{aligned}$$

By (F.1), (F.2), and $2p_e \le \text{ESB} \le \text{BNP}$, the proof of Lemma 5.1 is complete. ■

Proof of Lemma 5.2*:* Define $p_{e,0\leftrightarrow x}$ as the bit error probability of the MAP detector[1] given that the input X is uniformly distributed on $\{0, x\}$, namely,

$$p_{e,0\leftrightarrow x} := \mathsf{P}_{X\in_u\{0,x\}} \left(X \neq \hat{X}_{MAP}(Y) \right), \qquad \text{(F.3)}$$

where $\mathsf{P}_{X\in_u A}(\cdot)$ denotes the probability assuming X is evenly distributed on $A \subseteq \mathbb{Z}_m$. We note that $X \in_u \{0, x\}$ is equivalent to a binary-input channel with input alphabet $\{0, x\}$. Define $\text{BNP}(0 \leftrightarrow x) = \frac{1}{2}\text{BNP}(0 \to x) + \frac{1}{2}\text{BNP}(x \to 0)$ as the BNP value of the binary channel $\{0, x\} \mapsto \mathbf{Y}$. By the symmetry of BNP (5.5), we have $\text{BNP}(0 \to x) = \text{BNP}(x \to 0)$, which in turn implies $\text{BNP}(0 \leftrightarrow x) = \text{BNP}(0 \to x)$. By Lemma 5.1, we have

$$2p_{e,0\leftrightarrow x} \le \text{BNP}(0 \to x) \le 2\sqrt{p_{e,0\leftrightarrow x}(1-p_{e,0\leftrightarrow x})}.$$

As a result, proving Lemma 5.2 is equivalent to proving

$$\max_{x\in\mathbb{Z}_m\setminus\{0\}} \{p_{e,0\leftrightarrow x}\} \le p_e \le \sum_{x\in\mathbb{Z}_m\setminus\{0\}} p_{e,0\leftrightarrow x}. \qquad \text{(F.4)}$$

We need only to prove the result for the MSC case, and the proof for general MI-SO channels then follows by taking the probabilistic average of the constituent MSCs. For an MSC with the parameter vector $\mathbf{p} = \{p_0, \cdots, p_{m-1}\}$, we have

$$\begin{aligned}
p_{e,0\leftrightarrow x} &= \frac{1}{2} \sum_{y=0}^{m-1} \min(p_y, p_{y+x}) \qquad \text{(F.5)} \\
p_e &= 1 - \max(p_0, p_1, \cdots, p_{m-1}).
\end{aligned}$$

Without loss of generality, we may assume p_0 is the maximum entry in $\mathbf{p}$ and $p_e = 1 - p_0$. Then for any $x \neq 0$, we can rewrite $p_{e,0\leftrightarrow x}$ as

$$\begin{aligned}
p_{e,0\leftrightarrow x} &= \frac{1}{2} \left(p_x + \sum_{y=1}^{m-1} \min(p_y, p_{y+x}) \right) \le \frac{1}{2} \left(p_x + \sum_{y=1}^{m-1} p_y \right) \\
&\le \frac{1}{2} \sum_{y=1}^{m-1} 2p_y = 1 - p_0 = p_e,
\end{aligned}$$

[1]This MAP detector takes into account that X is uniformly distributed on $\{0, x\}$, and therefore it is a different MAP detector than that for uniformly distributed X.

and the first half of Ineq. (F.4) is proved. Also by Eq. (F.5) and the assumption that p_0 is the maximal component of $\mathbf{p}$, we have

$$p_{e,0\leftrightarrow x} \geq \frac{p_x + p_{-x}}{2}, \forall x \in \mathbb{Z}_m \backslash \{0\}.$$

Summing over all possible $x \in \mathbb{Z}_m \backslash \{0\}$, the second half of (F.4) is also proved. ■

F.2 Erasure Decomposition Lemma for BI-NSO Channels

Lemma F.1 (Erasure Decomposition Lemma for BI-NSO Channels) *Consider any BI-NSO channel F and let p_e denote the error probability of the MAP detector assuming uniform a priori distribution. Then F is physically degraded w.r.t. an erasure channel $F_{BEC,\epsilon}$ with $\epsilon = 2p_e$.*

Proof: By the BNSC decomposition lemma, we only need to prove the cases when F is a BNSC with crossover probabilities $(p_{0\to1}, p_{1\to0})$. Without loss of generality, we may further assume that $p_{0\to1} + p_{1\to0} \leq 1$, otherwise, we simply flip the output of F. The MAP error probability then becomes $\frac{p_{0\to1}+p_{1\to0}}{2}$. Define the range of BECs as $\{0, 1, E\}$ and the probability model of $F_{BEC,\epsilon}$ as follows.

$$F_{BEC,\epsilon}(y|x) = \begin{cases} (1-\epsilon) & \text{if } x = y = 0 \text{ or } x = y = 1 \\ \epsilon & \text{if } y = E \\ 0 & \text{otherwise} \end{cases}.$$

Construct an auxiliary channel $G : \{0, 1, E\} \mapsto \{0, 1\}$ as follows.

$$G(y|x) = \begin{cases} 1 & \text{if } x = y = 0 \text{ or } x = y = 1 \\ \frac{p_{1\to0}}{p_{0\to1}+p_{1\to0}} & \text{if } x = E \text{ and } y = 0 \\ \frac{p_{0\to1}}{p_{0\to1}+p_{1\to0}} & \text{if } x = E \text{ and } y = 1 \\ 0 & \text{otherwise} \end{cases}.$$

It can be easily verified that the concatenation of $F_{BEC,2p_e}$ and G becomes F with crossover probabilities $(p_{0\to1}, p_{1\to0})$. Therefore, F is physically degraded w.r.t. $F_{BEC,2p_e}$.

■

F.3 Erasure Decomposition Lemma for MI-SO Channels

An x-erasure MSC can be defined by specifying its parameter vector $\mathbf{p}$ as $p_0 = 1/2$, $p_x = 1/2$ and $p_z = 0, \forall z \notin \{0, x\}$. We will use $\mathbf{e}_x$ to denote this particular parameter $\mathbf{p}$. An x-erasure decomposition lemma is given as follows.

Lemma F.2 (x-erasure Decomposition) *Consider any MI-SO channel with pairwise MAP error $p_{e,0\leftrightarrow x}$ defined in (F.3). This MI-SO channel can be written as a degraded channel of a probabilistic combination of two MSCs, such that the probabilistic weight $dQ^\epsilon(\mathbf{q})$ satisfies*

$$dQ^\epsilon(\mathbf{q}) = \begin{cases} 1 - 2p_{e,0\leftrightarrow x} & \textit{if } \mathbf{q} = (1, 0, \cdots, 0) \\ 2p_{e,0\leftrightarrow x} & \textit{if } \mathbf{q} = \mathbf{e}_x \\ 0 & \textit{otherwise} \end{cases}.$$

Proof: We need to prove Lemma F.2 only for an MSC with arbitrary parameter $\mathbf{p}$. By taking the average over $dP(\mathbf{p})$, the same result holds for general MI-SO channels.

We first note that $dQ^{\epsilon}(\cdot)$ can be viewed as a $\mathbb{Z}_m \mapsto (\{0,1\} \times \mathbb{Z}_m)$ channel, where the first output component is 1 iff $\mathbf{q} = \mathbf{e}_x$. Let $\mathbf{p}$ denote the parameter of the original MSC. We would like to show that there exists another channel F such that after concatenating $dQ^{\epsilon}(\cdot)$ and F, we can reproduce the probability law $\mathbf{p}$ of the original MSC. To be more explicit, F is a $(\{0,1\} \times \mathbb{Z}_m) \mapsto \mathbb{Z}_m$ channel such that the concatenation $\mathbb{Z}_m \mapsto (\{0,1\} \times \mathbb{Z}_m) \mapsto \mathbb{Z}_m$ becomes an MSC with parameter $\mathbf{p}$. We prove the existence of F by explicitly specifying its probability law.

When the first component of the input of F is given, say 0 or 1, let the remaining $\mathbb{Z}_m \mapsto \mathbb{Z}_m$ channel be an MSC with parameter $\mathbf{r}$ or with parameter $\mathbf{s}$ (depending on the first component being 0 or 1). Define, $\forall i \in \mathbb{Z}_m$,

$$r_i = \frac{p_i - \frac{1}{2}\min(p_i, p_{i+x}) - \frac{1}{2}\min(p_i, p_{i-x})}{1 - \sum_{z \in \mathbb{Z}_m} \min(p_z, p_{z+x})}$$

$$s_i = \frac{\min(p_i, p_{i+x})}{\sum_{z \in \mathbb{Z}_m} \min(p_z, p_{z+x})}.$$

It is easy to check that both $\mathbf{r}$ and $\mathbf{s}$ are valid probability parameter vectors.

It is also easy to check that the end-to-end $dQ^{\epsilon} \circ F$ channel is an MSC. By noting that $2p_{e,0 \hookrightarrow x} = \sum_{z \in \mathbb{Z}_m} \min(p_z, p_{z+x})$, we can verify that the end-to-end channel has the same parameter $\mathbf{p}$ as the original MSC.

■

Appendix G

Proofs of Finite Dimensional Bounds

G.1 The Necessary Stability Condition for $\mathbb{Z}_m$ LDPC Codes

One necessary lemma for proving the necessary stability condition for $\mathbb{Z}_m$ LDPC codes is stated as follows.

Lemma G.1 (Monotonicity of $p_{e,0\leftrightarrow x}^{(l)}$) *Let $p_{e,0\leftrightarrow x}^{(l)}$ denote the pairwise error probability of the support tree channel of depth $2l$ (after l iterations). Then $p_{e,0\leftrightarrow x}^{(l)}$ is a non-increasing function of l. Furthermore, if $p_{e,0\leftrightarrow x}^{(l)} > 0$, then $p_{e,0\leftrightarrow x}^{(l+1)} > 0$.*

Proof: As l grows, the support tree gives more information by providing additional observations. As a result, the MAP error $p_{e,0\leftrightarrow x}^{(l)}$ is a non-increasing function of l.

For the second statement, we break one iteration into its check node part and its variable node part. Since a check node channel is a degraded channel with respect to each of its constituent channels, we have $p_{e,0\leftrightarrow x}^{(l+\frac{1}{2})} \geq p_{e,0\leftrightarrow x}^{(l)} > 0$, where $p_{e,0\leftrightarrow x}^{(l+\frac{1}{2})}$ is the pairwise error probability of the support tree of depth $2l+1$ (after incorporating the check node). For variable nodes, by the equation $\mathsf{BNP}_{var} = \mathsf{BNP}_{in,1} \bullet \mathsf{BNP}_{in,2}$, we have $p_{e,0\leftrightarrow x}^{(l+1)} = 0$ iff either $p_{e,0\leftrightarrow x}^{(0)} = 0$ or $p_{e,0\leftrightarrow x}^{(l+\frac{1}{2})} = 0$. Since both $p_{e,0\leftrightarrow x}^{(0)} \geq p_{e,0\leftrightarrow x}^{(l)} > 0$ and $p_{e,0\leftrightarrow x}^{(l+\frac{1}{2})} > 0$, it follows that $p_{e,0\leftrightarrow x}^{(l+1)} > 0$. ■

Proof of Theorem 5.3*:* Suppose Theorem 5.3 is false, namely, there exists a MI-SO channel such that $\lim_{l\to\infty} p_e^{(l)} = 0$ while there exists an $x_0 \in \mathbb{Z}_m \backslash \{0\}$ satisfying $\lambda_2 \rho'(1)\mathrm{BNP}(0 \to x_0) > 1$. By Lemmata 5.1, 5.2, and G.1, we have $\lim_{l\to\infty} p_{e,0\leftrightarrow x_0}^{(l)} = 0$ and $p_{e,0\leftrightarrow x_0}^{(l)} > 0$, $\forall l \in \mathbb{N}$. We first choose a small $\epsilon > 0$ and then find the value of l_0 such that $p_{e,0\leftrightarrow x_0}^{(l_0)} < \epsilon$. Without loss of generality, we may assume $p_{e,0\leftrightarrow x_0}^{(l_0)} = \epsilon > 0$ by Lemma G.1.

By Lemma F.2, we can replace the supporting tree channel of depth $2l_0$ by a probabilistic combination of a perfect channel and an x_0-erasure channel with weights $(1-2\epsilon, 2\epsilon)$. Consider the next iteration after this substitution, and denote the new pairwise MAP error probability by $q_{e,0\leftrightarrow x_0}^{(1)}$. We also say an x_0-erasure channel outputs $+0$ if the output is $Y = X + 0$. Similarly, an x_0-erasure channel outputs $+x_0$ if the output is $Y = X + x_0$. We

then have

$$\begin{aligned}
q^{(1)}_{e,0\leftrightarrow x_0} &= \mathsf{P}(X=0)\mathsf{P}\left(\hat{X}_{MAP}=x_0|X=0\right)+\mathsf{P}(X=x_0)\mathsf{P}\left(\hat{X}_{MAP}=0|X=x_0\right)\\
&\geq \frac{1}{2}\mathsf{P}(\text{one and only one check node constituent channel (CNCC) is } x_0\text{-erasure}\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{and outputs } +0|X=0)\\
&\qquad\cdot\mathsf{P}\left(\hat{X}_{MAP}=x_0|\text{one and only one CNCC is } x_0\text{-erasure and outputs } +0, X=0\right)\\
&\quad+\frac{1}{2}\mathsf{P}(\text{one and only one CNCC is } x_0\text{-erasure and outputs } +x_0|X=0)\\
&\qquad\cdot\mathsf{P}\left(\hat{X}_{MAP}=0|\text{one and only one CNCC is } x_0\text{-erasure and outputs } +x_0, X=0\right)\\
&\quad+\mathcal{O}(\epsilon^2)\\
&= \frac{1}{2}\left(\frac{2\epsilon\lambda_2\rho'(1)}{2}\right)\int_{m_{x_0}=-\infty}^{0}P_{x_0}(dm_{x_0})+\frac{1}{2}\left(\frac{2\epsilon\lambda_2\rho'(1)}{2}\right)\int_{m_{-x_0}=-\infty}^{0}P_{-x_0}(dm_{-x_0})+\mathcal{O}(\epsilon^2),
\end{aligned}$$

where $m_x=\log\frac{\mathsf{P}(Y=y|X=0)}{\mathsf{P}(Y=y|X=x)}$ is the LLR between $X=0$ and $X=x$, and $P_x(\cdot)$ is the density of the LLR message m_x given $X=0$. *Note:* If $x_0\neq -x_0$ in $\mathbb{Z}_m$, the inequality becomes an equality. If $x_0=-x_0$ in $\mathbb{Z}_m$, then $q^{(1)}_{e,0\leftrightarrow x_0}$ equals twice the right-hand side of the above expression.

By similar arguments, the second iteration gives

$$q^{(2)}_{e,0\leftrightarrow x_0} \geq \frac{1}{2}\epsilon\left(\lambda_2\rho'(1)\right)^2\left(\int_{m=-\infty}^{0}(P_{-x_0}\otimes P_{x_0})(dm)+\int_{m=-\infty}^{0}(P_{x_0}\otimes P_{-x_0})(dm)\right)+\mathcal{O}(\epsilon^2),$$

and after $2\Delta l$ iterations we have

$$q^{(2\Delta l)}_{e,0\leftrightarrow x_0} \geq \epsilon\left(\lambda_2\rho'(1)\right)^{2\Delta l}\int_{m=-\infty}^{0}(P_{-x_0}\otimes P_{x_0})^{\otimes\Delta l}(dm)+\mathcal{O}(\epsilon^2).$$

It can be shown that $(P_{-x_0}\otimes P_{x_0})$ becomes a symmetric distribution, i.e., $P(dm)=e^m P(-dm)$, and its Bhattacharyya noise parameter is $(\text{BNP}(0\to x_0))^2$. Choose $\delta>0$ such that $\lambda_2\rho'(1)(\text{BNP}(0\to x_0)-\delta)>1$. By the tightness of the Bhattacharyya noise parameter, we can lower bound $q^{(2\Delta l)}_{e,0\leftrightarrow x_0}$ for sufficiently large Δl by

$$q^{(2\Delta l)}_{e,0\leftrightarrow x_0} \geq \epsilon\left(\lambda_2\rho'(1)\right)^{2\Delta l}(\text{BNP}(0\to x_0)-\delta)^{2\Delta l}+\mathcal{O}(\epsilon^2).$$

Choose sufficiently large Δl such that $(\lambda_2\rho'(1)(\text{BNP}(0\to x_0)-\delta))^{2\Delta l}>2$ and sufficiently small ϵ, we have $q^{(2\Delta l)}_{e,0\leftrightarrow x_0}\geq 2\epsilon+\mathcal{O}(\epsilon^2)>\epsilon$. Since $p^{(l_0+2\Delta l)}_{e,0\leftrightarrow x_0}$ can be viewed as the pairwise MAP error probability from a degraded channel compared to that of $q^{(2\Delta l)}_{e,0\leftrightarrow x_0}$, we have

$$p^{(l_0+2\Delta l)}_{e,0\leftrightarrow x_0}\geq q^{(2\Delta l)}_{e,0\leftrightarrow x_0}>\epsilon=p^{(l_0)}_{e,0\leftrightarrow x_0},$$

which contradicts the monotonicity result in Lemma G.1. Using this contradiction, the proof of Theorem 5.3 is complete.

■

G.2 The Maximizing Distribution for Check Nodes with Two-Dimensional Constraints on $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$

Proof of Theorem 5.5*:* We take the approach of considering the marginal dP first and assuming that dQ is a point mass, i.e., dQ concentrates all its probability on a fixed q_0. To simplify the notation, we let $a := 2\sqrt{p(1-p)}$ and $b := 2\sqrt{q_0(1-q_0)}$, and drop the subscript 1 in $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$ to $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$. The original problem (5.27) becomes an infinite dimensional linear programming problem on dP:

$$\begin{aligned} \max \quad & \int \sqrt{a^2(1-b^2)+b^2}\, dP(a) \\ \text{subject to} \quad & \int dP(a) = 1 \\ & \int a\, dP(a) \le \mathrm{BNP}_{in} \\ & \int a^2\, dP(a) \le \mathrm{ESB}_{in} \\ & dP(a) \ge 0, \ \forall a \in [0,1]. \end{aligned}$$

The corresponding dual problem is

$$\begin{aligned} \min \quad & \xi := y_0 + y_1 \mathrm{BNP}_{in} + y_2 \mathrm{ESB}_{in} \\ \text{subject to} \quad & y_0 + a y_1 + a^2 y_2 \ge \sqrt{a^2(1-b^2)+b^2}, \ \forall a \in [0,1] \\ & y_1, y_2 \ge 0. \end{aligned}$$

Let

$$\begin{aligned} t &= \frac{\mathrm{ESB}_{in}}{\mathrm{BNP}_{in}}, \\ y_0^* &= b, \\ y_1^* &= \frac{2}{t}\left(\frac{t^2(1-b^2)+2b^2}{2\sqrt{t^2(1-b^2)+b^2}} - b\right) \\ y_2^* &= \frac{1}{t^2}\left(b - \frac{b^2}{\sqrt{t^2(1-b^2)+b^2}}\right). \end{aligned}$$

It is easy to check that both $y_1^*, y_2^* \ge 0$. By Lemma G.2 (stated at the end of this proof), $\mathbf{y}^* \stackrel{\Delta}{=} (y_0^*, y_1^*, y_2^*)$ is a valid solution. Let

$$dP^*(a) = \begin{cases} 1 - \frac{\mathrm{BNP}_{in}}{t} & \text{if } a = 0 \\ \frac{\mathrm{BNP}_{in}}{t} & \text{if } a = t \\ 0 & \text{otherwise} \end{cases}.$$

It can be verified that the duality gap between the two feasible solutions $\mathbf{y}^*$ and dP^* is zero. By the weak duality theorem of linear programming, dP^* is the maximizing distribution when dQ concentrates on q_0. Since dP^* does not depend on b (and thus does not depend on q_0), dP^* is the universal maximizer for general dQ. ■

Lemma G.2 *$y_0^* + a y_1^* + a^2 y_2^* \ge \sqrt{a^2(1-b^2)+b^2}$ for all $a, b \in [0,1]$.*

Proof: Let $f := y_0^* + ay_1^* + a^2 y_2^* - \sqrt{a^2(1-b^2)+b^2}$. By noting that

$$\begin{aligned} \frac{d^3 f}{da^3} &= \frac{3ab^2(1-b^2)^2}{\left(\sqrt{a^2(1-b^2)+b^2}\right)^5} \geq 0, \\ f(0) &= 0, \\ f(t) &= 0, \\ \text{and} \quad \left.\frac{df}{da}\right|_{a=t} &= 0, \end{aligned}$$

we conclude that $f \geq 0$ for all $a, b \in [0,1]$. ■

G.3 The Upper Bounding Distribution for Variable Nodes with Two-Dimensional Constraints on $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$

We take the approach of assuming dQ concentrates on a fixed q_0. Let $a := 2\sqrt{p(1-p)}$ and $b := 2\sqrt{q_0(1-q_0)}$ and drop the subscript 1 in $(\mathrm{BNP}_{in,1}, \mathrm{ESB}_{in,1})$ to write $(\mathrm{BNP}_{in}, \mathrm{ESB}_{in})$. The original problem (5.29) becomes a linear programming problem with the primal and dual representations as follows.

$$\begin{aligned} \max \quad & \zeta := \int \frac{a^2 b^2}{a^2(1-b^2)+b^2} dP(a) \\ \text{subject to} \quad & \int dP(a) = 1 \\ & \int a dP(a) \leq \mathrm{BNP}_{in} \\ & \int a^2 dP(a) \leq \mathrm{ESB}_{in} \\ & dP(a) \geq 0, \ \forall a \in [0,1], \end{aligned}$$

$$\begin{aligned} \min \quad & \xi := y_0 + y_1 \mathrm{BNP}_{in} + y_2 \mathrm{ESB}_{in} \\ \text{subject to} \quad & y_0 + a y_1 + a^2 y_2 \geq \frac{a^2 b^2}{a^2(1-b^2)+b^2}, \ \forall a \in [0,1] \\ & y_1, y_2 \geq 0. \end{aligned}$$

For convenience, we define $t := \mathrm{ESB}_{in}/\mathrm{BNP}_{in}$ and $r_b(a) := \frac{a^2 b^2}{a^2(1-b^2)+b^2}$.

Unlike the check node channel case, this time the optimal primal solution $dP^*(a)$ depends on the value of b, namely, to which of the three intervals, $\left[0, \sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}\right]$, $\left[\sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}, \sqrt{\frac{t^2}{1+t^2}}\right]$, and $\left[\sqrt{\frac{t^2}{1+t^2}}, 1\right]$, b belongs. These three different cases will be addressed in the following three propositions respectively.

Proposition G.1 *If* $b \in \left[0, \sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}\right]$, *the maximizing* $dP^*(a)$ *and the optimum values*

are as follows.

$$dP^*(a) = \begin{cases} 1 & \text{if } a = \text{BNP}_{in} \\ 0 & \text{otherwise} \end{cases}, \tag{G.1}$$
$$\zeta^* = \xi^* = \frac{\text{BNP}_{in}^2 b^2}{\text{BNP}_{in}^2(1-b^2)+b^2}.$$

Proof: It is easy to check that the specified $dP(a)$ is feasible. We then note that $r_b(0) = 0$ and $\left(\frac{1}{2}\sqrt{\frac{b^2}{1-b^2}}\right) \cdot a$ is the only tangent line of $r_b(a)$ passing through the origin (with the contact point $a_t = \sqrt{\frac{b^2}{1-b^2}}$). Furthermore, when $a \geq \sqrt{\frac{b^2}{1-b^2}}$, we have $\frac{d^2 r_b}{da^2} = 2b^4 \frac{b^2 - 3a^2(1-b^2)}{(a^2(1-b^2)+b^2)^3} \leq 0$ and $r_b(a)$ is thus a concave function in the interval $[\sqrt{\frac{b^2}{1-b^2}}, 1]$. From the above observations,

$$u_b(a) = \begin{cases} \left(\frac{1}{2}\sqrt{\frac{b^2}{1-b^2}}\right) \cdot a & \text{if } a \in [0, \sqrt{\frac{b^2}{1-b^2}}] \\ r_b(a) & \text{if } a \in [\sqrt{\frac{b^2}{1-b^2}}, 1] \end{cases} \tag{G.2}$$

is the convex hull of $r_b(a)$. By Jensen's inequality,

$$\int r_b(a) dP^*(a) \leq \int u_b(a) dP^*(a) \leq u_b\left(\int a dP^*(a)\right) = \frac{\text{BNP}_{in}^2 b^2}{\text{BNP}_{in}^2(1-b^2)+b^2}. \tag{G.3}$$

Since $dP^*(a)$ in (G.1) achieves the upper bound in (G.3), it is indeed the maximizing distribution. ■

Proposition G.2 *If* $b \in \left[\sqrt{\frac{\text{BNP}_{in}^2}{1+\text{BNP}_{in}^2}}, \sqrt{\frac{t^2}{1+t^2}}\right]$, *the maximizing* $dP^*(a)$ *and the optimum values are as follows.*

$$dP^*(a) = \begin{cases} CB_{in}\sqrt{\frac{1-b^2}{b^2}} & \text{if } a = \sqrt{\frac{b^2}{1-b^2}} \\ 1 - \text{BNP}_{in}\sqrt{\frac{1-b^2}{b^2}} & \text{if } a = 0 \\ 0 & \text{otherwise} \end{cases}, \tag{G.4}$$
$$\zeta^* = \xi^* = \frac{\text{BNP}_{in}}{2}\sqrt{\frac{b^2}{1-b^2}}$$

Proof: It is easy to check that the specified $dP^*(a)$ is feasible. By again invoking Jensen's inequality on $u_b(a)$ defined in (G.2), we have

$$\int r_b(a) dP^*(a) \leq \int u_b(a) dP^*(a) \leq u_b\left(\int a dP^*(a)\right) = \frac{\text{BNP}_{in}}{2}\sqrt{\frac{b^2}{1-b^2}}. \tag{G.5}$$

Since $dP^*(a)$ in (G.4) achieves the upper bound in (G.5), it is indeed the maximizing distribution. ■

Proposition G.3 *If* $b \in \left[\sqrt{\frac{t^2}{1+t^2}}, 1\right]$, *the maximizing* $dP^*(a)$ *and the optimum values are as follows.*

$$\begin{aligned} dP^*(a) &= \begin{cases} \frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}} & \text{if } a = t \\ 1 - \frac{\mathrm{BNP}_{in}^2}{\mathrm{ESB}_{in}} & \text{if } a = 0 \\ 0 & \text{otherwise} \end{cases}, \\ \zeta^* &= \xi^* = \frac{\mathrm{ESB}_{in} b^2}{t^2 (1-b^2) + b^2} \end{aligned}$$

Proof: It is easy to check that the specified $dP^*(a)$ is feasible. By choosing

$$\begin{aligned} y_0^* &= 0 \\ y_1^* &= \frac{2t^3 b^2 (1-b^2)}{\left(t^2(1-b^2)+b^2\right)^2} \geq 0 \\ y_2^* &= \frac{b^2 (b^2 - t^2(1-b^2))}{\left(t^2(1-b^2)+b^2\right)^2} \geq 0, \end{aligned}$$

we have $\xi^* = \frac{\mathrm{ESB}_{in} b^2}{t^2(1-b^2)+b^2} = \zeta^*$. So it remains to show that $\mathbf{y}^* = (y_0^*, y_1^*, y_2^*)$ is feasible for all $b \in [0,1]$. Let $g := a y_1^* + a^2 y_2^* - r_b(a)$. By the following observations

$$\begin{aligned} \frac{d^3 g}{da^3} &= \frac{24 a b^4 (1-b^2)}{\left(a^2(1-b^2)+b^2\right)^4} \left(b^2 - a^2(1-b^2)\right) \geq 0 \quad \text{if } a \in [0, \sqrt{\tfrac{b^2}{1-b^2}}] \\ g(0) &= 0 \\ g(t) &= 0 \\ \left.\frac{dg}{da}\right|_{a=t} &= 0, \end{aligned}$$

we have $g(a) \geq 0$ for all $b \in [0,1]$, $a \in [0, \sqrt{\frac{b^2}{1-b^2}}]$. We then consider the case $a \in [\sqrt{\frac{b^2}{1-b^2}}, 1]$. Also by the same observations, it can be implied that both $g\left(\sqrt{\frac{b^2}{1-b^2}}\right)$ and $g'\left(\sqrt{\frac{b^2}{1-b^2}}\right)$ are no less than zero. By noting that $g''(a) = 2y_2^* + 2b^4 \frac{3a^2(1-b^2)-b^2}{(a^2(1-b^2)+b^2)^3} \geq 0$ for all $a \in [\sqrt{\frac{b^2}{1-b^2}}, 1]$, we conclude that $g(a) \geq 0$ for all $a \in [\sqrt{\frac{b^2}{1-b^2}}, 1]$. From the above reasoning, we have $g(a) \geq 0$ for all $a, b \in [0,1]$, and thus $\mathbf{y}^*$ is feasible and the proposition follows. ■

From Propositions G.1 through G.3, we have the following tight upper bound:

$$\begin{aligned} \int r_b(a) dP^*(a) &= \int \frac{a^2 b^2}{a^2(1-b^2)+b^2} dP^*(a) \leq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b), \\ \text{where} \quad s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b) &:= \begin{cases} \frac{\mathrm{BNP}_{in}^2 b^2}{\mathrm{BNP}_{in}^2(1-b^2)+b^2} & \text{if } b \in \left[0, \sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}\right] \\ \frac{\mathrm{BNP}_{in}}{2}\sqrt{\frac{b^2}{1-b^2}} & \text{if } b \in \left[\sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}, \sqrt{\frac{t^2}{1+t^2}}\right] \\ \frac{\mathrm{ESB}_{in} b^2}{t^2(1-b^2)+b^2} & \text{if } b \in \left[\sqrt{\frac{t^2}{1+t^2}}, 1\right] \end{cases} . \end{aligned} \quad \text{(G.6)}$$

Hereafter, we will show that the b-value-independent $dP^{**}(a)$ in (5.30) is a loose upper bounding distribution, such that $dP^{**}(a)$ may not be feasible, but the resulting $\int r_b(a) dP^{**}(a)$ is no smaller than $s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$ for all $b \in [0,1]$.

Lemma G.3 $\int r_b(a)dP^{**}(a) \geq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$ *for all* $b \in [0, \sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}]$.

Proof: By the monotonicity of $r_b(a)$ as a function of a, we have

$$\begin{aligned}\int r_b(a)dP^{**}(a) &= f_{\mathrm{BNP}} r_b(\mathrm{BNP}_{in}) + f_{\mathrm{ESB}} r_b(\sqrt{\mathrm{ESB}_{in}}) + f_t r_b(t) \\ &\geq f_{\mathrm{BNP}} r_b(\mathrm{BNP}_{in}) + f_{\mathrm{ESB}} r_b(\mathrm{BNP}_{in}) + f_t r_b(\mathrm{BNP}_{in}) \\ &= r_b(\mathrm{BNP}_{in}) = s(CB_{in}, DCB_{in}, b).\end{aligned}$$

■

Lemma G.4 $\int r_b(a)dP^{**}(a) \geq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$ *for all* $b \in (\sqrt{\frac{t^2}{1+t^2}}, 1]$.

Proof: In this proof, we consider the variable $\alpha \triangleq a^2$ and rewrite $r_b(a) = r_b(\alpha) = \frac{\alpha b^2}{\alpha(1-b^2)+b^2}$. It can be shown that $\frac{d^2 r_b(\alpha)}{d\alpha^2} \leq 0$ and $r_b(\alpha)$ is a concave function of α. By noting that $\int \alpha dP^{**}(\alpha) = \mathrm{ESB}_{in}$ and the weights in $dP^{**}(\alpha)$ are concentrated only on three points $\alpha = \mathrm{BNP}_{in}^2$, ESB_{in}, and t^2 in an increasing order, we have $\int r_b(\alpha)dP^{**}(\alpha) \geq A$, where A is the intersection of the vertical line $\alpha = \mathrm{ESB}_{in}$ and the chord connecting $\big(\mathrm{BNP}_{in}^2, r_b(\mathrm{BNP}_{in}^2)\big)$ and $\big(t^2, r_b(t^2)\big)$.

We also notice that $s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$ is the intersection of the vertical line $\alpha = \mathrm{ESB}_{in}$ and the chord connecting $(0, r_b(0))$ and $\big(t^2, r_b(t^2)\big)$. By the concavity of $r_b(\alpha)$, we conclude $\int r_b(\alpha)dP^{**}(\alpha) \geq A \geq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$. ■

Lemma G.5 $\int r_b(a)dP^{**}(a) \geq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b)$ *for all* $b \in [\sqrt{\frac{\mathrm{BNP}_{in}^2}{1+\mathrm{BNP}_{in}^2}}, \sqrt{\frac{t^2}{1+t^2}}]$.

Proof: We prove this by directly applying calculus. By changing variables to $x := \sqrt{\frac{b^2}{1-b^2}}$ and using c as a shortcut of BNP_{in} (note that $\mathrm{ESB}_{in} = tc$), proving Lemma G.5 is equivalent to showing

$$\begin{aligned}\int r_b(a)dP^{**}(a) &= (1-f_{\mathrm{ESB}})\frac{t}{t+c}\frac{c^2x^2}{c^2+x^2} + f_{\mathrm{ESB}}\frac{tcx^2}{tc+x^2} \\ &\quad +(1-f_{\mathrm{ESB}})\frac{c}{t+c}\frac{t^2x^2}{t^2+x^2} \\ &= tcx^2\left(\frac{x^4+t^2c^2+x^2\left(t^2+c^2-(1-f_{\mathrm{ESB}})(t-c)^2\right)}{(x^2+t^2)(x^2+tc)(x^2+c^2)}\right) \\ &\geq s(\mathrm{BNP}_{in}, \mathrm{ESB}_{in}, b) \\ &= \frac{c}{2}x, \qquad \forall x \in [c,t] \subseteq [0,1].\end{aligned}$$

Multiplying the common denominator and changing the variable to $y := \frac{x}{\sqrt{tc}}$, the desired inequality becomes

$$\begin{aligned}&t^3c^3y^6 + t^2c^2(t^2+c^2+tc)y^4 + t^2c^2(t^2+c^2+tc)y^2 + t^3c^3 \\ &-2t\sqrt{tc}y(t^2c^2y^4 + t^2c^2 + tcy^2(t^2+c^2-(1-f_{\mathrm{ESB}})(t-c)^2)) \leq 0,\end{aligned}$$

for all $\frac{\sqrt{c}}{\sqrt{t}} < y \leq \frac{\sqrt{t}}{\sqrt{c}}$. By again changing the variable to $w := \sqrt{tc}(y + \frac{1}{y})$, we would like to prove that

$$\eta(w) - 2f_{\mathrm{ESB}}t(t-c)^2 \leq 0, \qquad \forall w \in [2\sqrt{tc}, (t+c)],$$

where $\eta(w)$ is defined in (5.31). By noting that $\eta(t+c) - 2f_{\text{ESB}}t(t-c)^2 = -2(t-c)(c(t+c) + f_{\text{ESB}}t(t-c)) \leq 0$ for all $f_{\text{ESB}} \in [0,1]$, we would like to show that there exists no root of $\eta(w) - 2f_{\text{ESB}}t(t-c)^2$ in $[2\sqrt{tc}, t+c]$. If $t - \sqrt{c(2t-c)} \leq 2\sqrt{tc}$, then by definition, $f_{\text{ESB}} = 0$. By simple calculus, there is no root in $[2\sqrt{tc}, t+c]$. If $2\sqrt{tc} - t + \sqrt{c(2t-c)} < 0$, there is one root of $\eta(w)$ in $[2\sqrt{tc}, t+c]$. By letting $f_{\text{ESB}} = \frac{\eta(w^*)}{2t(t-c)^2}$ where

$$w^* = \begin{cases} 2\sqrt{tc} & \text{if } \eta'(2\sqrt{tc}) \leq 0 \\ \frac{2t - \sqrt{4t^2 - 3(t-c)^2}}{3} & \text{otherwise} \end{cases},$$

we guarantee that $\eta(w) - 2f_{\text{ESB}}t(t-c)^2$, the shifted version of $\eta(w)$, has no root in $[2\sqrt{tc}, t+c]$. This completes the proof. ■

G.4 Explicit Expression of Φ_k

To write Φ_k explicitly, we need to define the following terms:

$$\begin{aligned}
t_{in} &= \text{ESB}_{in}/\text{BNP}_{in}, \\
t_0 &= \text{ESB}_0/\text{BNP}_0, \\
p_{in,\text{BNP}} &= \frac{1 - \sqrt{1 - (\text{BNP}_{in})^2}}{2}, \\
p_{in,\text{ESB}} &= \frac{1 - \sqrt{1 - \text{ESB}_{in}}}{2}, \\
p_{in,t} &= \frac{1 - \sqrt{1 - (t_{in})^2}}{2}, \\
p_{0,\text{BNP}} &= \frac{1 - \sqrt{1 - (\text{BNP}_0)^2}}{2}, \\
p_{0,\text{ESB}} &= \frac{1 - \sqrt{1 - \text{ESB}_0}}{2}, \\
p_{0,t} &= \frac{1 - \sqrt{1 - (t_0)^2}}{2},
\end{aligned}$$

and $(f_{in,\text{BNP}}, f_{in,\text{ESB}}, f_{in,t})$ and $(f_{0,\text{BNP}}, f_{0,\text{ESB}}, f_{0,t})$ are the $\left((1-f_{\text{ESB}})\frac{t}{t+\text{BNP}_{in}}, f_{\text{ESB}}, (1-f_{\text{ESB}})\frac{\text{BNP}_{in}}{t+\text{BNP}_{in}}\right)$ in (5.30) corresponding to $(\text{BNP}_{in}, \text{ESB}_{in})$ and $(\text{BNP}_0, \text{ESB}_0)$ respectively. Then

$$\begin{aligned}
&\text{ESB}_{out} \\
&\leq \Phi\left((\text{BNP}_0, \text{ESB}_0), (\text{BNP}_{in}, \text{ESB}_{in}), d_v - 1\right) \\
&= \sum_{\substack{i+j+k=d_v-1, \\ i,j,k \geq 0.}} 4\binom{d_v-1}{i,j,k} (f_{in,\text{BNP}})^i (f_{in,\text{ESB}})^j (f_{in,t})^k \left(\frac{(\text{BNP}_{in})^2}{4}\right)^i \left(\frac{t_{in}\text{BNP}_{in}}{4}\right)^j \left(\frac{(t_{in})^2}{4}\right)^k \\
&\quad \cdot \sum_{\hat{i}=0}^{i} \sum_{\hat{j}=0}^{j} \sum_{\hat{k}=0}^{k} \binom{i}{\hat{i}} \binom{j}{\hat{j}} \binom{k}{\hat{k}} \left(\frac{f_{0,\text{BNP}}\left(\frac{(\text{BNP}_0)^2}{4}\right)}{G_{i,j,k}(p_{0,\text{BNP}}, p_{in,\text{BNP}}, \hat{i}, p_{in,\text{ESB}}, \hat{j}, p_{in,t}, \hat{k})} \right. \\
&\quad \left. + \frac{f_{0,\text{ESB}}\left(\frac{t_0\text{BNP}_0}{4}\right)}{G_{i,j,k}(p_{0,\text{ESB}}, p_{in,\text{BNP}}, \hat{i}, p_{in,\text{ESB}}, \hat{j}, p_{in,t}, \hat{k})} + \frac{f_{0,t}\left(\frac{(t_0)^2}{4}\right)}{G_{i,j,k}(p_{0,t}, p_{in,\text{BNP}}, \hat{i}, p_{in,\text{ESB}}, \hat{j}, p_{in,t}, \hat{k})} \right),
\end{aligned}$$

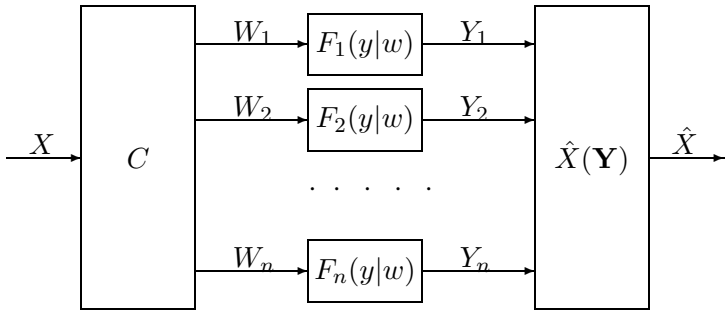

Figure G.1: General deterministic/randomized bit-to-sequence mapper with independent BI-SO observational channels.

where

$$\begin{aligned} G_{i,j,k}(x,y,\hat{i},z,\hat{j},w,\hat{k}) &:= xy^{\hat{i}}(1-y)^{i-\hat{i}}z^{\hat{j}}(1-z)^{j-\hat{j}}w^{\hat{k}}(1-w)^{k-\hat{k}} \\ &\quad +(1-x)(1-y)^{\hat{i}}y^{i-\hat{i}}(1-z)^{\hat{j}}z^{j-\hat{j}}(1-w)^{\hat{k}}w^{k-\hat{k}}. \end{aligned}$$

G.5 Proof of Theorem 5.7

We provide a proof of a more general theorem, which includes general error correcting codes and multiuser detection as special cases and is formally stated as follows.

As in Figure G.1, consider any *deterministic/randomized* sequence mapper[1] $C : \{0,1\} \mapsto \{0,1\}^n$ and $\mathbf{W} = C(X)$. Each coordinate W_i of $\mathbf{W}$ is passed through independent BI-SO channels $F_i(dy|w)$ to generate the observation Y_i, $i \in \{1,2,\cdots,n\}$. Let $\hat{X}(\mathbf{Y})$ be the MAP detector, and define

$$\begin{aligned} p_e(\{F_i\}) &:= \mathsf{P}_{X,\mathbf{Y}}(\hat{X}(\mathbf{Y}) \neq X) \\ \text{and } \ \mathrm{BNP}(\{F_i\}) &:= \mathsf{E}_{X,\mathbf{Y}}\left\{\sqrt{\frac{\mathsf{P}(\bar{X}|\mathbf{Y})}{\mathsf{P}(X|\mathbf{Y})}}\right\} \end{aligned}$$

as the error probability and BNP value of this $X \mapsto \mathbf{Y}$ vector channel given the conditional distributions $\{F_i\}$ of the observational channels. We then have the following theorem.

Theorem G.1 *For any uniform/nonuniform binary input distribution on X, we have*

$$\mathrm{BNP}(\{F_i\}) \leq \mathrm{BNP}(\{F_{BSC,\tilde{p}_i}\}),$$

where for any i, $\tilde{p}_i$ satisfies $4\tilde{p}_i(1-\tilde{p}_i) = \int 4p(1-p)dP_i(p)$. The integrator $dP_i(p)$ is the equivalent probabilistic weight in the BSC decomposition of channel F_i as described in Section 5.1.2.

Theorem 5.7 is a special case of Theorem G.1 obtained by letting $X \mapsto C(X)$ be the binary-input/vector-output support tree channel.

Note 1: In the setting of Theorem G.1, we only require all constituent channels to be of BI-SO. The bit-to-sequence mapper $X \mapsto C(X)$ does not need to possess any symmetric structure, which is different from the case of LDPC codes.

[1] X and $C(X)$ can be regarded as a binary-input/vector-output channel. Or $C(X)$ can be the subspace of codewords corresponding to information bit X.

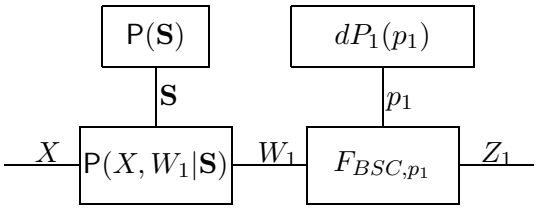

Figure G.2: The factor graph of the five random variables: $X, Z_1, W_1, \mathbf{S}$, and p_1.

Note 2: The definition of ESB in (5.2) is valid for general BI-NSO channels with arbitrary input distributions. However, with a non-uniform input distribution, $\mathrm{ESB}(F_{BSC,p}) \neq 4p(1-p)$. This is the reason why in Theorem G.1, we deliberately use $4\tilde{p}_i(1-\tilde{p}_i) = \int 4p(1-p)dP_i(p)$ instead of $\mathrm{ESB}(F_{BSC,\tilde{p}_i}) = \mathrm{ESB}(F_i)$.

Proof of Theorem G.1*:* By rewriting each BI-SO channel F_i as the probabilistic combination of BSCs with weights $dP_i(p_i)$, each observation y_i can be viewed as a pair $(z_i, p_i) \in \{0,1\} \times [0, 1/2]$, where z_i is the binary output of F_{BSC,p_i} and p_i is the side information specifying the crossover probability of the corresponding BSC. Taking the marginal approach, we will focus on $y_1 = (z_1, p_1)$ and treat all $y_2, y_3, y_4, \cdots, y_n$ as the side information $\mathbf{s}$. The conditional probability $\mathsf{P}(\cdot|\mathbf{S} = \mathbf{s}, p_1)$ can then be factored as

$$\begin{aligned} \mathsf{P}(X, W_1, Z_1|\mathbf{S} = \mathbf{s}, p_1) &= (p_1 + (1-2p_1)\delta(Z_1 - W_1))\,\mathsf{P}(X, W_1|\mathbf{S} = \mathbf{s}, p_1) \\ &= (p_1 + (1-2p_1)\delta(Z_1 - W_1))\,\mathsf{P}(X, W_1|\mathbf{S} = \mathbf{s}), \end{aligned} \tag{G.7}$$

where $\delta(x) = 1$ iff $x = 0$. To be able to substitute $\mathsf{P}(X, W_1|\mathbf{S}, p_1)$ with $\mathsf{P}(X, W_1|\mathbf{S})$, we use the fact that knowing what type of BSCs we are facing (namely, knowing p_1) provides no information[2] about the input distribution (W_1). This fact also implies that $dP_1(p)$ does not depend on the distribution of $\mathbf{S}$ either. As a result, we have

$$\mathsf{P}(\mathbf{S}, p_1) = \mathsf{P}(\mathbf{S})\mathsf{P}(p_1) = \mathsf{P}(\mathbf{S})dP_1(p_1). \tag{G.8}$$

By (G.7) and (G.8), the corresponding factor graph is drawn in Figure G.2. We can write the conditional distribution $\mathsf{P}(X, W_1|\mathbf{S})$ in the matrix form:

$$\begin{aligned} &\begin{pmatrix} \mathsf{P}(W_1 = 0, X = 0|\mathbf{S} = \mathbf{s}) & \mathsf{P}(W_1 = 0, X = 1|\mathbf{S} = \mathbf{s}) \\ \mathsf{P}(W_1 = 1, X = 0|\mathbf{S} = \mathbf{s}) & \mathsf{P}(W_1 = 1, X = 1|\mathbf{S} = \mathbf{s}) \end{pmatrix} \\ = &\begin{pmatrix} a & b \\ c & d \end{pmatrix}, \end{aligned}$$

where a, b, c, and d are functions of $\mathbf{s}$ satisfying $a, b, c, d \geq 0$ and $a + b + c + d = 1$. It is worth repeating that a, b, c, and d do not depend on p_1. The conditional input-output distribution $\mathsf{P}(X, Z_1|\mathbf{S}, p_1)$ then becomes

$$\begin{aligned} &\begin{pmatrix} \mathsf{P}(Z_1 = 0, X = 0|\mathbf{S} = \mathbf{s}, p_1) & \mathsf{P}(Z_1 = 0, X = 1|\mathbf{S} = \mathbf{s}, p_1) \\ \mathsf{P}(Z_1 = 1, X = 0|\mathbf{S} = \mathbf{s}, p_1) & \mathsf{P}(Z_1 = 1, X = 1|\mathbf{S} = \mathbf{s}, p_1) \end{pmatrix} \\ = &\begin{pmatrix} a(1-p_1) + cp_1 & b(1-p_1) + dp_1 \\ ap_1 + c(1-p_1) & bp_1 + d(1-p_1) \end{pmatrix}. \end{aligned}$$

[2] $dP(p)$ only depends on the channel distribution $F(dz|w)$, not on the *a priori* distribution of W. This is a special property of the BSC decomposition mentioned in Section 5.1.2, which does not hold for BNSC decomposition for general BI-NSO channels.

The value of BNP for the $X \mapsto \mathbf{Y}$ channel (or equivalently $X \mapsto (Z_1, p_1, \mathbf{S})$) becomes $\mathrm{BNP} = \mathsf{E}\left\{\sqrt{\frac{\mathsf{P}(\bar{X}|Z_1,p_1,\mathbf{S})}{\mathsf{P}(X|Z_1,p_1,\mathbf{S})}}\right\}$. Taking the expectation step by step, we have

$$\begin{aligned}
&\mathrm{BNP}_{X,Z_1|p_1,\mathbf{S}} \\
:= \;& \mathsf{E}_{X,Z_1|p_1,\mathbf{S}}\left\{\sqrt{\frac{\mathsf{P}(\bar{X}|Z_1,p_1,\mathbf{S})}{\mathsf{P}(X|Z_1,p_1,\mathbf{S})}}\right\} \\
= \;& 2\sqrt{(a(1-p_1)+cp_1)\,(b(1-p_1)+dp_1)} + 2\sqrt{(ap_1+c(1-p_1))\,(bp_1+d(1-p_1))}.
\end{aligned} \tag{G.9}$$

By Proposition G.4 (stated at the end of this proof), $\mathrm{BNP}_{X,Z_1|p_1,\mathbf{S}}$ is a concave function with respect to $\beta := 4p_1(1-p_1)$ for all valid a, b, c, and d. By Jensen's inequality, for any channel F_C,

$$\begin{aligned}
\mathrm{BNP}_{X,Z_1,p_1|\mathbf{S}} &:= \mathsf{E}_{X,Z_1,p_1|\mathbf{S}}\left\{\sqrt{\frac{\mathsf{P}(\bar{X}|Z_1,p_1,\mathbf{S})}{\mathsf{P}(X|Z_1,p_1,\mathbf{S})}}\right\} \\
&= \int_{p=0}^{1/2} \mathrm{BNP}_{X,Z_1|p_1,\mathbf{S}} dP_1(p_1) \\
&\le \mathrm{BNP}_{X,Z_1|\tilde{p}_1,\mathbf{S}},
\end{aligned} \tag{G.10}$$

where $\tilde{p}_1$ is the crossover probability such that $4\tilde{p}_1(1-\tilde{p}_1) = \int 4p_1(1-p_1)dP_1(p_1)$. By (G.10) and noting that $F_{BSC,\tilde{p}_1}$ is the universal maximizing distribution for any realization of $\mathbf{S}$, we obtain that

$$\begin{aligned}
\mathrm{BNP}(\{F_1, F_2, \cdots, F_n\}) &= \int \mathrm{BNP}_{X,Z_1,p_1|\mathbf{S}} d\mathsf{P}(\mathbf{S}) \\
&\le \int \mathrm{BNP}_{X,Z_1|\tilde{p}_1,\mathbf{S}} d\mathsf{P}(\mathbf{S}) \\
&= \mathrm{BNP}(\{F_{BSC,\tilde{p}_1}, F_2, F_3, \cdots, F_n\}).
\end{aligned}$$

By repeatedly applying this BNP-increasing channel replacement until all constituent channels F_i are replaced by $F_{BSC,\tilde{p}_i}$, the proof of Theorem G.1 is complete. ■

Proposition G.4 *For any constants $a, b, c, d \ge 0$ and $p \in [0, 1/2]$, we have*

$$\sqrt{(a(1-p)+cp)\,(b(1-p)+dp)} + \sqrt{(ap+c(1-p))\,(bp+d(1-p))}$$

is a concave function of $\beta := 4p(1-p)$.

Proof: This proof involves several changes of variables. It is worth noting that this proposition is a pure algebraic statement and the notations involved are irrelevant to those of the LDPC code problem.

We first let $X = 1-2p$, $A = \frac{a+c}{|a-c|}$, and $B = \frac{b+d}{|b-d|}$. Then the problem becomes to prove that both

$$f(X) = \sqrt{A+X}\sqrt{B+X} + \sqrt{A-X}\sqrt{B-X}$$

and

$$g(X) = \sqrt{A-X}\sqrt{B+X} + \sqrt{A+X}\sqrt{B-X}$$

are concave functions of $\beta = 1 - X^2$, for all $A, B \in [1, \infty]$ and $X \in [0, 1]$. We focus on the concavity of $f(X)$ first. Using the chain rule,

$$\frac{df(X)}{d\beta} = \frac{df(X)}{dX}\frac{dX}{d\beta} = \frac{1}{2X}\left(-\sqrt{\frac{A+X}{B+X}} - \sqrt{\frac{B+X}{A+X}} + \sqrt{\frac{A-X}{B-X}} + \sqrt{\frac{B-X}{A-X}}\right).$$

Since $\frac{d^2 f(X)}{d\beta^2} = -\frac{1}{2X}\frac{d}{dX}\frac{df(X)}{d\beta}$, showing the concavity of $f(X)$ as a function of β is equivalent to showing:

$$\begin{aligned}
&\frac{d}{dX}\frac{df(X)}{d\beta} \\
&= \frac{1}{2X^2}\left(\frac{1}{\sqrt{(A+X)(B+X)}}\left(A+B+X+\frac{X}{2}\left(\frac{A+X}{B+X}+\frac{B+X}{A+X}\right)\right)\right. \\
&\quad \left. + \frac{1}{\sqrt{(A-X)(B-X)}}\left(-A-B+X+\frac{X}{2}\left(\frac{A-X}{B-X}+\frac{B-X}{A-X}\right)\right)\right) \\
&\triangleq \frac{1}{2X^2} f_2(X) \geq 0.
\end{aligned}$$

To show $f_2(X) \geq 0$, we first note that $f_2(0) = 0$. Its first derivative is

$$\frac{df_2(X)}{dX} = \frac{3}{4}(A-B)^2 X\left(\frac{(A-X)+(B-X)}{\sqrt{(A-X)(B-X)}^5} - \frac{(A+X)+(B+X)}{\sqrt{(A+X)(B+X)}^5}\right).$$

By Lemma G.6 (stated at the end of this proof), we have $\frac{df_2(X)}{dX} \geq 0$. Thus $f_2(X) \geq 0$, which implies that $f(X)$ is concave as a function of β.

For $g(X)$, we have

$$\frac{dg(X)}{d\beta} = \frac{1}{2X}\left(\sqrt{\frac{A+X}{B-X}} - \sqrt{\frac{B-X}{A+X}} - \sqrt{\frac{A-X}{B+X}} + \sqrt{\frac{B+X}{A-X}}\right).$$

Since $\frac{d^2 g(X)}{d\beta^2} = -\frac{1}{2X}\frac{d}{dX}\frac{dg(X)}{d\beta}$, showing the concavity of $g(X)$ is equivalent to showing

$$\begin{aligned}
&\frac{d}{dX}\frac{dg(X)}{d\beta} \\
&= \frac{1}{2X^2}\left(\frac{1}{\sqrt{(A+X)(B-X)}}\left(-A+B-X+\frac{X}{2}\left(\frac{A+X}{B-X}+\frac{B-X}{A+X}\right)\right)\right. \\
&\quad \left. + \frac{1}{\sqrt{(A-X)(B+X)}}\left(A-B-X+\frac{X}{2}\left(\frac{A-X}{B+X}+\frac{B+X}{A-X}\right)\right)\right) \\
&\triangleq \frac{1}{2X^2} g_2(X) \geq 0.
\end{aligned}$$

To show $g_2(X) \geq 0$, we first note that $g_2(0) = 0$. Its first derivative is

$$\frac{dg_2(X)}{dX} = \frac{3}{4}(A+B)^2 X\left(\frac{(A+X)-(B-X)}{\sqrt{(A+X)(B-X)}^5} + \frac{(B+X)-(A-X)}{\sqrt{(A-X)(B+X)}^5}\right).$$

By Lemma G.6, we have $\frac{dg_2(X)}{dX} \geq 0$, which implies that $g(X)$ is concave as a function of β. This completes the proof of Proposition G.4. ■

Lemma G.6 *For all $a, b, c \geq 0$, we have*

$$\frac{a+b}{(ab)^{5/2}} \geq \frac{(a+c)+(b+c)}{((a+c)(b+c))^{5/2}}$$
$$\frac{(a+c)-b}{((a+c)b)^{5/2}} + \frac{(b+c)-a}{((b+c)a)^{5/2}} \geq 0.$$

Proof: By noting that $\frac{a+b}{ab} = \frac{1}{a}+\frac{1}{b} \geq \frac{1}{a+c}+\frac{1}{b+c} = \frac{(a+c)+(b+c)}{(a+c)(b+c)}$ and $0 \leq ab \leq (a+c)(b+c)$, we prove the first inequality.

For the second inequality, without loss of generality, we assume $a \leq b$. We then observe that

$$\begin{aligned} (a+c)b &\geq (b+c)a \geq 0 \\ \frac{(b+c)-a}{(b+c)a} &\geq 0 \\ \frac{(a+c)-b}{(a+c)b} + \frac{(b+c)-a}{(b+c)a} &= \frac{1}{b} - \frac{1}{a+c} + \frac{1}{a} - \frac{1}{b+c} \geq 0. \end{aligned} \tag{G.11}$$

Considering (G.11), after multiplying the non-negative second term $\frac{(b+c)-a}{(b+c)a}$ by a larger factor $\frac{1}{\sqrt{(b+c)a}^3}$ and the *possibly-negative* first term $\frac{(a+c)-b}{(a+c)b}$ by a smaller factor $\frac{1}{\sqrt{(a+c)b}^3}$, the new weighted sum is no less than zero, namely,

$$\frac{(a+c)-b}{((a+c)b)^{5/2}} + \frac{(b+c)-a}{((b+c)a)^{5/2}} \geq 0.$$

This completes the proof. ■

References

[1] K. Adam, "Learning while searching for the best alternative," *Journal of Economic Theory*, vol. 101, pp. 252–280, 2001.

[2] R. Agrawal, M. V. Hegde, and D. Teneketzis, "Asymptotically efficient adaptive allocation rules for the multiarmed bandit problem with switching cost," *IEEE Trans. Automat. Contr.*, vol. 33, no. 10, pp. 899–906, Oct. 1988.

[3] R. Agrawal, D. Teneketzis, and V. Anantharam, "Asymptotically efficient adaptive allocation schemes for controlled i.i.d. processes: Finite parameter space," *IEEE Trans. Automat. Contr.*, vol. 34, no. 3, pp. 258–267, Mar. 1989.

[4] ——, "Asymptotically efficient adaptive allocation schemes for controlled Markov chains: Finite parameter space," *IEEE Trans. Automat. Contr.*, vol. 34, no. 12, pp. 1249–1259, Mar. 1989.

[5] R. Ahlswede, "Group codes do not achieve Shannon's channel capacity for general discrete channels," *Ann. Math. Stat.*, vol. 42, no. 1, pp. 224–240, 1971.

[6] R. Ahlswede and J. Gemma, "Bounds on algebraic code capacities for noisy channels. i," *Information and Control*, vol. 19, no. 2, pp. 124–145, 1971.

[7] S. M. Aji and R. J. McEliece, "The generalized distributive law," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 325–343, March 2000.

[8] ——, "The generalized distributive law and free energy minimization," in *Proc. 39th Annual Allerton Conf. on Comm., Contr., and Computing.* Monticello, IL, USA, Oct. 2001.

[9] V. Anantharam, P. Varaiya, and J. Walrand, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part I: I.i.d. rewards," *IEEE Trans. Automat. Contr.*, vol. 32, no. 11, pp. 968–976, Nov. 1987.

[10] ——, "Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-part II: Markovian rewards," *IEEE Trans. Automat. Contr.*, vol. 32, no. 11, pp. 977–982, Nov. 1987.

[11] A. Banerjee, D. J. Costello, Jr., T. E. Fuja, and P. Massey, "Bit interleaved coded modulation using multiple turbo codes," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Lausanne, Switzerland, June 2002, p. 443.

[12] L. Bazzi, T. Richardson, and R. Urbanke, "Exact thresholds and optimal codes for the binary symmetric channel and Gallager's decoding algorithm A," *IEEE Trans. Inform. Theory*, vol. 50, no. 9, pp. 2010–2021, Sept. 2004.

[13] S. Benedetto and G. Montorsi, "Unveiling turbo codes: Some results on parallel concatenated coding schemes," *IEEE Trans. Inform. Theory*, vol. 42, no. 2, pp. 409–428, Mar. 1996.

[14] A. Bennatan and D. Burshtein, "Iterative decoding of LDPC codes over arbitrary discrete-memoryless channels," in *Proc. 41st Annual Allerton Conf. on Comm., Contr., and Computing*. Monticello, IL, USA, 2003, pp. 1416–1425.

[15] A. Bennatan and D. Burstein, "On the application of LDPC codes to arbitrary discrete-memoryless channels," *IEEE Trans. Inform. Theory*, vol. 50, no. 3, pp. 417–438, March 2004.

[16] J. Berkmann, "On turbo decoding of nonbinary codes," *IEEE Commun. Letters*, vol. 2, no. 4, pp. 94–96, April 1998.

[17] ——, "A symbol-by-symbol MAP decoding rule for linear codes over rings using the dual code," in *Proc. IEEE Int'l. Symp. Inform. Theory*. Cambridge, MA, 1998, p. 90.

[18] C. Berrou and A. Glavieux, "Near optimum error correcting coding and decoding: Turbo-codes," *IEEE Trans. Inform. Theory*, vol. 44, no. 10, pp. 1261–1271, Oct. 1996.

[19] D. A. Berry, "A Bernoulli two-armed bandit," *Ann. Math. Stat.*, vol. 43, no. 3, pp. 871–897, June 1972.

[20] D. A. Berry and B. Fristedt, *Bandit Problems, Sequential Allocation of Experiments*. London: Chapman and Hall, 1985.

[21] J. A. Bucklew, *Large Deviation Techniques in Decision, Simulation, and Estimation*. New York: Wiley, 1990.

[22] D. Burshtein and G. Miller, "Bounds on the performance of belief propagation decoding," *IEEE Trans. Inform. Theory*, vol. 48, no. 1, pp. 112–122, Jan. 2002.

[23] ——, "Asymptotic enumeration methods for analyzing LDPC codes," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1115–1131, June 2004.

[24] J. W. Byers, M. Luby, and M. Mitzenmacher, "A digital fountain approach to asynchronous reliable multicast," *IEEE J. Select. Areas Commun.*, vol. 20, no. 8, pp. 1528–1540, Oct. 2002.

[25] G. Caire, D. Burshtein, and S. Shamai, "LDPC coding for interference mitigation at the transmitter," in *Proc. 40th Annual Allerton Conf. on Comm., Contr., and Computing*. Monticello, IL, USA, Oct. 2002.

[26] G. Caire, S. Shamai, and S. Verdú, "Noiseless data compression with low-density parity-check codes," in *Advances in Network Information Theory,* P. Gupta, G. Kramer and A. J. van Wijngaarden, Eds.*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, vol. 66. American Mathematical Society, 2004, pp. 263–284.

[27] G. Caire, G. Taricco, and E. Biglieri, "Bit-interleaved coded modulation," *IEEE Trans. Inform. Theory*, vol. 44, no. 3, pp. 927–946, May 1998.

[28] H. Chernoff, *Sequential Analysis and Optimal Design.* Philadelphia: Society for Industrial and Applied Mathematics, 1972.

[29] S. Y. Chung, "On the construction of some capacity-approaching coding schemes," Ph.D. dissertation, MIT, 2000.

[30] S. Y. Chung, G. D. Forney, Jr., T. J. Richardson, and R. L. Urbanke, "On the design of low-density parity-check codes within 0.0045 dB of the Shannon limit," *IEEE Commun. Letters*, vol. 5, no. 2, pp. 58–60, Feb. 2001.

[31] S. Y. Chung, T. J. Richardson, and R. L. Urbanke, "Analysis of sum-product decoding of low-density parity-check codes using a Gaussian approximation," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 657–670, Feb. 2001.

[32] M. K. Clayton, "Covariate models for Bernoulli bandits," *Sequential Analysis*, vol. 8, no. 4, pp. 405–426, 1989.

[33] T. M. Cover and J. A. Thomas, *Elements of Information Theory.* New York: Wiley, 1991.

[34] M. C. Davey and D. J. C. MacKay, "Low-density parity check codes over GF(q)," *IEEE Commun. Letters*, vol. 2, no. 6, pp. 165–167, June 1998.

[35] A. Dembo and O. Zeitouni, *Large Deviation Techniques and Applications.* New York: Springer, 1998.

[36] C. Di, D. Proietti, E. Telatar, T. J. Richardson, and R. L. Urbanke, "Finite-length analysis of low-density parity-check codes on the binary erasure channel," *IEEE Trans. Inform. Theory*, vol. 48, no. 6, pp. 1570–1579, June 2002.

[37] D. Divsalar, H. Jin, and R. J. McEliece, "Coding theorems for "Turbo-Like" codes," in *Proc. 36th Annual Allerton Conf. on Comm., Contr., and Computing.* Monticello, IL, USA, 1998, pp. 210–220.

[38] R. L. Dobrushin, "Asymptotic optimality of group and systematic codes for some channels," *Theory of Probability and its Applications*, vol. 8, no. 1, pp. 47–60, 1963.

[39] P. Elias, "Coding for noisy channels," *IRE Conv. Rec.*, no. 4, pp. 37–46, March 1955.

[40] U. Erez and G. Miller, "The ML decoding performance of LDPC ensembles over $\mathbf{Z}_q$," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Yokohama, Japan, 2003, p. 86.

[41] B. J. Frey, R. Koetter, and A. Vardy, "Signal-space characterization of iterative decoding," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 766–781, Feb. 2001.

[42] E. Frostig and G. Weiss, "Four proofs of Gittins' multiarmed bandit theorem," *Applied Probability Trust*, Nov. 1999.

[43] R. G. Gallager, *Low-Density Parity-Check Codes*, ser. Research Monograph Series. Cambridge, MA: MIT Press, 1963, no. 21.

[44] ——, *Information Theory and Reliable Communication.* New York: John Wiley and Sons, 1968.

[45] J. Garcia-Frias, "Decoding of low-density parity check codes over finite-state binary Markov channels," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Washington, DC, 2001, p. 72.

[46] J. Garcia-Frias and W. Zong, "Approaching near shannon performance by iterative decoding of linear codes with low-density generator matrix," *IEEE Commun. Letters*, vol. 7, no. 6, pp. 266–268, June 2003.

[47] B. Ghosh and P. Sen, *Handbook of Sequential Analysis.* New York: Dekker, 1991.

[48] J. C. Gittins, "Bandit processes and dynamic allocation indices," *J. Royal Statistical Society. Series B (Methodological)*, vol. 41, no. 2, pp. 148–177, 1979.

[49] ——, "A dynamic allocation index for the discounted multiarmed bandit problem," *Biometrika*, vol. 66, no. 3, pp. 561–565, Dec. 1979.

[50] J. Hou, P. H. Siegel, and L. B. Milstein, "Performance analysis and code optimization of low density parity-check codes on Rayleigh fading channels," *IEEE J. Select. Areas Commun.*, vol. 19, no. 5, pp. 924–934, May 2001.

[51] J. Hou, P. H. Siegel, L. B. Milstein, and H. D. Pfister, "Capacity-approaching bandwidth-efficient coded modulation schemes based on low-density parity-check codes," *IEEE Trans. Inform. Theory*, vol. 49, no. 9, pp. 2141–2155, Sept. 2003.

[52] H. Jin and R. J. McEliece, "RA codes achieve AWGN channel capacity," in *Proc. 13th Int'l Symp. Applied Algebra, Algebraic Algorithms, and Error Correcting Codes.* Honolulu, HI, USA, 1999, pp. 10–18.

[53] H. Jin and R. McEliece, "Coding theorems for turbo code ensembles," *IEEE Trans. Inform. Theory*, vol. 48, no. 6, pp. 1451–1461, June 2002.

[54] H. Jin and T. J. Richardson, "Fast density evolution," in *Proc. 38th Conf. Inform. Sciences and Systems.* Princeton, NJ, USA, 2004.

[55] K. Kasai, T. Shibuya, and K. Sakaniwa, "Detailed representation of irregular LDPC code ensembles and density evolution," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Yokohama, Japan, 2003, p. 121.

[56] M. N. Katehakis and H. Robbins, "Sequential choice from several populations," in *Proc. Nat. Acad. Sci. USA*, vol. 92, Sept. 1995, pp. 8584–8585.

[57] A. Kavčić, X. Ma, and M. Mitzenmacher, "Binary intersymbol interference channels: Gallager codes, density evolution and code performance bound," *IEEE Trans. Inform. Theory*, vol. 49, no. 7, pp. 1636–1652, July 2003.

[58] A. Khandekar, "Graph-based codes and iterative decoding," Ph.D. dissertation, California Institute of Technology, 2002.

[59] A. Khandekar and R. J. McEliece, "A lower bound on the iterative decoding threshold of irregular LDPC code ensembles," in *Proc. 36th Conf. Inform. Sciences and Systems.* Princeton, NJ, USA, 2002.

[60] V. Krishnamurthy and R. J. Evans, "Hidden Markov model multiarm bandits: A methodology for beam scheduling in multitarget tracking," *IEEE Trans. Signal Processing*, vol. 49, no. 12, pp. 2893–2908, Dec. 2001.

[61] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[62] S. R. Kulkarni, "On bandit problems with side observations and learnability," in *Proc. 31st Allerton Conf. Commun. Contr. Comp.*, Sept. 1993, pp. 83–92.

[63] S. R. Kulkarni and G. Lugosi, "Finite-time lower bounds for the two-armed bandit problem," *IEEE Trans. Automat. Contr.*, vol. 45, no. 4, pp. 711–714, Apr. 2000.

[64] V. Kumar, O. Milenkovic, and K. Prakash, "On graphical representations of algebraic codes suitable for iterative decoding," in *Proc. 39th Conf. Inform. Sciences & Systems.* Baltimore, MD, USA, March 2005.

[65] B. M. Kurkoski, P. H. Siegel, and J. K. Wolf, "Joint message-passing decoding of LDPC codes and partial-response channels," *IEEE Trans. Inform. Theory*, vol. 48, no. 6, pp. 1410–1422, June 2002.

[66] T. L. Lai and H. Robbins, "Asymptotically optimal allocation of treatments in sequential experiments," in *Design of Experiments : Ranking and Selection,* Thomas J. Santner, Ajit C. Tamhane Eds. New York: Dekker, 1984.

[67] ——, "Asymptotically efficient allocation rules," *Adv. Appl. Math.*, vol. 6, no. 1, pp. 4–22, 1985.

[68] T. L. Lai and S. Yakowitz, "Machine learning and nonparametric bandit theory," *IEEE Trans. Automat. Contr.*, vol. 40, no. 7, pp. 1199–1209, July 1995.

[69] I. Land, P. A. Hoeher, S. Huettinger, and J. Huber, "Bounds on information combining," in *Proc. 3rd Int'l. Symp. Turbo Codes & Related Topics.* Brest, France, 2003, pp. 39–42.

[70] F. Lehmann, "Distance properties of irregular LDPC codes," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Yokohama, Japan, 2003, p. 85.

[71] G. Li, I. Fair, and W. Krzymień, "Analysis of nonbinary LDPC codes using Gaussian approximation," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Yokohama, Japan, 2003, p. 234.

[72] J. Li, K. R. Narayanan, E. Kurtas, and C. N. Georghiades, "On the performance of high-rate TPC/SPC codes and LDPC codes over partial response channels," *IEEE Trans. Commun.*, vol. 50, no. 5, pp. 723–734, May 2002.

[73] S. Litsyn and V. Shevelev, "On ensembles of low-density parity-check codes: Asymptotic distance distributions," *IEEE Trans. Inform. Theory*, vol. 48, no. 4, pp. 887–908, Apr. 2002.

[74] H. A. Loeliger, "Signal sets matched to groups," *IEEE Trans. Inform. Theory*, vol. 37, no. 6, pp. 1675–1681, Nov. 1991.

[75] M. G. Luby, "LT codes," in *Proc. 43rd Annu. IEEE Symp. Foundations of Compuer Science (FOCS'02).* Vancouver, BC, Canada, Nov. 2002, pp. 271–280.

[76] M. G. Luby, M. Mitzenmacher, M. A. Shokrollahi, and D. A. Spielman, "Analysis of low-density codes and improved designs using irregular graphs," in *Proc. 30th Annu. ACM Symp. Theory of Computing*, 1998, pp. 249–258.

[77] ——, "Efficient erasure correcting codes," *IEEE Trans. Inform. Theory*, vol. 47, pp. 569–584, Feb. 2001.

[78] D. J. C. MacKay, "Good error-correcting codes based on very sparse matrices," *IEEE Trans. Inform. Theory*, vol. 45, no. 2, pp. 399–431, Mar. 1999.

[79] E. E. Majani and H. Rumsey Jr., "Two results on binary-input discrete memoryless channels," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Budapest, Hungary, June 1991, p. 104.

[80] R. J. McEliece, "Are turbo-like codes effective on nonstandard channels?" *IEEE Inform. Theory Society Newsletter*, vol. 51, no. 4, Dec. 2001.

[81] R. J. McEliece, D. J. C. Mackay, and J. F. Cheng, "Turbo decoding as an instance of Pearl's "Belief Propagation" algorithm," *IEEE J. Select. Areas Commun.*, vol. 16, no. 2, pp. 140–152, Feb. 1998.

[82] G. Miller and G. Cohen, "The rate of regular LDPC codes," *IEEE Trans. Inform. Theory*, vol. 49, no. 11, pp. 2989–2992, Nov. 2003.

[83] R. Narayanaswami, "Coded modulation with low-density parity-check codes," Master's thesis, Texas A&M, 2001.

[84] A. Orlitsky, K. Viswanathan, and J. Zhang, "Stopping set distribution of LDPC code ensembles," in *Proc. IEEE Int'l. Symp. Inform. Theory.* Yokohama, Japan, 2003, p. 123.

[85] P. Oswald and A. Shokorollahi, "Capacity-achieving sequences for the erasure channel," *IEEE Trans. Inform. Theory*, vol. 48, no. 12, pp. 3017–3028, Dec. 2002.

[86] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Network of Plausible Inference.* San Mateo, CA: Morgan Kaufmann, 1988.

[87] L. Ping and K. Y. Wu, "Concatenated tree codes: A low-complexity, high-performance approach," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 791–799, Feb. 2001.

[88] R. Pyndiah, "Near-optimum decoding of product codes: Block turbo codes," *IEEE Trans. Commun.*, vol. 46, no. 8, pp. 1003–1010, Aug. 1998.

[89] E. Ratzer and D. MacKay, "Sparse low-density parity-check codes for channels with cross-talk," in *Proc. IEEE Inform. Theory Workshop.* Paris, France, March 31 – April 4 2003.

[90] T. J. Richardson, M. A. Shokrollahi, and R. L. Urbanke, "Design of capacity-approaching irregular low-density parity-check codes," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 619–637, Feb. 2001.

[91] T. J. Richardson and R. L. Urbanke, "Multi-edge type LDPC codes," personal communications.

[92] ——, "The capacity of low-density parity-check codes," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 599–618, Feb. 2001.

[93] ——, "Efficient encoding of low-density parity-check codes," *IEEE Trans. Inform. Theory*, vol. 47, no. 2, pp. 638–656, Feb. 2001.

[94] H. Robbins, "Some aspects of the sequential design of experiments," *Bull. Am. Math. Soc.*, vol. 58, pp. 527–535, 1952.

[95] P. Robertson and T. Wörz, "Bandwidth-efficient turbo trellis-coded modulation using punctured component codes," *IEEE Trans. Inform. Theory*, vol. 16, no. 2, pp. 206–218, Feb. 1998.

[96] J. Sarkar, "One-armed bandit problems with covariates," *Ann. Statist.*, vol. 19, no. 4, pp. 1978–2002, 1991.

[97] S. Shamai and I. Sason, "Variations on the Gallager bounds, connections, and applications," *IEEE Trans. Inform. Theory*, vol. 48, no. 12, pp. 3029–3051, Dec. 2002.

[98] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, July and Oct. 1948.

[99] N. Shulman and M. Feder, "The uniform distribution as a universal prior," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1356–1362, June 2004.

[100] M. Sipser and D. A. Spielman, "Expander codes," *IEEE Trans. Inform. Theory*, vol. 42, no. 6, pp. 1710–1722, Nov. 1996.

[101] D. A. Spielman, "Linear-time encodable and decodable error-correcting codes," *IEEE Trans. Inform. Theory*, vol. 42, no. 6, pp. 1723–1731, Nov. 1996.

[102] R. Storn and K. Price, "Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces," *Journal of Global Optimization*, vol. 11, pp. 341–359, 1997.

[103] I. Sutskover, S. Shamai, and J. Ziv, "Extremes of information combining," in *Proc. 41st Annual Allerton Conf. on Comm., Contr., and Computing.* Monticello, IL, USA, 2003.

[104] W. Tan and J. R. Cruz, "Signal-to-noise ratio mismatch for low-density parity-check coded magnetic recording channels," *IEEE Trans. Magn.*, vol. 40, no. 2, pp. 498–506, Mar. 2004.

[105] S. ten Brink, "Convergence behavior of iteratively decoded parallel concatenated codes," *IEEE Trans. Commun.*, vol. 49, no. 10, pp. 1727–1737, Oct. 2001.

[106] A. Thangaraj and S. W. McLaughlin, "Thresholds and scheduling for LDPC-coded partial response channels," *IEEE Trans. Magn.*, vol. 38, no. 5, pp. 2307–2309, Sep. 2002.

[107] G. Ungerboeck, "Channel coding with multilevel/phase signals," *IEEE Trans. Inform. Theory*, vol. 28, no. 1, pp. 55–66, Jan. 1982.

[108] U. Wachsmann, R. F. H. Fischer, and J. B. Huber, "Multilevel codes: Theoretical concepts and practical design rules," *IEEE Trans. Inform. Theory*, vol. 45, no. 5, pp. 1361–1391, July 1999.

[109] C. C. Wang, S. R. Kulkarni, and H. V. Poor, "Bandit problems with side observations," in *Proc. 41st IEEE Conf. on Decision and Contr.*, vol. 4. Las Vegas, USA, Dec. 2002, pp. 3988–3993.

[110] ——, "Bandit problems with arbitrary side observations," in *Proc. 42nd IEEE Conf. on Decision and Contr.*, vol. 3. Maui, USA, Dec. 2003, pp. 2948–2953.

[111] ——, "Density evolution for asymmetric memoryless channels," in *Proc. Int'l. Symp. Turbo Codes & Related Topics*. Brest, France, Sept. 2003, pp. 121–124, to appear in *IEEE Trans. Inform. Theory*.

[112] ——, "On finite-dimensional bounds for LDPC-like codes with iterative decoding," in *Proc. Int'l Symp. Inform. Theory & its Applications*. Parma, Italy, Oct. 2004, submitted to *IEEE Trans. Inform. Theory*.

[113] ——, "Arbitrary side observations in bandit problems," *Advances in Applied Mathematics*, vol. 34, no. 4, pp. 903–938, May 2005, special issue dedicated to Dr. David P. Robbins.

[114] ——, "Bandit problems with side observations," *IEEE Trans. Automat. Contr.*, vol. 50, no. 5, pp. 338–355, May 2005.

[115] ——, "On the typicality of the linear code among the LDPC coset code ensemble," in *Proc. 39th Conf. Inform. Sciences and Systems*. Baltimore, USA, March 2005.

[116] M. Woodroofe, "A one-armed bandit problem with a concomitant variable," *J. Amer. Stat. Assoc.*, vol. 74, no. 368, pp. 799–806, Dec 1979.

[117] M. Yang and W. E. Ryan, "Lowering the error-rate floors of moderate-length high-rate irregular LDPC codes," in *Proc. IEEE Int'l. Symp. Inform. Theory*. Yokohama, Japan, 2003, p. 237.

[118] J. S. Yedida, W. T. Freeman, and Y. Weiss, "Bethe free energy, Kikuchi approximations, and belief propagation algorithms," Mitsubishi Electric Research Laboratories, Technical Report TR2001-16, 2001.

[119] T. Zoubeidi, "Optimal allocations in sequential tests involving two populations with covariates," *Commun. Statist.: Theory and Methods*, vol. 23, no. 4, pp. 1215–1225, 1994.

Index