DISSERTATION

INFORMATION THEORETIC PROBLEMS IN NETWORKS AND ACTIVE SENSING

SUBMITTED BY

HUA LI

ELECTRICAL AND COMPUTER ENGINEERING

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

COLORADO STATE UNIVERSITY

FORT COLLINS, COLORADO

FALL 2007

COLORADO STATE UNIVERSITY

October 16, 2007

WE HEREBY RECOMMEND THAT THE DISSERTATION PREPARED UNDER OUR SUPERVISION BY HUA LI ENTITLED INFORMATION THEORETIC PROBLEMS IN ACTIVE SENSING AND NETWORKS BE ACCEPTED AS FULFILLING IN PART REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY.

Committee on Graduate Work

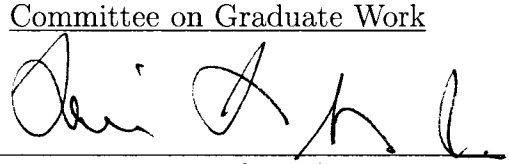

Louis L. Scharf

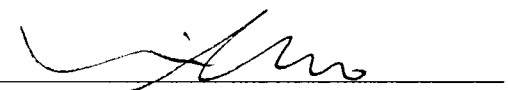

J. Rockey Luo

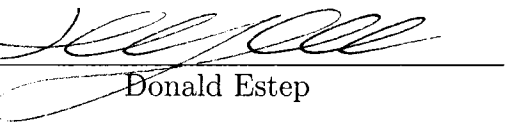

Donald Estep

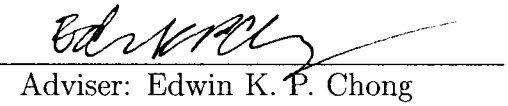

Adviser: Edwin K. P. Chong

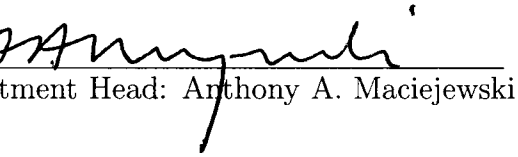

Department Head: Anthony A. Maciejewski

ABSTRACT OF DISSERTATION

INFORMATION THEORETIC PROBLEMS IN NETWORKS AND ACTIVE SENSING

This dissertation covers three related topics. They are on network information theory, search, and tracking respectively.

In the first part of the first topic, we propose a group theoretic model for information. Exploiting this formalization, we identify a comprehensive both qualitative and quantitative parallelism between information lattices and subgroup lattices. As a consequence of this fundamental relation, we show that any continuous law holds in general for the entropies of information elements if and only if the same law holds in general for the log-indices of subgroups. By constructing subgroup counterexamples we find surprisingly that common information obeys neither the submodularity nor the supermodularity law. Our mathematic model for information is conceptually significant.

In the second part of the first topic, we show that none of the three extra conditional mutual information terms in the Zhang-Yeung inequality can be dropped for the inequality to remain valid and that quasi-Hamilton groups satisfy the Ingleton inequality. This is the first time that certain classes of non-abelian groups are found to satisfy the Ingleton inequality.

In the second topic, we study a class of search problems. Both bounded and unbounded search problems are considered. We show that the Bounded Discrete Linear Search Problem is quadratic-time solvable but that the graph search problem is NP-complete, derive bounds for the erroneous BDLSP, and show that optimal policies for an Unbounded Discrete Linear Search Problem (UBDLSP) exist if and only if the double-sided mean of its underlying distribution is finite. We propose a provably effective procedure approximating optimal values and optimal

policies for UBDLSPs and show that the increments of the optimal policies for UBDLSPs with heavy-tailed distributions are necessarily unbounded.

In the third topic, we study the fundamental limits of trackability using information theoretic approach. We show that for a target to be trackable the minimum query quota required is no less than the entropy rate H of the Markov chain of the target but no more than $\lceil H + 1 \rceil$. Subsequently, we propose an adaptive strategy to learn the target motion law and track the target simultaneously. It is remarkable that the extra burden of learning the motion law sacrifices no loss of tracking performance in the asymptotic region.

Hua Li

Electrical and Computer Engineering Department

Colorado State University

Fort Collins, CO 80523

Fall 2007

ACKNOWLEDGMENTS

My foremost thanks go to my dissertation advisor Dr. Edwin Chong. Without him, this dissertation would not have been possible. I am grateful for his polished academic guidance and precious trust in my instinct for choosing research topics. His constant quest for deeper understanding of fundamentals and numerous insightful comments and criticisms shaped my research, opened my mind, and permanently changed my way of looking at this world. His infinite patience and always sincere encouragement carried me through difficult times. His influence on me went beyond my research--he showed me as a role-model how to be a better person. I still have difficulties to express myself when it comes to saying thanks to him even though he had taught me how to articulate, perhaps because my deep gratitude towards Dr. Chong has gone "intrinsically" beyond what words can express.

I am grateful to Dr. Louis Scharf, Dr. Donald Estep and Dr. Rockey Luo for their help and good suggestions for my dissertation work.

I would like to thank Dr. Louis Scharf, Dr. Peter Brockwell. Dr. Donald Estep, and many other professors in Electrical and Computer Engineering, Mathematics, and Statistics departments who have taught me classes. Many of them let me sit in their classes for free. Without this good learning environment, it would have been much more difficult or even impossible for me to find the problem, formulate the problem, and finish my dissertation at this level.

I would like to thank all my labmates Dr. Yun Li, Patricia Barbosa, Zhi Zhang, Ramin Zahadi, Lucas Krakow, Sowmya Lolla, Dr. Ying He, Dr. Gang Wu, Dr. Jung-Min Park, Dr. Jeffrey Herdtner, Dr. Dong-Won Shin, James Smith, Vladimir Shestak, Ye Hong, and Luis Briceno for the joyful moments they brought to me.

I wish to thank my parents and grandparents, who are always there when I need them. I feel deeply sorry that my grandmother, whose love I can always indulge in when I was a kid, passed away during my Ph.D. study.

Lastly and most importantly, I wish to thank my lovely wife, Zhifei Fan, on whose constant encouragement, trust, and love I have relied throughout the years.

TABLE OF CONTENTS

Page

Page

Page

Page

LIST OF FIGURES

CHAPTER 1

Introduction

This dissertation covers three related topics. They are on network information theory, search, and tracking respectively.

1.1 A Group Theoretic Model for Information

Recently, we have uncovered an obscure paper written by Shannon [1], in which the notions of information elements and information lattices were proposed to build a general theory for multi-terminal network communication. In Chapter 2, we formalize these two notions and establish isomorphisms between information lattices and certain subgroup lattices. Exploiting this formalization, we identify a comprehensive parallelism between information lattices and subgroup lattices. Qualitatively, we demonstrate isomorphisms between information lattices and subgroup lattices. Quantitatively, we establish a decisive approximation relation between the entropy structures of information lattices and the log-index structures of the corresponding subgroup lattices. This approximation extends the approximation for joint information carried out previously by Chan and Yeung [2]. As a consequence of our approximation result, we show that any continuous law holds in general for the entropies of information elements if and only if the same law holds in general for the log-indices of subgroups. As an application, by constructing subgroup counterexamples we find surprisingly that common information, unlike joint information, obeys neither the submodularity nor the supermodularity law. We emphasize that the notion of information elements is conceptually significant—formalizing it helps to reveal the deep connection between information theory and group theory. The parallelism established here admits an appealing group-action explanation and provides useful insights into the intrinsic structure among infor-

mation elements from a group-theoretic perspective. Part of the material in this chapter has been published in [3] and submitted in [4].

1.2 On the Entropy Function and the Ingleton Inequality

In Chapter 3, we focus on characterizing the range of the entropy function, essential to multi-terminal information theory. We disprove certain potentially valid non-Shannon-type inequalities with a computer-aided search for counterexamples and show that none of the three extra conditional mutual information terms in the Zhang-Yeung inequality can be dropped for the inequality to remain valid. Appealing to the fundamental "bridge" we build between information theory and group theory, a general condition, subsuming all the previously known conditions, is obtained for the Ingleton inequality – we show that quasi-Hamilton groups, including Hamiltonian groups as a subclass, satisfy the Ingleton inequality. To our best knowledge, this is the first time that certain classes of non-abelian groups are found to satisfy the Ingleton inequality. Part of the material in this chapter has been published in [5].

1.3 Search on Lines and Graphs

In Chapter 4, we study a class of search problems. The general setup models the ubiquitous situation where a searcher aims to find an immobile target with minimum expected latency and the location of the target is a discrete random variable distributed over a set of possible locations. Both bounded and unbounded search problems are considered. We first consider the Bounded Discrete Linear Search Problem (BDLSP) – the distribution of the target location has a finite support on the integer line – and show that the BDLSP is quadratic-time solvable by formulating it as a Markov Decision Problem (MDP). However, the graph search problem (GSP) – the target is located on the vertices of a graph – is shown to be NP-

complete. Then we consider the erroneous BDLSP (EBDLSP) the searcher may miss the target with a non-zero probability when the target location is visited and derive lower and upper bounds on the performance loss for the EBDLSP in terms of that of its error-free counterpart BDLSP. In the second part, we consider the Unbounded Discrete Linear Search Problem (UBDLSP) the distribution of the target location has an infinite support on the integer line. Theoretically, we show that an optimal search policy exists if and only if the double-sided mean of the distribution is finite. Then, we focus on symmetric UBDLSPs and establish the expanding property of optimal policies for symmetric UBDLSPs. Algorithmically, we propose a procedure effectively approximating the optimal value and the optimal policy given that the optimal policy is unique. To investigate the convergence rate of the procedure, we study the growth rate of the turning points of optimal policies for symmetric UBDLSP. We show that the increment sequences of the optimal policies for symmetric UBDLSPs with heavy-tailed distributions are necessarily unbounded. Part of the material in this chapter has been submitted in [6].

1.4 On the Fundamental Limits of Target Trackability

In Chapter 5, we propose an information theoretic formulation for the problem of target tracking via sensor querying. Our goal is to study the fundamental limits of trackability. The target motion is modeled by a finite state-space Markov chain. We show that for a target to be trackable the minimum query quota required is no less than the entropy rate H of the Markov chain, but no more than $\lceil H + 1 \rceil$. Subsequently, we consider the adaptive case where the target motion law, namely the probability transition function of the Markov chain, is unknown to the tracker a priori. In this case, the tracker is expected to learn the target motion law and track the target simultaneously. It turns out that the extra burden of learning

the motion law for the tracker sacrifices no loss of tracking performance in the asymptotic region we show that for a target to be universal-trackable no more than $\lceil H+1 \rceil$ number of queries at each time step is required. Part of the material in this chapter has been published in [7] and submitted in [8].

1.5 Other Work on Networks

My other Ph.D. work on networks has been published in [9 12].

List of References

[1] C. E. Shannon, "The lattice theory of information," *IEEE Transactions on Information Theory*, vol. 1, no. 1, pp. 105 107, Feb. 1953.

[2] T. H. Chan and R. W. Yeung, "On a relation between information inequalities and group theory," *IEEE Transactions on Information Theory*, vol. 48, no. 7, pp. 1992 1995, July 2002.

[3] H. Li and E. K. P. Chong, "Information lattices and subgroup lattices: Isomorphisms and approximations," in *Proceedings of the 45th Annual Allerton Conference on Communication, Control and Computing*, Monticello, Illinois, Sept. 26-28 2007.

[4] H. Li and E. K. P. Chong, "A group theoretic model for information," submitted to IEEE Transactions on Information Theory.

[5] H. Li and E. K. P. Chong, "On connections between group homomorphisms and the Ingleton inequality," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 24 29 2007, pp. 1996 2000.

[6] H. Li and E. K. P. Chong, "Search on lines and groups," submitted to Mathematics of Operations Research.

[7] P. R. Barbosa, H. Li, E. K. P. Chong, J. Hannig, and S. R. Kulkarni, "Zero-error target tracking through limited querying of binary sensors," in *Proceedings of the 44th Annual Allerton Conference on Communication, Control and Computing*, Monticello, Illinois, September 27 29 2006, pp. 1424 1431.

[8] H. Li, P. R. Barbosa, E. K. P. Chong, J. Hannig, and S. R. Kulkarni, "Zero-error target tracking with limited communication," submitted to IEEE Journal on Selected Areas in Communications, Control and Communications.

[9] M. Veeraraghavan, X. Zheng, W. C. Feng, H. Lee, E. K. P. Chong, and H. Li, "Scheduling and transport for file transfers on high-speed optical circuits," *Journal of Grid Computing, special issue on High Performance Networking*, vol. 1, no. 4, pp. 395 405, 2003.

[10] M. Veeraraghavan, X. Zheng, W. C. Feng, H. Lee, E. K. P. Chong, and H. Li, "Scheduling and transport for file transfers on high-speed optical circuits," in *Proceedings of the Second International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet 2004)*, Argonne National Laboratory, Argonne, Illinois, February 16 17 2004.

[11] M. Veeraraghavan, H. Lee, H. Li, and E. K. P. Chong, "Lambda scheduling algorithm for file transfers on high-speed optical circuits," in *Proceedings of the Workshop on Grids and Advanced Networks (GAN04), part of the IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2004)*, Chicago, Illinois, April 19–22 2004, pp. 617 624.

[12] M. Veeraraghavan, H. Lee, E. K. P. Chong, and H. Li, "A varying-bandwidth list scheduling heuristic for file transfers," in *Proceedings of the 2004 International Conference on Communications (ICC 2004)*, Paris, France, June 20 24 2004, pp. 1050 1054.

CHAPTER 2

A Group Theoretic Model for Information

2.1 Summary

In this work we formalize the notions of information elements and information lattices, first proposed by Shannon. Exploiting this formalization, we identify a comprehensive parallelism between information lattices and subgroup lattices. Qualitatively, we demonstrate isomorphisms between information lattices and subgroup lattices. Quantitatively, we establish a decisive approximation relation between the entropy structures of information lattices and the log-index structures of the corresponding subgroup lattices. This approximation extends the approximation for joint entropies carried out previously by Chan and Yeung. As a consequence of our approximation result, we show that any continuous law holds in general for the entropies of information elements if and only if the same law holds in general for the log-indices of subgroups. As an application, by constructing subgroup counterexamples we find surprisingly that common information, unlike joint information, obeys neither the submodularity nor the supermodularity law. We emphasize that the notion of information elements is conceptually significant—formalizing it helps to reveal the deep connection between information theory and group theory. The parallelism established in this work admits an appealing group-action explanation and provides useful insights into the intrinsic structure among information elements from a group-theoretic perspective.

2.2 Introduction

Information theory was born with the celebrated entropy formula measuring the *amount* of information for the purpose of communication. However, a suitable mathematical model for *information itself* remained elusive over the last

sixty years. It is reasonable to assume that information theorists have had certain intuitive conceptions of information, but in this work we seek a mathematic model for such a conception. In particular, building on Shannon's work [1], we formalize the notion of *information elements* to capture the syntactical essence of *information*, and identify information elements with σ-algebras and *sample-space-partitions*. As we shall see in the following, by building such a mathematical model for information and identifying the lattice structure among information elements, the seemingly surprising connection between information theory and group theory, established by Chan and Yeung [2], is revealed via isomorphism relations between information lattices and subgroup lattices. Consequently, a fully-fledged and decisive approximation relation between the entropy structure of information lattices and the subgroup-index structure of corresponding subgroup lattices is obtained.

We first motivate our formal definition for the notion of information elements.

2.2.1 Informationally Equivalent Random Variables

Recall the profound insight offered by Shannon [3] on the essence of communication: "the fundamental problem of communication is that of reproducing at one point exactly or approximately a message selected at another point." Consider the following motivating example. Suppose a message, *in English*, is delivered from person A to person B. Then, the message is translated and delivered *in German* by person B to person C (perhaps because person C does not know English). Assuming the translation is faithful, person C should receive the message that person A intends to convey. Reflecting upon this example, we see that the message (information) assumes two different "representations" over the process of the entire communication – one in English and the other in German, but the message (information) itself remains the same. Similarly, coders (decoders), essential components of communication systems, perform the similar function of "translating" one repre-

sentation of the same information to another one. This suggests that "information" itself should be defined in a translation invariant way. This "translation-invariant" quality is precisely how we seek to characterize information.

To introduce our formal definition for information elements to capture the essence of information itself, we note that information theory is built within the probabilistic framework, in which one-time information sources are usually modeled by random variables. Therefore, we start in the following with the concept of *informational equivalence* between random variables and develop the formal concept of information elements from first principles.

Recall that, given a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ and a measurable space $(\mathrm{S}, \mathcal{S})$, a random variable is a measurable function from Ω to S. The set S is usually called the state space of the random variable, and $\mathcal{S}$ is a σ-algebra on S. The set Ω is usually called the *sample space*; $\mathcal{F}$ is a σ-algebra on Ω, usually called the *event space*; and $\mathbf{P}$ denotes a probability measure on the measurable space $(\Omega, \mathcal{F})$.

To illustrate the idea of *informational equivalence*, consider a random variable $X : \Omega \to \mathrm{S}$ and another random variable $X' = f(X)$, where the function $f : \mathrm{S} \to \mathrm{S}'$ is bijective. Certainly, the two random variables X and X' are *technically different* for they have different codomains. However, it is intuitively clear that that they are "equivalent" in some sense. In particular, one can infer the exact state of X by observing that of X', and vice versa. For this reason, we may say that the two random variables X and X' carry the same piece of information. Note that the σ-algebras induced by X and X' coincide with each other. In fact, two random variables such that the state of one can be inferred from that of the other induce the same σ-algebra. This leads to the following definition for *information equivalence*.

Definition 2.2.1. *We say that two random variables X and X' are* informationally equivalent, *denoted $X \cong X'$, if the σ-algebras induced by X and X' coincide.*

It is easy to verify that the "being-informational-equivalent" relation is an equivalence relation. The definition reflects our intuition, as demonstrate in the previous motivating examples, that two random variables carry the same piece information if and only if they induce the same σ-algebra. This motivates the following definition for *information elements* to capture the syntactical essence of information itself.

Definition 2.2.2. *An* information element *is an equivalence class of random variables with respect to the "being-informationally-equivalent" relation.*

We call the random variables in the equivalent class of an information element m *representing random variables* of m. Or, we say that a random variable X *represents* m.

We believe that our definition of information elements reflects Shannon's original intention [1]:

> Thus we are led to define the actual information of a stochastic process as that which is common to all stochastic processes which may be obtained from the original by reversible encoding operations.

Intuitive (also informal) discussion on identifying "information" with σ-algebras surfaces often in probability theory, martingale theory, and mathematical finance. In probability theory, see for example [4], the concept of conditional probability is usually introduced with discussion of treating the σ-algebras conditioned on as the "partial information" available to "observers." In martingale theory and mathematical finance, see for example [5, 6], *filtrations* — increasing sequences of σ-algebras — are often interpreted as records of the information available over time.

A Few Observations

Proposition 2.2.3. *If* $X \cong X'$, *then* $H(X) = H(X')$.

(Throughout the chapter, we use $H(X)$ to denote the entropy of random variable X.)

The converse to Proposition 2.2.3 fails two random variables with a same entropy do not necessarily carry the same information. For example, consider two binary random variables $X, Y : \Omega \to \{0, 1\}$, where $\Omega = \{a, b, c, d\}$ and $\mathbf{P}$ is uniform on Ω. Suppose $X(\omega) = 0$ if $\omega = a, b$ and 1 otherwise, and $Y(\omega) = 0$ if $\omega = a, c$ and 1 otherwise. Clearly, we have $H(X) = H(Y) = 1$, but one can readily agree that X and Y do *not* carry the same information. Therefore, the notion of "informationally-equivalent" is stronger than that of "identically-distributed."

On the other hand, we see that the notion of "informationally-equivalent" is weaker than that of "being-equal."

Proposition 2.2.4. *If $X = X'$, then $X \cong X'$.*

The converse to Proposition 2.2.4 fails as well, since two informationally equivalent random variable X and X' may have totally different state spaces, so that it does not even make sense to say $X = X'$.

As shown in the following proposition, the notion of "informational equivalence" characterizes a kind of state space invariant "equalness."

Proposition 2.2.5. *Two random variables X and Y with state spaces $\mathcal{X}$ and $\mathcal{Y}$, respectively, are informationally equivalent if and only if there exists a one-to-one correspondence $f : \mathcal{X} \to \mathcal{Y}$ such that $Y = f(X)$.*

Remark: Throughout the chapter, we fix a probability space unless otherwise stated. For ease of presentation, we confine ourselves in the following to finite discrete random variables. However, most of the definitions and results can be applied to more general settings without significant difficulties.

2.2.2 Identifying Information Elements via σ-algebras and Sample-Space-Partitions

Since the σ-algebras induced by informationally equivalent random variables are the same, we can unambiguously identify information elements with σ-algebras. Moreover, because we deal with finite discrete random variables exclusively in this work, we can afford to discuss σ-algebras more explicitly as follows.

Recall that a *partition* Π of a set A is a collection $\{\pi_i : i \in [k]\}$ of disjoint subsets of A such that $\cup_{i\in[k]}\pi_i = A$. (Throughout the chapter, we use the bracket notation $[k]$ to denote the generic index set $\{1, 2, \cdots, k\}$.) The elements of a partition Π are usually called the *parts* of Π. It is well known that there is a natural one-to-one correspondence between partitions of the sample space and the σ-algebras — any given σ-algebra of a sample space can be generated uniquely, via union operation, from the atomic events of the σ-algebra, while the collection of the atomic events forms a partition of the sample space. For example, for a random variable $X : \Omega \to \mathcal{X}$, the atomic events of the σ-algebra induced by X are $X^{-1}(\{x\}), x \in \mathcal{X}$. For this reason, from now on, we shall identify an information element by either its σ-algebra or its corresponding sample space partition.

It is well known that the number of distinct partitions of a set of size n is the nth Bell number and that the Stirling number of the second kind $S(n, k)$ counts the number of ways to partition a set of n elements into k nonempty parts. These two numbers, crucial to the remarkable results obtained by Orlitsky et al. in [7], suggest a possibly interesting connection between the notion of information elements discussed in this work and the "patterns" studied in [7].

2.2.3 Shannon's Legacy

As we mentioned before, the notion of *information elements* was originally proposed by Shannon in [1]. In the same paper, Shannon also proposed a partial

order for information elements and a lattice structure for collections of information elements. We follow Shannon and call such lattices *information lattices* in the following.

Abstracting the notion of information elements out of their representations random variables is a conceptual leap, analogous to the leap from the concrete calculation with matrices to the study of abstract vector spaces. To this end, we formalize both the ideas of information elements and information lattices. By identifying information elements with sample-space-partitions, we are equipped to establish a comprehensive parallelism between information lattices and subgroup lattices. Qualitatively, we demonstrate isomorphisms between information lattices and certain subgroup lattices. With such isomorphisms established, quantitatively, we establish an approximation for the entropy structure of information lattices, consisting of joint, common, and many other information elements, using the log-index structures of their counterpart subgroup lattices. Our approximation subsumes the approximation carried out only for joint information elements by Chan and Yeung [2]. Building on [2], the parallelism identified in this work reveals an intimate connection between information theory and group theory and suggests that group theory may provide suitable mathematical language to describe and study laws of information.

The full-fledged parallelism between information lattices and subgroup lattices established in this work is one of our main contributions. With this intrinsic mathematical structure among multiple information elements being uncovered, we anticipate more systematic attacks on certain network information problems, where a better understanding of intricate internal structures among multiple information elements is in urgent need. Indeed, the ideas of information elements and information lattices were originally motivated by network communication problems in [1],

Shannon wrote:

> The present note outlines a new approach to information theory which is aimed specifically at the analysis of certain communication problems in which there exist a number of sources simultaneously in operation.

and

> Another more general problem is that of a communication system consisting of a large number of transmitting and receiving points with some type of interconnecting network between the various points. The problem here is to formulate the best system design whereby, in some sense, the best overall use of the available facilities is made.

It is not hard to see that Shannon was attempting to solve now-well-known network coding capacity problems.

Certainly, we do not claim that all the ideas in this work are our own. For example, as we pointed out previously, the notions of information elements and information lattices were proposed in the 1950s by Shannon [1]. However, this work of Shannon's is not widely known, perhaps owing to the abstruseness of the ideas. Formalizing these ideas and connecting them to current research is one of the primary goals of this work. For all other results and ideas that have been previously published, we separate them from those of our own by giving detailed references to their original sources. ,

2.2.4 Organization

The chapter is organized as follows. In Section 2.3, we introduce a "being-richer-than" partial order between information elements and study the information lattices induced by this partial order. In Section 2.4, we formally establish isomorphisms between information lattices and subgroup lattices. Section 2.5 is devoted to the quantitative aspects of information lattices. We show that the entropy structure of information lattices can be approximated by the log-index structure

of their corresponding subgroup lattices. As a consequence of this approximation result, in Section 2.6, we show that any continuous law holds for the entropies of common and joint information if and only if the same law holds for the log-indices of subgroups. As an application of this result, we show a result, which is rather surprising, that unlike joint information neither the submodularity nor the supermodularity law holds for common information in general. We conclude the chapter with a discussion in Section 2.7.

2.3 Information Lattices

2.3.1 "Being-richer-than" Partial Order

Recall that every information element can be identified with its corresponding sample-space-partition. Consider two sample-space-partitions Π and Π'. We say that Π is *finer than* Π', or Π' is *coarser than* Π, if each part of Π is contained in some part of Π'.

Definition 2.3.1. *For two information elements m_1 and m_2, we say that m_1 is* richer than *m_2, or m_2 is* poorer than *m_2, if the sample-space-partition of m_1 is finer than that of m_2. In this case, we write $m_1 \geq m_2$.*

It is easy to verify that the above defined "being-richer-than" relation is a partial order.

We have the following immediate observations:

Proposition 2.3.2. *$m_1 \geq m_2$ if and only if $H(m_2|m_1) = 0$.*

As a corollary to the above proposition, we have

Proposition 2.3.3. *If $m_1 \geq m_2$, then $H(m_1) \geq H(m_2)$.*

The converse of Proposition 2.3.3 does not hold in general.

With respect to representative random variables of information elements, we have

Proposition 2.3.4. *Suppose random variables X_1 and X_2 represent information elements m_1 and m_2 respectively. Then, $m_1 \geq m_2$ if and only if $X_2 = f(X_1)$ for some function f.*

A similar result to Proposition 2.3.4 was previously observed by Renyi [8] as well.

The "being-richer-than" relation is very important to information theory, because it characterizes a universal information-theoretic constraint put on all deterministic coders (decoders) – the input information element of any coder is always richer than the output information element. For example, partially via this principle, Yan et al. recently characterized the capacity region of general acyclic multi-source, multi-sink networks [9]. Harvey et al. [10] obtained an improved computable outer bound for general network coding capacity regions by applying this same principle under a different name called *information dominance* – the authors of the paper stated: "...information dominance plays a key role in our investigation of network capacity."

2.3.2 Information Lattices

Recall that a lattice is a set endowed with a partial order in which any two elements have a unique supremum and a unique infimum with respect to the partial order. Conventionally, the supremum of two lattice elements x and y is also called the *join* of x and y; the infimum is also called the *meet*. In our case, with respect to the "being-richer-than" partial order, the supremum of two information elements m_1 and m_2, denoted $m_1 \vee m_2$, is the poorest among all the information elements that are richer than both m_1 and m_2. Conversely, the infimum of m_1 and m_2, denoted $m_1 \wedge m_2$, is the richest among all the information elements that are poorer than both m_1 and m_2. In the following, we also use m^{12} to denote the join of m_1 and m_2, and m_{12} the meet.

Definition 2.3.5. *An* information lattice *is a set of information elements that is closed under the join* $\vee$ *and meet* $\wedge$ *operations.*

Recall the one-to-one correspondence between information elements and sample-space-partitions. Consequently, each information lattice corresponds to a partition lattice (with respect to the "being-finer-than" partial order on partitions), and vice versa. This formally confirms the assertions made in [1]: "they (information lattices) are at least as general as the class of finite partition lattices."

Since the collection of information lattices could be as general as that of partition lattices, we should not expect any special lattice properties to hold generally for all information lattices, because it is well-known that any finite lattice can be embedded in a finite partition lattice [11]. Therefore, it is not surprising to learn that information lattices are in general not distributive, not even modular.

2.3.3 Joint Information Element

The *join* of two information elements is straightforward. Consider two information elements m_1 and m_2 represented respectively by two random variables X_1 and X_2. It is easy to check that the joint random variable (X_1, X_2) represents the join m^{12}. For this reason, we also call m^{12} (or $m_1 \vee m_2$) the *joint information element* of m_1 and m_2. It is worth pointing out that the joint random variable (X_2, X_1) represents m^{12} equally well.

2.3.4 Common Information Element

In [1], the meet of two information elements is called *common information.* More than twenties years later, the same notion of common information was independently proposed and first studied in detail by Gács and Körner [12]. For the first time, it was demonstrated that common information could be far less than mutual information. ("Mutual information" is rather a misnomer because it

does not correspond naturally to any information element [12].) Unlike the case of joint information elements, characterizing common information element via their representing random variables is much more complicated. See [12, 13] for details.

In contrast to the all-familiar joint information, common information receives far less attention. Nonetheless, it has been shown to be important to cryptography [14 17], indispensable for characterizing of the capacity region of multi-access channels with correlated sources [18], useful in studying information inequalities [19, 20], and relevant to network coding problems [21].

2.3.5 Previously Studied Lattices in Information Theory

Historically, at least three other lattices [22 24] have been considered in attempts to characterize certain ordering relations between information elements. Two of them, studied respectively in [22] and [24], are subsumed by the information lattices considered in this work.

2.4 Isomorphisms between Information Lattices and Subgroup Lattices

In this section, we discuss the qualitative aspects of the parallelism between information lattices generated from sets of information elements and subgroup lattices generated from sets of subgroups. In particularly, we establish isomorphism relations between them.

2.4.1 Information Lattices Generated by Information Element Sets

It is easy to verify that both the binary operations "$\vee$" and "$\wedge$" are *associative* and *commutative*. Thus, we can readily extend them to cases of more than two information elements. Accordingly, for a given set $\{m_i : i \in [n]\}$ of information elements, we denote the joint information element of the subset $\{m_i : i \in \alpha\}$, $\alpha \subseteq [n]$, of information elements by m^α and the common information element by m_α.

Definition 2.4.1. *Given a set* $\mathbf{M} = \{m_i : i \in [n]\}$ *of information elements. the* information lattice generated by $\mathbf{M}$, *denoted* $\mathrm{L}_{\mathbf{M}}$, *is the smallest information lattice that contains* $\mathbf{M}$. *We call* $\mathbf{M}$ *a* generating set *of the lattice* $\mathrm{L}_{\mathbf{M}}$.

It is easy to see that each information element in $\mathrm{L}_{\mathbf{M}}$ can be obtained from the information elements in the generating set $\mathbf{M}$ via a sequence of join and meet operations. Note that the set $\{m_\alpha : \alpha \subseteq [n]\}$ of information elements forms a meet semi-lattice and the set $\{m^\beta : \beta \subseteq [n]\}$ forms a join semi-lattice. However, the union $\{m_\alpha, m^\beta : \alpha, \beta \subseteq [n]\}$ of these two semi-lattices does *not* necessarily form a lattice. To see this, consider the following example constructed with partitions (since partitions are in one-to-one correspondence with information elements). Let $\{\pi_i : i = [4]\}$ be a collection of partitions on the set $\{1, 2, 3, 4\}$ where $\pi_1 = 12|3|4$, $\pi_2 = 14|2|3$, $\pi_3 = 23|1|4$, and $\pi_4 = 34|1|2$. See Figure 1 for the Hasse diagram of the lattice generated by the collection $\{\pi_i : i = [4]\}$. It is easy to see $(\pi_1 \vee \pi_2) \wedge (\pi_3 \vee \pi_4) = 124|3 \wedge 234|1 = 24|1|3$, but $24|1|3 \notin \{\pi_\alpha, \pi^\beta : \alpha, \beta \in [4]\}$. Similarly, we have $(\pi_1 \vee \pi_3) \wedge (\pi_2 \vee \pi_4) = 13|2|4 \notin \{\pi_\alpha, \pi^\beta : \alpha, \beta \in [4]\}$.

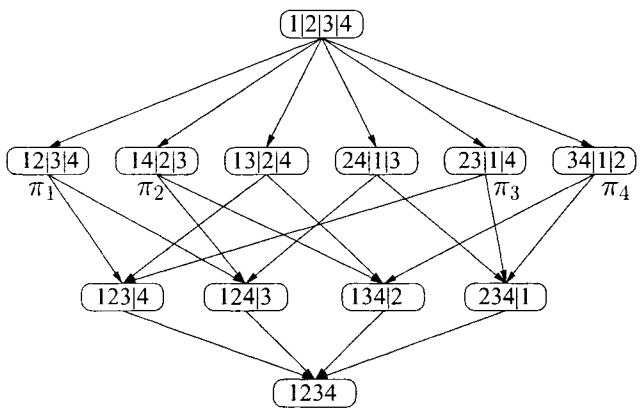

Figure 1. Lattice generated by $\{\pi_i : i = [4]\}$

2.4.2 Subgroup Lattices

Consider the binary operations on subgroups intersection and union. We know that the intersection $G_1 \cap G_2$ of two subgroups is again a subgroup. However, the union $G_1 \cup G_2$ does *not* necessarily form a subgroup. Therefore, we consider the subgroup *generated* from the union $G_1 \cup G_2$, denoted G^{12} (or $G_1 \vee G_2$). Similar to the case of information elements, the intersection and "$\vee$" operations on subgroups are both associative and commutative. Therefore, we readily extend the two operations to the cases with more than two subgroups and, accordingly, denote the intersection $\cap_{i\in[n]}G_i$ of a set of subgroups $\{G_i : i \in [n]\}$ by $G_{[n]}$ and the subgroup generated from the union by $G^{[n]}$. It is easy to verify that the subgroups $G_{[n]}$ and $G^{[n]}$ are the infimum and the supremum of the set $\{G_i : i \in [n]\}$ with respect to the "being-a-subgroup-of" partial order. For notation consistency, we also use "$\wedge$" to denote the intersection operation.

Note that, to keep the notation simple, we "overload" the symbols "$\vee$" and "$\wedge$" for both the join and the meet operations with information elements and the intersection and the "union-generating" operations with subgroups. Their actual meaning should be clear within context.

Definition 2.4.2. *A* subgroup lattice *is a set of subgroups that is closed under the* $\wedge$ *and* $\vee$ *operations.*

For example, the set of all the subgroups of a group forms a lattice.

Similar to the case of information lattices generated by sets of information elements, we consider in the following *subgroup lattices generated by a set of subgroups.*

Definition 2.4.3. *Given a set* $\mathbf{G} = \{G_i : i \in [n]\}$ *of subgroups, the* subgroup lattice generated by $\mathbf{G}$, *denoted* $\mathrm{L}_{\mathbf{G}}$, *is the smallest lattices that contains* $\mathbf{G}$. *We call* $\mathbf{G}$ *a* generating set *of* $\mathrm{L}_{\mathbf{G}}$.

Note that the set $\{G_\alpha : \alpha \subseteq [n]\}$ forms a semilattice under the meet $\wedge$ operation and the set $\{G^\beta : \beta \subseteq [n]\}$ forms a semilattice under the join $\vee$ operation. However, as in the case of information lattices, the union $\{G_\alpha, G^\beta : \alpha, \beta \subseteq [n]\}$ of the two semilattices does *not* necessarily form a lattice.

In the remainder of this section, we relate information lattices generated by sets of information elements and subgroup lattices generated by collections of subgroups and demonstrate isomorphism relations between them. For ease of presentation, as a special case we first introduce an isomorphism between information lattices generated by sets of coset-partition information elements and their corresponding subgroup lattices.

2.4.3 Special Isomorphism Theorem

We endow the sample space with a *group* structure — the sample space in question is taken to be a group G. For any subgroup of G, by Lagarange's theorem [25], the collection of its cosets forms a partition of G. Certainly, the coset-partition, as a sample-space-partition, uniquely defines an information element. A collection $\mathbf{G} = \{G_i : i \in [n]\}$ of subgroups of G, in the same spirit, identifies a set $\mathbf{M} = \{m_i : i \in [n]\}$ of information elements via this subgroup coset-partition correspondence.

Remark: throughout the chapter, groups are taken to be multiplicative, and cosets are taken to be right cosets.

It is clear that, by our construction, the information elements in $\mathbf{M}$ and the subgroups in $\mathbf{G}$ are in one-to-one correspondence via the subgroup coset-partition relation. It turns out that the information elements on the entire information lattice $\mathrm{L}_\mathbf{M}$ and the subgroups on the subgroup lattice $\mathrm{L}_\mathbf{G}$ are in one-to-one correspondence as well via the same subgroup coset-partition relation. In other words, both the join and meet operations on information lattices are faithfully "mirrored"

by the join and meet operations on subgroup lattices.

Theorem 2.4.4. *(Special Isomorphism Theorem) Given a set* $\mathbf{G} = \{G_i : i \in [n]\}$ *of subgroups, the subgroup lattice* $\mathrm{L}_{\mathbf{G}}$ *is isomorphic to the information lattice* $\mathrm{L}_{\mathbf{M}}$ *generated by the set* $\mathbf{M} = \{m_i : i \in [n]\}$ *of information elements, where* m_i, $i \in [n]$, *are accordingly identified via the coset-partitions of the subgroups* G_i, $i \in [n]$.

The theorem is shown by demonstrating a mapping, from the subgroup lattice $\mathrm{L}_{\mathbf{G}}$ to the information lattice $\mathrm{L}_{\mathbf{M}}$, such that it is a lattice-morphism, i.e., it honors both join and meet operations, and is bijective as well. Naturally, the mapping $\phi : \mathrm{L}_{\mathbf{G}} \rightarrow \mathrm{L}_{\mathbf{M}}$ assigning to each subgroup $G_i \in \mathrm{L}_{\mathbf{G}}$ the information element identified by the coset-partition of the subgroup G_i is such a morphism. Since this theorem and its general version, Theorem 2.4.7, are crucial to our later results—Theorems 2.5.4 and 2.6.17—and certain aspects of the reasoning are novel, we include a detailed proof for it in Appendix 2.8.1.

2.4.4 General Isomorphism Theorem

The information lattices considered in Section 2.4.3 are rather limited—by Lagrange's theorem, coset-partitions are all equal partitions. In this subsection, we consider arbitrary information lattices—we do not require the sample space to be a group. Instead, we treat a general sample-space-partition as an *orbit-partition* resulting from some group-action on the sample space.

Group-Actions and Permutation Groups

Definition 2.4.5. *Given a group* G *and a set* A, *a group-action of* G *on* A *is a function* $(g, a) \mapsto g(a)$, $g \in G$, $a \in A$, *that satisfies the following two conditions:*

- $(g_1 g_2)(a) = \big(g_1(g_2(a))\big)$ *for all* $g_1, g_2 \in G$ *and* $a \in A$;
- $e(a) = a$ *for all* $a \in A$, *where* e *is the identity of* G.

We write (G, A) to denote the group-action.

Now, we turn to the notions of *orbits* and *orbit-partitions*. We shall see that every group-action (G, A) induces unambiguously an equivalence relation as follows. We say that x_1 and x_2 are *connected* under a group-action (G, A) if there exists a $g \in G$ such that $x_2 = g(x_1)$. We write $x_1 \overset{G}{\sim} x_2$. It is easy to check that this "being-connected" relation $\overset{G}{\sim}$ is an equivalence relation on A. By the fundamental theorem of equivalence relations, it defines a partition on A.

Definition 2.4.6. *Given a group-action (G, A), we call the equivalence classes with respect to the equivalence relation $\overset{G}{\sim}$, or the parts of the induced partition of A, the* orbits *of the group-action. Accordingly, we call the induced partition the* orbit-partition *of (G, A) .*

Sample-Space-Partition as Orbit-Partition

In fact, starting with a partition Π of a set A, we can go in the other direction and unambiguously define a group action (G, A) such that the orbit-partition of (G, A) is exactly the given partition Π. To see this, note the following salient feature of group-actions: For any given group-action (G, A), associated with every element g in the group is a mapping from A to itself and any such mappings must be bijective. This feature is the direct consequence of the group axioms. To see this, note that every group element g has a unique inverse g^{-1}. According to the first defining property of group-actions, we have $(gg^{-1})(x) = g\left(g^{-1}(x)\right) = e(x) = x$ for all $x \in A$. This requires that the mappings associated with g and g^{-1} to be invertible. Clearly, the identity e of the group corresponds to the identity map from A to A.

With the observation that under group-action (G, A) every group element corresponds to a permutation of A, we can treat every group as a collection of permutations that is closed under permutation composition. Specifically, for a

given partition Π of a set A, it is easy to check that all the permutations of A that permute the elements of the parts of Π *only* to the elements of the same parts form a group. These permutations altogether form the so-called *permutation representation* of G (with respect to A). For this reason in the following, without loss of generality, we treat all groups as permutation groups. We denote by G_Π the permutation group corresponding as above to a partition Π—G_Π acts naturally on the set A by permutation, and the orbit partition of (G_Π, A) is exactly Π.

From group theory, we know that this orbit-partition permutation-group-action relation is a one-to-one correspondence. Since every information element corresponds definitively to a sample-space-partition, we can identify every information element by a permutation group. Given a set $\mathbf{M} = \{m_i : i \in [n]\}$ of information elements, denote the set of the corresponding permutation groups by $\mathbf{G} = \{G_i : i \in [n]\}$. Note that all the permutations in the permutation groups G_i, $i \in [n]$, are permutations of the same set, namely the sample space. Hence, all the permutation groups G_i, $i \in [n]$, are subgroups of the symmetric group $S_{|\Omega|}$, which has order $2^{|\Omega|}$. Therefore, it makes sense to take intersection and union of groups from the collection $\mathbf{G}$.

From Coset-Partition to Orbit-Partition—From Equal Partition to General Partition

In fact, the previously studied coset-partitions are a special kind of orbit-partitions. They are orbit-partitions of group-actions defined by the native group multiplication. Specifically, given a subgroup G_1 of G, a group-action (G_1, G) is defined such that $g_1(a) = g_1 \circ a$ for all $g_1 \in G_1$ and $a \in G$, where "$\circ$" denotes the native binary operation of the group G. The orbit-partition of such a group-action is exactly the coset-partition of the subgroup G_1. Therefore, by taking a different kind of group-action — permutation rather than group multiplication — we are freed

from the "equal-partition" restriction so that we can correspond arbitrary information elements identified with arbitrary sample-space-partitions to subgroups. It turns out information lattices generated by sets of information elements and subgroup lattices generated by the corresponding sets of permutation groups remain isomorphic to each other. Thus, the isomorphism relation between information lattices and subgroup lattices holds in full generality.

Isomorphism Relation Remains Between Information Lattices and Subgroup Lattices

Similar to Section 2.4.3, we consider a set $\mathbf{M} = \{m_i, i \in [n]\}$ of information element. Unlike in Section 2.4.3, the information elements m_i, $i \in [n]$ considered here are arbitrary. As we discussed in the above, with each information element m_i we associate a permutation group G_i according to the orbit-partition permutation-group-action correspondence. Denote the set of corresponding permutation groups by $\mathbf{G} = \{G_i, i \in [n]\}$.

Theorem 2.4.7. *(General Isomorphism Theorem) The information lattice* $\mathrm{L}_{\mathbf{M}}$ *is isomorphic to the subgroup lattice* $\mathrm{L}_{\mathbf{G}}$.

The arguments for Theorem 2.4.7 are similar to those for Theorem 2.4.4 we demonstrate that the orbit-partition permutation-group-action correspondence is a lattice isomorphism between $\mathrm{L}_{\mathbf{M}}$ and $\mathrm{L}_{\mathbf{G}}$.

2.5 An Approximation Theorem

From this section on, we shift our focus to the quantitative aspects of the parallelism between information lattices and subgroup lattices. In the previous section, by generalizing from coset-partitions to orbit-partitions, we successfully established an isomorphism between general information lattices and subgroup lattices. In this section, we shall see that not only is the qualitative structure pre-

served, but also the quantitative structure — the entropy structure of information lattices — is essentially captured by their isomorphic subgroup lattices.

2.5.1 Entropies of Coset-partition Information Elements

We start with a simple and straightforward observation for the entropies of coset-partition information elements on information lattices.

Proposition 2.5.1. *Let $\{G_i : i \in [n]\}$ be a set of subgroups of group G and $\{m_i : i \in [n]\}$ be the set of corresponding coset-partition information elements. The entropies of the joint and common information elements on the information lattice, generated from $\{m_i : i \in [n]\}$, can be calculated from the subgroup-lattice, generated from $\{G_i : i \in [n]\}$, as follows*

$$h(m^{[n]}) = \log \frac{|G|}{|\wedge_{i\in[n]} G_i|}. \tag{1}$$

and

$$h(m_{[n]}) = \log \frac{|G|}{|\vee_{i\in[n]} G_i|}. \tag{2}$$

Proposition 2.5.1 follows easily from the isomorphism relation established by Theorem 2.4.7.

Note that the right hand sides of both Equation (1) and (2) are the logarithms of the indices of subgroups. In the following, we shall call them, in short, *log-indices*.

Proposition 2.5.1 establishes a quantitative relation between the entropies of the information elements on coset-partition information lattices and the log-indices of the subgroups on the isomorphic subgroup lattices. This quantitative relation is *exact*. However, the scope of Proposition 2.5.1 is rather restrictive — it applies only to certain special kind of "uniform" information elements, because, by Lagrange's theorem, all coset-partitions are equal partitions.

In Section 2.4, by generalizing from coset-partitions to orbit-partitions we successfully removed the "uniformness" restriction imposed by the coset-partition structure. At the same time, we established a new isomorphism relation, namely orbit-partition permutation-group-action correspondence, between information lattices and subgroup lattices. It turns out that this generalization maintains an "rough" version of the quantitative relation established in Proposition 2.5.1 between the entropies of information lattices and the log-indices of their isomorphic permutation-subgroup lattices. As we shall see in the next section, the entropies of the information elements on information lattices can be approximated, up to arbitrary precision, by the log-indices of the permutation groups on their isomorphic subgroup lattices.

2.5.2 Subgroup Approximation Theorem

To discuss the approximation formally, we introduce two definitions as follows.

Definition 2.5.2. *Given an information lattice* $\mathrm{L}_{\mathbf{M}}$ *generated from a set* $\mathbf{M} = \{m_i, i \in [n]\}$ *of information elements, we call the real vector*

$$\big(H(m) : m \in \mathrm{L}_{\mathbf{M}}\big),$$

whose components are the entropies of the information elements on the information lattice $\mathrm{L}_{\mathbf{M}}$ *generated by* $\mathbf{M}$*, listed according to a certain prescribed order, the* entropy vector of $\mathrm{L}_{\mathbf{M}}$*, denoted* $h(\mathrm{L}_{\mathbf{M}})$.

The entropy vector $h(\mathrm{L}_{\mathbf{M}})$ captures the informational structure among the information elements of $\mathbf{M}$.

Definition 2.5.3. *Given a subgroup lattice* $\mathrm{L}_{\mathbf{G}}$ *generated from a set* $\mathbf{G} = \{G_i, i \in [n]\}$ *of subgroups of a group* G*, we call the real vector*

$$\frac{1}{|G|}\left(\log\frac{|G|}{|G'|} : G' \in \mathrm{L}_{\mathbf{G}}\right),$$

whose components are the normalized log-indices of the subgroups on the subgroup lattice $\mathrm{L}_{\mathbf{G}}$ *generated by* $\mathbf{G}$*, listed according to a certain prescribed order, the* normalized log-index vector of $\mathrm{L}_{\mathbf{G}}$*, denoted* $l(\mathrm{L}_{\mathbf{G}})$.

In the following, we assume that $l(\mathrm{L}_{\mathbf{G}})$ and $h(\mathrm{L}_{\mathbf{M}})$ are accordingly aligned.

Theorem 2.5.4. *Let* $\mathbf{M} = \{m_i, i \in [n]\}$ *be a set of information elements. For any* $\epsilon > 0$ *there exists an* $N > 0$ *and a set* $\mathbf{G}^N = \{G_i : i \in [n]\}$ *of subgroups of the symmetry group* S_N *of order* 2^N *such that*

$$\|h(\mathrm{L}_{\mathbf{M}}) - l(\mathrm{L}_{\mathbf{G}^N})\| < \epsilon. \tag{3}$$

where "$\|\cdot\|$*" denotes the norm of real vectors.*

Theorem 2.5.4 subsumes the approximation carried out by Chan and Yeung in [2], which is limited to joint entropies. The approximation procedure we carried out to prove Theorem 2.5.4 is similar to that of Chan and Yeung [2] both use Stirling's approximation formula for factorials. But, with the group-action relation between information elements and permutation groups being exposed, and the isomorphism between information lattices and subgroup lattices being revealed, the approximation procedure becomes transparent and the seemingly surprising connection between information theory and group theory becomes mathematically natural. For these reasons, we included a detailed proof in Appendix 2.8.2.

2.6 Parallelism between Continuous Laws of Information Elements and those of Subgroups

As a consequence of Theorem 2.5.4, we shall see in the following that if a continuous law holds in general for information elements, then the same law must hold for the log-indices of subgroups, and vice versa.

In the following, for reference and comparison purposes, we first review the known laws concerning the entropies of joint and common information elements.

These laws, usually expressed in the form of *information inequalities*, are deemed to be fundamental to information theory [26].

2.6.1 Laws for Information Elements

Non-Negativity of Entropy

Proposition 2.6.1. *For any information element m, we have $H(m) \geq 0$.*

Laws for Joint Information

Proposition 2.6.2. *Given a set $\{m_i, i \in [n]\}$ of information elements, if $\alpha \subseteq \beta$, $\alpha, \beta \subseteq [n]$, then $H(m^\alpha) \leq H(m^\beta)$.*

Proposition 2.6.3. *For any two sets of information elements $\{m_i : i \in \alpha\}$ and $\{m_j : j \in \beta\}$, the following inequality holds:*

$$H(m^\alpha) + H(m^\beta) \geq H(m^{\alpha \cup \beta}) + H(m^{\alpha \cap \beta}).$$

This proposition is mathematically equivalent to the following one.

Proposition 2.6.4. *For any three information elements m_1, m_2, and m_3, the following inequality holds:*

$$H(m^{12}) + H(m^{23}) \geq H(m^{123}) + H(m^3).$$

Note that $H(m^3) = H(m_3)$.

Proposition 2.6.3 (or equivalently 2.6.4) is usually called the submodularity law for entropy function. Proposition 2.6.1, 2.6.2, and 2.6.3 are known, collectively, as the *polymatroidal axioms* [27, 28]. Up until very recently, these are the only known laws for entropies of joint information elements.

In 1998, Zhang and Yeung discovered a new information inequality, involving four information elements [28].

Proposition 2.6.5. *(Zhang-Yeung Inequality) For any four information elements* m_i, $i = 1, 2, 3,$ *and* 4, *the following inequality holds:*

$$\begin{aligned} & 3H(m^{13}) + 3H(m^{14}) + H(m^{23}) + H(m^{24}) + 3H(m^{34}) \\ & \geq H(m^1) + 2H(m^3) + 2H(m^4) \\ & \quad + H(m^{12}) + 4H(m^{134}) + H(m^{234}). \end{aligned} \tag{4}$$

This newly discovered inequality, classified as a *non-Shannon type information inequality* [26], proved that our understanding on laws governing the quantitative relations between information elements is incomplete. Recently, six more new four-variable information inequalities were discovered by Dougherty et al. [29].

Information inequalities such as those presented above were called "laws of information" [26, 30]. Seeking new information inequalities is currently an active research topic [19, 28, 31, 32]. In fact, they should be more accurately called "laws of joint information", since these inequalities involves only joint information only. We shall see below laws involving common information.

Common Information v.s. Mutual Information

In contrast to joint information, little research has been done to laws involving common information. So far, the only known non-trivial law involving both joint information and common information is stated in the following proposition, discovered by Gács and Körner [12].

Proposition 2.6.6. *For any two information element* m_1 *and* m_2, *the following inequality holds:*

$$H(m_{12}) \leq I(m_1; m_2) = H(m^1) + H(m^2) - H(m^{12}).$$

Note that $m^1 = m_1$ and $m^2 = m_2$.

Laws for Common Information

Dual to the non-decreasing property of joint information, it is immediately clear that entropies of common information are non-increasing.

Proposition 2.6.7. *Given a set* $\{m_i, i \in [n]\}$ *of information elements, if* $\alpha \subseteq \beta$ $\alpha, \beta \subseteq [n]$, *then* $H(m_\alpha) \geq H(m_\beta)$.

Comparing to the case of joint information, one may naturally expect, as a dual counterpart of the submodularity law of joint information, a supermodularity law to hold for common information. In other words, we have the following conjecture.

Conjecture 2.6.8. *For any three information elements* m_1, m_2, *and* m_3, *the following inequality holds:*

$$H(m_{12}) + H(m_{23}) \leq H(m_{123}) + H(m_2). \tag{5}$$

We see this conjecture as natural because of the intrinsic duality between the join and meet operations of information lattices. Due to the combinatorial nature of common information [12], it is not obvious whether the conjecture holds. With the help of our approximation results established in Theorem 2.5.4 and 2.6.17, we find, surprisingly, that neither the conjecture nor its converse holds. In other words, common information observes neither the submodularity nor the supermodularity law.

2.6.2 Continuous Laws for Joint and Common Information

As a consequence of Theorem 2.5.4, we shall see in the following that if a continuous law holds for information elements, then the same law must hold for the log-indices of subgroups, and vice versa. To convey this idea, we first present the simpler case involving only joint and common information elements. To state our result formally, we first introduce two definitions.

Definition 2.6.9. *Given a set* $\mathbf{M} = \{m_i : i \in [n]\}$ *of information elements, consider the collection* $\mathcal{M} = \{m_\alpha, m^\beta : \alpha, \beta \subseteq [n]\}$ *of join and meet information elements generated from* $\mathbf{M}$. *We call the real vector*

$$\Big(H(m^\alpha). H(m_\beta) : \alpha, \beta \subseteq [n], \alpha, \beta \neq \Phi\Big),$$

whose components are the entropies of the information elements of $\mathcal{M}$, *the* entropy vector of $\mathcal{M}$, *denoted by* $h_{\mathcal{M}}$.

Definition 2.6.10. *Given a set* $\mathbf{G} = \{G_i : i \in [n]\}$ *of subgroups of a group* G, *consider the set* $\mathcal{G} = \{G_\alpha, G^\beta : \alpha, \beta \subseteq [n]\}$ *of the subgroups generated from* $\mathbf{G}$. *We call the real vector*

$$\frac{1}{|G|}\Big(\log \frac{|G|}{|G^\alpha|}. \log \frac{|G|}{|G_\beta|} : \alpha, \beta \subseteq [n], \alpha, \beta \neq \Phi\Big),$$

whose components are the normalized log-indices of the subgroups in $\mathcal{M}$, *the* normalized log-index vector of $\mathcal{G}$, *denoted by* $l_{\mathcal{G}}$.

In this context, we assume that the components of both $l_{\mathcal{G}}$ and $h_{\mathcal{M}}$ are listed according to a common fixed order. Moreover, we note that both the vectors $h_{\mathcal{M}}$ and $l_{\mathcal{G}}$ have dimension $2^{n+1} - n - 2$.

Theorem 2.6.11. *Let* $f : \mathbb{R}^{2^{n+1}-n-2} \to \mathbb{R}$ *be a continuous function. Then,* $f(h_{\mathcal{M}}) \geq 0$ *holds for all sets* $\mathbf{M}$ *of* n *information elements if and only if* $f(l_{\mathcal{G}}) \geq 0$ *holds for all sets* $\mathbf{G}$ *of* n *subgroups of any group.*

Theorem 2.6.11 is a special case of Theorem 2.6.17.

Theorem 2.6.11 and its generalization Theorem 2.6.17 extend the result obtained by Chan and Yeung in [2] in the following two ways. First, Theorem 2.6.11 and 2.6.17 apply to all continuous laws, while only linear laws were considered in [2]. Even though so far we have not yet encountered any nonlinear law for

entropies, it is highly plausible that nonlinear information laws may exist given the recent discovery that at least certain part of the boundary of the entropy cones involving at least four information elements are curved [33]. Second, our theorems encompass both common information and joint information, while only joint entropies were considered in [2]. For example, laws such as Propositions 2.6.6 and 2.6.7 cannot even be expressed in the setting of [2]. In fact, as we shall see later in Section 2.6.4, the laws of common information depart from those of joint information very early —unlike joint information, which obeys the submodularity law, common information admits neither submodularity nor supermodularity. For these reasons, we believe that our extending the subgroup approximation to common information is of interest in its own right.

2.6.3 Continuous Laws for General Lattice Information Elements

In this section, we extend Theorem 2.6.11 to all the information elements in information lattices, not limited to the "pure" joint and common information elements. In the following, we introduce some necessary machinery to formally present the result in full generality.

Note that an element from the lattice generated from a set X has its expression built from the generating elements of the lattice in the similar way that *terms* are built from *literals* in mathematical logic. In particular, we define *lattice-terms* as follows:

Definition 2.6.12. *An expression E is called a* lattice-term *formed from a set* X *of literals if either E is a literal from* X *or E is formed from two lattice-terms with either the* join *or the* meet *symbols: $E = x \,\mathrm{OP}\, y$, where x and y are lattice-terms and* OP *is either the join symbol $\vee$ or the meet symbol $\wedge$.*

Definition 2.6.13. *Suppose that E_i, $i \in [k]$, are lattice-terms generated from a*

literal set of size n: $\mathrm{X} = \{x_1, \cdots, x_n\}$. *We call an expression of the form*

$$f\big(H(E_1), \cdots, H(E_k)\big),$$

where f represents a function from $\mathbb{R}^k$ *to* $\mathbb{R}$ *and H represents the entropy function, an n-variable* generalized information expression.

We *evaluate* an n-variable generalized information expression $f\big(H(E_1), \cdots, H(E_k)\big)$ against a set $\mathrm{M} = \{m_i : i \in [n]\}$ of information elements by substituting x_i with m_i respectively, calculating the entropy of the information elements obtained by evaluating the lattice-terms E_i according to the semantics of the join and meet operations on information elements, and then obtaining the corresponding function value. We denote this value by

$$f\big(H(E_1), \cdots, H(E_k)\big)\big|_{\mathrm{M}}.$$

Definition 2.6.14. *If an n-variable generalized information expression* $f\big(H(E_1), \cdots, H(E_k)\big)$ *is evaluated non-negatively for any set of n information elements, i.e.,*

$$f\big(H(E_1), \cdots, H(E_k)\big)\big|_{\mathrm{M}} \geq 0, \text{ for all } \mathrm{M},$$

then we call

$$f\big(H(E_1), \cdots, H(E_k)\big) \geq 0$$

an n-variable information law.

Similar to generalized information expressions, we define *generalized log-index expression* as follows.

Definition 2.6.15. *we call an expression of the form*

$$f\big(L(E_1), \cdots, L(E_k)\big),$$

where f represents a function from $\mathbb{R}^k$ *to* $\mathbb{R}$ *and L represents the normalized log-index function of subgroups, an n-variable* generalized log-index expression.

We *evaluate* an n-variable generalized log-index expression $f\big(L(E_1),\cdots,L(E_k)\big)$ against a set $\mathbf{G}=\{G_i : i\in[n]\}$ of subgroups of a group G by substituting x_i with G_i respectively, calculating the log-index of the subgroups obtained by evaluating the lattice-terms E_i according to the semantics of the join and meet operations on subgroups, and then obtaining the corresponding function value. We denote this value by

$$f\big(L(E_1),\cdots,L(E_k)\big)\big|_{\mathbf{G}}.$$

Definition 2.6.16. *If an n-variable generalized log-index expression $f\big(H(E_1),\cdots,H(E_k)\big)$ is evaluated non-negatively for any set of n subgroups of any group, i.e.,*

$$f\big(L(E_1),\cdots,L(E_k)\big)\big|_{\mathbf{G}}\geq 0,\ \textit{for all}\ \mathbf{G},$$

then we call

$$f\big(L(E_1),\cdots,L(E_k)\big)\geq 0$$

an n-variable subgroup log-index law.

With the above formalism and corresponding notations, we are ready to state our equivalence result concerning the generalized information laws.

Theorem 2.6.17. *Suppose that f is continuous. Then an n-variable information law*

$$f\big(H(E_1),\cdots,H(E_k)\big)\geq 0$$

holds if and only if the corresponding n-variable subgroup log-index law

$$f\big(L(E_1),\cdots,L(E_k)\big)\geq 0$$

holds.

Proof. To see one direction, namely that $f\big(L(E_1),\cdots,L(E_k)\big)\geq 0$ implies that $f\big(H(E_1),\cdots,H(E_k)\big)\geq 0$, assume that there exists a set $\mathbf{M}$ of information elements such that $f\big(H(E_1),\cdots,H(E_k)\big)\big|_{\mathbf{M}} = a$ for some $a<0$. By the continuity of the function f and Theorem 2.5.4, we are guaranteed to be able to construct, from the information lattice generated from $\mathbf{M}$, some subgroup lattice $\mathrm{L}_{\mathbf{G}}$ such that the value of the function f at the normalized log-indices of the correspondingly constructed subgroups is arbitrarily close to $a<0$. This contradicts the assumption that $f\big(L(E_1),\cdots,L(E_k)\big)\big|_{\mathbf{G}}\geq 0$ holds for all sets $\mathbf{G}$ of n subgroups of any group.

On the other hand, for any normalized log-indices of the subgroups from subgroup lattices, it can be readily interpreted as the entropies of information elements by taking permutation representation for the subgroups on the subgroup lattice and then producing an information lattice, according to the orbit-partition–permutation-group-action correspondence. Therefore, that $f\big(H(E_1),\cdots,H(E_k)\big)\big|_{\mathbf{M}}\geq 0$ holds for all sets $\mathbf{M}$ implies that $f\big(L(E_1),\cdots,L(E_k)\big)\big|_{\mathbf{G}}\geq 0$ holds for all sets $\mathbf{G}$. □

2.6.4 Common Information Observes Neither Submodularity Nor Supermodularity Laws

As discussed in the above, appealing to the duality between the join and the meet operations, one might conjecture, dual to the well-known submodularity of joint information, that common information would observe the supermodularity law. It turns out that common information observes neither the submodularity (6) nor the supermodularity (7) law neither of the following two inequalities holds in general:

$$h(m_{12}) + h(m_{23}) \geq h(m_{123}) + h(m_2) \tag{6}$$

$$h(m_{12}) + h(m_{23}) \leq h(m_{123}) + h(m_2). \tag{7}$$

Because common information is combinatorial in flavor—it depends on the "zero pattern" of joint probability matrices [12] — it is hard to directly verify the validity of (6) and (7). However, thanks to Theorem 2.6.17, we are able to construct subgroup counterexamples to invalidate (6) and (7) indirectly.

To show that (7) fails, it suffices to find three subgroups G_1, G_2, and G_3 such that

$$|G_1 \vee G_2||G_2 \vee G_3| < |G_1 \vee G_2 \vee G_3||G_2|. \tag{8}$$

Consider $G = S_5$, the symmetry group of order 2^5, and its subgroups $G_1 = \langle(12345)\rangle$, $G_2 = \langle(12)(45)\rangle$, and $G_3 = \langle(12543)\rangle$. The subgroup G_1 is the permutation group generated by permutation (12345), G_2 by (12)(45), and G_3 by (12543). (Here, we use the standard cycle notation to represent permutations.) Consequently, we have $G_1 \vee G_2 = \langle(12345), (12)(45)\rangle$, $G_2 \vee G_3 = \langle(12543), (12)(45)\rangle$, and $G_1 \vee G_2 \vee G_3 = \langle(12345), (12)(45), (12543)\rangle$. It is easy to see that both $G_1 \vee G_2$ and $G_2 \vee G_3$ are dihedral groups of order 10 and that $G_1 \vee G_2 \vee G_3$ is the alternative group A_5, hence of order 60. The order of G_2 is 2. Therefore, we see that the subgroups G_1, G_2, and G_3 satisfy (8). By Theorem 2.6.17, the supermodularity law (7) does not hold in general for common information. (Thank to Professor Eric Moorhouse for contributing this counterexample.)

Similar to the case of supermodularity, the example with $G_2 = \{e\}$ and $G_1 = G_3 = G$, $|G| \neq 1$, invalidates the group version of (6). Therefore, according to Theorem 2.6.17, the submodularity law (6) does not hold in general for common information either.

2.7 Discussion

This work builds on some of Shannon's little-recognized legacy and adopts his interesting concepts of information elements and information lattices. We formalize all these concepts and clarify the relations between random variables and information elements, information elements and σ-algebras, and, especially, the one-to-one correspondence between information elements and sample-space-partitions. We emphasize that such formalization is conceptually significant. As demonstrated in this work, beneficial to the formalization carried out, we are able to establish a comprehensive parallelism between information lattices and subgroup lattices. This parallelism is mathematically natural and admits intuitive group-action explanations. It reveals an intimate connection, both structural and quantitative, between information theory and group theory. This suggests that group theory might serve a promising role as a suitable mathematical language in studying deep laws governing information.

Network information theory in general, and capacity problems for network coding specifically, depend crucially on our understanding of intricate structures among multiple information elements. By building a bridge from information theory to group theory, we can now access the set of well-developed tools from group theory. These tools can be brought to bear on certain formidable problems in areas such as network information theory and network coding. Along these lines, by constructing subgroup counterexamples we show that neither the submodularity nor the supermodularity law holds for common information, neither of which is obvious from traditional information theoretic perspectives.

2.8 Appendix

2.8.1 Proof of Theorem 2.4.4

Proof. To show two lattices are isomorphic, we need to demonstrate a mapping, from one lattice to the other, such that it is a lattice-morphism – it honors both join and meet operations – and bijective as well. Instead of proving that $\mathrm{L}_{\mathbf{G}}$ is isomorphic to L_G directly, we show that the dual of $\mathrm{L}_{\mathbf{G}}$ is isomorphic to $\mathrm{L}_{\mathbf{M}}$. Figuratively speaking, the dual of a lattice L is the lattice obtained by flipping L upside down. Formally, the dual lattice L′ of a lattice L is the lattice defined on the same set with the partial order reversed. Accordingly, the join operation of the prime lattice L corresponds to the meet operation for the dual lattice L′ and the meet operation of L to the join operation for L′. In the other words, we show that L_G is isomorphic to $\mathrm{L}_{\mathbf{M}}$ by demonstrating a bijective mapping $\phi : \mathrm{L}_{\mathbf{G}} \rightarrow \mathrm{L}_{\mathbf{M}}$ such that

$$\phi(G \vee G') = \phi(G) \wedge \phi(G'), \tag{9}$$

and

$$\phi(G \wedge G') = \phi(G) \vee \phi(G'), \tag{10}$$

hold for all $G, G' \in \mathrm{L}_{\mathbf{G}}$.

Note that each subgroups on the subgroup lattice $\mathrm{L}_{\mathbf{G}}$ is obtained from the set $\mathbf{G} = \{G_i : i \in [n]\}$ via a sequence of join and meet operations and each information element on the information lattice $\mathrm{L}_{\mathbf{M}}$ is obtained similarly from the set $\mathbf{M} = \{m_i : i \in [n]\}$. Therefore, to show that $\mathrm{L}_{\mathbf{G}}$ is isomorphic to $\mathrm{L}_{\mathbf{M}}$, according to the induction principle, it is enough to demonstrate a bijective mapping ϕ such that

- $\phi(G_i) = m_i$, for all $G_i \in \mathbf{G}$ and $m_i \in \mathbf{M}$;
- For any $G, G' \in \mathrm{L}_{\mathbf{G}}$, if $\phi(G) = m$ and $\phi(G') = m'$, then

$$\phi(G \vee G') = m \wedge m', \text{ and} \tag{11}$$

$$\phi(G \wedge G') = m \vee m'. \tag{12}$$

Naturally, we take $\phi : \mathrm{L_G} \rightarrow \mathrm{L_M}$ to be the mapping that assigns to each subgroup $G \in \mathrm{L_G}$ the information element identified by the coset-partition of the subgroup G. Thus, the initial step of the induction holds by assumption. On the other hand, it is easy to see that the mapping ϕ so defined is bijective simply because different subgroups always produce different coset-partitions and vice versa. Therefore, we are left to show that Equation (11) and (12) holds.

We first show that ϕ satisfies Equation (11). In other words, we show that the coset-partition of the intersection subgroup $G \cap G'$ is the coarsest among all the sample-space-partitions that are finer than both the coset-partitions of G and G'. To see this, let Π be a sample-space-partition that is finer than both the coset-partitions of G and G' and π be a part of Π. Since Π is finer than the coset-partitions of G, π must be contained in some coset C of G. For the same reason, π must be contained in some coset C' of G' as well. Consequently, $\pi \subseteq C' \cap C'$ hold. Realizing that $C \cap C'$ is a coset of $G \cap G'$, we conclude that the coset-partition of $G \cap G'$ is coarser than Π. Since Π is chosen arbitrary, this proves that the coset-partition of the intersection subgroup $G \cap G'$ is the coarsest among all the sample space partitions that are finer than both the coset-partitions of G and G'. Therefore, Equation (11) holds for ϕ.

The proof for Equation (12) is more complicated. We use an idea called "transitive closure". Similarly, we need to show that the coset-partition of the subgroup $G \vee G'$ generated from the union of G and G' is the finest among all the sample-space-partitions that are coarser than both the coset-partitions of G and G'. Let Π be a sample-space-partition that is coarser than both the coset-partitions of G and G'. Denote the coset partition of the subgroup $G \vee G'$ by $\bar{P}i$. Let $\bar{\pi}$ be a part of $\bar{\Pi}$. It suffices to show that $\bar{\pi}$ is contained in some part of Π. Pick an

element x from $\bar{\pi}$. This element x must belong to some part π of Π. It remains to show $\bar{\pi} \subseteq \pi$. In other words, we need to show that $y \in \pi$ for any $y \neq x, y \in \pi^{ij}$. Note that π is a part of the coset-partition of the subgroup $G_i \vee G_j$. In other words, π is a coset of $G_i \vee G_j$. The following reasoning depends on the following fact from group theory [25].

Proposition 2.8.1. *Two elements g_1 and g_2 belong to a same (right) coset of a subgroup if and only if $g_1 g_2^{-1}$ belongs to the subgroup.*

Since x and y belong to a same coset π of the subgroup $G \vee G'$, we have $yx^{-1} \in G \vee G'$. Note that any element g from $G \vee G'$ can be written in the form of $g = a_1 b_1 a_1 b_2 \cdots a_K b_K$ where $a_k \in G$ and $b_k \in G'$ for all $k \in [K]$. Suppose $yx^{-1} = g = a_1 b_1 a_1 b_2 \cdots a_K b_K$. We have

$$y = a_1 b_1 a_2 b_2 \cdots a_K b_K x.$$

In the following we shall show that y belongs to $\bar{\pi}$ by induction on the sequence $a_1 b_1 \cdots a_K b_K$.

First, we claim $b_K x \in \bar{\pi}$. To see this, note that $x \in \bar{\pi}$. Since $(b_K x)x^{-1} = b_K \in G'$, by Proposition 2.8.1, we know that $b_K x$ and x belong to a same coset C_K of G'. By assumption, the partition Π is coarser than the coset-partition of G', the coset C_K must be contained in $\bar{\pi}$, since it already contains an element x of C_K.

For the same reason, with $b_K x \in \bar{\pi}$ showed, we can see that $a_K b_K x$ belongs to $\bar{\pi}$ as well, because $(a_K b_K x)(b_K x)^{-1} = a_K \in G$ implies $a_K b_K x$ and $b_K x$ belong to a same coset of G.

Continuing the above argument inductively on the sequence $a_1 b_1 \cdots a_K b_K$, we can finally have $a_1 b_1 \cdots a_K b_K x \in \bar{\pi}$. Therefore, we have $y \in \bar{\pi}$. This concludes the proof. □

2.8.2 Proof of Theorem 2.5.4

Proof. The approximation process is decomposed into three steps. The first step is to "dilate" the sample space such that we can turn a non-uniform probability space into a uniform probability space. The sample space partitions of the information elements are accordingly "dilated" as well. After dilating the sample space, depending on the approximation error tolerance, i.e., ϵ, we may need to further "amplify" the sample space. Then, we follow the same procedure as in Section 2.4.4 and construct a subgroup lattice using the orbit-partition permutation-group-action correspondence.

We assume the probability measure $\mathbf{P}$ on the sample space are rational. In other words, the probabilities of the elementary event $\mathrm{p}_i = \Pr\{\omega_i\}$, $\omega_i \in \Omega$ are all rational numbers, namely $\mathrm{p}_i = \frac{p_i}{q_i}$ for some $p_i, q_i \in \mathbb{N}$. This assumption is reasonable, because any finite dimensional real vector can be approximated, up to an arbitrary precision, by some rational vector.

Let M be the least common multiple of the set $\{q_i\}$ of denominators. We "split" each sample point in Ω into $\frac{Mp_i}{q_i}$ points. Note that $\frac{Mp_i}{q_i}$ is integral. We need to accordingly "dilate" the sample space partitions of the information elements. Specifically, for each part π of the partition of every information element m_i, its "dilated" partition π', in the dilated sample space $\hat{\Omega}$, contains exactly all the sample points that are "split" from the sample points in π. The dilated sample space $\hat{\Omega}$ has size of $\sum_{\omega_i \in \Omega} \frac{Mp_i}{q_i}$. To maintain the probability structure, we assign to each sample point in the dilated sample space $\hat{\Omega}$ probability $\frac{1}{|\hat{\Omega}|}$. In other words, we equip the dilated sample space with a uniform probability measure. It is easy to check that the entire (quantitative) probability structure remains the same. Thus, we can consider all the information elements as if defined on the dilated probability space.

If necessary, depending on the approximation error tolerance ϵ, we may further "amplify" the dilated sample space $\hat{\Omega}$ by K times by "splitting" each of its sample points into to K points. At the same time, we scale the probability of each sample point in the post-amplification sample space down by K times to $\frac{1}{K|\hat{\Omega}|}$. By abusing of notation, we still use $\hat{\Omega}$ to denote the post-amplification sample space. Similar to the "dilating" process, all the partitions are accordingly amplified.

Before we move to the third step, we compute entropies for information elements in terms of the cardinality of the parts of its dilated sample space partition. Consider an information element m_i. Denote its pre-dilation sample space partition by $\Pi_i = \{\pi_i^j, j \in [J]\}$ and its post-amplification sample space partition by $\hat{\Pi}_i = \{\hat{\pi}_i^j, j \in [J]\}$. It is easy to see that the entropy $H(m_i)$ can be calculated as follows:

$$\begin{aligned} H(m_i) &= -\sum_{j\in[J]} \Pr\{\pi_i^j\} \log \Pr\{\pi_i^j\} \\ &= -\sum_{j\in[J]} \Pr\{\hat{\pi}_i^j\} \log \Pr\{\hat{\pi}_i^j\} \\ &= -\sum_{j\in[J]} \frac{|\hat{\pi}_j^i|}{|\hat{\Omega}|} \log \frac{|\hat{\pi}_j^i|}{|\hat{\Omega}|}. \end{aligned} \tag{13}$$

All the entropies of the other information elements, including the joint and common information elements, on the entire information lattices can be computed in the exactly same way in terms of the cardinalities of the parts of their dilated sample space partitions.

In the third step, we follow the same procedure as in Section 2.4.4, and construct, based on the orbit-partition-permutation-group-action correspondence, a subgroup lattice that isomorphic to the information lattice generated by the set of information elements $\{m_i : i \in [n]\}$. More specifically, the subgroup lattice is constructed according to their "post-amplification" sample space partitions.

Suppose, on the constructed subgroup lattice, the permutation groups G_i cor-

responds to the information element m_i. As in the above, the "post-amplification" sample space partition of m_i is $\hat{\Pi}_i = \{\hat{\pi}_i^j, j \in [J]\}$. Then, the cardinality of the permutation group is simply

$$|G_i| = \prod_{j \in J} \hat{\pi}_i^j!.$$

According to the isomorphism relation established in Theorem 2.4.7, the above calculations remain valid for all the subgroups on the subgroup lattices.

Recall that all the groups on the subgroup lattice are permutation groups and are all subgroups of the symmetry group of order $|\hat{\Omega}|$. So the log-index of G_i, corresponding to m_i, is

$$\log \frac{|\hat{\Omega}|!}{|G_i|} = \log \frac{|\hat{\Omega}|!}{\prod_{j \in J} \hat{\pi}_i^j!}. \tag{14}$$

As we see from Equation (1) and (2) of Proposition 2.5.1, the entropies of the coset-partition information elements on information lattices equal *exactly* the log-indices of their subgroups on subgroup lattices. However, for the information lattice generated from general information elements, namely information elements with non-equal sample space partitions, as we see from Equation (13) and (14), the entropies of the information elements on the information lattice does not equal the log-indices of their corresponding permutation groups on the subgroup lattices exactly any more. But, as we can shall see, the entropies of the information elements are well *approximated* by the log-indices of their corresponding permutation groups. Recall the following Stirling's approximation formula for factorials:

$$\log n! = n \log n - n + o(n). \tag{15}$$

"Normalizing" the log-index in Equation (14) by a factor $\frac{1}{|\hat{\Omega}|}$ and then substituting the factorials with the above Stirling approximation formula, we get

$$\frac{1}{|\hat{\Omega}|} \log \frac{|\hat{\Omega}|}{|G_i|} = \frac{1}{|\hat{\Omega}|} \Big(|\hat{\Omega}| \log |\hat{\Omega}| - |\hat{\Omega}| - \Big(\sum_{j \in [J]} |\hat{\pi}_i^j| \log |\hat{\pi}_i^j| - |\hat{\pi}_i^j| \Big) + o(|\hat{\Omega}|) \Big).$$

Note that in the above substitution process, we combined some finite $o(|\hat{\Omega}|)$ terms "into" one $o(|\hat{\Omega}|)$ term.

It is clear that $\sum_{j\in[J]} |\hat{\pi}_i^j| = |\hat{\Omega}|$, since $\{\hat{\pi}_i^j : j \in [J]\}$ forms a partition of $\hat{\Omega}$. Therefore, we get

$$\begin{aligned}\frac{1}{|\hat{\Omega}|}\log\frac{|\hat{\Omega}|}{|G_i|} &= \frac{1}{|\hat{\Omega}|}\Big(|\hat{\Omega}|\log|\hat{\Omega}| - \sum_{j\in[J]} |\hat{\pi}_i^j|\log|\hat{\pi}_i^j| + o(\hat{\Omega})\Big)\\ &= h(m_i) + \frac{o(|\hat{\Omega}|)}{|\hat{\Omega}|}.\end{aligned}$$

So, the difference between the entropy $H(m_i)$ and the normalized log-index of its corresponding permutation subgroup G_i diminishes for $\hat{\Omega}$ large.

Since both the entropy vector $h_{\mathbf{M}}$ and the log-index vector $l_{\mathbf{G}^N}$ are of finite dimension, it follows easily

$$\left\| h_{\mathbf{M}} - \frac{l_{\mathbf{G}^N}}{N} \right\|_1 = \frac{o(|\hat{\Omega}|)}{|\hat{\Omega}|} \to 0,$$

with

$$N = |\hat{\Omega}| = K \sum_{\omega_i\in\Omega} \frac{Mp_i}{q_i} \to \infty, \text{ by taking } K \to \infty.$$

This concludes the proof. □

List of References

[1] C. E. Shannon, "The lattice theory of information," *IEEE Transactions on Information Theory*, vol. 1, no. 1, pp. 105 107, Feb. 1953.

[2] T. H. Chan and R. W. Yeung, "On a relation between information inequalities and group theory," *IEEE Transactions on Information Theory*, vol. 48, no. 7, pp. 1992 1995, July 2002.

[3] C. E. Shannon, "A mathematical theory of communication." *Bell System Technical Journal*, vol. 27, pp. 379 423 and 623 656, July and October 1948.

[4] P. Billingsley, *Probability and Measure*, 3rd ed. John Wiley & Sons, 1995.

[5] S. E. Shreve, *Stochastic Calculus for Finance I: The Binomial Asset Pricing Model.* Springer, 2005.

[6] S. Ankirchner, S. Dereich, and P. Imkeller, "The Shannon information of filtrations and the additional logarithmic utility of insiders," *The Annals of Probability*, vol. 34, pp. 743–778, 2006.

[7] A. Orlitsky, N. P. Santhanam, and J. Zhang, "Universal compression of memoryless sources over unknown alphabets," *IEEE Transactions on Information Theory*, vol. 50, no. 7, pp. 1469–1481, July 2004.

[8] A. Renyi, *Foundations of Probability*. Holden-Day Inc., 1970.

[9] X. Yan, R. W. Yeung, and Z. Zhang, "The capacity region for multi-source multi-sink network coding," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 2007, pp. 116–120.

[10] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2345–2364, June 2006.

[11] P. Pudlák and J. Tůma, "Every finite lattice can be embedded in a finite partition lattice," *Algebra Universalis*, vol. 10, pp. 74–95, 1980.

[12] P. Gács and J. Körner, "Common information is far less than mutual information," *Problems of Control and Information Theory*, vol. 2, pp. 149–162, 1973.

[13] H. S. Witsenhausen, "On sequences of pairs of dependent random variables," *SIAM Journal on Applied Mathematics*, vol. 28, pp. 100–113, 1975.

[14] R. Ahlswede and I. Csiszàr, "Common randomness in information theory and cryptography – part I: Secret sharing," *IEEE Transactions on Information Theory*, vol. 39, pp. 1121–1132, 1993.

[15] R. Ahlswede and I. Csiszàr, "Common randomness in information cryptography – part II: CR capacity," *IEEE Transactions on Information Theory*, vol. 44, pp. 225–240, 1998.

[16] I. Csiszàr and P. Narayan, "Common randomness and secret key generation with a helper," *IEEE Transactions on Information Theory*, vol. 46, pp. 344–366, 2000.

[17] S. Wolf and J. Wullschleger, "Zero-error information and application in cryptography," in *Proceedings of the 2004 IEEE Information Theory Workshop (ITW 2004)*, 2004.

[18] T. Cover, A. E. Gamal, and M. Salehi, "Multiple access channels with arbitrarily correlated sources," *IEEE Transactions on Information theory*, vol. 26, pp. 648–657, 1980.

[19] Z. Zhang, "On a new non-Shannon type information inequality," *Communications in Information and Systems*, vol. 3, no. 1, pp. 47 60, June 2003.

[20] D. Hammer, A. Romashchenko, A. Shen, and N. Vereshchagin, "Inequalities for Shannon entropy and Kolmogorov complexity," *Journal of Computer and System Sciences*, vol. 60, no. 2, pp. 442 464, April 2000.

[21] U. Niesen, C. Fragouli, and D. Tuninetti, "On capacity of line networks," submitted to IEEE Transactions on Information Theory.

[22] S. Fujishige, "Polymatroidal dependence structure of a set of random variables," *Information and Control*, vol. 39, pp. 55 72, 1978.

[23] F. Cicalese and U. Vaccaro, "Supermodularity and subadditivity properties of entropy on the majorization lattice," *IEEE Transactions on Information Theory*, vol. 48, no. 4, pp. 933–938, Apr. 2002.

[24] A. Chernov, A. Muchnik, A. Romashchenko, A. Shen, and N. Vereshchagin, "Upper semilattice of binary strings with the relation 'x is simple conditional to y'," *Theoretical Computer Science*, vol. 271, no. 1, pp. 69 95, Jan. 2002.

[25] D. S. Dummit and R. M. Foote, *Abstract Algebra*. 3rd ed. Wiley, 2003.

[26] R. W. Yeung, *A First Course in Information Theory*. Kluwer Academic/Plenum Publishers, 2002.

[27] J. G. Oxley, *Matroid Theory*. Oxford University Press, 1992.

[28] Z. Zhang and R. W. Yeung, "On characterization of entropy function via information inequalities," *IEEE Transactions on Information Theory*, vol. 44, no. 4, pp. 1440 1452, July 1998.

[29] R. Dougherty, C. Freiling, and K. Zeger, "Six new non-Shannon information inequalities," in *Proceedings of the 2006 IEEE International Symposium on Information Theory*, 2006, pp. 233 236.

[30] N. Pippenger, "What are the laws of information theory," in *1986 Special Problems on Communication and Computation Conference*, Palo Alto, California, Sept. 3-5 1986.

[31] F. Matúš, "Piecewise linear conditional information inequality," *IEEE Transactions on Information Theory*, vol. 52, no. 1, pp. 236 238, Jan. 2006.

[32] H. Li and E. K. P. Chong, "On connections between group homomorphisms and the Ingleton inequality," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 24 29 2007, pp. 1996 2000.

[33] F. Matúš, "Infinitely many information inequalities," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 24-29 2007, pp. 41 44.

CHAPTER 3

On the Entropy Function and the Ingleton Inequality

3.1 Summary

We investigate, in Section 3.3, a set of potentially valid 4-variable non-Shannon-type information inequalities constructed by adding at most three conditional mutual information terms to the Ingleton inequality. In particular, we attempt to disprove them by computer-aided searching for counterexamples among sets of joint distributions of 4 binary random variables. We find that all except three of them are invalidated. One direct consequence of this result is that none of the three extra conditional mutual information terms in the inequality discovered by Zhang and Yeung can be dropped for the inequality to remain valid.

Then, we show in Section 3.4 that random variables mapped under group homomorphisms from a uniformly distributed background random variable satisfy the Ingleton inequality. As corollaries, we recover two previous known results. The first is that the network throughput of linear network codes is, in general, constrained by the Ingleton Inequality. The second and related result is that the network throughput of abelian-group network codes—group network codes that are restricted to abelian groups—is also constrained by the Ingleton inequality.

Further along the line of questing more general conditions for the Ingleton inequality, in Section 3.5 we identify a general group-theoretic condition for the Ingleton inequality, subsuming all previously known conditions, including those obtained in Section 3.4. Specifically, we show that quasi-Hamiltonian groups satisfy the Ingleton inequality. Quasi-Hamiltonian groups include as a subclass the well-known Hamiltonian groups, which are non-abelian. To our best knowledge, this is the first time that the Ingleton inequality is found to hold for certain classes of non-abelian groups.

3.2 Introduction

Ahlswede et al. [1] first demonstrated that network coding is capable of achieving better information flow throughput for single-source multicast networks than the current practice of storing-and-forwarding. Research interest in network coding has flourished in both information theory and networking communities, especially after Li et al. [2] showed that network codes as simple as linear block codes can achieve the maximum multicast throughput promised for general network codes in [1].

However, the capacity problem for general multi-source, multi-sink networks proves to be difficult. So far, besides multicast networks, only the capacity region of the special single-source two-sink network has been successfully characterized similarly in terms of the maximum-flows (minimum-cuts) from the source to the two sinks [3–5]. Various computable or explicit outer bounds were derived by Yeung [6], Kramer and Savari [7] and Harvey et al. [8]. Recently, Yan et al. [9] obtained an exact but implicit characterization, in terms of the fundamental regions of the entropy function, for the network capacity regions of general acyclic multi-source, multi-sink networks. Notably, this is the first exact characterization established for general network capacity regions. However, the fundamental regions themselves remain elusive.

Current research on the range of the entropy function focuses on the method of *information inequalities*, linear inequalities of joint entropies. Information inequalities are referred to as "laws of information theory" [6, 10]. For decades, the only known laws for the entropy function were the *polymatroidal axioms*. Information inequalities that can be written as linear combinations of the instances of the polymatriodal axioms are called *basic* or *Shannon-type* information inequalities – these laws were essentially established by Shannon in his founding paper [11]. It turns

out that for the cases involving four or more variables there exist other laws that are *not* the consequence of the polymatriodal axioms — in 1998, Zhang and Yeung discovered the first non-Shannon type information inequality [12]. Recently six more were found by Dougherty et al. [13]. Various other results on non-Shannon-type information inequalities have been reported by Lněnička [14], Makarychev et al. [15], Zhang and Yeung [16], and Matúš [17].

It is well-known that the polymatroidal axioms fully characterize the rank function of *matroids* [18]. The rank function of *vector matroids* — a special class of matroids that can be represented with sets of vectors — were found to be further constrained by the Ingleton inequality [19]. However, it was found that the entropy function in general does *not* respect the Ingleton inequality [12, 20]. With similar reasoning in establishing the capacity region characterization for general multi-source, multi-sink networks [9], we can argue that linear network codes may sacrifice certain throughput performance compared with general network codes: see the counterexamples constructed by Dougherty et al. [21] and the recent theoretical evidence provided by Chan and Grant [22].

3.2.1 The Entropy Function

Throughout this chapter, we consider only discrete random variables over finite sample spaces, and all the logarithms for entropies are taken to be base 2 unless mentioned otherwise.

For a set of random variables, for example $\{X_1, X_2, ..., X_n\}$, we denote its joint entropy by $H(\{X_1, X_2, ..., X_n\})$. For simplicity, we may write $H(X_1, X_2, ..., X_n)$ instead when no confusion arises. We first introduce our definition for *entropy functions*.

Definition 3.2.1. *Let n be a positive integer. The* entropy function $\mathcal{H}_n$ *is a vector-valued function which maps every set Ω of n random variables to a $(2^n - 1)$-*

dimensional real vector, denoted $(H_n(\omega_1), \cdots, H(\omega_i), \cdots, H(\omega_{2^n-1}))$, *where* ω_i, $i = 1, 2, \cdots, 2^n - 1$, *are subsets of* Ω *and are indexed according to a certain pre-specified order.*

We call such a (2^n-1)-dimensional real vector an *entropy vector*. Note that our definition for *entropy functions* is slightly different from what has been proposed in the past, such as in [16].

Entropy functions are of central importance in information theory. It is of both theoretical and practical interest to characterize entropy functions, especially their ranges, because the boundaries of theirs ranges separate what is possible and what is impossible with respect to coding freedom, one of the essential kinds of design freedom granted in information theory. Denote the range of $\mathcal{H}_n$ by $\mathbf{H}_n$. Note that $\mathbf{H}_n$ lies in $\mathbb{R}^{2^n-1}$. We call $\mathbf{H}_n$ the *n-variable entropy range*, or simply *entropy range*. Because of the discreteness of the underlying random variables, we would expect that $\mathbf{H}_n$ has "irrational holes" in general. But it has been shown in [16] that those "irrational holes" are *asymptotically* achievable by sequences of n-random variable sets. Therefore, it is justified, also for the sake of mathematical convenience, to focus on its closure, denoted $\overline{\mathbf{H}}_n$. It turns out that $\overline{\mathbf{H}}_n$ is a pointed convex cone, lying in the nonnegative orthant of $\mathbb{R}^{2^n-1}$ [16]. For this reason, we call it the n-variable *entropy cone*, or simply *entropy cone*. Since this geometric view was first taken by Zhang and Yeung in [23] and [24], some progress has been made towards characterizing $\overline{\mathbf{H}}_n$.

3.2.2 Shannon Cones

Throughout the chapter, we refer to linear combinations of (joint) entropies, for example $H(X) + H(Y) - H(X, Y)$, as *information expressions* and associated inequalities of the form $H(X) + H(Y) - H(X, Y) \geq 0$ as *information inequalities*. We say that an information inequality is *valid* if it holds for any joint distribution of

the random variables involved. Because all other types of the *Shannon information measures*, namely mutual information, conditional entropy, and conditional mutual information, are all linear combinations of (joint) entropies, to shorten expressions, we may write some of the information inequalities encountered later in terms of "Shannon information measures."

Note that a superficially new information inequality can be obtained by permuting the variables from an old inequality. For example, from $I(A;B|C) \geq 0$, we obtain a "new" inequality $I(C;B|A) \geq 0$.(Throughout this chapter, we use $I(A;B|C)$ to denote the conditional mutual information between the random variables A and B conditioning on C.) To avoid such trivialities, throughout the chapter we refer to the class of all the essentially equivalent information inequalities as *one* inequality.

In terms of information inequalities, the so-called *polymatroidal axioms* satisfied by entropy functions can be presented as follows:

Axiom 3.2.2. *(Polymatroidal Axioms [18])*

1. *(Non-negative)* $H(\alpha) \geq 0$,

2. *(Non-decreasing)* $H(\beta) - H(\alpha) \geq 0$ *if* $\alpha \subseteq \beta$,

3. *(Submodular)* $H(\alpha) + H(\beta) - H(\alpha \cup \beta) - H(\alpha \cap \beta) \geq 0$,

where α *and* β *are sets of random variables.*

Information inequalities that can be expressed as non-negative linear combinations of instances of the above three polymatroidal axioms were classified in the textbook [6] as *basic* or *Shannon-type* information inequalities these laws were essentially established by Shannon in his founding paper [11]. Altogether, the Shannon-type information inequalities involving at most n random variables de-

fine a convex cone in $\mathbb{R}^{2^n-1}$. We call this cone the *Shannon cone*, denoted $\mathbf{S}_n$. Clearly, the Shannon cone $\mathbf{S}_n$ is an outer bound for the entropy cone $\overline{\mathbf{H}}_n$.

3.2.3 Improving the Outer Bound by Finding More non-Shannon-type Information Inequalities

Surprisingly, for decades the foregoing polymatroidal axioms remained as the only known laws for entropy functions. It remained unknown whether $\mathbf{S}_n$ and $\overline{\mathbf{H}}_n$ coincide until Zhang and Yeung discovered in 1998 [12] the first 4-random variable non-Shannon-type information inequality, an inequality that cannot be expressed as a non-negative linear combination of Shannon-type inequalities. Consequently, it is now known that $\mathbf{S}_n$ is *strictly* larger than $\overline{\mathbf{H}}_n$ for $n \geq 4$. In the same paper, it was shown that $\mathbf{H}_2 = \mathbf{S}_2$ and $\overline{\mathbf{H}}_3 = \mathbf{S}_3$ ($\mathbf{H}_3 \neq \mathbf{S}_3$ though because of "irrational holes.")

Inequality 3.2.3. *(Zhang-Yeung Inequality [12]) For any four random variables A, B, C, and D, the following inequality holds:*

$$\begin{aligned} &I(C;D|A) + I(C;D|B) + I(A;B) - I(C;D) \\ &\quad + I(A;D|C) + I(A;C|D) + I(C;D|A) \geq 0. \end{aligned} \tag{16}$$

Towards characterizing entropy cones $\overline{\mathbf{H}}_n$ for $n \geq 4$, the 4-variable Zhang-Yeung inequality was subsequently generalized to the cases involving arbitrary $n > 4$ random variables [15]. Recently six new non-Shannon-type information inequalities were discovered by Dougherty et al. [13] by generalizing Zhang and Yeung's original proof technique in a significant way. Some other results on non-Shannon-type information inequalities have been reported by Lněnička [14], Zhang and Yeung [16], and Matúš [17].

Obviously, by finding more and more non-Shannon-type inequalities we get closer and closer, from the outside, to the boundary of the entropy cone $\overline{\mathbf{H}}_n$, $n \geq 4$. However, there is a subtle point that we should point out: It is not clear whether

$\overline{\mathbf{H}}_n$, $n \geq 4$, is polyhedral at all, even though, as discussed previously, $\overline{\mathbf{H}}_2$ and $\overline{\mathbf{H}}_3$ happen to be so. Therefore, theoretically, there is a chance that the entropy cone $\overline{\mathbf{H}}_n$, $n \geq 4$, may not be fully characterized by a *finite* set of information inequalities, which are linear by definition. It turns out that this is indeed the case—Matúš [25] recently showed that at least some part of the boundary of entropy cone $\overline{\mathbf{H}}_n$, $n \geq 4$, is curved. Thus, we may eventually need to seek non-linear information inequalities to succinctly characterize such entropy cones (even though we have not encountered one yet).

3.2.4 Ingleton Cones

In the quest for a full characterization of $\overline{\mathbf{H}}_n$, the above approach of information inequalities is the most popular one so far, and indeed it proved to be successful in the cases of $n = 2$ and 3. Information inequalities in themselves are very important, for they "essentially govern the impossibility in information theory" [6]. However, it seems very difficult to find new information inequalities – we lack systematic tools to find them. To tackle the problem, we may explore from the opposite direction, namely, from the "inside" of $\overline{\mathbf{H}}_n$. Ingleton cones, defined by the Ingleton inequality together with Shannon-type inequalities, seem to serve as a good starting point. We denote the Ingleton cone involving n random variables by $\mathbf{I}_n$, $n \geq 4$.

Inequality 3.2.4. *(Ingleton Inequality)*

$$I(C;D|A) + I(C;D|B) + I(A;B) - I(C;D) \geq 0,$$

where A, B, C, and D denote four random variables.

The Ingleton inequality holds for rank functions of *vector matroids* [19], but not for entropy functions in general [12, 20, 26]. Conversely, it was shown in [20] that all the linear inequalities satisfied by entropy functions hold for rank functions

of vector matroids. Therefore, the sets defined by the linear inequalities satisfied by vector rank functions are *strictly* smaller than entropy cones, with Ingleton cones sandwiched in between. (To our best knowledge, we do not yet have a full characterization for rank functions of vector matroids either.)

Since the start of the quest towards characterizing $\overline{\mathbf{H}}_n$, most work has focused on finding new information inequalities, for whenever a new information inequality is established, a better outer bound is obtained. However, little work has been done from the inside of $\overline{\mathbf{H}}_n$. In the next section, we try to approach the boundary of $\overline{\mathbf{H}}_n$ from its interior. In particular, we concentrate our attention on the case of 4 random variables – this is the simplest unresolved case. We investigate a set of potentially valid non-Shannon-type information inequalities involving 4 random variables, and try to disprove them by searching, with the aid of computers, for counterexamples. All the counterexamples discovered along the search serve as "landmarks" in our adventure towards the boundary of $\overline{\mathbf{H}}_n$. Fig. 2 shows the relation among the Shannon cone, the entropy cone, and the Ingleton cone.

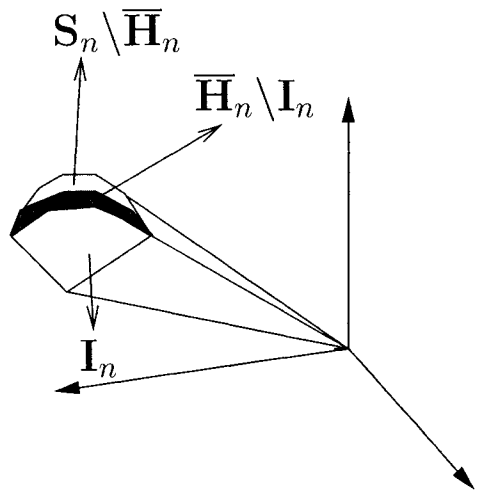

Figure 2. Shannon cone, entropy cone, and Ingleton cone

3.3 Exploring the Space between the Ingleton Cone $\mathbf{I}_4$ and the Shannon Cone $\mathbf{S}_4$

From now on, we restrict ourselves to the case involving only 4 random variables. In this case, we explore the territory between $\mathbf{S}_4$ and $\mathbf{I}_4$, denoted $\mathbf{S}_4 \backslash \mathbf{I}_4$.

3.3.1 Adding Conditional Mutual Information Terms to the Ingleton Inequality

To explore $\mathbf{S}_4 \backslash \mathbf{I}_4$, we need a certain sense of "direction" to navigate. Observe that the Zhang-Yeung inequality (16) can be written as the linear combination of the Ingleton inequality and three extra conditional mutual information terms, which are non-negative. We naturally wonder whether any of these three extra conditional mutual information terms can be dropped or replaced. More generally, we ask whether there are any valid non-Shannon information inequalities, other than the Zhang-Yeung inequality, that are linear combinations of the Ingleton inequality and one, two, or three conditional mutual information terms (multiplicity allowed). They take the following form

$$\begin{aligned} &[I(C;D|A) + I(C;D|B) + I(A;B) - I(C;D)] \\ &+ \alpha_1 I(X_1^1; X_2^1|X_3^1) + \alpha_2 I(X_1^2; X_2^2|X_3^2) \\ &+ \alpha_3 I(X_1^3; X_2^3|X_3^3) \geq 0, \end{aligned} \tag{17}$$

where X_j^i, $i, j = 1, 2, 3$, are random variables taken from $\{A, B, C, D\}$ and $\alpha_i = 0, 1, 2$ such that $\sum_i \alpha_i \leq 3$. Furthermore, for each i, X_1^i, X_2^i, and X_3^i are taken to be distinct. There are totally 122 such inequalities, excluding the Zhang-Yeung and the Ingleton ones.

Note that, according to [27], all those inequalities are *balanced*. In the same paper, it was shown that every (discrete) information inequality can be written as the linear combination of its "balanced" counterpart and a set of "residual weights," and that every (discrete) information inequality is valid if and only if

its balanced counterpart is valid. In this sense, these candidate inequalities are minimal, for they contain no "residual weights."

3.3.2 Identifying "Pass-through" Inequalities

We are going to use these information inequalities to provide us some sense of "direction" when we explore the territory between $\mathbf{I}_4$ and $\mathbf{S}_4$. As a first step, we aim to identify, among the 122 candidates, those which "pass through" $\mathbf{S}_4 \backslash \mathbf{I}_4$, explained next.

For each information inequality, for example the Ingleton inequality $I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D) \geq 0$, associated with it is a hyperplane defined by the associated equation $I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D)=0$. We call an inequality a "pass-through" inequality if its associated hyperplane "passes through" $\mathbf{S}_4 \backslash \mathbf{I}_4$, i.e., in $\mathbf{S}_4 \backslash \mathbf{I}_4$ there is at least one point at which the inequality holds and at least one point at which the inequality fails.

We know from geometry that a convex polyhedral cone can be represented in two different ways. It can be either represented as the intersection of a finite number of half-spaces via a set of linear inequalities or as a non-negative linear combination of its extreme points, those lying on its extreme rays [28].

Relevant to our case here, $\mathbf{I}_4$ can be represented by its 35 extreme points and $\mathbf{S}_4$ can be represented by the 35 extreme points of $\mathbf{I}_4$ and six extra extreme points. See [20] for a detailed account.

With the help of the six extra extreme points of $\mathbf{S}_4$, we can easily identify the "pass-through" inequalities, for the hyperplane of an inequality passes through $\mathbf{S}_4 \backslash \mathbf{I}_4$ if and only if the inequality fails at one or more of the six extra extreme points of $\mathbf{S}_n$. (The inequality is guaranteed to hold by its construction at the 35 extreme points of $\mathbf{I}_n$.) It is easy to see that a "pass-through" inequalities cannot be written as a non-negative linear combination of Shannon-type information

inequalities. So, they are potentially valid non-Shannon-type information inequalities. For this reason, we can also use the software Information Theoretic Inequality Prover (ITIP) [29] to identify them. However, in this 4-variable case with $\mathbf{S}_4$ and $\mathbf{I}_4$ so explicitly represented, the above extreme point method seems more straightforward and intuitive. Using the extreme point procedure as described above, among the 122 candidate inequalities, we identify 50 such "pass-through" inequalities, or potentially valid non-Shannon-type inequalities. Owing to space limitations, we do not give a full account for them, but instead we give an example as follows:

$$\begin{aligned} &I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D) \\ &+I(A;B|D)+I(C;D|B)+I(C;D|A)\geq 0. \end{aligned} \tag{18}$$

3.3.3 Searching For Counterexamples

Because very little about entropy functions is understood, when we face a set of potentially valid non-Shannon-type information inequalities, it is unclear which ones are actually valid and which ones are not: we do not have a systematic way to construct counterexamples to invalidate a potential inequality nor a rich set of tools to prove it either. For example, a highly non-trivial counterexample, based on certain structures of projective planes, was constructed in [12] to demonstrate the invalidity of the Ingleton inequality. Currently all the published (unconditional) non-Shannon-type information inequalities (e.g. [15] and [13]) were proved essentially using the original method of Zhang and Yeung. We lack both proof and disproof techniques for non-Shannon-type inequalities. Even worse, we do not have an efficient way to discover them. Our approach here is an improvisation at best.

Facing the difficulties coming from both sides, we start from what we believe a relatively easy side, namely attempting to disprove a potentially non-Shannon-type information inequality by constructing counterexamples for it. Since we do

not have much theoretical support for constructing counterexamples, we organize a computer-aided search for counterexamples for the 50 potentially valid non-Shannon-type inequalities. But before launching the time-consuming search, we test these inequalities using the foregoing mentioned example constructed to disprove the Ingleton inequality. It turns out that none of them can be invalidated by the example.

Clearly, the value of an inequality at a 4-random variable set depends on its entropy vector, which can be computed from its joint distribution. Consider a set of 4 binary random variables (A, B, C, D). Its joint distribution is fully determined through a 16-dimensional *probability vector* of the form

$$(p_{0000}, p_{0001}, \cdots, p_{1111}), \tag{19}$$

where $p_{abcd} = \Pr\{A = a, B = b, C = c, D = d\}$, $a, b, c, d \in \{0, 1\}$, and $\sum_{a,b,c,d} p_{abcd} = 1$. Note that any non-negative 16-dimensional real vector can be turned into a probability vector by normalizing it. For example, the nonnegative vector $(0, 2, 0, 0, 0, 0, 0, 0, 1, 2, 0, 0, 0, 0, 0, 1)$ can be turned into the probability vector $(0, \frac{1}{3}, 0, 0, 0, 0, 0, 0, \frac{1}{6}, \frac{1}{3}, 0, 0, 0, 0, 0, \frac{1}{6})$. So, for convenience, we may refer to non-negative real vectors as probability vectors in the following when no confusion arises.

It is clear that when we confine ourselves to sets of 4 binary random variables, all the probability vectors form a simplex in $\mathbb{R}^{16}$. We aim to search a certain number of points belonging to this simplex. In other words, we need to discretize the simplex in some manner. However, we do not have an obvious way, say with respect to a particular information inequality in question, to do this discretization so that we can quickly find a counterexample to invalidate the inequality, given that the inequality is indeed invalid. On the other hand, an inequality that is valid for all the binary distributions is not guaranteed to be valid in general, because it

is possible that it may fail at some other distributions with sample spaces other than $\{0,1\}^4$.

We confine our project to the case with binary random variables because of the excessive computational burden incurred in higher dimensional cases. Consider the case of ternary random variables for example. In this case, the probability vector is of dimension $3^4 = 81$. It is clear that any brute-force enumeration of the points induced by a meaningful discretizing of this 81-dimensional probability simplex is computationally prohibitive.

As a first attempt, we search for counterexamples among the probability vectors from the following set

$$\{(p_{0000}, \cdots, p_{1111}) : p_{abcd} = 0, 1, 2, 3, \text{ for } a, b, c, d \in \{0, 1\}\}$$

The size of this set is 4^{16}. Note that each component of the probability vector from the above set takes 4 possible values, namely 0, 1, 2, and 3. We say that the components have "dynamic range" of 4. For each probability vector from the set, we compute its entropy vector and test the 50 candidate non-Shannon inequalities. The search is completed in a reasonable time, and 43 of them were invalidated in the end.

Let us call a probability vector that invalidates a candidate non-Shannon-type inequality a "witness" probability vector for the inequality and the entropy vector computed from the probability vector a "witness" entropy vector. For example, inequality (18) fails at the "witness" probability vector $(0, 0, 0, 1, 0.0.0, 1, 0, 0, 0, 2, 2.1, 2, 0)$. Owing to space limitations, we do not list here all the witness probability vectors for the inequalities found to be invalid.

At this point, we are left with 7 unknown inequalities.

We observe that all the remaining 7 inequalities have three extra conditional mutual information terms. This is not surprising, because those inequalities with

three extra conditional mutual information terms, which are always non-negative, are the hardest (among the 50 inequalities) to invalidate.

One consequence of the above observation is that we now know that *none of the extra conditional mutual information terms of the Zhang-Yeung inequality can be dropped for it to remain valid.*

By increasing the "dynamic range" of the components of probability vector from 4 to 5, we successfully invalidate 3 more from the remaining 7 inequalities. This search takes considerably longer than the previous one.

We see that by increasing the "dynamic range" of the components of probability vectors we get finer constructions of entropy vectors so that more inequalities are invalidated. However, we quickly hit the "exponential wall": a set with "dynamic range" of 5 is already as large as $5^{16} \approx 1.5 \times 10^{11}$. It becomes clear that searching larger probability vector spaces by uniformly increasing the "dynamic range" of their components does not seem to be a viable solution.

At the same time, we observe that all the "witness" probability vectors discovered share a pattern: roughly half (7–8) of the components of each "witness" probability vector are zero. By keeping the zero components of a "witness" probability vector fixed, we can afford to increase the "dynamic range" of the non-zero components while its zero-pattern is maintained. In particular, for a few selected "witness" probability vectors whose zero-patterns appear "typical," we increase the "dynamic range," from 5 to 10, of their non-zero components in the hope of finding "witness" probability vectors for the remaining 4 inequalities. Indeed, we find a number of "witness" probability vectors to invalidate one of the remaining 4 inequalities.

Thus, we are left with 3 unresolved potentially valid non-Shannon-type in-

equalities:

$$I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D) \\ +I(A;B|D)+I(A;C|D)+I(A;D|C)\geq 0, \tag{20}$$

$$I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D) \\ +I(A;B|D)+I(A;C|D)+I(B;C|D)\geq 0, \tag{21}$$

$$I(C;D|A)+I(C;D|B)+I(A;B)-I(C;D) \\ +I(A;C|D)+I(A;D|C)+I(B;C|D)\geq 0. \tag{22}$$

Unfortunately, we are unable to prove any of them either. Hence, they remain as plausible conjectures.

3.3.4 Expanding the Frontier Using "Witness" Entropy Vectors as Landmarks

In the above search, we get at least one "witness" entropy vector for each invalid inequality. The set of all "witness" entropy vectors together with the 35 extreme points of $\mathbf{I}_4$ define a convex cone in $\mathbb{R}^{15}$. By the convexity of the entropy function [16], the convex cone is contained in $\overline{\mathbf{H}}_4$. On the other hand, it contains $\mathbf{I}_4$ *strictly*, because each "witness" entropy vector lies outside of $\mathbf{I}_4$, for otherwise the inequalities it invalidates would hold at the "witness" entropy vector. Figuratively speaking, we have pushed back our "frontier," from the known boundary of $\mathbf{I}_4$, using the "witness" entropy vectors as landmarks, and got closer towards the unknown the boundary of $\overline{\mathbf{H}}_4$ from its interior. In other words, The convex cone constructed using these "witness" entropy vectors provides a better inner bound for $\overline{\mathbf{H}}_4$. It can also be used to construct better inner bounds for entropy cones $\overline{\mathbf{H}}_n$, $n > 4$, by representing it with a finite set of information inequalities. Since an improved inner bound for $\overline{\mathbf{H}}_n$ can, in principle, help to get better inner bounds for network coding capacities, the inner bound we have so obtained may find future

applications in improving the inner bound for network coding capacities upon those usually obtained using the Ingleton inequality alone.

3.3.5 Discussion

Owing to a lack of understanding of entropy functions, it is usually very hard to prove non-Shannon-type information inequalities. In the first place, we are in urgent need of methods to discover new non-Shannon-type information inequalities.

Our approach of constructing candidate non-Shannon-type information inequalities in this section is heuristic rather than systematic. The search we carry out is merely an experiment to familiarize us with the problem of characterizing the entropy cone $\overline{\Gamma}_4$.

So far, among the 50 candidate non-Shannon inequalities we investigate, three remain unresolved. As a direct consequence, we now know that none of the three extra conditional mutual information terms in the Zhang-Yeung inequality can be dropped for it to remain valid.

3.4 Random Variables Satisfying the Ingleton Inequality

In this section, we show that random variables mapped under group homomorphisms from a uniformly distributed background random variable satisfy the Ingleton inequality. As corollaries, we recover two previous known results. The first is that the network throughput of linear network codes is, in general, constrained by the Ingleton Inequality. The second and related result is that the network throughput of abelian-group network codes group network codes that are restricted to abelian groups– is also constrained by the Ingleton inequality.

With the aim of characterizing entropy functions, partially driven by a desire for better network codes and better bounds for network coding capacity, Chan and Yeung [30] first connected group theory to information inequalities. Remark-

ably, they established a one-to-one correspondence between information inequalities and group inequalities. Based on this theoretical finding, it was demonstrated that *group network codes* have the full potential to achieve network coding capacities [31]. In the same paper, it was claimed that abelian-group network codes group network codes that are restricted to abelian groups are in general not powerful enough to achieve network coding capacity based on the conclusion that the Ingleton inequality has to be honored for abelian-group network codes.

In this section, we attempt to characterize the random variables that satisfy the Ingleton inequality, i.e., the inverse image of the Ingleton cone $\mathcal{H}^{-1}(\mathbf{I}_n)$. See Fig. 3 for an illustration of the entropy function $\mathcal{H}_n$ mapping from the random variable space to the entropy space. We treat random variables as results of some intermediate transformations from a "simple" background random variable. In particular, we show that random variables mapped under group homomorphisms from a uniformly distributed background random variable satisfy the Ingleton inequality. As corollaries, we recover two previous known results. The first is that the network throughput of linear network codes is, in general, constrained by the Ingleton Inequality. The second and related result is that the network throughput of abelian-group network codes —group network codes that are restricted to abelian groups— is also constrained by the Ingleton inequality. Our group-theoretic treatment of random variables satisfying the Ingleton inequality provides a more general and unifying view of those two results.

We further make the following remarks. We consider, throughout this section, only discrete random variables on finite sample spaces and all the logarithms for entropies are taken to be base 2 unless mentioned otherwise. Groups are assumed to be multiplicative and 1 is taken to be the generic group unit.

The rest of the section is organized as follows: In Section 3.4.1, we state our

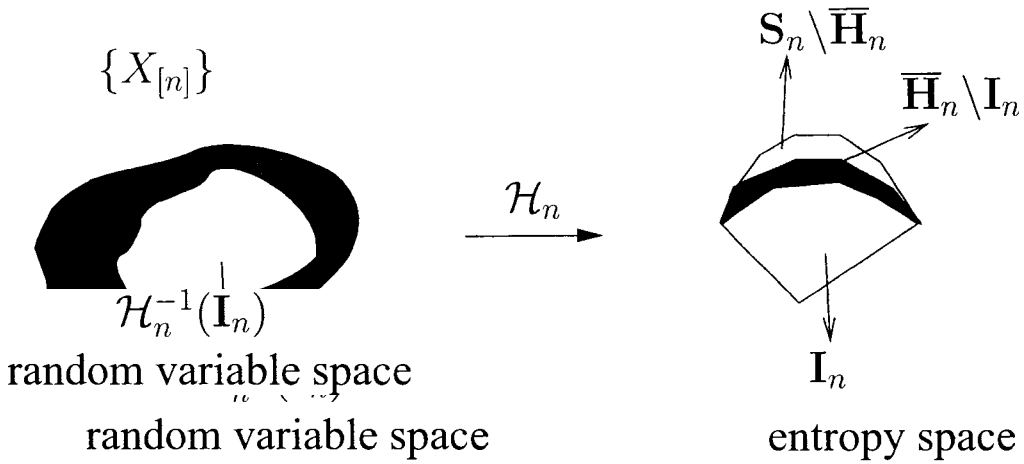

Figure 3. Mapping from random variable space to entropy space

main result and give a self-contained proof for it. In Section 3.4.2, we recover two known results as corollaries of the main result. We conclude the section in Section 3.4.3 with a discussion.

3.4.1 Group-homomorphism Random Variables Satisfy the Ingleton Inequality

Theorem 3.4.1. *Let G be a finite* group *and $\psi_1, \cdots, \psi_4$ group homomorphisms from G to groups $G'_1, \cdots, G'_4$, respectively. Let X be a random variable uniformly distributed on G. Then, for the random variables $\psi_1(X), \cdots, \psi_4(X)$, distributed on $G'_1, \cdots, G'_4$, respectively, the following inequality holds:*

$$\begin{aligned} &I(\psi_1(X); \psi_2(X)|\psi_3(X)) + I(\psi_1(X); \psi_2(X)|\psi_4(X)) \\ &+ I(\psi_3(X); \psi_4(X)) - I(\psi_1(X); \psi_2(X)) \geq 0. \end{aligned} \tag{23}$$

Inequality (23) is called the Ingleton inequality, first discovered for *rank functions* of *vector matroids* [19].

Remark: In general, $\psi_1, \cdots, \psi_4$ may not be surjective. Therefore, it is possible that for some $i = 1, \cdots, 4$ and $y \in G'_i$ $\Pr\{G'_i(X) = y\} = 0$. In this case, we take the convention that $0 \log 0 = 0$.

Before proving Theorem 3.4.1, we introduce the following simple lemma.

Lemma 3.4.2. *Let ψ be a group homomorphism from a finite group G to a group G' and X a random variable uniformly distributed on G. Then, the random variable $\psi(X)$ is uniformly distributed on the range $\psi(G)$ of ψ.*

Note that Lemma 3.4.2 generalizes Lemma 5.8 in [32].

Proof. It is a direct consequence of the Lagrange theorem: the collection of the pre-images of a homomorphism ψ, namely $\{\psi^{-1}(y) : y \in \psi(G)\}$, forms an equal partition of G. In particular, the sizes of the pre-images are all equal to $\ker \psi$. Therefore,

$$\Pr\{\psi(X) = y\} = \frac{|\ker \psi|}{|G|},$$

for all $y \in \psi(G)$. □

With Lemma 3.4.2 established, we are now ready to prove Theorem 3.4.1.

Proof of Theorem 3.4.1. Denote the kernel of a group homomorphism ψ by $\ker \psi$. By the *fundamental theorem on homomorphisms*, we have

- $\ker \psi$ is a *normal* subgroup of G, and
- $\psi(G)$ is isomorphic to the quotient group $G/\ker \psi$.

Furthermore, $|\psi(G)| = |G/\ker \psi| = \frac{|G|}{|\ker \psi|}$ holds. By Lemma 3.4.2, it is easy to see that $H(\psi(X)) = \log \frac{|G|}{|\ker \psi|}$.

Next, we calculate the joint entropy $H(\psi_1(X), \psi_2(X))$. We can treat (ψ_1, ψ_2) so formed as another group homomorphism from G to $G_1' \times G_2'$, where G_1' and G_2' are co-domains of the homomorphisms ψ_1 and ψ_2 respectively. It is easy to see that $\ker(\psi_1, \psi_2) = \ker \psi_1 \cap \ker \psi_2$, for $(\psi_1(g), \psi_2(g)) = 1$ holds if and only if $\psi_1(g) = 1$ and $\psi_2(g) = 1$ hold, $g \in G$. Therefore, we have

$$H(\psi_1(X), \psi_2(X)) = \log \frac{|G|}{|\ker \psi_1 \cap \ker \psi_2|}.$$

By induction, we have

$$H(\psi_1(X), \psi_2(X), \cdots, \psi_n(X)) = \log \frac{|G|}{\cap_{i=1}^{n} \ker \psi_i}.$$

Then, we calculate

$$I(\psi_1(X); \psi_2(X)) = \log \frac{|G|}{\frac{|\ker \psi_1||\ker \psi_2|}{|\ker \psi_1 \cap \ker \psi_2|}}$$

and

$$I(\psi_1(X); \psi_2(X)|\psi_3(X)) = \log \frac{|\ker \psi_1 \cap \ker \psi_2 \cap \ker \psi_3||\ker \psi_3|}{|\ker \psi_1 \cap \ker \psi_3||\ker \psi_2 \cap \ker \psi_3|}.$$

Thus, we can transform the Ingleton inequality (23) into an inequality on the cardinalities of the subgroups generated through intersection of the kernels $\ker \psi_1, \cdots, \ker \psi_4$.

Invoke the following inequality:

Inequality 3.4.3. *[20]*

$$2H(E|A) + 2H(E|B) + I(A; B|C) + I(A; B|D) + I(C; D) \geq H(E) \tag{24}$$

where A, B, C, D, and E are five random variables. This inequality was proved in [20]. We omit its proof here, for it is of Shannon-type and can be readily verified, manually or using the software Information Theoretic Inequality Prover (ITIP) [29].

Substitute A, B, C, and D in (24) with $\psi_1(X)$, $\psi_2(X)$, $\psi_3(X)$, and $\psi_4(X)$, respectively. We get

$$\begin{aligned} H(E) \leq\ & 2H(E|\psi_1(X)) + 2H(E|\psi_2(X)) \\ & + I(\psi_1(X); \psi_2(X)|\psi_3(X)) + I(\psi_1(X); \psi_2(X)|\psi_4(X)) \\ & + I(\psi_3(X); \psi_4(X)). \end{aligned} \tag{25}$$

Next, we construct random variable E. Recall that there is a one-to-one correspondence between group homomorphisms and normal subgroups via the

homomorphism-kernel relation. In other words, a homomorphism from group G to G' can be defined through specifying a normal subgroup of G as its kernel. From group theory, we know that the *direct product* of two normal subgroups is again a normal subgroup. Hence, the direct product $\ker\psi_1 \ker\psi_2$ is a normal subgroup, because both $\ker\psi_1$ and $\ker\psi_2$ are normal subgroups. Take $\ker\psi_1 \ker\psi_2$ as a kernel and let homomorphism $\psi_{1;2}$ be the homomorphism defined by $\ker\psi_1 \ker\psi_2$. Let E be $\psi_{1;2}(X)$.

In the following, we consider random variable E, or equivalently $\psi_{1;2}(X)$, and derive a few properties of it. First, we have

$$H(\psi_{1;2}(X)) = \log \frac{|G|}{|\ker\psi_1 \ker\psi_2|}.$$

We know from group theory that for any two subgroups K_1 and K_2, the following equality holds:

$$|K_1K_2| = \frac{|K_1||K_2|}{|K_1 \cap K_2|}.$$

It follows that

$$|\ker\psi_1 \ker\psi_2| = \frac{|\ker\psi_1||\ker\psi_2|}{|\ker\psi_1 \cap \ker\psi_2|}.$$

Therefore, we have

$$\begin{aligned} H(\psi_{1;2}(X)) &= \log \frac{|G|}{|\ker\psi_1 \ker\psi_2|} \\ &= \log \frac{|G|}{\frac{|\ker\psi_1||\ker\psi_2|}{|\ker\psi_1 \cap \ker\psi_2|}} \\ &= I(\psi_1(X); \psi_2(X)). \end{aligned}$$

Second, we show that random variable $\psi_{1;2}(X)$ *functionally* depends on the random variables $\psi_1(X)$ and $\psi_2(X)$. In particular, we show that $\psi_{1;2}$ *functionally* depends on $\psi_1(X)$ by demonstrating a homomorphism from $\psi_1(G)$ to $\psi_{1;2}(G)$. Then a similar result follows for $\psi_2(X)$. Note that $\ker\psi_1$ is a normal subgroup of G and a subgroup of $\ker\psi_{1;2} = \ker\psi_1 \ker\psi_2$. By the *fundamental theorem of*

homomorphisms, there exists a unique homomorphism $\phi : \psi_1(G) \to \psi_{1;2}(G)$ such that $\psi_{1;2} = \phi \circ \psi_1$, where "$\circ$" denotes homomorphism composition. In other words, we have $\psi_{1;2}(g) = \phi(\psi_1(g))$ for all $g \in G$. Hence, random variable $\psi_{1;2}(X)$, or E, functionally depends on random variable $\psi_1(X)$ through homomorphism ϕ. Consequently, we have $H(E|\psi_1(X)) = 0$ and, similarly, $H(E|\psi_2(X)) = 0$.

Substitute E in (25) with $\psi_{1;2}(X)$. We get

$$I(\psi_1(X);\psi_2(X)|\psi_3(X)) + I(\psi_1(X);\psi_2(X)|\psi_4(X)) + I(\psi_3(X);\psi_4(X)) - I(\psi_1(X);\psi_2(X)) \geq 0.$$

This concludes the proof. □

3.4.2 Two Corollaries

In this section, we recover two previously known results as corollaries of Theorem 3.4.1. The first one is on linear network codes.

By specializing Theorem 3.4.1 to vector spaces, we recover the original Ingleton inequality obtained for *rank functions* of vector matroids. Its information theoretical counterpart can be stated as follows.

Corollary 3.4.4. *Let F be a finite* field *and V a finite-dimensional vector space over F. Suppose $L_1, \cdots, L_4$ are four linear transformations from V to the vector spaces $V_1', \cdots, V_4'$ respectively. Let X be a random variable uniformly distributed on V. Then, for random variables $L_1(X), \cdots, L_4(X)$, distributed on $V_1', \cdots, V_4'$ respectively, the following information inequality holds:*

$$I(L_1(X);L_2(X)|L_3(X)) + I(L_1(X);L_2(X)|L_4(X)) + I(L_3(X);L_4(X)) - I(L_1(X);L_2(X)) \geq 0. \tag{26}$$

The result stated in Corollary 3.4.4 was implied in [20] in the context of showing that every information inequality is satisfied by rank functions of vector

matroids. In [33], Corollary 3.4.4 has been used to characterize the pre-images of Ingleton cones under entropy functions in a slightly different form.

Corollary 3.4.4 essentially shows that the network throughput of linear network codes is constrained by the Ingleton inequality. Therefore, non-linear network codes are, in general, expected to be able to provide better throughput.

From Theorem 3.4.1, another result on group inequalities, first stated in [31], can easily be recovered:

Corollary 3.4.5. *[31] Suppose G is an* abelian *group and $K_1, \cdots, K_4$ are subgroups of G. For any non-empty subset α of $\{1,2,3,4\}$, let $g_\alpha = \log \frac{|G|}{|\cap_{i\in\alpha} K_i|}$. Then*

$$g_{1,2} + g_{1,3} + g_{1,4} + g_{2,3} + g_{2,4} - g_1 - g_2 - g_{3,4} - g_{1,2,3} - g_{1,2,4} \geq 0. \qquad (27)$$

For abelian groups, all subgroups are *normal*. By the one-to-one correspondence between normal subgroups and group homomorphisms, subgroup $K_1, \cdots, K_4$ define four group homomorphisms by taking $K_1, \cdots, K_4$ as their kernels respectively. Then inequality (27) follows easily.

The implication of Corollary 3.4.5 is that abelian-group network codes, group network codes that are restricted to *abelian* groups, in general are not powerful enough to achieve network coding capacity.

3.4.3 Discussion

It is a difficult task to fully characterize the entropy function. The approach of information inequalities has proven to be successful in the cases involving 2 and 3 random variables. However, the case with 4 random variables remains far from clear so far. Finding and proving new information inequalities are notoriously hard. Currently, only techniques to construct and prove (unconditional) non-Shannon-type information inequalities are the original method [12] and its various generalizations. On the other hand, first studying entropy functions for some

special classes of random variables may be a more fruitful approach in the hope that finding more "conditional" laws may enable us to eventually discover global "unconditional" laws.

Thanks to [30], we now know that to study entropy functions, it is enough to confine ourselves to certain properties of subgroups. In fact, the subgroup characterization for entropy functions described therein is decisive. The group-theoretic conditions provided in this section are only sufficient conditions, and we do not have a full characterization for random variables satisfying the Ingleton inequality. This section shows that, to "leave" the relatively familiar region where the law of the Ingleton inequality applies, we need to turn our attention to *non-abelian* groups and *non-normal* subgroups of general groups.

3.5 A General Group-theoretic Condition for Ingleton Inequality

In this section, appealing to the fundamental relation we established previously in Chapter 2, we identify a general group-theoretic condition for the Ingleton inequality, subsuming all previously known conditions. Specifically, we show that quasi-Hamiltonian groups satisfy the Ingleton inequality. Quasi-Hamiltonian groups include as a subclass the well-known Hamiltonian groups, which are non-abelian. To our best knowledge, this is the first time that the Ingleton inequality is found to hold for certain classes of non-abelian groups.

This section is organized as follows: In Section 3.5.1, based on the approximation relation established in [30, 34] between group inequalities and information inequalities, we show that quasi-Hamiltonian groups satisfy the Ingleton inequality via a hybrid approach using results from both information theory and group theory. In Section 3.5.4, we briefly review certain features of quasi-Hamiltonian groups and outline a high-level characterization for them. The section is concluded by Section 3.5.5.

Before we close this section, a few remarks on our convention are in order. Throughout the section, groups are taken to be multiplicative. We write, the multiplication of two elements g_1, g_2 of a group G as $g_1 g_2$.

3.5.1 The Ingleton Inequality Holds for Quasi-Hamiltonian Groups

Thus far, partially driven by the desire for a better understanding of the limitation of various classes of network coding and partially driven by the goal of characterizing the range of the entropy function, various group-theoretic conditions have been obtained for the Ingleton inequality [6, 35, 36]. In particular, it was shown that abelian groups satisfy the Ingleton inequality. In this section, we advance in the same direction and identify a more general group-theoretic condition, subsuming all those previously obtained, for the Ingleton inequality. Specifically, we show that the Ingleton inequality holds for a class of groups called *quasi-Hamiltonian groups*. In addition to abelian groups, the class of quasi-Hamiltonian groups includes certain classes of non-abelian groups such as the well-known Hamiltonian groups. To our best knowledge, this is the first time that the Ingleton inequality is shown to hold for certain non-abelian groups. More usefully, quasi-Hamiltonian groups can be readily identified through the characterizations for nilpotent groups and modular groups.

Ingleton Inequality

The Ingleton inequality was first discovered by Ingleton [19] during his quest for characterizing vector representable matroids.

Proposition 3.5.1 (Ingleton inequality for vector matroids [19]). *Let V be a finite dimensional vector space and $A, B, C, D \subseteq V$ be four sets of vectors from V. Then*

the following inequality holds:

$$\begin{aligned} &\dim(A) + \dim(B) + \dim(A \cup B \cup C) \\ &\quad + \dim(A \cup B \cup D) + \dim(C \cup D) \\ &\leq \dim(A \cup B) + \dim(A \cup C) + \dim(A \cup D) \\ &\quad + \dim(B \cup C) + \dim(B \cup D), \end{aligned} \tag{28}$$

where $\dim(A \cup B \cup C)$ *denotes the dimension of the subspace spanned by the vectors in the union of* A, B, *and* C.

The information theoretic version of the Ingleton inequality is usually written in the following form:

$$\begin{aligned} &H(X_1) + H(X_2) + H(X_{123}) + H(X_{124}) + H(X_{34}) \\ &\leq H(X_{12}) + H(X_{13}) + H(X_{14}) + H(X_{23}) + H(X_{24}), \end{aligned} \tag{29}$$

where $H(X_{123})$ denotes the joint entropy $H(X_1, X_2, X_3)$. The inequality can be written more concisely in terms of (conditional) mutual information:

$$I(X_1; X_2|X_3) + I(X_1; X_2|X_4) + I(X_3; X_4) - I(X_1; X_2) \geq 0.$$

In [36], we established that the Ingleton inequality holds for group-homomorphism random variables, or equivalently for *normal subgroups.* As an immediate consequence, all abelian groups satisfy the Ingleton inequality. It is interesting to determine whether there are non-abelian groups satisfying the Ingleton inequality. From group theory, we know that all the subgroups of the so-called Hamiltonian groups are normal. Consequently, Hamiltonian groups satisfy the Ingleton inequality. It turns out, as we shall see in the sequel, that the Ingleton inequality holds for a more general class of groups called quasi-Hamiltonian groups.

3.5.2 Quasi-Hamiltonian Groups

Given two subgroups H and K of a group G, we call the set $\{hk : k \in K, h \in H\}$ the *product* of H and K, denoted HK. In general, $HK = KH$ fails to hold.

The following group-theoretic definitions are from [37].

Definition 3.5.2 (Mutually permutable subgroups). *We say that two subgroups H and K of a group G are* mutually permutable *if $HK = KH$ holds.*

Definition 3.5.3 (Quasinormal subgroups). *A subgroup of a group G is called* quasinormal *if it is mutually permutable with any other subgroup of G.*

Definition 3.5.4 (Quasi-Hamiltonian groups). *A group is called* quasi-Hamiltonian *if every of its subgroups are quasinormal.*

3.5.3 The Ingleton Inequality Holds for Quasi-Hamiltonian Groups

In the following, we provide a general condition in terms of permutabilities of the subgroups for the Ingleton inequality to hold. Then, we obtain, as a corollary, the desired result that the Ingleton inequality holds for quasi-Hamiltonian groups.

We use the square bracket notation $[n]$ to denote the generic index set $\{1, \cdots, n\}$. Give a set $\{G_i : i \in [n]\}$ of subgroups of a group G and a subset $\alpha \subseteq [n]$, we denote the intersection $\cap_{i\in\alpha} G_i$ by G_α. Similarly, given a set $\{m_i : i \in [n]\}$ of random variables and a subset $\alpha \in [n]$, we denote by m^α the joint random variable $(a_i : i \in \alpha)$.

Theorem 3.5.5. *Let G_1, G_2, G_3, and G_4 be four subgroups of a group G. If G_1 and G_2 are mutually permutable, then the following inequality holds:*

$$|G_1||G_2||G_{123}||G_{124}||G_{34}| \geq |G_{12}||G_{13}||G_{23}||G_{14}||G_{24}|. \tag{30}$$

Inequality (30) is the group version of the Ingleton inequality. It can be translated directly from the information theoretic version (29) by replacing each joint entropy term with the log-index of its corresponding intersection subgroups, according to the relation, shown in Proposition 3.5.8, between group inequalities and information inequalities

The proof of the theorem uses two simple facts from group theory [38, Prop. 3.13, 3.14] and the previously mentioned connection established between information inequalities and log-index inequalities for subgroups [30] and [34, Thm. 4]. They are collected as follows.

Proposition 3.5.6. *For two subgroups H and K of a group G, we have*

$$|HK| = \frac{|H||K|}{|H \cap K|}.$$

Proposition 3.5.7. *The product HK of two subgroups H and K of a group G is a subgroup if and only if $HK = KH$ holds, i.e., H and K are mutually permutable.*

Proposition 3.5.8. *Let $\mathbf{M} = \{m_i : i \in [n]\}$ be a set of n random variables and*

$$h_{\mathbf{M}} = \big(H(m^\alpha) : \alpha \subseteq [n]\big)$$

be the entropy vector *whose components are the joint entropies $H(m^\alpha)$, $\alpha \subseteq [n]$. Let $\mathbf{G} = \{G_i : i \in [n]\}$ be a set of n subgroups of a group G and*

$$l_{\mathbf{G}} = \big(\frac{1}{|G|} \log \frac{|G|}{|G_\alpha|} : \alpha \subseteq [n]\big)$$

be the scaled log-index vector *for the subgroups generated from the intersections of the subsets of subgroups. Suppose the components of the entropy vector $h_{\mathbf{M}}$ and those of the scaled log-index vector $l_{\mathbf{G}}$ are arranged in a same order. Let $f : \mathbb{R}^{2^{n+1}-1} \to \mathbb{R}$ be a continuous function. Then, the inequality $f(h_{\mathbf{M}}) \geq 0$ holds for all sets $\mathbf{M}$ of n random variables if and only if the inequality $f(l_{\mathbf{G}}) \geq 0$ holds for all sets $\mathbf{G}$ of n subgroups of any group.*

Note that Proposition 3.5.8 is an abridged version (only joint entropies involved) of Theorem 4 of [34]. We state such an abridged version because in the following proof we consider only joint entropies. Moreover, because information inequalities are linear, only linear functions are considered.

Proof. (of Theorem 3.5.5) We start with the following Shannon-type information inequality first proved in [20, Th. 8]:

$$2H(E|A) + 2H(E|B) + I(A;B|C) + I(A;B|D) + I(C;D) \geq H(E). \tag{31}$$

For completeness, we include its proof here. To see that the inequality holds, note that the following equality holds:

$$H(E) = H(E|A) + H(E|B) + I(A;B) - H(E|A,B) - I(A;B|E).$$

Then, we have

$$H(E) \leq H(E|A) + H(E|B) + I(A;B), \tag{32}$$

because both $H(E|A,B)$ and $I(A;B|E)$ are non-negative. Taking "conditional" on the both sides of the above inequality, we get

$$H(E|C) \leq H(E|A,C) + H(E|B,C) + I(A;B|C),$$

and

$$H(E|D) \leq H(E|A,D) + H(E|B,D) + I(A;B|D).$$

Combining the last two inequalities, we get

$$\begin{aligned} H(E|C) + H(E|D) \leq & H(E|A,C) + H(E|B,C) \\ & + I(A;B|C) + H(E|A,D) + H(E|B,D) + I(A;B|D). \end{aligned}$$

Then, we have

$$\begin{aligned} H(E|C) + H(E|D) + I(C;D) \leq & H(E|A,C) \\ & + H(E|B,C) + I(A;B|C) + H(E|A,D) \\ & + H(E|B,D) + I(A;B|D) + I(C;D). \end{aligned}$$

By (32), we have

$$H(E) \leq H(E|C) + H(E|D) + I(C;D).$$

Therefore, the following inequality holds:

$$\begin{aligned} H(E) \leq & H(E|A,C) + H(E|B,C) + I(A;B|C) \\ & + H(E|A,D) + H(E|B,D) + I(A;B|D) + I(C;D). \end{aligned}$$

Because of

$$\begin{aligned} H(E|A,C) &\leq H(E|A), \\ H(E|B,C) &\leq H(E|B), \\ H(E|A,D) &\leq H(E|A), \text{ and} \\ H(E|B,D) &\leq H(E|B), \end{aligned}$$

the above inequality can be simplified to

$$H(E) \leq 2H(E|A) + 2H(E|B) + I(A;B|C) + I(A;B|D) + I(C;D).$$

Thus, (31) holds.

Next, we write (31) in terms of joint entropies so that we can convert it to an inequality of log-indexes of subgroups:

$$\begin{aligned} & H(A,C) + H(B,C) - H(A,B,C) - H(C) \\ & \quad + H(A,D) + H(B,D) - H(A,B,D) - H(D) + H(C) \\ & \quad + H(D) - H(C,D) \\ & \geq H(E) - 2\big(H(E,A) - H(A)\big) \\ & \quad - 2\big(H(E,B) - H(B)\big). \end{aligned} \tag{33}$$

According to Proposition 3.5.8, we obtain the following group inequality holds for subgroups G_i, $i = [5]$, of any group G by replacing A with G_1, B with G_2, C with

G_3, D with G_4, and E with G_5:

$$\begin{aligned}
&\log\frac{|G|}{|G_1\cap G_3|}+\log\frac{G}{|G_2\cap G_3|}-\log\frac{|G|}{|G_1\cap G_2\cap G_3|}\\
&-\log\frac{|G|}{|G_3|}+\log\frac{|G|}{|G_1\cap G_4|}+\log\frac{G}{|G_2\cap G_4|}\\
&-\log\frac{|G|}{|G_1\cap G_2\cap G_4|}-\log\frac{|G|}{|G_4|}+\log\frac{|G|}{|G_3|}\\
&+\log\frac{|G|}{|G_4|}-\log\frac{|G|}{|G_3\cap G_4|}\\
&\geq\log\frac{|G|}{|G_5|}-2\log\frac{|G_1|}{|G_1\cap G_5|}-2\log\frac{|G_2|}{|G_2\cap G_5|}.
\end{aligned}\tag{34}$$

By assumption, $G_1G_2=G_2G_1$ holds. Consequently, according to Proposition 3.5.7, G_1G_2 is a subgroup of G. Choose $G_5=G_1G_2$. Then, we see $G_1\cap G_5=G_1$ and $G_2\cap G_5=G_2$. By Proposition 3.5.6, we have

$$\log\frac{|G|}{|G_5|}=\log\frac{|G|}{|G_1G_2|}=\log\frac{|G||G_1\cap G_2|}{|G_1||G_2|}=\log\frac{|G|}{|G_1|}+\log\frac{|G|}{|G_2|}-\log\frac{|G|}{|G_1\cap G_2|}.$$

Simplifying inequality (34) accordingly, we get

$$\begin{aligned}
&|G_1||G_2||G_1\cap G_2\cap G_3||G_1\cap G_2\cap G_4||G_3\cap G_4|\\
&\geq|G_1\cap G_2||G_1\cap G_3||G_2\cap G_3||G_1\cap G_4||G_2\cap G_4|.
\end{aligned}$$

In short notation, it is written as

$$|G_1||G_2||G_{123}||G_{124}||G_{34}|\geq|G_{12}||G_{13}||G_{23}||G_{14}||G_{24}|$$

This concludes the proof. □

As a consequence of Theorem 3.5.5, the Ingleton inequality holds for any group whose subgroups are all quasinormal. Such groups are simply the *quasi-Hamiltonian* groups.

Corollary 3.5.9. *If G_1, G_2, G_3, and G_4 are four subgroups of a quasi-Hamiltonian group G, then the following inequality holds:*

$$|G_1||G_2||G_{123}||G_{124}||G_{34}|\geq|G_{12}||G_{13}||G_{23}||G_{14}||G_{24}|.$$

3.5.4 On Quasi-Hamiltonian groups

In this section, we provide a brief overview of quasi-Hamiltonian groups and an outline for identifying them via the well-established characterizations for nilpotent groups and modular groups.

Recall that all the subgroups of a group G forms a lattice. This lattice is usually called the subgroup lattice of G. Note that subgroup lattices thus defined are different from those defined in [34] that are generated from a set of subgroups of a group.

Definition 3.5.10. *A subgroup of G is called* modular *if it is modular on the subgroup lattice of G. Accordingly, a group G is called* modular *if all of its subgroups are modular.*

Proposition 3.5.11. *[37. Th. 5.1.1] A subgroup is* quasinormal *if and only if it is modular and* subnormal.

Since finite groups whose subgroups are all subnormal are exactly *nilpotent groups*, we have the following characterization for quasi-Hamiltonian groups.

Proposition 3.5.12. *A group is quasi-Hamiltonian if and only if it is nilpotent and modular.*

The class of quasi-Hamiltonian includes Hamiltonian groups as a subclass. Furthermore, it was shown that there exist non-Hamiltonian quasi-Hamiltonian groups [37]. Group theorists have fully characterized both nilpotent groups and modular groups. Specifically, a group is nilpotent if and only if it is the direct product of its Sylow subgroups [38]. Modular groups have been fully characterized as well, but the characterization is more complicated than that of nilpotent groups. See [37, Ch. 2] for a detailed account.

3.5.5 Discussion

In this section, we generalize the previously known conditions for the Ingleton inequality. Specifically, we show that the Ingleton inequality holds for a class of groups called *quasi-Hamiltonian* groups. This condition subsumes all those previously known for the Ingleton inequality. Besides abelian groups, the class of quasi-Hamiltonian groups includes as a subclass Hamiltonian groups, which are well known to be non-abelian. To our best knowledge, this is the first time that certain classes of non-abelian groups are found to satisfy the Ingleton inequality. More usefully, we point out that quasi-Hamiltonian groups can be identified via the well-established characterizations for nilpotent groups and modular groups.

By identifying more general conditions for the Ingleton inequality, we provide better guidelines for designing network codes that aim to achieve better throughput, since in general the performance of the network codes subject to certain constraints may potentially fail to achieve the maximum throughput (network coding capacity) that can be achieved by general unconstrained codes. According to the conditions obtained in this section for the Ingleton inequality, similar to linear codes and abelian group codes, group network codes limited to quasi-Hamiltonian groups are bounded to be constrained by the Ingleton inequality, and hence they are expected to potentially suffer, on general multi-source, multi-sink networks, the same throughput inefficiency as linear codes and abelian group codes.

List of References

[1] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Transactions on Information theory*, vol. 46, no. 4, pp. 1204 1216, July 2000.

[2] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Transactions on Information theory*, vol. 49, no. 2, pp. 371 381, Feb. 2003.

[3] E. Erez and M. Feder, "Capacity region and network codes for two receivers multicast with private and common data," in *Workshop on Coding, Cryptog-*

raphy and Combinatorics, 2003.

[4] C. K. Ngai and R. W. Yeung, "Multisource network coding with two sinks," in *Communications, Circuits and Systems, 2004. ICCCAS 2004. 2004 International Conference on*, 2004.

[5] A. Ramamoorthy and R. D. Wesel, "The single source two terminal network with network coding," in *Canadian Workshop on Information Theory*, 2005.

[6] R. W. Yeung, *A First Course in Information Theory.* Kluwer Academic/Plenum Publishers, 2002.

[7] G. Kramer and S. A. Savari, "Edge-cut bounds on network coding rates," *Jornal of Network and Systems Management*, vol. 14, no. 1, pp. 49 67, March 2006.

[8] N. J. A. Harvey, R. Kleinberg, and A. R. Lehman, "On the capacity of information networks," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2345–2364, June 2006.

[9] X. Yan, R. W. Yeung, and Z. Zhang, "The capacity region for multi-source multi-sink network coding," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 2007, pp. 116 120.

[10] N. Pippenger, "What are the laws of information theory," in *1986 Special Problems on Communication and Computation Conference*, Palo Alto, California, Sept. 3-5 1986.

[11] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379 423 and 623 656, July and October 1948.

[12] Z. Zhang and R. W. Yeung, "On characterization of entropy function via information inequalities," *IEEE Transactions on Information Theory*, vol. 44, no. 4, pp. 1440 1452, July 1998.

[13] R. Dougherty, C. Freiling, and K. Zeger, "Six new non-Shannon information inequalities," in *Proceedings of the 2006 IEEE International Symposium on Information Theory*, 2006, pp. 233 236.

[14] R. Lněnička, "On the tightness of the Zhang-Yeung inequality for Gaussian vectors," *Communications in Information and Systems*, vol. 3, no. 1, pp. 41 46, June 2003.

[15] K. Makarychev, Y. Makarychev, A. Romashchenko, and N. Vereshchagin, "A new class of non-Shannon-type inequalities for entropies," *Communications in Information and Systems*, vol. 2, no. 2, pp. 147 166, December 2002.

[16] Z. Zhang and R. W. Yeung, "A non-Shannon-type conditional inequality of information quantities," *IEEE Transactions on Information theory*, vol. 43, no. 6, pp. 1982 1986, Nov. 1997.

[17] F. Matúš, "Piecewise linear conditional information inequality," *IEEE Transactions on Information Theory*, vol. 52, no. 1, pp. 236 238, Jan. 2006.

[18] D. J. A. Welsh, *Matroid Theory*. New York: Academic, 1976.

[19] A. W. Ingleton, "Representation of matroids," in *Combinatorial mathematics and its applications*, D. Welsh, Ed. London: Academic Press, 1971, pp. 149 167.

[20] D. Hammer, A. Romashchenko, A. Shen, and N. Vereshchagin, "Inequalities for Shannon entropy and Kolmogorov complexity," *Journal of Computer and System Sciences*, vol. 60, no. 2, pp. 442 464, April 2000.

[21] R. Dougherty, C. Freiling, and K. Zeger, "Insufficiency of linear coding in network information flow," *IEEE Transactions on Information theory*, vol. 51, no. 8, pp. 2745 2579, August 2005.

[22] T. Chan and A. Grant, "Entropy vectors and network codes." [Online]. Available: http://arxiv.org/abs/cs.IT/0702063v1

[23] R. W. Yeung, "A new outlook on Shannon's information measures," *IEEE Transactions on Information theory*, vol. 37, pp. 466 474, 1991.

[24] R. W. Yeung, "A framework for linear information inequalities," *IEEE Transactions on Information theory*, vol. 43, pp. 1924 1934, 1997.

[25] F. Matúš, "Infinitely many information inequalities," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 24-29 2007, pp. 41 44.

[26] F. Matúš and M. Studeny, "Conditional independences among four random variables I," *Combinatorics, Probability & Computing*, vol. 4, pp. 269 278, 1995.

[27] T. H. Chan, "Balanced information inequalities," *IEEE Transactions on Information theory*, vol. 49, pp. 3261 3267, 2003.

[28] G. M. Ziegler, *Lectures on Polytopes*, J. Ewing, F. Gehring, and P. Halmos, Eds. Springer-Verlag, 1995.

[29] R. W. Yeung and Y.-O. Yan, "Information theoretic inequality prover." [Online]. Available: http://home.ie.cuhk.edu.hk/ ITIP/

[30] T. H. Chan and R. W. Yeung, "On a relation between information inequalities and group theory," *IEEE Transactions on Information Theory*, vol. 48, no. 7, pp. 1992 1995, July 2002.

[31] T. H. Chan, "On the optimality of group network codes," in *Proceedings. International Symposium on Information Theory*, 2005.

[32] R. Dougherty, C. Freiling, and K. Zeger, "Matroids, networks, and non-Shannon information inequalities," *IEEE Transactions on Information theory*, submitted.

[33] H. Li and E. K. P. Chong, "Disproving certain potentially valid non-Shannon-type information inequalities," manuscript.

[34] H. Li and E. K. P. Chong, "Information lattices and subgroup lattices: Isomorphisms and approximations," manuscript, submitted.

[35] T. H. Chan, "Group characterizable entropy functions," in *2007 IEEE International Symposium on Information Theory*, submitted.

[36] H. Li and E. K. P. Chong, "On connections between group homomorphisms and the Ingleton inequality," in *Proceedings of the 2007 IEEE International Symposium on Information Theory*, Nice, France, June 24 29 2007, pp. 1996 2000.

[37] R. Schmidt, *Subgroup Lattices of Groups*. Walter De Gruyter, 1994.

[38] D. S. Dummit and R. M. Foote, *Abstract Algebra*, 3rd ed. Wiley, 2003.

CHAPTER 4

Search on Lines and Graphs

4.1 Summary

In this work we investigate discrete linear search and graph search problems. We first formulate the Bounded Discrete Linear Search Problem (BDLSP) as a Markov Decision Problem (MDP) and show that BDLSPs admit strongly polynomial-time solutions. In contrast, we show that the Graph Search Problem (GSP) is NP-complete by revealing the equivalence between GSP and the Weighted Minimum Latency Problem (WMLP). We then consider the Erroneous BDLSP (EBDLSP) and obtain lower and upper bounds on the optimal cost in terms of that of it error-free counterpart BDLSP. In the second part of the work, we investigate the Unbounded Discrete Linear Search Problem (UBDLSP). We first establish that for an optimal policy to exist for a general UBDLSP it is both necessary and sufficient for its double-sided mean of the underlying distribution to be finite. Then, we consider a special class of UBDLSPs—symmetric UBDLSPs and prove the expanding property of optimal policies for symmetric UBDLSPs. Based on the expanding property, we devise a procedure to approximate, by solving a sequence of finite-truncated BDLSPs, the optimal costs and the optimal policies for symmetric UBDLSPs. We prove that the sequence of approximated optimal costs converges to the true optimal cost and that if the optimal policy is unique, then the sequence of partial policies converges to the true optimal policy. Furthermore, we investigate the increments of the turning-point sequence of the optimal policy for symmetric UBDLSPs with heavy-tailed distributions. It turns out that, in contrast to the boundness result obtained earlier for the typical thin-tailed distribution (Gaussian distribution), the increment sequence for heavy-tailed distributions is necessarily unbounded.

4.2 Introduction

Using minimum effort to locate an item on a line or over an area is a ubiquitous problem. For example, in the linear case, an item is assumed to be located on a number of possible positions with certain probabilities and a searcher aims to find the item with the minimum expected traveling distance. It turns out that such seemingly simple problems are generally hard. To provide some background to our work on these problems, we review in the following a history of the search problem.

4.2.1 War-time Efforts

During the Second World War, the operations research pioneer Koopman and his collaborators [1] first started a systematic and extensive investigation on the search problem. The angle they took then was to devise algorithms maximizing the success probability subject to a constraint on the search effort. Results obtained [1] were kept confidential until the mid fifties and published in a series of three papers [2–4]. This line of research was collected into a book by Stone [5], a Lanchester prize winner. See [6] for a summary of recent developments taking this approach.

4.2.2 The Continuous Linear Search Problem

In 1963, another operations research pioneer Bellman [7] posed the following carefully formulated search problem in *SIAM Review*:

> Suppose that we know that a particle is located in the interval $(x, x+dx)$ somewhere along the real line $-\infty < x < \infty$ with a probability density function $g(x)$. We start at some initial point x_0 and can move in either direction. What policy minimizes the expected time required to find the particle, assuming a uniform velocity and
>
> (a) assuming that the particle will be recognized when we pass x, or
>
> (b) assuming that there is a probability $p > 0$ of missing the particle as we go past it?
>
> Also, what would be the optimum start point x_0?

The same problem was independently formulated by Beck [8] around the same time:

> The senior author (Beck) introduced the problem in 1963 by private communication at the same time it was independently proposed by Richard Bellman in the *SIAM Review*.

After Bellman's formulation, Franck started working on this problem, partly as the topic of his doctoral dissertation, and two years later published the first paper [9] on this problem. On the other hand, Beck and his collaborators started a four-decade quest towards understanding the problem [8, 10–16]. In [10], Beck coined the name for the problem — *linear search problem* (LSP). The journey proved to be rough — even characterizing the condition for an optimal policy to exist resulted in two slightly different outcomes [9, 10]. The condition offered by Franck was followed by an immediate correction [17] and one counterexample twenty years later [18]. This seemingly simple problem turned out to be surprisingly difficult—both analytically and numerically.

A colorful version of the LSP, under the name of the "Gaussian cookie problem," surfaced in the discussion of the sci.math newsgroup around the late eighties, even after Rousseeuw [19] attempted to numerically approximate the solutions to the LSPs with various well-known distributions such as normal, student, logistic, and Laplace distributions. As it was shown in [9], if distribution functions are "nice," the turning-points of optimal policies admit an iterative relation in the form of Equation (35). In particular, suppose we write the turning-point sequence of an optimal policy as:

$$(\cdots, a_{2i}, a_{2i-2}, \cdots, a_2, 0, a_1, a_3, \cdots, a_{2i-1}, a_{2i+1}, \cdots).$$

If the cumulative distribution function F is absolutely continuous in some neighborhood of turning-point a_i and if the density function f is continuous at a_i, then

the following iterative relation holds:

$$|a_{i+1}| = -|a_i| + \frac{\Pr\{X \notin [a_{i-1}, a_i]\}}{f(a_i)}. \tag{35}$$

However, the crux is the first turning-point, a_1. In [19], Rousseeuw's efforts were mainly devoted to approximating this first turning-point. For normal distributions, the first turning-point is approximated around 1.44084, resembling none of the well-known constants such as e, π, and $\sqrt{2}$. Three years later, Beck and Beck [8] tried numerical approaches to the problem as well and noticed that their approximation algorithm was numerically unstable. But, in a following paper [14], they managed to prove the convergence of their algorithm.

In 1974, Fristedt and Heath considered the problem with various generalized cost functions [20]. In 1987, Balkhi [21] considered the problem of choosing optimal starting points for LSPs, the final question in Bellman's problem statement. Bruss and Robertson, in 1988, published a comprehensive survey on the problem [22]. In [22], the dynamic-programming approach was first proposed. After the survey, Beck and his collaborators published two more papers [15, 16] addressing the LSP with generalized cost functions. In the middle of the nineties, Washburn [23] started to consider a variant of the LSP called the backpacker's linear search problem and proposed a dynamic-programming approach to the discrete version of the problem.

It is worth pointing out that Rousseeuw has created an entry dedicated to the linear search problem in the *Online Encyclopedia of Mathematics from SpringerLink* [24].

In this work, we focus on the discrete version of the LSP. In Section 4.2.2, we review the history of its continuous counterpart partly to provide a background to the linear search problem in general and partly to introduce certain sources from which we shall use results in the sequel.

In the rest of this introductory section, we introduce, respectively, the *bounded discrete linear search problem* (BDLSP), the *unbounded discrete linear search problem* (UBDLSP), and the *graph search problem* (GSP). For the purpose of bridging the relevant research of the two relatively segregated communities the operations research community and the computer science community we give a comprehensive review of the related work from both communities.

4.2.3 The Discrete Linear Search Problem BDLSP

Suppose that a target is known to a searcher to be located at one of $2n+1$ possible locations on

$$[-n, n] = \{-n, -n+1, \cdots, 1, 0, 1, \cdots . n-1, n\},$$

and that with probability p_i the target is at location i, $i = [-n, n]$. (In the following, we use "location" and "position" interchangeably.) The searcher starts with some initial location and travels on the line segment $[-n, n]$ to find the target (with no jumps allowed). The presence of the target can be detected only when the searcher visits the location of the target. Our goal is to devise an *optimal search policy* for the searcher to find the target with minimum traveled distance. Specifically, for a given policy π, if we denote the traveled distance under π to first visit a location i by $l^{\pi}(i)$, then the expected distance traveled c^{π} of the policy π (to find the target) can be computed as:

$$c^{\pi} := \sum_{i \in [-n,n]} l^{\pi}(i) p_i.$$

A policy π^* is called *optimal* if $c^{\pi^*} \leq c^{\pi}$ for all π. We call c^{π} the *value* of π and the minimum value

$$c^* := c^{\pi^*}$$

the *optimal value*.

UBDLSP

In the above BDLSP, the search domain is bounded. Assuming, instead, a probability distribution with unbounded support, for example $\{p_i, i \in \mathbb{Z}\}$, we convert the BDLSP to an *unbounded discrete linear search problem* (UBDLSP). The value of a policy π is now calculated as an infinite sum:

$$c^{\pi} := \sum_{i \in \mathbb{Z}} p_i l^{\pi}(i).$$

Similarly, a policy π^* is called *optimal* if $c^{\pi^*} \leq c^{\pi}$ for all π, and the minimum value

$$c^* := c^{\pi^*}$$

the *optimal value*.

UBDLSP with Worst-case Cost Criteria

Kao and Littman [25] first considered the BDLSP, under the name of "informed cow-path problem." On the other hand, computer scientists started to consider unbounded search problems [26] in the seventies. The unbounded linear search problem was first considered by Baeza-Yates et al. and Kao et al. [27, 28]. However, they formulated the problem without assuming a probability distribution on the target locations but sought policies with best worst-case *competitive ratio*. In other words, the performance of a policy π is measured by its worst-cast competitive ratio, defined as

$$r^{\pi} = \limsup_{n} \frac{l^{\pi}(n)}{|n|},$$

This min-max or game theoretic formulation is typical in the computer science community. In fact, the continuous linear search problem thus formulated was studied as early as in the seventies by Gal [29], as the topic of his doctoral dissertation, and by Beck and Newman in [12]. A number of results were obtained [30 32] and accumulated into a book [33]. In particular, the so-called *doubling policy* was

discovered to be optimal with the best competitive ratio 9. Historically, this result was first discovered by Beck and Newman in [12] and by Gal [31], and rediscovered later by Baeza-Yates [27] and Kao et al. [28].

In fact, in [25] under the name of "informed cow-path problem," Kao and Littman were mainly concerned with the bounded discrete linear search problem with the target location distribution $\mathbf{p}$ known to belong to a set $\mathbf{P} = \{\mathbf{p}^i : i \in [k]\}$ (throughout the chapter, we use the bracket notation $[k]$ to denote the generic index set $\{1, 2, \cdots, k\}$) of possible distributions, and they aimed there to find optimal policies π^* with the minimum worst-cast cost, namely,

$$\pi^* = \operatorname*{argmin}_{\pi} \max_{\mathbf{p} \in \mathbf{P}} c^{\pi}(\mathbf{p}),$$

where $c^{\pi}(\mathbf{p})$ is the cost of policy π when the target location distribution is taken to be $\mathbf{p}$. The bounded discrete linear search problem with a *fixed* probability distribution, i.e., the BDLSP as formulated in Section 4.2.3, was investigated there as a stepping-stone.

4.2.4 Search on a Plane

The unbounded linear search problem was subsequently generalized to the cases with multiple rays [28, 34, 35] and to various two-dimensional scenarios [27, 36, 37]. In fact, a version of the two-dimensional search problem was proposed as an editorial note to the original Bellman linear search problem as follows [7]:

> A related class of two dimension search problem are the following "swimming in a fog" problems. A person has been shipwrecked in a fog and wishes to determine the optimal path of swimming to get to shore (in the least expected time—assuming a uniform rate of swimming). The boundary conditions can be any of the following:
>
> 1. The ocean is a half-plane,
> 2. Condition (1) plus the knowledge that the initial distance to shore is $\leq D$ (with a uniform distribution),
> 3. The ocean boundary is a given curve, i.e., a circle, rectangle, or possibly not closed (a parabola),
> 4. Condition (2) and (3), etc.

Two years later, Bellman himself proposed the following two-dimensional search problem in the *Bulletin of the American Mathematical Society* [38]:

> We are given a region R and a random point P within the region. Determine the paths which
>
> (a) Minimize the expected time to reach the boundary, or
>
> (b) Minimize the maximum time required to reach the boundary.
>
> Consider, in particular, the cases
>
> (a) R is the region between two parallel lines at a known distance d apart.
>
> (b) R is the semi-infinite plane and we are given the distance d from the point P to the bounding line.

The two dimensional search problem is usually rephrased more vividly as the "lost in a forest" or "lost at sea" problem. Most researchers considered the problem with respect to the min-max criterion and sought for the best worst-case "escape policy." The problem in its full generality is deemed to be still open [39] and was included by Williams in his recent list of "Million Buck Problems" [40]. Nonetheless, several special cases have been solved along the way. The case with "an infinite-long straight strip with known width," the first case proposed in Bellman's formulation above, was solved by Zalgaller [41] in 1961. In 1957 Isbell [42] solved the second "half-plane with known distance" case.

It was recently realized that the two dimensional search problem is closely related to Moser's well-known "worm problem" [43, 44]. Using a result from the research on the "worm problem," Finch and Wetzel showed that if regions are of "flat" shapes, then best escape policies are line segments [39]. "Flat" regions include rectangles, regular n-gons ($n > 3$), and circular sectors with angle $\theta \geq \frac{\pi}{3}$. Surprisingly, the seemingly humble triangle case remains open. See [39] for a recent survey.

Little is known for the general unbounded and uninformed cases such as search for a line on a plane with unknown distance and slope. Baeza-Yates et al. conjectured that a logarithmic spiral is optimal to find a line on a plane at an unknown distance away and with unknown slope [27]. See [45] for a discussion on the conjecture.

Little work has been done on the two-dimensional search problem with respect to the expected escape distance criterion. Even less is known on the cases with missing probabilities.

4.2.5 Organization and Contributions of the Chapter

Section 4.3 is dedicated to two bounded discrete search problems: the BDLSP and the GSP. We formulate the BDLSP as a Markov Decision Problem (MDP) and show that it admits a strongly polynomial-time solution even when the searcher is required to optimize the starting location. Then, we generalize the BDLSP to the *Graph Search Problem* (GSP). It turns out the GSP is much harder than the BDLSP – it is *NP-complete*.

In Section 4.4, we consider the erroneous BDLSP (EBDLSP) where the target is missed with a certain probability whenever it is visited. We bound the optimal value of the EBDLSP in terms of that of its error-free counterpart BDLSP and show that if the missing probabilities are very small, or search is much less expensive than traveling, then the policy optimal for the corresponding BDLSP is almost optimal for the EBDLSP.

Section 4.5 is dedicated to the UBDLSP. We first provide a necessary and sufficient condition for an optimal policy to exist and show an expanding property of optimal policies for symmetric UBDLSPs. Then, using the *principle of optimality*, we show that the optimal values of symmetric UBDLSPs can be successively approximated by solving sequences of BDLSPs and that the optimal policy can be

similarly approximated as long as the optimal policy is unique (up to symmetric). At the end of the section, we study the growth rate of the turning-points of optimal policies for symmetric UBDLSPs. The growth rate is a factor important to the efficiency of our approximation method. It turns out that the growth rate of the optimal policy for heavy-tailed UBDLSPs are at least *superlinear*. This result suggests that the optimal values of heavy-tailed UBDLSPs are in general hard to approximate efficiently.

In Section 4.6, we discuss future directions on search problems and conclude the chapter.

4.3 Search on Bounded Lines and Graphs

In this section, we study two bounded discrete linear search problems the BDLSP and the GSP. We first formulate the BDLSP as an MDP and show that optimal policies for BDLSPs can be found in strongly polynomial-time. Next, we generalize the problem and consider the GSP. This multi-dimensional search problem can be formulated as an MDP as well, but with an exponentially larger state space. It turns out that the exponentially larger state space of the MDP formulation is most likely impossible to simplify we show that the GSP is NP-complete.

4.3.1 BDLSP

We can see through the following two examples, as shown in Fig. 4 and 5, that optimal policies for BDLSPs are numerically "sensitive" to the underline distribution.

In Fig. 4, the underlying distribution $(p_{-3}, p_{-2}, p_{-1}, p_0, p_1, p_2, p_3) = (0.1, 0.25.0.05, 0.4, 0.05, 0.10, 0.05)$ and the optimal policy is to start from location 0, walk left up to the left end, and then turn around to search the other

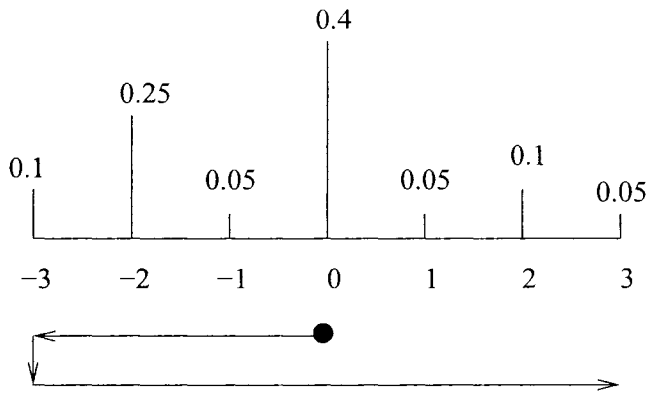

Figure 4. BDLSP example 1

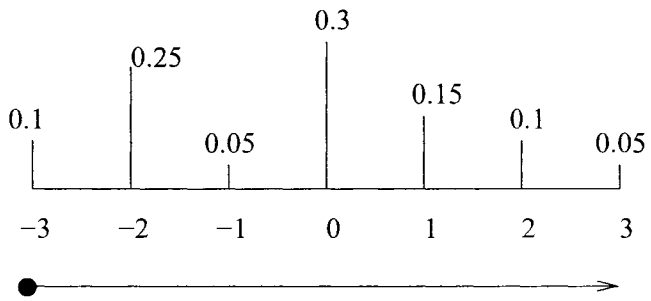

Figure 5. BDLSP example 2

side. With the distribution slightly changed to $(p_{-3}, p_{-2}, p_{-1}, p_0, p_1, p_2, p_3) = (0.1, 0.25.0.05, 0.3, 0.15, 0.10, 0.05)$ as shown in Fig. 5, the searcher should start from the left-end and walk straight towards the right-end to be optimal.

MDP Basics

An MDP is typically specified by four ingredients—a state space $\mathcal{S}$, an action space $\mathcal{A}$, a state-transition probability law $p^{\alpha}_{s_k,s_{k+1}}$, $s_k, s_{k+1} \in \mathcal{S}$ and $\alpha \in \mathcal{A}$, and a (one-step) cost function $g(s_k, \alpha)$. At a typical stage k, the system is at state s_k, an action α is taken, and a cost of $g(s_k, \alpha)$ is incurred. Subsequently, the system transits to a next state s_{k+1} with probability $p^{\alpha}_{s_k,s_{k+1}}$ according to the state-transition law.

A policy π is usually defined to be a sequence $\{u_k\}$ of state to action mappings $u_k : \mathcal{S} \to \mathcal{A}$. When a policy is given, the system evolves as a Markov chain with actions taken, stage by stage, according to the policy. Suppose the MDP is of finite horizon $[1, K]$. For a fixed policy π, the cost-to-go at a generic stage k is given by

$$J_k^\pi(s_k) = \mathrm{E}\left[\sum_{i=k}^{K} g(s_i, u(s_i))\middle| s_k\right],$$

and is a function of the current state s_k. The cost-to-go function satisfies the following recursive formula:

$$J_k^\pi(s_k) = g\big(s_k, u(s_k)\big) + \sum_{s_{k+1}} p_{s_k, s_{k+1}}^{u_k(s_k)} J_{k+1}^\pi(s_{k+1}). \tag{36}$$

The total cost incurred for a given policy π depends on the initial state and is hence denoted by $J_0^\pi(s_0)$. A policy π^* is called *optimal* for s_0 if $J_0^{\pi^*}(s_0) = \inf_\pi J_0(\pi, s_0)$.

The recursive formula (36) for cost-to-go functions suggests the following recursive formulas characterizing the optimal policies:

$$J_k^{\pi^*}(s_k) = \min_{u_k(s_k)} \left[g\big(s_k, u_k(s_k)\big) + \sum_{s_{k+1}} p_{s_k, s_{k+1}}^{u_k(s_k)} J_{k+1}^{\pi^*}(s_{k+1})\right], \text{ for all } k \text{ and } s_k, \tag{37}$$

and

$$u_k^*(s_k) = \operatorname*{argmin}_{\alpha \in \mathcal{A}} \left[g\big(s_k, \alpha\big) + \sum_{s_{k+1}} p_{s_k, s_{k+1}}^{\alpha} J_{k+1}^{\pi^*}(s_{k+1})\right], \text{ for all } k \text{ and } s_k, \tag{38}$$

where the sequence $\{u_k^*\}_{k=0}^K$ constitutes the optimal policy π^*.

The recursive formulas (37) and (38), usually called Bellman equations, are the cornerstones of MDP theory.

The consequence of Bellman equations is the following *dynamic-programming algorithm*. Suppose the state space, the action space, and the horizon of the MDP in question are all finite. Then Equations (37) and (38) imply the following sequences of "rules" to choose optimal actions, running backwards from the terminal stage K to the initial stage 0.

- Stage $K-1$: choose an action $u^*_{K-1}(s_{K-1})$ for each possible penultimate state s_{K-1} to minimize $g(s_{K-1}, \alpha)$. The optimal cost-to-go at stage $K-1$ for each possible state s_{K-1} is given by $J^{\pi^*}_{K-1}(s_{K-1}) = g(s_{K-1}, u^*_{K-1}(s_{K-1}))$.

- Stage $K-2$: choose an action $u^*_{K-2}(s_{K-2})$ for each possible state s_{K-2} to minimize

$$g(s_{K-2}, \alpha) + \sum_{s_{K-1}} p^{\alpha}_{s_{K-2}, s_{K-1}} J^{\pi^*}_{K-1}(s_{K-1})$$

In other words, we choose the action

$$u^*(s_{K-2}) = \underset{\alpha \in \mathcal{A}}{\operatorname{argmin}} \left[g(s_{K-2}, \alpha) + \sum_{s_{K-1}} p^{\alpha}_{s_{K-2}, s_{K-1}} J^{\pi^*}_{K-1}(s_{K-1}) \right],$$

and the optimal cost-to-go is given by

$$J^{\pi^*}_{K-2}(s_{K-2}) = \min_{\alpha \in \mathcal{A}} \left[g(s_{K-2}, \alpha) + \sum_{s_{K-1}} p^{\alpha}_{s_{K-2}, s_{K-1}} J^{\pi^*}_{K-1}(s_{K-1}) \right].$$

Note that the optimal costs-to-go for stage $K-1$, $J^{\pi^*}_{K-1}(s_{K-1})$, is known from the previous step.

- Stage k: choose action $u^*_k(s_k)$ for each possible state s_k to minimize

$$g(s_k, \alpha) + \sum_{s_{k+1}} p^{\alpha}_{s_k, s_{k+1}} J^{\pi^*}_{k+1}(s_{k+1}).$$

In other words, we choose the action

$$u^*(s_k) = \underset{\alpha \in \mathcal{A}}{\operatorname{argmin}} \left[g(s_k, \alpha) + \sum_{s_k} p^{\alpha}_{s_k, s_{k+1}} J^{\pi^*}_{K-1}(s_{k+1}) \right],$$

and the optimal cost-to-go is given by

$$J^{\pi^*}_k(s_k) = \min_{\alpha \in \mathcal{A}} \left[g(s_k, \alpha) + \sum_{s_{k+1}} p^{\alpha}_{s_k, s_{k+1}} J^{\pi^*}_{K-1}(s_{k+1}) \right].$$

Again, the optimal costs-to-go for stage $k+1$, $J^{\pi^*}_{k+1}(s_{k+1})$, is known from the previous step.

The above backward inductive derivation can be carried out accordingly up to the initial state s_0. During this backward induction, we obtained an optimal policy, namely the sequence of state-to-action mapping $\{u_k^*(s_k)\}_{k=0}^K$, and finally the cost of the optimal policy $J_0^{\pi^*}(s_0)$. Therefore, we have solved the problem.

The number of numerical operations carried out during this backward derivation is $O(|K||S|^2|A|)$, because at each stage k for each state s_k and each action the algorithm takes up to $|\mathcal{S}|$ multiplications and $|\mathcal{S}|-1$ additions. Furthermore, the algorithm needs to take up to $|\mathcal{A}|-1$ comparisons to choose the optimal action for each state. Therefore, the total number of operations required at each stage is upper bounded by $|\mathcal{S}|\big(|\mathcal{A}|(2|\mathcal{S}|-1)+|\mathcal{A}|-1\big)$. Hence, the total number of operations for the entire algorithm is upper bounded by $|K|[|\mathcal{S}|(|\mathcal{A}|(2|\mathcal{S}|-1)+|\mathcal{A}|-1)]$, which is $O(|K||\mathcal{S}|^2|\mathcal{A}|)$. Furthermore, it is clear that the total number of numerical operations involved is independent of the size of the numerical values of the parameters of the problem – the bit-length of the transition probabilities $p_{i,j}^{\alpha}$. Thus, we say the dynamical-programming algorithm is a *strongly polynomial* algorithm for finite state-space, finite action-space MDPs of finite horizon.

Formulate the BDLSP as an MDP

As we mentioned in Section 4.2.3, the same problem was considered by Kao and Littman [25] under the name of "informed cow-path problem." There, they essentially took the MDP approach as well. For the purpose of completeness and comparison with the following graph search case, we include the MDP formulation and its complexity analysis here.

As pointed out by Kao and Littman [25], this problem could be formulated naturally into a finite-indefinite-horizon POMDP by taking the set of all the possible locations of the target as the state space and the search results at each time

as observations. By formulating the problem in such a general form, we risk falling "naturally" into the trap of believing that the problem is intrinsically hard, since it was well-known that general POMDPs are PSPACE-hard [46] and even *undecidable* in the cases of infinite and indefinite horizon [47]. However, the BDLSP turns out to be special It has distinct structures to exploit. In particular, the problem can be formulated as an MDP by taking the searched regions, which are intervals, as the states.

In the following, we formulate the BDLSP as an MDP. We follow the notation and assumptions established in Section 4.2.3. However, instead of fixing the searcher's starting position to location 0, we include the starting position as part of the policy design.

Suppose that a searcher chooses a starting position, say k, and that it then moves either to the left $(k-1)$ or the right position $(k+1)$. After this movement, the situation is characterized by two integers n_l and n_r, $-n \leq n_l < n_r \leq n$, and an indicator $\mathrm{p} \in \{l, r\}$ which takes the value of l(eft) if the searcher is at position n_l and r(ight) if at n_r. This characterization, namely, the triple (n_l, n_r, p), suffices to describe various situations that can occur during the search process. Therefore, the set $\{(n_l, n_r, \mathrm{p}) : -n \leq n_l \leq n_r \leq n, \mathrm{p} \in \{l, r\}\}$ of triples together with two special (virtual) states, namely the initial state $\mathbf{i}$ and the terminal state $\mathbf{t}$, forms the state space $\mathcal{S} = \{\mathbf{i}, \mathbf{t}\} \cup \{(n_l, n_r, \mathrm{p}) : -n \leq n_l \leq n_r \leq n, \mathrm{p} \in \{l, r\}\}$.

At the stage 0 when the system is in the (virtual) initial state $\mathbf{i}$, the searcher chooses a starting position k (action) and then enters the first stage state (k, k, p), $k \in [-n, n]$ and p immaterial here. We denote these choices by the integers from the set $[-n, n]$. At the stage 1, the searcher has two choices either go left or right. We denote these two possible actions for the first stage by **lf** and **rt** respectively. In the middle of the search process, at a typical state (n_l, n_r, p), the searcher has

two choices as well – it can either continue in the previous direction and move a step further or turn around, pass the other end of the searched region, and explore one step further. We call these two actions "continuing" and "turning around," denoted **c** and **r** respectively. Therefore, the set $[-n, n]$ of feasible actions at stage 0 and that at stage 1, together with that of typical middle states, form the action space $\mathcal{A} = \{-n, \cdots, n, \mathbf{lf}, \mathbf{rt\,c}, \mathbf{r}\}$.

Next, we describe the state-transition law p_{ij}^{α}, $i, j \in \mathcal{S}$ and $\alpha \in \mathcal{A}$. We first describe the transition law for the stages 0 and 1. The system transits from the initial state **i**, under an action k, with probability p_k to the terminal state **t** – meaning the target is found at position k – and with probability $1 - p_k$ to the state (k, k, r) – the target is not found at position k. (Note that for states of the form (k, k, p) the position indicator p is degenerate.) Then, at stage 1, the system transits from a state (k, k, r), $k \in (-n, n)$, under the action of **rt**, with probability $p_{k+1}/(1 - p_k)$ to the terminal state **t** – the target is found at position $k + 1$ – and with probability $1 - p_{k+1}/(1 - p_k)$ to the state $(k, k+1, r)$ – the target is not found at position $k+1$. Similar for the case with the **lf** action, the system transits from a state (k, k, r), $k \in (-n, n)$, under the action of **lf**, with probability $p_{k-1}/(1 - p_k)$ to the terminal state **t** – the target is found at position $k - 1$ – and with probability $1 - p_{k-1}/(1 - p_k)$ to the state $(k - 1, k, l)$ – the target is not found at position $k - 1$. Finally, for the state $(-n, -n, r)/(n, n, r)$, the only feasible action is **rt** / **lf**. The system transits from $(-n, -n, r)$ to **t** with probability $p_{-n+1}/(1 - p_{-n})$ and to $(-n, -n + 1, r)$ with probability $p_{-n+1}/(1 - p_{-n})$, and similarly from (n, n, r) to **t** with probability $p_{n-1}/(1 - p_n)$ and to $(n - 1, n, l)$ with probability $p_{n-1}/(1 - p_n)$.

For the system evolution past stage 1, the law is given by the following tran-

sition probability table:

p_{ij}^{α}	$\mathbf{t}$	(n_l, n_r+1, r)	(n_l-1, n_r, l)
$(n_l, n_r, r), \mathbf{c}$	$\frac{p_{n_r+1}}{1-\sum_{i\in[n_l,n_r]} p_i}$	$1-\frac{p_{n_r+1}}{1-\sum_{i\in[n_l,n_r]} p_i}$	
$(n_l, n_r, l), \mathbf{c}$	$\frac{p_{n_l-1}}{1-\sum_{i\in[n_l,n_r]} p_i}$		$1-\frac{p_{n_l-1}}{1-\sum_{i\in[n_l,n_r]} p_i}$
$(n_l, n_r, r), \mathbf{r}$	$\frac{p_{n_l-1}}{1-\sum_{i\in[n_l,n_r]} p_i}$		$1-\frac{p_{n_l-1}}{1-\sum_{i\in[n_l,n_r]} p_i}$
$(n_l, n_r, l), \mathbf{r}$	$\frac{p_{n_r+1}}{1-\sum_{i\in[n_l,n_r]} p_i}$	$1-\frac{p_{n_r+1}}{1-\sum_{i\in[n_l,n_r]} p_i}$	

We assume $-n \le n_l < n_r \le n$ in the table. It reads from the column index to the row index as "the system transits from the state, say (n_l, n_r, r), under the action $\mathbf{c}$, to the next state $\mathbf{t}$ with probability $p_{n_r+1}/(1-\sum_{i\in[n_l,n_r]} p_i)$ and to next state (n_l, n_r+1, r) with probability $1-p_{n_r+1}/(1-\sum_{i\in[n_l,n_r]} p_i)$."

It remains to describe the cost function $g(i,\alpha)$, $i \in \mathcal{S}$ and $\alpha \in \mathcal{A}$. Similar to the state-transition law, we consider three cases. First, we assume the searcher, at the initial stage, is allowed to "jump" into any position with zero cost, i.e., $g(\mathbf{i},\alpha) = 0$, for all $\alpha \in [-n, n]$. For states of the form (k,k,r), we have $g\big((k,k,r),\alpha\big) = 1$, for all $k \in [-n, n]$ and $\alpha \in \{\mathbf{lf}, \mathbf{rt}\}$. For all the typical middle states of the form (n_l, n_r, p), $n_l < n_r$, we have $g\big((n_l, n_r, \mathrm{p}), \mathbf{c}\big) = 1$ and $g\big((n_l, n_r, \mathrm{p}), \mathbf{r}\big) = |n_r - n_l| + 1$.

The BDLSP is thus formulated as an MDP. Automatically, it admits the solution of the dynamic-programming algorithm we discuss in Section 4.3.1 and is solved by working backwards from the terminal state $\mathbf{t}$ to the initial state s_0.

It is easy to see that the number of states of the resulting MDP is $O(n^2)$, the action space is $O(2n)$, and the horizon is $O(n)$. Hence, according to the complexity discussion following the dynamic-program algorithm in 4.3.1, we know that the time-complexity of the dynamic-programming algorithm for the BDLSP is $O(n^4)$, independent of the sizes of the probability entries p_i. Therefore, we conclude that the bounded linear search problem is strongly polynomial.

However, taking a closer look at the transition diagram, we can easily see that the above estimation for the time complexity is too conservative – it overestimates the time-complexity. To see this, first note that in our case at each state there are

only two next states for each action. Therefore, for each state and each action, the involved numerical operations are 2 multiplications and 1 addition. To choose an optimal action, at the initial state, the algorithms takes $2n$ comparisons, and for all other states, at most 1 comparison is required since there are at most two choices at a time. Finally, note that every state was visited by the algorithm at most once. Therefore, the totally number of numerical operations is of the order of the system state space, i.e., $\Theta(n^2)$, rather than $O(n^4)$. It is interesting to note that the additional burden of choosing an optimal starting point dose *not* increase the complexity in the asymptotic region, since, as pointed out by Kao and Littman [25] as well, the time-complexity of the linear search problem with fixed starting points is also quadratic in n.

4.3.2 Search on Graphs—the GSP

The search domain we consider in Section 4.3.1 is the simplest kind of graph a linear network. It is natural to consider the case of search on more general graphs. It turns out that the complexity increases dramatically we shall show that the GSP is NP-complete.

Consider an undirected graph $G = (V, E)$. Denote the probability that the target locates at a vertex $v \in V$ of the graph by p_v. Without loss of generality, we assume that associated with each edge e is a cost l_e bearing the meaning of physical distance for the searcher to travel over the edge. The goal for the searcher is to find the target on the graph with minimum travel distance. Clearly, this problem is a natural extension of the forgoing linear search problem. In fact, GSP is equivalent to the so-called the *weighted minimum latency problem* (WMLP) [48, 49]. The *minimum latency problem* (MLP) is usually stated as follows: Suppose we are given a graph (V, E) such that associated with each edge is a non-negative cost. The goal is to construct a tour that minimizes the total latency $\sum_{v \in V} l(v)$, where

$l(v)$ is the latency to visit vertex v. In the case of WMLP, associated with each node is a non-negative weight w_v and the goal is to construct a tour that minimizes the total *weighted* latency $\sum_{v \in V} w_v l(v)$. By taking the distance traveled before first visiting v as latency and the target location probability as weight, it is clear that GSP is equivalent to WMLP. In the computer science literature, the MLP is also termed the *delivery man problem*, *traveling repairman problem* [48, 50, 51] (in contrast to the well-known *traveling salesman problem* (TSP)), and *school-bus driver problem* [52].

In fact, the BDLSP we discuss in Section 4.3.1 was studied as a special case of MLP as well, and the fact that it admits polynomial time-complexity solutions using dynamic-programming was discovered as early as 1986 [50]. Recently, Garcia et al. [49] even discovered a linear time-complexity algorithm by identifying a Monge matrix structure of the problem.

Theorem 4.3.1. *The GSP is NP-complete.*

Proof. First, note that by taking probability $p_v = \frac{w_v}{\sum_{v \in V} w_v}$, we convert the WMLP to the GSP. In particular, the optimal search policy for the corresponding GSP gives the optimal weight minimum latency tour for WMLP and vice versa. Next, each MLP is a special case of WMLP with equal weights. Therefore, any deterministic polynomial-time algorithm to the GSP would give a deterministic polynomial-time algorithm to WMLP and in turn to MLP. Because MLP is NP-complete [48, 53], the GSP is NP-complete. □

It is worth pointing out that the above proof also establishes the fact that GSP cannot even be ϵ-approximated in polynomial-time unless P = NP because it was shown in [53] and [48] that the MLP cannot be ϵ-approximated in polynomial-time unless P = NP.

4.4 Erroneous Bounded Linear Search Problem (EBDLSP)

In this section, we investigate a variation of the finite linear search problem the *erroneous bounded discrete linear search problem* (EBDLSP). When the searcher visits a location k at which the target indeed is, it may miss the target with a probability $\epsilon_k > 0$. Our goal is still to devise an optimal policy that minimizes its expected travel distance to find the target. In the following, we assume that the searcher always starts with position 0. Search problems with errors or misdetection are notoriously hard to solve. See [54] for a recent survey on the twenty-question-with-a-liar problem. To our best knowledge, this work is the first to address the linear search problem with errors.

As before, we assume the target is located on the integer line segment $[-n, n]$ and at a particular position $i \in [-n, n]$ with probability p_i. (Since the searcher always starts with position 0, we assume $p_0 = 0$.) Given that the target is at location i, for each inspection of i, the searcher may miss the target with probability ϵ_i, $0 < \epsilon_i < 1$, independent of both the target location distribution and the search history.

4.4.1 Costs for both Searching and Traveling

It is clear that if the searcher is not charged for each inspection of a location, the searcher can keep searching a location many times so that the missing probability at each location can be made arbitrary small and then continue to visit a next location. In this way, the total search cost can be made arbitrarily close to that of the case without errors. To avoid such trivialities, we could either impose a restriction that the searcher must leave the location once a location is visited, or (more naturally) we can charge the searcher a cost q_i for each inspection at location i while the searcher is allowed to stay at a location and keep searching it as many times as such a decision is desirable. In the latter case, we assume the cost

of traveling is commensurable with that of inspection – this is reasonable for many practical situations. For example, we may take q_i to be the time spent on each inspection of position i and the traveling cost d_i to be the time needed to travel between the two consecutive locations i to $i+1$. For the sake of model richness, in the following we take the latter approach – the searcher is allowed to inspect location i many times without leaving but for a cost of q_i for each inspection.

In the following, we investigate the performance penalty caused by errors and bound the optimal value of EBDLSP in terms of that of its error-free counterpart BDLSP.

4.4.2 Bound on Performance Loss for Misdetection

For mathematical convenience, we consider the search problem with homogeneous error probabilities, i.e., $\epsilon_i = \epsilon$ for all $i \in [-n, n]$. For an EBDLSP, denoted $\mathcal{P}_\epsilon$, we denote its error-free counterpart by $\mathcal{P}$. With a slight modification of the cost structure of the BDLSP discussed previously in Section 4.3.1, the error-free problem $\mathcal{P}$ can be solved equally efficiently using dynamic-programming.

To state our main result, we denote by $c^*(\mathcal{P}_\epsilon)$ and $c^*(\mathcal{P})$ the optimal costs of $\mathcal{P}_\epsilon$ and $\mathcal{P}$ respectively, and by C' the "total one-pass cost" – the total cost incurred for the searcher to travel from one end to the other end of the line segment $[-n, n]$ and inspect each location once. Furthermore, we denote by l the cost for the optimal policy of $\mathcal{P}$ to traverse the entire segment, i.e.,

$$l = \max_{i \in [-n,n]} l^{\pi^*}(i),$$

where π^* is the optimal policy for the BDLSP $\mathcal{P}$ and $l^{\pi^*}(i)$ is the cumulative cost incurred under the policy π^* up to the time when the location i is first visited and inspected.

Theorem 4.4.1. *The optimal cost of an EBDLSP $\mathcal{P}_\epsilon$ is bounded as follows:*

$$c^*(\mathcal{P}) \leq c^*(\mathcal{P}_\epsilon) \leq c^*(\mathcal{P})(1-\epsilon) + l\frac{\epsilon}{1-\epsilon} + C\frac{\epsilon}{(1-\epsilon)^2}.$$

This theorem suggests and justifies using the error-free optimal policy in the case where the detection is highly reliable, i.e., the error probability ϵ is small, because it is easy to see that:

$$c^*(\mathcal{P})(1-\epsilon) + l\frac{\epsilon}{1-\epsilon} + C\frac{\epsilon}{(1-\epsilon)^2} \to c^*(\mathcal{P}) \text{ as } \epsilon \to 0.$$

In fact, it also partially confirms our intuition that in the cases where search costs are relatively lower than traveling costs, repeating several searches at each location, whenever visited, could most likely be efficient since by doing so the missing probability at each visited location is geometrically reduced with only a linear increase of the small search costs.

The lower-bound part is obvious the optimal cost $c^*(\mathcal{P}_\epsilon)$ of an EBDLSP cannot be lower than that of its error-free counterpart.

To see the upper bound, consider the following multi-stage search policy, denoted π_m: at the first stage, the searcher starts from the location 0, the searcher traverses the line segment according to the optimal policy π of the error-free BDLSP $\mathcal{P}$ but searches the locations only when they are first visited. The searcher necessarily ends at one end of the line segment, say location n. At the next stage, the searcher travels straight from one end to the other end (from location n to $-n$) and inspects each visited location once. Then, the searcher travels back from $-n$ to n and inspects each of visited location once again. The searcher travels back and forth in this fashion along the line segment until the target is located. It can be show that the expected cost of this multi-stage search policy π_m is upper bounded as follows:

Lemma 4.4.2. *For an EBDLSP $\mathcal{P}_\epsilon$ with error probability ϵ, the cost $c^{\pi_m}(\mathcal{P}_\epsilon)$ of the multi-stage search policy is upper bounded as follows:*

$$c^{\pi_m}(\mathcal{P}_\epsilon) \leq c^*(\mathcal{P})(1-\epsilon) + l\frac{\epsilon}{1-\epsilon} + C\frac{\epsilon}{(1-\epsilon)^2}.$$

The proof of Lemma 4.4.2 is relegated to Appendix 4.7.1.

4.5 Search on Unbounded Lines

In this section, we investigate the problem of search on an unbounded line. Here, instead of considering computational complexities, we shift our focus to characterizing optimal policies, since the size of the problem and the description of optimal policies are typically unbounded. On the other hand, instead of seeking exact solutions, we consider approximation methods to these problems.

4.5.1 Introduction

Recall that an instance of the unbounded discrete liner search problem (UB-DLSP) is characterized by the underlying target location probability mass function with an unbounded support, $\mathrm{P} = \{p_i : i \in \mathbb{Z}\}$. The goal is to minimize the expected distance traveled to find the target. In this section, we assume that the searcher always starts from position 0. Furthermore, we assume that the probability mass function has unbounded support on both sides. This assumption can be more explicitly represented in terms of the following convenient notation $F(i)$, $i \in \mathbb{Z}$, $i \neq 0$, for "tail" probabilities:

$$F(i) = \begin{cases} \sum_{k>i} p_k, & i > 0 \\ \sum_{k<i} p_k, & i < 0 \end{cases} \tag{39}$$

We use $F(0_+)$ and $F(0_-)$ to denote $\sum_{i>0} p_i$ and $\sum_{i<0} p_i$ respectively. Clearly, the double-sided unbounded support assumption is equivalent to

$$F(i) > 0, \text{ for all } i \in \mathbb{Z}.$$

This assumption is used in the proof for Theorem 4.5.2.

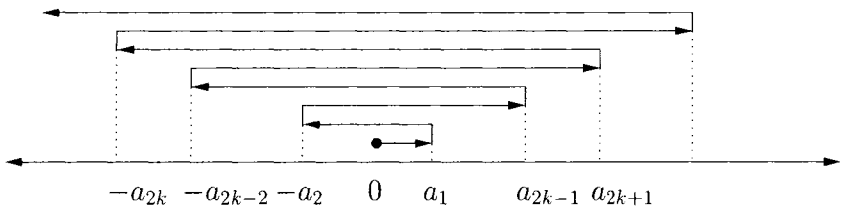

Figure 6. Zig-zag policy for UBDLSPs

Clearly, for a UBDLSP with double-sided unbounded support, any policy being able to locate the target with probability 1 can be described by a double-sided infinite sequence of integers, representing the turning-points, in either of following two forms depending on the first-step direction:

$$\{\cdots, a_{2k}, a_{2k-2}, \cdots, a_2, 0, a_1, \cdots, a_{2k-1}, a_{2k+1}, \cdots\} \tag{40}$$

and

$$\{\cdots, a_{2k+1}, a_{2k-1}, \cdots, a_1, 0, a_2, \cdots, a_{2k}, a_{2k+2}, \cdots\}, \tag{41}$$

where $|a_n| \leq |a_{n+2}|$ for all $n \in \mathbb{N}$. With policies of the first form as shown in Fig. 6, the searcher starts to search the right-hand side first, turns around at location a_1 towards the left-hand side, searches the locations from -1 up to a_2, turns around again towards to the right-hand side, and so on. For policies of the second form, the searcher starts its search on the left-hand side first and then follows the similar zig-zag procedure. We call the turning-point sequences (40) and (41) the *turning-point representation* of policies.

It is not hard to see that any good search policy must satisfy the following two conditions:

(1) $|a_{n+2}| > |a_n|$ for all $n \in \mathbb{N}$, and

(2) $p_{a_n} > 0$ for all $n \in \mathbb{N}$.

We call policies satisfying these conditions *reasonable policies*, and correspondingly refer to the two conditions as *reasonability conditions*. Denote the set of all *reasonable policies* by Π.

For a given policy $\pi = \{a_i : i \in \mathbb{Z}\}$, we denote its expected cost by c^π.

Definition 4.5.1. *A policy π^* is called an* optimal policy *if $c^{\pi^*} = c^* = \inf_{\pi\in\Pi} c^\pi$.*

Denote by $l^\pi(i)$ the traveled distance for the searcher, under the policy π, to first visit location i, $i \in \mathbb{Z}$. The cost of π can be calculated:

$$c^\pi = \sum_{i\in\mathbb{Z}} l^\pi(i)p_i.$$

In terms of the turning-point representation (40) or (41), we can write $l^\pi(i)$ out:

$$l^\pi(i) = 2\sum_{k=1}^{n} |a_k| + |i|, \text{ for } i > 0 \text{ and } i \in (a_{k-1}, a_{k+1}] \text{ or } i < 0 \text{ and } i \in [a_{k+1}, a_{k-1}).$$

It is not hard to see that the cost c^π can be alternatively written in terms of the tail probability notation $F(i)$:

$$c^\pi = 2\sum_{k\in\mathbb{N}} |a_k|[F(a_{k-1}) + F(a_k)] + \sum_{i\in\mathbb{Z}} |i|p_i,$$

where we take the convention $a_0 = 0_+$ if $a_1 < 0$ and $a_0 = 0_-$ if $a_1 > 0$. See Equation (52) in Appendix 4.7.3 and the detailed calculation therein.

Note that the second term is the double-sided mean of the underlying distribution, which is policy-invariant. Denoting this mean by M, we can simplify the expression for c^π:

$$c^\pi = 2\sum_{i\in\mathbb{N}} |a_k|[F(a_{k-1}) + F(a_k)] + M.$$

In the following, we characterize optimal policies. But, before doing so, we first address a subtle issue—the existence of optimal policies.

4.5.2 The Existence of Optimal Policies

We first clarify the meaning of "existence" for optimal policies. Recall that a policy π^* is called optimal if $c^{\pi^*} = c^* = \inf_{\pi \in \Pi} c^\pi$. Note that the set $\{c^\pi : \pi \in \Pi\}$ may not have a finite lower bound. In such cases, it makes little sense to take the infimum of the set. Second, note that the set Π of reasonable policies contains infinitely many, in fact uncountably many, elements. Hence, even if the set $\{c^\pi : \pi \in \Pi\}$ is lower bounded, the infimum may not be achieved—the set may be open at the bottom. For these considerations, we say that an optimal policy exists to mean that the infimum is finite and achievable.

It turns out that the existence issue is technically non-trivial in general. A necessary and sufficient condition for continuous cases was established by Beck [10] and another slightly different one by Franck [9]. A great deal of effort was spent there on tackling the subtlety of possible "infinitely small starting movement" for policies to strive to be "optimal" for certain distributions. Nonetheless, this subtlety does not arise in the discrete cases we address here. For this reason, the existence condition can be cleanly stated as follows.

Theorem 4.5.2. *Optimal policies for UBDLSP exist if and only if the double-sided mean of the underlying distribution is finite.*

We point out that the main idea of the following proof for Theorem 4.5.2 is similar to that of [9] and [10]. But, thanks to the discrete setting, their convolved arguments can be significantly simplified. For the sake of this simplification and to set the stage for the later discussion as well, we include in the following a self-contained proof.

Proof. First we claim:

Claim 4.5.3. *The infimum c^* is finite if and only if the double-sided mean M is finite. Specifically, we have $M \leq c^* \leq 9M$.*

To see the claim, we start with a straightforward observation.

Observation 4.5.4.

$$c^{\pi} \geq M, \text{ for all } \pi.$$

The observation holds because the travel distance $l_{\pi}(i)$ to visit a position i cannot be less than $|i|$ under any policy π. Therefore, we have

$$c^{\pi} = \sum_{i \in \mathbb{Z}} l_{\pi}(i) p_i \geq \sum_{i \in \mathbb{Z}} |i| p_i = M.$$

Therefore, if M is infinite, then c^* must be infinite as well. (It is easy to check that the distribution $\{p_0 = 0, p_i = p_{-i} = \frac{1}{2n^2} : i \in \mathbb{N}\}$ is such an example that has an infinite double-sided mean.)

Next, we prove the following converse to Observation 4.5.4:

Observation 4.5.5.

$$c^* \leq 9M.$$

This inequality can be easily shown by considering the performance of the best worst-cast policy we mentioned before, the so-called "doubling policy":

$$\pi_d = \{\cdots, 2^{2k+1}, \cdots, -8, -2, 0, 1, 4, \cdots, -2^{2k}, \cdots\}.$$

It is easy to see that the travel distance $l^{\pi_d}(i)$ to first visit any location i, is upper bounded by $9|i|$. Thus, the result follows.

With Claim 4.5.3 established, we are left to show that if M is finite, then an optimal policy exists. We start with a sequence $\{\pi_n : n \in \mathbb{N}\}$ of policies such that

$$c_n^{\pi} \to c^*.$$

Given a small positive number δ, there exists an $N \in \mathbb{N}$ such that $c^{\pi_n} < c^* + \delta$ for all $n > N$. Instead of working with the original policy sequence, we focus on the following subsequence of policies starting from π_{N+1}:

$$\{\pi_{N+n} : n \in \mathbb{N}\}.$$

For notation simplicity, we re-index this tail sequence (with little danger of confusion) directly as $\{\pi_n : n \in \mathbb{N}\}$. Next, we "extract" an optimal policy from the policy sequence using the "subsequence" argument.

Recall that a policy π_n is specified via the sequence of its turning-points $\{a_k^n : k \in \mathbb{N}\}$. First, we show that there exists a sequence $\{b_k : k \in \mathbb{Z}\}$ such that for each k the inequality $|a_k^n| < b_k$ holds for all n. We prove this by induction (on k). Recall the cost calculation formula:

$$c^\pi = 2\sum_{i\in\mathbb{Z}} |a_k|[F(a_{k-1}) + F(a_k)] + M. \tag{42}$$

Clearly, depending on whether a_1^n is positive or negative, we have either $2|a_1^n|F(0_-) \leq c^{\pi_n}$ or $2|a_1^n|F(0_+) \leq c^{\pi_n}$. Since $c^{\pi_n} < c^* + \delta$ holds for all n, we see that

$$|a_1^n| \leq b_1 = \frac{c^* + \delta}{2\min(F(0_+), F(0_-))} \text{ for all } n.$$

By our double-sided unboundness assumption, the denominator is non-zero. Therefore, we see that the first turning-point sequence $\{a_1^n : n \in \mathbb{N}\}$ is bounded by b_1.

Next, we assume that

$$|a_k^n| \leq b_k \text{ for all } n.$$

We are left to show

$$|a_{k+1}^n| \leq b_{k+1} \text{ for all } n,$$

for some finite b_{k+1}. Note that

$$F(a_k^n) \geq \min(F(b_k), F(-b_k)).$$

Furthermore, by our double-sided unboundedness assumption again, we have $F(b_k), F(-b_k) > 0$. Similar to the case of $k = 0$, by resorting to the cost calculation formula (42), we have $|a_{k+1}^n|F(a_k^n) \leq 2c^{\pi_n} \leq c^* + \delta$, for all n. Therefore, we get

$$|a_{k+1}^n| \leq b_{k+1} = \frac{c^* + \delta}{2\min(F(b_k), F(-b_k))}, \text{ for all } n.$$

Thus, we proved there exists a real sequence $\{b_k : k \in \mathbb{Z}\}$ such that for each k the inequality $|a_k^n| \leq b_k$ holds for all n.

Now, we are ready to apply the "subsequence" procedure to extract an optimal policy from the policy sequence $\{\pi_n : n \in \mathbb{N}\}$. Starting with the sequence of the first turning-points $\{a_1^n : n \in \mathbb{N}\}$ of the policy sequence, we know that there exists a convergent subsequence $\{a_1^{n_i^1} : i \in \mathbb{N}\}$ since the sequence is bounded by b_1. Denote its limit by a_1^*. Next, consider the subsequence $\{a_2^{n_i^1} : i \in \mathbb{N}\}$ of the second turning-points. For the same reason, there exists a convergent subsequence $\{a_2^{n_i^2} : i \in \mathbb{N}\}$. Similarly, we denote its limit by a_2^*. In the same spirit, we can carry this "subsequence extracting" procedure inductively for all k and obtain a sequence:

$$\{a_k^* : k \in \mathbb{N}\}$$

of limit turning-points. We denote by π^* the policy corresponding to this limit turning-point sequence.

It is left to show that the policy π^* specified by the above sequence of "limiting" turning-points is an optimal policy. In other words, we need to show

$$c^{\pi^*} = 2\sum_{k\in\mathbb{N}} |a_k^*|[F(a_{k-1}^*) + F(a_k^*)] + M = c^*.$$

Recall that the convergent subsequence extracted from the kth turning-point sequence is indexed as $\{a_k^{n_i^k} : i \in \mathbb{N}\}$. It is easy to see that, for any k, there exists an $N_k > 0$ such that

$$\left||a_j^*| - |a_j^{n_i^k}|\right| < \frac{1}{2}$$

holds for all $i > N_k$ and all $j = 1, \cdots, k$. Since all the turning-points are integers, the limiting points must be integers as well. Consequently, we have

$$\left||a_j^*| - |a_j^{n_i^k}|\right| = 0, \tag{43}$$

for all $i > N_k, j = 1, \cdots, k$. On the other hand, for each $k > 0$, there exists an $N'_k > 0$ such that

$$|c^{\pi_n} - c^*| < \frac{1}{k} \tag{44}$$

holds for all $n > N'_k$. We can find an $N(k) > 0$ such that $n^k_{N(k)} > \max\{n^k_{N_k}, N'_k\}$ holds. Then, we have that (43) and (44) hold simultaneously for all $i > N(k)$.

Denote by S^π_k the partial sum $2\sum_1^k |a_k|[F(a_{k-1}) + F(a_k)] + M$ of the infinite sum calculating the cost of policy π. From (43) we have

$$|S_k^{\pi_{n^k_i}} - S_k^{\pi^*}| = 0, \text{ for all } i > N(k). \tag{45}$$

Combining (44) and (45), we get

$$\begin{aligned} |c^{\pi^*} - c^*| &\leq |c^{\pi^*} - S_k^{\pi^*}| + |S_k^{\pi^*} - S_k^{\pi_{n^k_i}}| + |S_k^{\pi_{n^k_i}} - c^{\pi_{n^k_i}}| + |c^{\pi_{n^k_i}} - c^*| \\ &\leq |c^{\pi^*} - S_k^{\pi^*}| + |S_k^{\pi_{n^k_i}} - c^{\pi_{n^k_i}}| + \frac{1}{k}, \text{ for all } i > N(k). \end{aligned}$$

Note that as k goes to infinity, the term $|c^{\pi^*} - S_k^{\pi^*}|$ vanishes. On the other hand, we have that the diagonalized term $|S_k^{\pi_{n^k_k}} - c^{\pi_{n^k_k}}|$ vanishes as well as k goes infinity. Therefore, by taking k and i large, we can make $|c^{\pi^*} - c^*|$ arbitrarily small. Therefore, we conclude that $|c^{\pi^*} - c^*| = 0$ and that the policy π^* is optimal. □

4.5.3 Uniqueness of Optimal Policies

It is natural to be curious about the uniqueness of an optimal policy. Clearly, uniqueness is not the case in general. For example, when the underlying distribution is symmetric, any optimal policy starting from one side has a "mirroring" counterpart that starts from the other side. However, it is unclear so far, up to such a symmetric multiplicity, whether optimal policies are generally unique. To our best knowledge, the uniqueness (up to symmetric multiplicity) of the optimal policy was investigated and proved by Beck and Beck [14] only for the continuous case with the underlying distribution being normal. Their proof exploits certain

special property of the normal distribution. We show later that the uniqueness is crucial to establishing the convergence of our approximation procedure proposed in Section 4.5.6.

4.5.4 The Expanding Property of Optimal Policies for Symmetric UBDLSPs

From this section on, we focus on UBDLSPs with symmetric distributions. First, we establish the following property of optimal policies for symmetric UBDLSPs.

Theorem 4.5.6. *If the underlying distribution of a UBDLSP is symmetric and has a finite double-sided mean, then its optimal policies must be* expanding*:*

$$|a_{k+1}| > |a_k|, k \in \mathbb{N}.$$

Note that the *reasonability condition* (1) says that any "reasonable" policy, certainly including optimal policies, must be "one-sided" expanding. Theorem 4.5.6 says that the optimal policies for symmetric distributions must be "double-sided" expanding as well.

We first observed this phenomena during computer experiments and later found that it can be rigorously established. At that time, we were unaware of the whole body of work by Beck and Franck, including the fact that the same expanding property was early observed for continuous symmetric distributions by Beck and Beck [8] during their investigation. In fact, the first proof for the progressively expanding property for UBDLSP with continuous distributions given in [8] turned out faulty and later corrected by themselves in [14]. We find the main idea of our proof similar to theirs. However, because of the discreteness of our settings, the details are significantly different. We include a self-contained proof in Appendix 4.7.2.

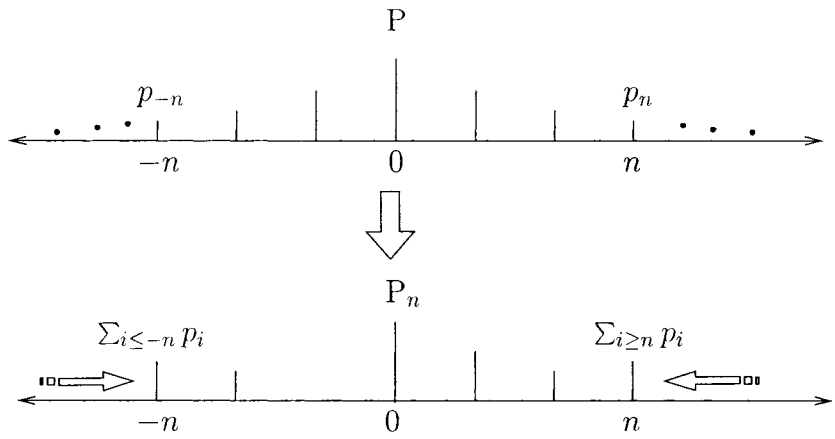

Figure 7. n-truncated BDLSP for UBDLSP

4.5.5 Approximating Optimal Values for Symmetric UBDLSP

In this section and the next, we approximate the optimal value and the optimal policy for symmetric UBDLSP by solving a sequence of finitely-truncated BDLSPs obtained from the original UBDLSP. Given a symmetric distribution $\mathrm{P} = \{p_i : i \in \mathbb{N}\}$ in which $p_i = p_{-i}$ for all i, consider the n-truncated BDLSP as shown in Fig. 7 with its underlying distribution

$$\big(F(n) + p_n, p_{n-1}, \cdots, p_1, p_0, p_1, \cdots, p_{n-1}, p_n + F(n)\big),$$

where $F(n)$ is the tail probability $\sum_{i>n} p_i$ of the original unbounded distribution. Since the original unbounded distribution is symmetric, the left tail probability $F(-i)$ equals the right tail probability $F(i)$ for all i. We denote the n-truncated distribution by P_n and with a little abuse of notation we call the corresponding n-truncated UBDLSP P_n as well.

As we show in Section 4.3.1, the optimal cost and the optimal policy of the n-truncated BDLSP can be computed efficiently. Denote by c_n^* the optimal value of the n-truncated BDLSP P_n. We have the following result.

Theorem 4.5.7. *Suppose that the optimal policy for a symmetric UBDLSP exists. The sequence $\{c_n^*\}_{n\in\mathbb{N}}$ of the optimal values for the finitely-truncated BDLSPs converges to the optimal value c^* of the UBDLSP.*

The theorem is not surprising and confirms our intuition. However, the result is not immediately obvious either, considering that we have few tools accessible and useful in such a primitive analysis setting and hence can hardly say anything definite about the optimal policy. Our proof exploits the *principle of optimality* [55] and takes advantage of the expanding property for symmetric UBDLSPs established in Theorem 4.5.6.

Proof. The proof is carried out in two steps. First, we show that the sequence $\{c_n^*\}$ converges. Then, we show the sequence converges to c^* by demonstrating that one of its subsequences converges to c^*.

We show the convergence of $\{c_n^*\}$ by proving that it is monotone increasing and bounded from above. We first show that it is monotone increasing by contraction. Suppose that $c_{n+1}^* < c_n^*$ holds for some n, and c_{n+1}^* is achieved by optimal policy π_{n+1}^*. Without loss of generality, suppose that the policy π_{n+1}^* has the following one-sided turning-point representation:

$$\pi_{n+1}^* = (a_1, \cdots, a_k. n+1, n+1),$$

where $a_i > 0$ for all $i \in [k]$ and $a_i > a_j$ for $k \geq i > j = 1$. To avoid cluttering the main text, we relegate the detailed calculation for the costs of policies in both finite truncated and unbounded cases to Appendix 4.7.3. The cost of π_{n+1}^* is given by:

$$c^{\pi_{n+1}^*} = 2\sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)] + \sum_{i=-n-1}^{n+1} |i|p_i + 2(n+1)F(a_k) + 2(n+1)F(n+1).$$

From optimal policy π_{n+1}^* for the $(n+1)$-truncated distribution P_{n+1}, we construct a policy π_n' for the n-truncated problem P_n and show that $c^{\pi_n'} < c_n^*$ to get the

contradiction. Specifically, π'_n is constructed as follows:

$$\pi'_n = (a_1, a_2, \cdots, a_{k-1}, n, n), \text{ if } a_k = n, \text{ or } \pi'_n = (a_1, a_2, \cdots, a_k, n, n), \text{ if } a_k < n.$$

In the first case, the cost of π'_n can be computed:

$$\begin{aligned} c^{\pi'_n} &= \sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)] + \sum_{i=-n}^{n} |i|p_i \\ &= c^{\pi^*_{n+1}} - 2\big[(n+1)F(a_k) + (n+1)F(n+1) + (n+1)p_{n+1}\big] \\ &\leq c^{\pi^*_{n+1}} = c^*_{n+1} < c^*_n. \end{aligned}$$

In the second case, the cost can be computed:

$$\begin{aligned} c^{\pi'_n} &= \sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)] + \sum_{i=-n}^{n} |i|p_i + 2nF(a_k) + 2nF(n) \\ &= c^{\pi^*_{n+1}} - 2\big[(n+1)p_{n+1} + (n+1)F(n+1) + (n+1)F(a_k) - nF(n) - nF(a_k)\big] \\ &= c^{\pi^*_{n+1}} - 2\big[F(a_k) + F(n+1) + p_{n+1}\big] \\ &\leq c^{\pi^*_{n+1}} = c^*_{n+1} < c^*_n. \end{aligned}$$

Therefore, in both cases, we have $c^{\pi'_n} < c^*_n$, contradicting to the minimality of c^*_n. Hence, we conclude that the sequence $\{c^*_n\}$ is non-decreasing.

Next, we show that the sequence is bounded. By the same argument as that for Observation 4.5.5, it is easy to see that the cost of the following doubling policy for the n-truncated BDLSP P_n,

$$\pi^n_d = (-n, \cdots, -2^{2k+1}, \cdots, -2, 0, 1, 4, 2^{2k}, \cdots, n),$$

is bounded by $9M_n$, where M_n is the double-sided mean for the n-truncated distribution P_n. Note that

$$M_n = 2\sum_{i=1}^{n} |i|p_i + 2nF(n) = 2\sum_{i=1}^{n} |i|p_i + 2\sum_{i=n+1}^{\infty} np_i \leq 2\sum_{i=1}^{\infty} |i|p_i = M.$$

Therefore, the sequence $\{c^*_n\}_{n\in\mathbb{N}}$ is bounded above by $9M$. We conclude that the sequence $\{c^*_n\}$ converges.

It remains to show that the sequence converges to the optimal value c^* of the original UBDLSP. Suppose that an optimal policy π has one-sided turning-point representation:

$$\pi = \{a_1, a_2, \cdots, a_k, \cdots\}.$$

Without loss of generality, we assume the policy starts from the right side – the searcher turns at locations a_1, $-a_2$, a_3, $-a_4$, and so on. According to the expanding property, we have $a_{k+1} > a_k$ for all $k \in \mathbb{N}$. We claim that the subsequence $\{c^*_{a_k}\}_{k\in\mathbb{N}}$ of the sequence $\{c^*_n\}_{n\in\mathbb{N}}$ converges to the optimal cost c^*. To see this, consider the a_k-truncated BDLSP P_{a_k}. We claim that the following partial policy obtained by "truncating" the optimal policy π,

$$\pi_{a_k} = (a_1, a_2, \cdots, a_k, a_k),$$

is optimal for P_{a_k}. In other words, we have

$$c^{\pi_{a_k}} = c^*_{a_k}.$$

To see this, note that the cost of the optimal policy π^* is

$$c^{\pi^*} = 2\sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)] + 2\sum_{i=k+1}^{\infty} a_i[F(a_{i-1}) + F(a_i)] + 2\sum_{i=1}^{\infty} ip_i.$$

Since a_k is fixed, the first term

$$2\sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)]$$

depends only on the partial policy sequence $(a_1, \cdots, a_{k-1})$, while the second term is independent of this partial sequence and the last term is constant. By the principle of optimality, the partial sequence $(a_1, \cdots, a_k)$ must minimize the first term for the overall policy to be optimal, for otherwise the partial policy sequence $(a_1, \cdots, a_{k-1})$ could be replaced with a better alternative to obtain a lower overall cost and a better overall policy, contradicting to the optimality of π^*.

Note that the cost of policy π_{a_k} for the a_k-truncated BDLSP is

$$c^{\pi_k} = 2\sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)] + 2\sum_{i=1}^{a_k} ip_i.$$

By the optimality of the sequence $(a_1, \cdots, a_{k-1})$ for the term $2\sum_{i=1}^{k} a_i[F(a_{i-1}) + F(a_i)]$, we see that $c^{\pi_{a_k}}$ must be minimized by the sequence $(a_1, \cdots, a_{k-1})$ as well, because the second term is fixed. Therefore, we have $c^*_{a_k} = c^{\pi_{a_k}}$ for all $k = 1, 2, \cdots$. Furthermore, we have

$$c^{\pi_{a_k}} = 2\sum_{i=1}^{k} a_i[F(a_{i-1})+F(a_i)]+2\sum_{i=1}^{a_k} ip_i \rightarrow c^* = 2\sum_{i=1}^{\infty} a_i[F(a_{i-1})+F(a_i)]+2\sum_{i=1}^{\infty} ip_i,$$

as k goes to infinity. Hence $c^*_{a_k}$ converges to c^*. Therefore, we conclude that the sequence $\{c^*_n : n \in \mathbb{N}\}$ converges to c^*. The proof is completed. □

4.5.6 Approximating Optimal Policies for Symmetric UBDLSPs

We know from the previous section that the optimal value sequence $\{c^*_n\}$ of the sequence $\{\mathrm{P}_n\}$ of n-truncated BDLSPs converges to the optimal value c^* of the original symmetric UBDLSP. It is natural to ask whether the sequence π_n of the optimal policies obtained for the n-truncated BDLSPs P_n converges to the optimal policy of the original UBDLSP. In the proof of Theorem 4.5.7, we showed that the optimal policy for the a_k-truncated BDLSPs, where a_k is a turning-point of an optimal policy for the symmetric UBDLSP, coincides with an optimal policy for the UBDLSP for the first k turning-points. In particular, suppose that the following policy

$$\pi^* = (a_1, a_2, \cdots, a_k, \cdots)$$

is optimal for a given symmetric UBDLSP. Then, for each $k \in \mathbb{N}$, the following policy obtained from π^* by truncating-then-padding,

$$\pi_{a_k} = (a_1, a_2, \cdots, a_k, a_k),$$

is optimal for the a_k-truncated BDLSP. Suppose the optimal policy for the n-truncated BDLSP is $\pi_n = (a_1^n, a_2^n, \cdots, a_k^n, n, n)$. In the following, we investigate whether the kth turning-point sequence $\{a_k^n : n \in \mathbb{N}\}$ converges to the kth turning-points of the optimal policy for all k.

To be precise about convergence of policies, we state the following definition:

Definition 4.5.8. *We say that a sequence of (infinite) policies*

$$\pi_n = (a_1^n, a_2^n, \cdots, a_k^n, \cdots)$$

converges to the policy

$$\pi^* = (a_1^*, a_2^*, \cdots, a_k^*, \cdots),$$

denoted $\pi_n \to \pi^*$, *if*

$$a_k^n \to a_k^*, \text{ as } n \to \infty, \text{ for all } k \in \mathbb{N}.$$

To align with the above definition, we "pad" the optimal policy (which is finite)

$$\pi_n = (a_1^n, a_2^n, \cdots, a_k^n)$$

for the n-truncated BDLSP P_n and form the corresponding "infinite" policy

$$\pi_n' = (a_1^n, a_2^n, \cdots, a_k^n, 0, 0, \cdots)$$

with $a_j^n = 0$ for all $j > k$.

It turns out that if the optimal policy π^* of the original UBDLSP is unique, then the sequence $\{\pi_n'\}$ of the padded optimal policy obtained by solving the n-truncated BDLSPs converges to the optimal policy π^* for the original UBDLSP.

Theorem 4.5.9. *Suppose that a symmetric UBDLSP has a unique optimal policy* π *(up to symmetry duplication). Then, the sequence of padded optimal policies* $\{\pi_n'\}$ *of the* n*-truncated BDLSP sequence converges to* π.

It is possible that generally the turning-point sequence, for example $\{a_k^n\}_{n\in\mathbb{N}}$ of the kth turning-points, obtained during our approximation process may not settle down but oscillates among several numbers infinitely. Theorem 4.5.9 ensures that as long as the optimal policy for the original UBDLSP is unique the turning-points obtained during our approximation process eventually settle down to a fixed value and hence proves the effectiveness of our approximation procedure in computing the optimal policy π^*.

Before we prove Theorem 4.5.9, we provide a "weaker" result as follows:

Theorem 4.5.10. *Given a symmetric UBDLSP, let $\{\pi'_n\}_{n\in\mathbb{B}}$ be the sequence of the padded optimal policies for the n-truncated BDLSPs. Consider one of its subsequences $\{\pi'_{n_i}\}_{i\in\mathbb{N}}$. If*

$$\pi'_{n_i} \to \pi,\ i \to \infty$$

then π is optimal for the original UBDLSP.

This is a corollary of Theorem 4.5.7 because $\{c^*_{n_i}\}$, as a subsequence of $\{c^*_n\}$, converges to c^*. This theorem implies that if the "partial" optimal policy sequence oscillates "consistently" among several candidates, then these candidates are all in fact optimal. Clearly, such a situation can happen only when the UBDLSP has multiple optimal policies.

Proof of Theorem 4.5.9. Denote by Π_k the set of policies for the UBDLSP whose first k turning-points coincide with those of the optimal policy π^*. In other words, we have

$$\pi = (a_1, a_2, \cdots, a_k \cdots) \in \Pi_k \iff a_i = a_i^* \text{ for all } i = 1, \cdots, k.$$

Clearly, we have $c^* = \inf\{c^\pi : \pi \in \Pi_k^*\}$. Now, since the optimal policy is unique, for fixed k we have

$$c'_k := \inf\{c^\pi : \pi \notin \Pi_k\} > c^*.$$

Let $\delta_k := c'_k - c^* > 0$.

Next, we show that there exists an $N(k) > 0$ such that for any partial policy π^n, $n > N(k)$, if its padded policy π'_n does not belong to Π_k, then it cannot be optimal for the n-truncated BDLSP. Suppose $\pi^n = (a_1, a_2, \cdots, a_k, a_{k+1}, \cdots, a_j, n, n)$ is a search policy for the BDLSP P_n and $\pi'_n \notin \Pi_k$. The cost of the policy π_n is

$$c^{\pi_n} = 2\sum_{i=1}^{n} ip_i + 2n[F(a_j) + F(n)] + 2\sum_{i=1}^{j-1} a_i[F(a_{i-1}) + F(a_i)].$$

Consider the following search policy $\bar{\pi}_n$ for the original UBDLSP constructed from π_n by "patching" it with doubling search policy in the tail region:

$$\bar{\pi}_n = (a_1, a_2, \cdots, a_k, \cdots, a_j, n, n, 2n, 4n, \cdots, 2^i n, \cdots).$$

It is easy to see that the cost for $\bar{\pi}_n$ is

$$\begin{aligned} c^{\bar{\pi}_n} &= 2\sum_{i=1}^{n} ip_i + 2n[F(a_j) + F(n)] + 2\sum_{i=1}^{j-1} a_i[F(a_{i-1}) + F(a_i)] + \sum_{i\notin[-n,n]} l(i)p_i \\ &= c^{\pi_n} + \sum_{i\notin[-n,n]} l(i)p_i, \end{aligned}$$

The last term $\sum_{i\notin[-n,n]} l(i)p_i$ is the extra expected cost incurred for $\bar{\pi}_n$, where $l(i)$ is the extra travel distance to first visit location $i \in (-\infty, -n) \cup (n, \infty)$ after the searcher finishes searching the region $[-n, n]$. The doubling policy employed for the tail region ensures that $l(i) < 9|i| + n < 10|i|$ for all $i \in (-\infty, -n) \cup (n, \infty)$. Therefore, we get

$$c^{\bar{\pi}_n} - c^{\pi_n} = \sum_{i\notin[-n,n]} l(i)p_i \le 2[\sum_{i=n+1}^{\infty} (10i)p_i] = 20\sum_{i=n+1}^{\infty} ip_i.$$

Because the underlying distribution has finite double-sided mean, the term $\sum_{i=n+1}^{\infty} ip_i$ vanishes as n goes to infinity. Therefore, there exists an $N(k)$ such that $20\sum_{i=n+1}^{\infty} ip_i < \frac{\delta_k}{2}$ for all $n > N(k)$. Then, we have

$$c^{\pi_n} > c^{\bar{\pi}_n} - \frac{\delta_k}{2} \ge c'_k - \frac{\delta_k}{2} = c^* + \delta_k - \frac{\delta_k}{2} > c^*, n > N(k).$$

Therefore, for any $n > N(k)$, if $\pi'_n \notin \Pi_k$, then π_n cannot be optimal for the n-truncated BDLSP P_n since from the previous theorem we know that the optimal cost satisfies $c^*_n \leq c^*$ for all n. Hence, we conclude that the first k turning-points of the optimal policies π^*_n for the BDLSPs P_n must eventually settle to $(a^*_1, a^*_2, \cdots, a^*_k)$, the first k turning-points of the optimal policy π^* for the original UBDLSP. This completes the proof. □

4.5.7 The Increment Sequence of Optimal Policies for Symmetric UBDLSPs with Heavy-tailed Distributions

In the previous section, we established the effectiveness of our approximation approach. It is natural to further investigate its efficiency – the convergence rate of the approximation procedure. It is clear that the convergence rate depends on the tail behavior of the underlying distribution. Specifically, the tail property affects the approximation in two different but related ways. One is that the heavier the tail, the less accurate our approximation with the BDLSP obtained at a fixed n. The other is, as we observed in our computer experiments, that for heavy-tailed distributions, the "strides" of optimal policies – the increment between the consecutive turning-points of the optimal policy – tend to diverge quickly. Both effects suggest that we need to solve an n-truncated BDLSP with some large n to estimate accurately the optimal cost of heavy-tailed distributions. Nonetheless, we have no good estimation so far for the combination of these two effects. As a step towards understanding the convergence rate of our approximation approach, in the following, we formally establish some boundedness results for the strides of optimal policies for UBDLSPs with heavy-tailed distributions.

We quickly review here the related work by Beck and Beck [8] on the normal distribution. It was shown that the increments between consecutive turning-points of the optimal policy for the exceptionally thin-tailed normal distributions is upper

bounded:

$$|a_{k+1}| - |a_k| < 2.5, \text{ for all } k \in \mathbb{N}.$$

In fact, the increment sequence tends to vanish as we can see from a more powerful result obtained in the same paper for the asymptotic behavior of the optimal policy for the normal distribution:

$$\lim_{n\to\infty} \frac{a_n}{\sqrt{2n \ln n}} = 1.$$

We say that a positive function $f(x) : \mathbb{R}_+ \to \mathbb{R}_+$ is asymptotically x^α, denoted $f(x) \sim x^\alpha$, if there exist constants $a, b \in \mathbb{R}$ and an x_0 such that

$$bx^\alpha \geq f(x) \geq ax^\alpha \text{ for all } x > x_0.$$

Definition 4.5.11. *A distribution is called* fat-tailed *if both of its right tail* $F_+(x) = \Pr\{X > x\}$ *and left tail* $F_-(x) = \Pr\{X < -x\}$ *are asymptotically* $x^{-\alpha}$ *for some* $\alpha > 0$. *If the exponent parameter* $\alpha < 2$, *then the fat-tailed distribution is called* heavy-tailed.

Theorem 4.5.12. *Suppose that the underlying distribution of a symmetric UB-DLSP is heavy-tailed and that an optimal policy exists. Then, for any optimal policy*

$$\pi^* = (\cdots, a_{2k}, a_{2k-2}, \cdots, a_2, 0, a_1, \cdots, a_{2k-1}, a_{2k+1}, \cdots),$$

the increment sequence

$$\{|a_{k+1}| - |a_{k-1}| : k \in \mathbb{N}\}$$

is unbounded.

The unboundness of the increment sequence suggests that the optimal policies aggressively explore the uncharted region when the underlying distribution is heavy-tailed. Theorem 4.5.12 and the boundness result for normal distribution, which is a typical thin-tailed distribution, form a clear dichotomy.

Proof. By the heavy-tail assumption, there exist $N > 0$, $c > 0$, and $\alpha \in (1, 2)$ such that $F(n) = F(-n) > cn^{-\alpha}$ for all $n \geq N$. The reason that α must be greater than 1 is for the distribution to have a finite double-sided mean.

We prove the theorem by contradiction. Suppose that the increment sequence of an optimal policy is bounded, i.e., there exist a $d > 0$ such that

$$|a_{k+1}| - |a_{k-1}| < d$$

holds for all k. (Note that by the reasonability condition (2), it holds that $|a_{k+1}| - |a_{k-1}| > 0$ for all k.)

Without loss of generality, we focus on the right side of the turning-points, namely the subsequence

$$\{0, a_1, \cdots, a_{2k-1}, a_{2k+1}, \cdots\}.$$

Fix a large $m > 0$. Consider the set of turning-points $\{a_{2k+i} : a_{2k+i} \in (N, N + m], i \in \mathbb{N}\}$. Clearly, the set must contain at least $\lfloor \frac{m}{d} \rfloor$ turning-points. Rewrite the set as

$$\{a_{n_0+i} : i = 0, \cdots, j-1\},$$

where $j \geq \lceil \frac{m}{d} \rceil$ and $a_{n_0} > N$ is the first turning-point of the set. Recall the cost calculation formula

$$c^\pi = 2 \sum_{k=1}^{\infty} a_k \left[F(a_{k-1}) + F(a_k) \right] + M.$$

By the heavy-tail assumption, we have

$$c^\pi \geq 2 \sum_{k=1}^{\infty} a_k F(a_k) \geq 2 \sum_{i=0}^{j-1} a_{n_0+i} F(a_{n_0+i}) \geq 2c \sum_{i=0}^{j-1} a_{N'+i} (a_{N'+i})^{-\alpha}.$$

Moreover, noting that the function $x^{-(\alpha-1)}$, $\alpha > 1$, is monotone decreasing, we can bound the cost

$$c^\pi \geq 2c \left[j \, a_{N'+j-1}^{-(\alpha-1)} \right] \geq 2j(N+m)^{-(\alpha-1)} \geq 2\frac{m}{d}(N+m)^{-(\alpha-1)} \sim m^{-(\alpha-2)}.$$

Since $\alpha < 2$, by taking m to infinity, we can make the lower bound on c^{π} arbitrarily large. This contradicts the finiteness of the optimal cost. Therefore, we conclude that the increment sequence must be unbounded. □

4.6 Discussion

In this work we focus on discrete linear search and graph search problems. We take the MDP approach to the BDLSP and see that the BDLSP is easy to solve—efficient algorithms exist for BDLSP. But the generalization from simple linear graph to general graph makes the problem much harder GSP turns out to be NP-complete. In fact, it is MAX-SNP-hard as we show that it is equivalent to the WMLP, which is known to be MAX-SNP-hard [48, 53]. Similar to taking the MDP approach to the BDLSP, it would be interesting to approach the *erroneous* BDLSP using a Partially Observable MDP formulation.

We show in the above the *unboundness* of the increment sequence of optimal policies for symmetric UBDLSPs with heavy-tailed distributions. We offer in the following a *boundness* conjecture for symmetric UBDLSPs with *light-tailed* distributions:

Conjecture 4.6.1. *If the underlying distribution of a symmetric UBDLSP is light-tailed (i.e., the tail probability $F(n) = F(-n) \sim \alpha^n$, $0 < \alpha < 1$), then the increment sequence of the turning-points of the optimal policy is bounded.*

This conjecture has been validated in the continuous case with Gaussian distributions by Beck and Beck [8]. Nonetheless, it appears to be difficult to established this result in general.

As we discussed in Section 4.2, although several special cases have been solved, the general two-dimensional search problem is considerably harder than the one-dimensional case and is still widely open. The most interesting topic there is the logarithmic spiral conjecture for the planar search problem where a line lies at an

unknown distance away and with unknown slope. It is also of practical interest to investigate the erroneous two-dimensional search problem this is a completely uncharted research topic.

4.7 Appendix

4.7.1 Proof of Lemma 4.4.2

We analyze the proposed multi-stage search policy π_m stage by stage as follows. Denote by $l^k(i)$, $k \geq 1$, the cumulative cost for the searcher to first arrive and inspect location i at the kth stage. Note that $l = \max_{i\in[-n,n]} l^1(i)$. According to our multi-stage policy, we see that the cumulative cost satisfies $l^k(i) < l + (k-1)C$ for $k \geq 2$.

Denote by p_i^k the probability that the target is found at location i when the searcher first visits and inspects it during the kth stage. Clearly, we have $p_i^1 \leq p_i(1-\epsilon)$. Denote by p^k the probability that the target is found during the kth stage. It is easy to see that the following inequality holds:

$$\sum_{i\in[-n,n]} p_i^k \leq p^k.$$

Because at each stage each location is inspected at least once, the conditional probability that the target is not found at the kth stage given the target is not found at the $(k-1)$th stage is less than ϵ. Therefore, the probability that the target is not found at any of the first k stages is less than ϵ^k. Since the target is found at the kth stage only when the search at the first $k-1$ stages fails, we have

$$p^k \leq \epsilon^{k-1}.$$

Now we can bound the expected cost of the multi-stage search policy π_m as follows:

$$
\begin{aligned}
c^{\pi_m}(\mathcal{P}_\epsilon) &= \sum_{k=1}^{\infty} \sum_{i\in[-n,n]} l^k(i) p_i^k \\
&= \sum_{i\in[-n,n]} l^1(i) p_i^1 + \sum_{k\geq 2} \sum_{i\in[-n,n]} l^k(i) p_i^k \\
&\leq \sum_{i\in[-n,n]} l(i) p_i (1-\epsilon) + \sum_{k\geq 2} \sum_{i\in[-n,n]} [l+(k-1)C]\, p_i^k \\
&= \sum_{i\in[-n,n]} l(i) p_i (1-\epsilon) + \sum_{k\geq 2} [l+(k-1)C] \sum_{i\in[-n,n]} p_i^k \\
&\leq \sum_{i\in[-n,n]} l(i) p_i (1-\epsilon) + \sum_{k\geq 2} [l+(k-1)C]\, \epsilon^{k-1} \\
&= c^*(\mathcal{P})(1-\epsilon) + l\frac{\epsilon}{1-\epsilon} + C\sum_{k\geq 1} k\epsilon^k \\
&= c^*(\mathcal{P})(1-\epsilon) + l\frac{\epsilon}{1-\epsilon} + C\frac{\epsilon}{(1-\epsilon)^2}
\end{aligned}
$$

This concludes the proof.

4.7.2 Proof of Theorem 4.5.6

The condition for finite double-sided means is solely for the purpose of the existence of optimal policies so that it makes sense for us to discuss the properties of them. In fact, this condition will not be used explicitly in the following.

Because of the symmetry of the underlying distribution, any optimal policy starting from right side has a symmetric counterpart that starts from the left side and equally achieves optimality, and vice versa. For this reason, we shall consider only the policies that start from the right side. To facilitate discussion, instead of using double-sided infinite sequence representations, we specify policies by sequences of positive integers $\{a_1, a_2, \cdots, a_k \cdots\}$, meaning that the searcher travels to location a_1 first, then turns around and travels to location $-a_2$, then turns around again and travels to location a_3, and so on.

In this symmetry situation, the tail probability function $F(i)$ is symmetric as well, namely, $F(k) = F(-k)$, $k \in \mathbb{N}$. Correspondingly, the cost calculation formula

is simplified to (see Appendix 4.7.3, Equation (52)):

$$c^{\pi} = 2\sum_{k=1}^{\infty} a_k\big[F(a_{k-1}) + F(a_k)\big] + M.$$

We prove the desired result by contradiction. Suppose $\pi^* = \{a_k : k \in \mathbb{N}\}$ is optimal and the policy π^* fails to satisfy the expanding property. Specifically, we assume it first fails at $(k+2)$th turning-point, namely, $a_{k+1} > a_{k+2}$ and $a_i < a_{i+1}$ for all $i < k+1$. We show in the following that the policy π^* can be improved by a new policy, denoted π', that is obtained by switching a_{k+1} and a_{k+2} and adjusting other turning-points accordingly as follows:

$$\pi' = \{a_1, \cdots, a_k, a_{k+2}, a_{k+1}, a_{k+j-1}, a_{k+j}, a_{k+j+2}, \cdots\},$$

where $j = \min\{2n : a_{k+2n} \geq a_{k+1}, n \in \mathbb{N}\}$. Note that j is even and must be greater than 2, because $a_{k+2} < a_{k+1}$ by our assumption. In other words, j must be greater than or equal to 4.

Next, we show that $c^{\pi'} < c^{\pi^*}$. To see this, we write the cost of π' out:

$$\begin{aligned} c^{\pi'} =& M + \sum_{i=1}^{k} a_i\big[F(a_{i-1}) + F(a_i)\big] \\ &+ a_{k+2}\big[F(a_k) + F(a_{k+2})\big] + a_{k+1}\big[F(a_{k+2}) + F(a_{k+1})\big] \\ &+ a_{k+j-1}\big[F(a_{k+1}) + F(a_{k+j-1})\big] + \sum_{i=k+j}^{\infty} a_i\big[F(a_{i-1}) + F(a_i)\big]. \end{aligned}$$

Compare the above with c^{π^*} similarly written out:

$$\begin{aligned} c^{\pi^*} =& M + \sum_{i=1}^{k} a_i\big[F(a_{i-1}) + F(a_i)\big] \\ &+ a_{k+1}\big[F(a_k) + F(a_{k+1})\big] + a_{k+2}\big[F(a_{k+1}) + F(a_{k+2})\big] \\ &+ a_{k+3}\big[F(a_{k+2}) + F(a_{k+3})\big] + \cdots + a_{k+j-2}\big[F(a_{k+j-3}) + F(a_{k+j-2})\big] \\ &+ a_{k+j-1}\big[F(a_{k+j-2}) + F(a_{k+j-1})\big] \\ &+ \sum_{i=k+j}^{\infty} a_i\big[F(a_{i-1}) + F(a_i)\big] \end{aligned} \tag{46}$$

As we just discussed in the above, the number j must be greater than or equal to 4. It turns out that the proofs for case of $j = 4$ and that of $j > 4$ are slightly different.

We consider the case of $j = 4$ first. Note that in this case the third line of Equation (46) vanishes. Hence, the difference is

$$\begin{aligned} & c^{\pi^*} - c^{\pi'} \\ &= a_{k+1}\big[F(a_k) + F(a_{k+1})\big] + a_{k+2}\big[F(a_{k+1}) + F(a_{k+2})\big] + a_{k+3}\big[F(a_{k+2}) + F(a_{k+3})\big] \\ &\quad - \Big(a_{k+2}\big[F(a_k) + F(a_{k+2})\big] + a_{k+1}\big[F(a_{k+2}) + F(a_{k+1})\big] \\ &\qquad + a_{k+3}\big[F(a_{k+1}) + F(a_{k+3})\big]\Big). \end{aligned} \tag{47}$$

Simplifying, we get

$$\begin{aligned} & c^{\pi^*} - c^{\pi'} \\ &= a_{k+1}[F(a_k) - F(a_{k+2})] - a_{k+2}[F(a_k) - F(a_{k+1})] + a_{k+3}[F(a_{k+2}) - F(a_{k+1})] \\ &= (a_{k+1} - a_{k+2})[F(a_k) - F(a_{k+2})] + (a_{k+3} - a_{k+2})[F(a_{k+2}) - F(a_{k+1})]. \end{aligned}$$

By the reasonability condition (1), we have $a_{k+3} > a_{k+1}$ and $a_{k+2} > a_k$. Therefore, $a_{k+3} > a_{k+2}$ holds by our assumption $a_{k+1} > a_{k+2}$. Together with the reasonability condition (2), we have $[F(a_k) - F(a_{k+2})] \geq p_{a_{k+2}} > 0$, and $[F(a_{k+2}) - F(a_{k+1})] \geq p_{a_{k+1}} > 0$. Therefore, it holds that $c^{\pi^*} > c^{\pi'}$, contradicting the optimality of π^*.

For the case $j > 4$, the difference is more complicated:

$$\begin{aligned}
c^{\pi^*} - c^{\pi'} = & + a_{k+1}\big[F(a_k) + F(a_{k+1})\big] + a_{k+2}\big[F(a_{k+1}) + F(a_{k+2})\big] \\
& + a_{k+3}\big[F(a_{k+2}) + F(a_{k+3})\big] + \cdots + a_{k+j-2}\big[F(a_{k+j-3}) + F(a_{k+j-2})\big] \\
& + a_{k+j-1}\big[F(a_{k+j-2}) + F(a_{k+j-1})\big] \\
& - \Big(a_{k+2}\big[F(a_k) + F(a_{k+2})\big] \\
& + a_{k+1}\big[F(a_{k+2}) + F(a_{k+1})\big] + a_{k+j-1}\big[F(a_{k+1}) + F(a_{k+j-1})\big]\Big) \\
= & (a_{k+1} - a_{k+2})[F(a_k) - F(a_{k+2})] \\
& + a_{k+j-1}\big[F(a_{k+j-2}) - F(a_{k+1})\big] \\
& + a_{k+3}\big[F(a_{k+2}) + F(a_{k+3})\big] + \cdots + a_{k+j-2}\big[F(a_{k+j-3}) + F(a_{k+j-2})\big] \\
& - a_{k+2}[F(a_{k+2}) - F(a_{k+1})] \\
= & (a_{k+1} - a_{k+2})[F(a_k) - F(a_{k+2})] \\
& + a_{k+j-1}\big[F(a_{k+j-2}) - F(a_k)\big] \\
& + a_{k+3}F(a_{k+3}) + \cdots + a_{k+j-2}\big[F(a_{k+j-3}) + F(a_{k+j-2})\big] \\
& a_{k+3}F(a_{k+2}) - a_{k+2}[F(a_{k+2}) - F(a_{k+1})]
\end{aligned} \tag{48}$$

Notice the last line of Equation (48) is strictly positive, because $a_{k+3} > a_{k+1} > a_{k+2}$ and $F(a_{k+2}) > F(a_{k+2}) - F(a_{k+1})$ by the second reasonability condition. The term $a_{k+j-1}\big[F(a_{k+j-2}) - F(a_{k+1})\big]$ is strictly positive as well because $a_{k+j-2} < a_{k+1}$ according to the definition of j. For the same reason as in the case of $j = 4$, the term $(a_{k+1} - a_{k+2})[F(a_k) - F(a_{k+2})]$ is strictly positive. Furthermore, the term $a_{k+3}F(a_{k+3}) + \cdots + a_{k+j-2}\big[F(a_{k+j-3}) + F(a_{k+j-2})\big]$ is obviously positive. Therefore, we see that $c^{\pi^*} > c^{\pi'}$, contradicting the optimality of π^*.

Hence, we conclude an optimal policy for a UBDLSP with symmetric distribution, if it exists, must satisfy the expanding property $a_{k+1} > a_k$ for all $k \in \mathbb{N}$.

4.7.3 Calculating the Costs of Policies

We start with the general n-truncated distribution:

$$\mathrm{P} = \big(F(-m) + p_{-m}, p_{-(m-1)}, \cdots, p_{-1}, p_0, p_1, \cdots, p_{n-1}, p_n + F(n) \big).$$

Without loss of generality, consider a policy π for the distribution P:

$$\pi = \{-a_{2k+2}, -a_{2k}, -a_{2k-2}, \cdots, -a_2, 0, a_1, a_3, \cdots, a_{2k-1}, a_{2k+1}, a_{2k+3}\},$$

where a_{2k+2} and a_{2k+3} equal n. The cost of policy π can be calculated:

$$c^{\pi} = \sum_{i=-a_{2k+2}}^{a_{2k+3}} l^{\pi}(i) p_i,$$

where $l^{\pi}(i)$ is the traveled distance for the searcher, under the policy π, to first visit location i. It is easy to see that

$$l^{\pi}(i) = 2S_{2j} + |i|, \text{ for } i \in (a_{2j-1}, a_{2j+1}], j = 1, 2, \cdots, k+1,$$

$$l^{\pi}(i) = 2S_{2j-1} + |i|, \text{ for } i \in [-a_{2j}, -a_{2j-2}), j = 1, 2, \cdots, k+1,$$

where

$$S_i = \sum_{j=1}^{i} |a_j|, \ i = 1, 2, \cdots, 2k+1.$$

Now, the cost is calculated as:

$$
\begin{aligned}
c^{\pi} &= \sum_{i=-n}^{n} l_{\pi}(i)p_i \\
&= (2S_{2k+1} + a_{2k+2})(p_{a_{2k+2}} + F(-a_{2k+2})) \\
&\quad + \sum_{i\in(-a_{2k+2},-a_{2k})} (2S_{2k+1} + |i|)p_i \\
&\quad + \sum_{i\in[-a_{2k},-a_{2k-2})} (2S_{2k-1} + |i|)p_i \\
&\quad + \sum_{i\in[-a_{2k-2},-a_{2k-4})} (2S_{2k-3} + |i|)p_i \\
&\quad \cdots \\
&\quad + \sum_{i\in[-a_2,0)} (2S_1 + |i|)p_i \\
&\quad + \sum_{i\in(0,a_1]} (2S_0 + |i|)p_i \\
&\quad + \sum_{i\in(0,a_3]} (2S_2 + |i|)p_i \\
&\quad \cdots \\
&\quad + \sum_{i\in(a_{2k-3},a_{2k-1}]} (2S_{2k-2} + |i|)p_i \\
&\quad + \sum_{i\in(a_{2k-1},a_{2k+1}]} (2S_{2k} + |i|)p_i \\
&\quad + \sum_{i\in(a_{2k+1},a_{2k+3})} (2S_{2k+2} + |i|)p_i \\
&\quad + (2S_{2k+2} + a_{2k+3})(p_{a_{2k+3}} + F(a_{2k+3})).
\end{aligned}
$$

Expanding the terms out, we get:

$$
\begin{aligned}
c^{\pi} = {} & 2\left(\sum_{i=1}^{2k+1} a_i\right)\left[p_{-a_{2k+2}} + F(a_{2k+2})\right] + a_{2k+2}p_{-a_{2k+2}} + a_{2k+2}F(-a_{2k+2}) \\
& + 2\left(\sum_{i=1}^{2k+1} a_i\right)\left[\sum_{i\in(-a_{2k+2},-a_{2k})} p_i\right] + \sum_{i\in(-a_{2k+2},-a_{2k})} |i|p_i \\
& + 2\left(\sum_{i=1}^{2k-1} a_i\right)\left[\sum_{i\in[-a_{2k},-a_{2k-2})} p_i\right] + \sum_{i\in[-a_{2k},-a_{2k-2})} |i|p_i \\
& + 2\left(\sum_{i=1}^{2k-3} a_i\right)\left[\sum_{i\in[-a_{2k-2},-a_{2k-4})} p_i\right] + \sum_{i\in[-a_{2k-2},-a_{2k-4})} |i|p_i \\
& \cdots \\
& + 2\left(\sum_{i=1}^{1} a_i\right)\left[\sum_{i\in[-a_2,0)} p_i\right] + \sum_{i\in[-a_2,0)} |i|p_i \\
& + \sum_{i\in(0,a_1)} |i|p_i \\
& + 2\left(\sum_{i=1}^{2} a_i\right)\left[\sum_{i\in(a_1,a_3])} p_i\right] + \sum_{i\in(a_1,a_3]} |i|p_i \\
& \cdots \\
& + 2\left(\sum_{i=1}^{2k-2} a_i\right)\left[\sum_{i\in(a_{2k-3},a_{2k-1}]} p_i\right] + \sum_{i\in(a_{2k-3},a_{2k-1}]} |i|p_i \\
& + 2\left(\sum_{i=1}^{2k} a_i\right)\left[\sum_{i\in(a_{2k-1},a_{2k+1})} p_i\right] + \sum_{i\in(a_{2k-1},a_{2k+1})} |i|p_i \\
& + 2\left(\sum_{i=1}^{2k+2} a_i\right)\left[\sum_{i\in(a_{2k+1},a_{2k+3})} p_i\right] + \sum_{i\in(a_{2k+1},a_{2k+3})} |i|p_i \\
& + 2\left(\sum_{i=1}^{2k+2} a_i\right)\left[p_{a_{2k+3}} + F(a_{2k+3})\right] + a_{2k+3}p_{a_2k+3} + a_{2k+3}F(a_{2k+3}).
\end{aligned}
$$

Organizing the terms with respect to a_i, $i = 1, \cdots, 2k+3$, we obtain:

$$
\begin{aligned}
c^{\pi} = & \sum_{i=-a_{2k+2}}^{a_{2k+3}} |i| p_i \\
& + a_{2k+2} F(-a_{2k+2}) + a_{2k+3} F(a_{2k+3}) \\
& + 2a_{2k+1} \left(\left[\sum_{i \in [-a_{2k+2}, -a_{2k})} p_i + F(-a_{2k+2}) \right] + \left[\sum_{i \in (a_{2k+1}, a_{2k+3}]} p_i + F(a_{2k+3}) \right] \right) \\
& + 2a_{2k-1} \left(\left[\sum_{i \in [-a_{2k+2}, -a_{2k-2})} p_i + F(-a_{2k+2}) \right] + \left[\sum_{i \in (a_{2k-1}, a_{2k+3}]} p_i + F(a_{2k+3}) \right] \right) \\
& + 2a_{2k-3} \left(\left[\sum_{i \in [-a_{2k+2}, -a_{2k-4})} p_i + F(-a_{2k+2}) \right] + \left[\sum_{i \in (a_{2k-3}, a_{2k+3}]} p_i + F(a_{2k+3}) \right] \right) \\
& \\
& + 2a_1 \left(\left[\sum_{i \in [-a_{2k+2}, 0)} p_i + F(-a_{2k+2}) + \right] + \left[\sum_{i \in [a_1, a_{2k+2})} p_i + F(a_{2k+3}) \right] \right) \\
& + 2a_2 \left(\left[\sum_{i \in (a_1, a_{2k+3}]} p_i + F(a_{2k+3}) \right] + \left[\sum_{i \in [-a_{2k+2}, -a_2)} p_i + F(-a_{2k+2}) \right] \right) \\
& \cdots \\
& + 2a_{2k-2} \left(\left[\sum_{i \in (a_{2k-3}, a_{2k+3}]} p_i + F(a_{2k+3}) \right] + \left[\sum_{i \in [-a_{2k+2}, -a_{2k-2})} p_i + F(-a_{2k+2}) \right] \right) \\
& + 2a_{2k} \left(\left[\sum_{i \in (a_{2k-1}, a_{2k+3}]} p_i + F(a_{2k+3}) \right] + \left[\sum_{i \in [-a_{2k+2}, -a_{2k})} p_i + F(-a_{2k+2}) \right] \right) \\
& + 2a_{2k+2} F(a_{2k+1}).
\end{aligned}
$$

Simplifying and substituting a_{2k+2} and a_{2k+3} with n, we get

$$\begin{aligned} c^{\pi} = & \sum_{i=-a_{2k+2}}^{a_{2k+3}} |i| p_i \\ & + 2nF(a_{2k+1}) + nF(-n) + nF(n) \\ & + 2\sum_{j=1}^{k} a_{2j} \left[F(a_{2j-1}) + F(-a_{2j})\right] \\ & + 2\sum_{j=0}^{k} a_{2j+1} \left[F(-a_{2j}) + F(a_{2j+1})\right]. \end{aligned} \tag{49}$$

For n-truncated symmetric distributions, the cost formula (49) turns into

$$c^{\pi} = 2\sum_{i=1}^{n} |i| p_i + 2n \left[F(a_{2k+1}) + F(n)\right] + 2\sum_{j=1}^{2k+1} a_j \left[F(a_{j-1}) + F(a_j)\right]. \tag{50}$$

From the formula (49) for bounded distributions, we can easily obtain the following formula for unbounded distributions by taking both the left and the right boundaries to infinity:

$$\begin{aligned} c^{\pi} = & \sum_{i=-\infty}^{\infty} |i| p_i \\ & + 2\sum_{k=1}^{\infty} a_{2k} \left[F(a_{2k-1}) + F(-a_{2k})\right] \\ & + 2\sum_{k=0}^{\infty} a_{2k+1} \left[F(-a_{2k}) + F(a_{2k+1})\right]. \end{aligned} \tag{51}$$

In the case of *signed* turning-point representation for the policy:

$$(\cdots, a_{2k}, a_{2k-2}, \cdots, a_2, 0, a_1, a_3, \cdots, a_{2k-1}, a_{2k+1}, \cdots),$$

where $a_i < 0$ for i odd $a_i > 0$ for i even, the cost formula can be simplified to:

$$c^{\pi} = \sum_{i=-\infty}^{\infty} |i| p_i + 2\sum_{k=1}^{\infty} |a_k| \left[F(a_{k-1}) + F(a_k)\right]. \tag{52}$$

List of References

[1] B. O. Koopman, "Search and screening," Center for Naval Analysis, Tech. Rep. Operations Evaluation Group 56, 1946.

[2] B. O. Koopman, "The theory of search I. kinematic bases," *The Journal of the Operations Research*, vol. 4, pp. 324 346, 1956.

[3] B. O. Koopman, "The theory of search II. target detection," *The Journal of the Operations Research*, vol. 4, pp. 503 531, 1956.

[4] B. O. Koopman, "The theory of search III. the optimum distribution of searching effort," *The Journal of the Operations Research*, vol. 5, pp. 613 626, 1957.

[5] L. D. Stone, *Theory of Optimal Search.* Academic Press, 1975.

[6] L. D. Stone, "What's happened in search theory since the 1975 Lanchester prize," *Operations Research*, vol. 37, no. 3, pp. 501 506, 1989.

[7] R. Bellman, "Problem 63-9, an optimal search." *SIAM review*, vol. 5, no. 3, p. 274, 1963.

[8] A. Beck and M. Beck, "Son of the linear search problem," *Israel Journal of Mathematics*, vol. 48, pp. 109 122, 1984.

[9] W. Franck, "On an optimal search problem," *SIAM review*, vol. 7, pp. 503–512, 1965.

[10] A. Beck, "On the linear search problem," *Israel Journal of Mathematics*, vol. 2, pp. 221 228, 1964.

[11] A. Beck, "More on the linear search problem," *Israel Journal of Mathematics*, vol. 3, pp. 61 70, 1965.

[12] A. Beck and D. Newman, "Yet more on the linear search problem," *Israel Journal of Mathematics*, vol. 8, pp. 419 429, 1970.

[13] A. Beck and P. Warren, "The return of the linear search problem," *Israel Journal of Mathematics*, vol. 14, pp. 169 183, 1973.

[14] A. Beck and M. Beck, "The linear search problem rides again," *Israel Journal of Mathematics*, vol. 53, pp. 365 372, 1986.

[15] A. Beck and M. Beck, "The revenge of the linear search problem," *SIAM Journal on Control and Optimization*, vol. 30, pp. 112 122, 1992.

[16] V. Baston and A. Beck, "Generalizations in the linear search problem," *Israel Jornal of Mathematics*, vol. 90, pp. 301 323, 1995.

[17] W. Franck, "Errata: An optimal search problem," *SIAM Rev.*, vol. 8, p. 524, 1965.

[18] A. Beck, "An optimal search: Problem 63-9," *SIAM review*, vol. 27, no. 3, pp. 447 448, 1985.

[19] P. J. Rousseeuw, "Optimal search paths for random variables," *Journal of Computational and Applied Mathematics*, vol. 9, pp. 279–286, 1983.

[20] B. Fristedt and D. Heath, "Searching for a particle on the real line," *Advances in Applied Probability*, vol. 6, pp. 79 102, 1974.

[21] Z. T. Balkhi, "The generalized linear search problem, existence of optimal search paths," *Journal of the Operations Research Society of Japan*, vol. 30, pp. 399 421, 1987.

[22] F. T. Bruss and J. B. Robertson, "A survey of the linear-search problem," *Math. Scientist*, vol. 13, pp. 75 89, 1988.

[23] A. Washburn, "Dynamic programming and the backpacker's linear search problem," *Journal of Computational and Applied Mathematics*, vol. 60, pp. 357 365, 1995.

[24] P. J. Rousseeuw, "Search problem (linear)." [Online]. Available: http://eom.springer.de/S/s083630.htm

[25] M.-Y. Kao and M. L. Littman, "Algorithms for informed cows," in *Working Notes AAAI'97 Workshop on Online-Search*. 1997.

[26] J. L. Bentley and A. C.-C. Yao, "An almost optimal algorithm for unbounded searching," *Information Processing Letters*, vol. 5, pp. 82–87, 1976.

[27] R. A. Baeza-Yates, J. C. Culberson, and G. J. E. Rawlins, "Searching in the plane," *Information and Computation*, vol. 106, pp. 234 252, 1993.

[28] M.-Y. Kao, J. H. Reif, and S. R. Tate, "Searching in an unknown environment: An optimal randomized algorithm for the cow-path proble," *Information and Computation*, vol. 131, pp. 63–79, 1996.

[29] S. Gal, "Minimax solution for certain search problems," Ph.D. dissertation, Hebrew University, 1972.

[30] S. Gal, "A general search game," *Israel Journal of Math*, vol. 12, pp. 32 45, 1972.

[31] S. Gal, "Minimax solutions for linear search problems," *SIAM Journal of Applied Math*, vol. 27, pp. 17–30, 1974.

[32] S. Gal and D. Chazan, "On the optimality of the exponential functions for some minimax problems," *SIAM Journal of Applied Math*, vol. 30, pp. 324 348, 1976.

[33] S. Gal, *Search Games*. Academic Press, 1980.

[34] A. López-Ortiz and S. Schuierer, "The ultimate strategy to search on m rays?" *Lecture Notes in Computer Science*, vol. 1449, pp. 75 84, 1998.

[35] P. Jaillet and M. Stafford, "Online searching," *Operations Research*, vol. 49, pp. 501 515, 2001.

[36] A. López-Ortiz, "On-line searching on bounded and unbounded domain," Ph.D. dissertation, University of Waterloo, 1996.

[37] S. Schuierer, "Lower bounds in on-line geometric searching," *Computational Geometry: Theory and Applications*, vol. 18, pp. 37 53, 2001.

[38] R. Bellman, "Minimization problem," *Bulletin of the American Mathematical Society*, vol. 62, p. 270, 1956.

[39] S. R. Finch and J. E. Wetzel, "Lost in a forest," *The American Mathematical Monthly*, vol. 111, pp. 645 654, 2004.

[40] S. W. Williams, "Million buck problems," *Mathematical Intelligencer*, vol. 24, pp. 17 20, 2002.

[41] V. A. Zalgaller, "How to get out of the woods? on a problem of Bellman," *Matematicheskoe Prosveshchenie*, vol. 6, pp. 191 195. 1961.

[42] J. R. Isbell, "An optimal search problem," *Naval Res. Logist. Quart*, vol. 4, pp. 357 359, 1957.

[43] W. O. J. Moser, "Problems, problems, problems," *Discrete Applied Mathematics*, vol. 31, pp. 201 225, 1991.

[44] J. E. Wetzel, "Fits and covers," *Mathematics Magazine*, vol. 76, pp. 349 363, 2003.

[45] S. R. Finch, "The logarithmic spiral conjecture," 2005, arXiv. [Online]. Available: http://arxiv.org/abs/math.OC/0501133

[46] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of Markov decision processes," *Mathematics of Markov Decision Processes*, vol. 12, pp. 441 450, 1987.

[47] O. Madani, S. Hanks, and A. Condon, "On the undecidability of probabilistic planning and related stochastic optimization problems," *Artificial Intelligence*, vol. 147, pp. 5 34, 2003.

[48] A. Blum, P. Chalasani, D. Coppersmith, B. Pulleyblank, P. Raghavan, and D. Sudan, "The minimum latency problem," in *26th ACM Symposium on the Theory of Computing (STOC'94)*, 1994.

[49] A. García, P. Jodrá, and J. Jejel, "A note on the traveling repairman problem," *Networks*, vol. 40, pp. 27–31, 2002.

[50] F. Afrati, S. Cosmadakis, C. H. Papadimitriou, G. Papageorgiou, and N. Papakostantinou, "The complexity of the travelling repairman problem," *Theoretical Informatics and Applications*, vol. 20, pp. 79 87, 1986.

[51] J. N. Tsitsiklis, "Special cases of traveling salesman and repairman problems with time windows," *Networks*, vol. 1992, pp. 263 282, 22.

[52] T. G. Will, "Extremal results and algorithms for degree sequences of graphs," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1993.

[53] S. Sahni and T. Gonzalez, "P-complete approximation problems," *Journal of the Association for Computing Machinery*, vol. 23, pp. 555 565, 1976.

[54] A. Pelc, "Searching games with errors—fifty years of coping with liars," *Theoretical Computer Science*, vol. 270, pp. 71 109, 2002.

[55] M. J. Atallah, *Algorithms and Theory of Computation Handbook*. CRC Press LLC, 1999.

CHAPTER 5

On the Fundamental Limits of Target Trackability

5.1 Summary

We consider the problem of tracking a target that moves according to a Markov chain. A tracker queries a set of sensors to obtain tracking information. We are interested in finding the minimum number of queries per time step such that a target is trackable. We consider both the cases where the motion law, i.e., the transition probability function of the Markov chain, is known or unknown to the tracker a priori. In each case, three scenarios are analyzed. First we investigate the case where the tracker is required to know the exact location of the target at each time step. We then relax this requirement and explore the case where the tracker may lose track of the target at some time step, but it is able to "catch-up," regaining up-to-date information about the target's track at some later time step. Finally, we consider the case where tracking information is only known after a delay of d time steps. We provide necessary and sufficient conditions on the number of queries per time step needed to track in these three scenarios for each case (known or unknown motion law). These conditions are stated in terms of the entropy rate of the target's Markov chain. The work presented in this chapter is a joint effort with Patricia Barbosa.

5.2 Introduction

The problem of searching by asking questions has been the subject of extensive research for many years as we discussed in Chapter 4. Its origins can be traced back to Ulam [1] and Rényi [2], who introduced variations of the famous "twenty questions problem." Since then, several other formulations of this two-person game have been considered in the literature [3–6]. In this work, our goal is to study

the fundamental limits of target trackability and derive theoretical bounds on the number of queries per time step a tracker is required to ask a set of sensors to track a target.

The problem of target tracking through limited querying is of particular interest in the sensor network setting. Sensor networks have emerged as one of the most promising technologies in recent years. While much of the research done in this area explores networking issues like time synchronization [7, 8], sensor localization [9, 10], and routing [11, 12], additional communications problems such as data compression and message complexity have become increasingly important as the number of networked sensing devices continues to grow. For the majority of existing sensor networks, these small and inexpensive devices impose serious energy constraints affecting the network lifetime by having to transmit sensing information (over possibly long communication channels) to a remote monitoring station (estimator) [13]. Moreover, the reliability and the capacity of the channel available for communication with the estimator lead to restrictions on the volume of data sent over such networks. As a consequence, the estimator needs to make judicious decisions when selecting sensors to send data, so that communication with the sensor network is kept to a minimum.

It is within this setting that we propose a sensor model in which sensors are capable of sending only one-bit messages to an estimator. These messages are used to gather tracking information about a moving target. In the literature, one-bit-message sensor networks are called *binary sensor networks* and have been previously considered for target tracking [14–16]. In [17], Evans et al. analyzed the problem of optimal sensor selection; however their approach is to formulate the problem as a partially observed stochastic control problem, where sensors are not constrained to one-bit messages, and the estimator also controls the channel data

rate so that mean squared errors are bounded.

The remainder of this chapter is organized as follows. Section 5.3 formalizes the tracking problem under three different definitions. In Section 5.4, we study the case where the motion law is known to the tracker and present the necessary and sufficient conditions for followability, trackability, and d-trackability. In Section 5.5, we study the adaptive case where the motion law is unknown to the tracker a priori and present the necessary and sufficient conditions respectively for universal-followability, universal-trackability, and universal-d-trackability. The results are later proved in Section 5.6. Finally, Section 5.7 concludes this chapter.

5.3 Problem Formulation

Consider a target moving around an area. Suppose that the area is partitioned into a number of non-overlapping regions, referred to as locations. At each location, a sensor is deployed to monitor the motion of the target. We model the motion of the target by a discrete time finite state-space Markov chain

$$\{X_\iota : \iota \in \mathbb{N}\}$$

and take the set $\mathcal{X}$ of the indices of the locations as the state space of the Markov chain. We assume that the Markov chain is ergodic and time-homogeneous and denote the one-step transition probabilities by

$$p_{x,y} = \Pr\{X_\iota = y | X_{\iota-1} = x\},\ x, y \in \mathcal{X}.$$

We denote the history of the target motion up to time t by

$$X^t = (X_1, X_2, \cdots, X_t),$$

and call it the *target track* (up to time t).

Fig. 8 illustrates such a tracking sensor network, where the target track $x^5 = (4, 5, 1, 6, 2)$. Noting that the locations and the sensors are in one-to-one

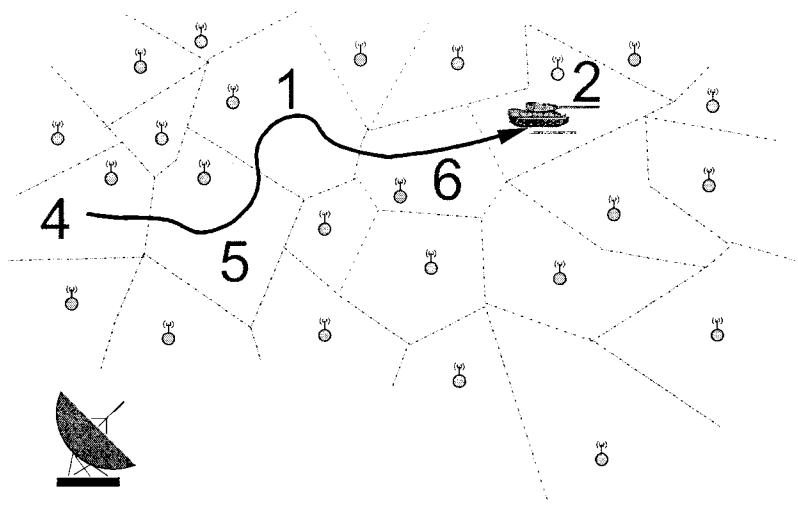

Figure 8. A tracking sensor network

correspondence, we index both of them by the same index set $\mathcal{X}$ and use the terms target location and sensor as synonymous in the following.

By querying the sensors, a tracker aims to track the target. At each time step, the tracker is allowed to query the sensors a number of times. We denote the ith query of time step t by $q_{t,i}$. Furthermore, each query, sent by the tracker to the sensors, consists of a number of *questions*, each of which addresses a particular sensor with a specific time stamp. Formally, we write

$$q_{t,i} = \{(s_{t,i}^j, \tau_{t,i}^j) : j \in J_{t,i}, \tau_{t,i}^j \leq t, s_{t,i}^j \in \mathcal{X}\}$$

where the pair $(s_{t,i}^j, \tau_{t,i}^j)$ denotes the question "has the sensor $s_{t,i}^j$ detected the target at time $\tau_{t,i}^j$?". In response to the query $q_{t,i}$, the tracker receives a binary response $r_{t,i} \in \{0, 1\}$.

We assume that the target detection of the sensors are flawless and non-overlapping, and the communication between the tracker and the sensors is error-free. Under these assumptions, the response $r_{t,i}$ to the query $q_{t,i}$ amounts to a

binary random variable as follows:

$$r_{t,i} = \begin{cases} 1, & \text{if } X^j_{\tau^j_{t,i}} = s^j_{t,i} \text{ for some } j \in J_{t,i}; \\ 0, & \text{otherwise.} \end{cases}$$

Denote the query-response history, up to the ith query-response round of time step t, by

$$Q_{t,i} = \Big((q_{1,1,}, r_{1,1}), \cdots, (q_{1,j_1}, r_{1,k_1}), \tag{53}$$

$$(q_{2,1,}, r_{2,1}), \cdots, (q_{2,j_1}, r_{2,k_2}), \tag{54}$$

$$\cdots$$

$$(q_{t-1,1,}, r_{t-1,1}), \cdots, (q_{t-1,k_{t-1}}, r_{t-1,k_{t-1}}), \tag{55}$$

$$(q_{t,1,}, r_{t,1}), \cdots, (q_{t,i}, r_{t,i}) \Big), \tag{56}$$

where k_t denotes the total number of queries issued at time step t. Note that (53) collects the k_1 queries of time step 1, (54) the k_2 queries of time step 2, (55) the k_{t-1} queries of time step $t-1$, and (56) the first i queries of time step t.

Definition 5.3.1. *A* policy Π *is a sequence of mappings from query and response history to next queries and updated track estimates:*

$$\pi_{t,i} : Q_{t,i} \mapsto (q_{t,i+1}, \hat{\mathrm{X}}^t_i) \text{ for } i < k_t \quad \text{and} \quad \pi_{t,k_t} : Q_{t,k_t} \mapsto (q_{t+1,1}, \hat{\mathrm{X}}^t_{k_t}),$$

where $\hat{\mathrm{X}}^t_i$ *denotes the* i*th estimate of the target track* X^t.

Because track estimates depend on policies, we denote the ith estimate of target track X^t under a specific policy Π by $\hat{\mathrm{X}}^t_i(\Pi)$. To make the word "tracking" precise, we consider the following three distinct degrees of "tracking," which we call *following*, *tracking*, and *d-tracking*.

Definition 5.3.2. *A policy* Π *is called a* following *policy if* $\hat{X}^t_{k_t}(\Pi) = X^t$ *holds for all* $t \in \mathbb{N}$ *almost surely.*

Definition 5.3.3. *A policy* Π *is called a* tracking *policy if* $\hat{X}^t_{k_t}(\Pi) = X^t$ *holds for infinitely many* $t \in \mathbb{N}$ *almost surely.*

Definition 5.3.4. *A policy* Π *is called a* d-tracking *policy if* $\hat{X}^t_{k_{t+d}}(\Pi) = X^t$ *holds for infinitely many* $t \in \mathbb{N}$ *almost surely, where* $\hat{X}^t_{k_{t+d}}(\Pi)$ *denotes the "partial track estimates" whose components are the first* t *components of the* $(t+d)$*-long vector of the track estimate* $\hat{X}^{t+d}_{k_{t+d}}(\Pi)$.

We consider in this work the situation where the tracker is allowed to query at most C times at each time step. We call the number C the *query quota*. Note that C is an integer. Corresponding to the above three distinct degrees of "tracking," we have the following three distinct degrees of "trackability."

Definition 5.3.5. *We say that a target is* followable *if there exist a following policy* Π *such that the number of queries used by* Π *at time step* t*, denoted* $k_t(\Pi)$*, is no more than* C *for all* $t \in \mathbb{N}$.

Definition 5.3.6. *We say that a target is* trackable *if there exists a tracking policy such that* $k_t(\Pi) \leq C$ *for all* $t \in \mathbb{N}$.

Definition 5.3.7. *We say that a target is* d-trackable *if there exists a* d*-tracking policy* Π *such that* $k_t(\Pi) \leq C$ *for all* $t \in \mathbb{N}$.

5.4 Tracking with Known Motion Law

In this section, we consider the case where the motion law – the transition probabilities of the Markov Chain – are known to the tracker. In this case, the mappings $\pi_{t,i}$, $t \in \mathbb{N}$, $i \in [k_t]$, of policies are allowed to be dependent on the transition probabilities $p_{x,y}$, $x, y \in \mathcal{X}$. Denote the entropy rate of the Markov chain by H, which is calculated as [18]:

$$H = \sum_{x \in \mathcal{X}} \pi_x \sum_{y \in \mathcal{X}} -p_{x,y} \log p_{x,y}. \tag{57}$$

We have the following results on the "trackability." Their proofs are relegated to Section 5.6 after we state their counterparts in the case of unknown motion law.

Theorem 5.4.1 (Followability). *A target is followable if and only if* $C \geq \max_{x \in \mathcal{X}} \log |N_x|$, *where* N_x *denotes the set* $\{y \in \mathcal{X} : p_{x,y} > 0\}$ *of the possible next states of* $x \in \mathcal{X}$.

Theorem 5.4.2 (Trackability).

a) If $C \geq H + 1$, *then the target is trackable.*

b) If the target is trackable, then $C \geq H$.

Theorem 5.4.3 (d-trackability).

a) If $C \geq H + \frac{1}{d}$, *then the target is* d*-trackable.*

b) If the target is d*-trackable for some* $d > 0$, *then* $C \geq H$.

5.5 Tracking with Unknown Motion Laws

In this section, we consider the case where the motion law of the target the transition probabilities of the Markov chain is *unknown* to the tracker a priori. In this case, the mappings $\pi_{t,i}$, $t \in \mathbb{N}$, $i \in [k_t]$, of policies are prohibited from depending on the transition probabilities $p_{x,y}$, $x, y \in \mathcal{X}$. We call such policies *universal* policies, following the terminology of information theory for universal coding [18]. Correspondingly, we have the following definitions for universal-following, universal-tracking, and universal-d-tracking.

Definition 5.5.1. *A universal policy* Π *is called a* universal-following *policy if* $\hat{X}^t_{k_t}(\Pi) = X^t$ *holds for all* $t \in \mathbb{N}$ *almost surely.*

Definition 5.5.2. *A universal policy* Π *is called a* universal-tracking *policy if* $\hat{X}^t_{k_t}(\Pi) = X^t$ *holds for infinitely many* $t \in \mathbb{N}$ *almost surely.*

Definition 5.5.3. *A universal policy* Π *is called a* universal-d-tracking *policy if* $\hat{X}_{k_{t+d}}^{t+d}(\Pi) = X^t$ *holds for infinitely many* $t \in \mathbb{N}$ *almost surely.*

Definition 5.5.4. *We say that a target is* universally followable *if there exists a universal-following policy* Π *such that* $k_t(\Pi) \leq C$ *for all* $t \in \mathbb{N}$.

Definition 5.5.5. *We say that a target is* universally trackable *if there exists a universal-tracking policy such that* $k_t(\Pi) \leq C$ *for all* $t \in \mathbb{N}$.

Definition 5.5.6. *We say that a target is* universally d-trackable *if there exists a universal-d-tracking policy* Π *such that* $k_t(\Pi) \leq C$ *for all* $t \in \mathbb{N}$.

As we shall see in the following, it is remarkable that the universality required for universal-tracking and universal-d-tracking strategies does NOT incur extra queries comparing with the case where motion law is known a priori. This resembles the well known result that there exist universal source encoders to code information source *optimally* even without knowing the statistics of the sources. The proofs for these results are relegated to Section 5.6 after their counterparts in the case of known motion law.

Theorem 5.5.7 (Universal-followability). *A target is universal-followable if and only if* $C \geq \log |\mathcal{X}|$.

Theorem 5.5.8 (Universal-trackability).

a) If $C \geq H + 1$, *then the target is universal-trackable.*

b) If the target is universal-trackable, then $C \geq H$.

Theorem 5.5.9 (Universal-d-trackability).

a) If $C \geq H + \frac{1}{d}$, *then the target is universal-d-trackable.*

b) If the target is universal-d-trackable for some $d > 0$, *then* $C \geq H$.

The part (a) of Theorem 5.5.9 can be proved based on Theorem 5.5.8 using the block-coding idea as presented in the proof of Theorem 5.4.3. The converse, i.e., the part (b), follows trivially from the part (b) of Theorem 5.5.8. Therefore, we omit its detailed proof in Section 5.6.

5.6 Proofs

5.6.1 Proof of Theorem 5.4.1

We first prove the necessity, by contradiction. Let $X_1 = x_1 \in \mathcal{X}$. Assuming $C < \max_{x\in\mathcal{X}} \log |N_x|$, there exists a state $x^* \in \mathcal{X}$ such that $C < \log |N_{x^*}|$. Since the Markov chain $\{X_t : t \in \mathbb{N}\}$ is ergodic, thus irreducible, there exists a time step $t^* \in \mathbb{N}$ such that the t^*-step transition probability from state x_1 to state x^* is strictly positive, i.e., $p^{(t^*)}_{x_1,x^*} > 0$ (where $p^{(1)}_{x_1x^*} = p_{x_1,x^*}$). From the definition of query quota, we know that at most C bits per time step can be transmitted to the tracker. Therefore, the number of choices for estimating X_{t+1} is at most 2^C at time step $t+1$. But it is clear that no policy is able to identify all possible choices for $X_{t+1} \in N_{x^*}$ with at most C queries with $2^C < |N_{x^*}|$. Therefore, we have

$$\mathrm{P}\left\{\hat{X}^{t+1}_{k_{t+1}} \neq X^{t+1} \middle| X_t = x^*\right\} \geq \min_{y\in\mathcal{X}} p_{x^*,y} > 0.$$

Hence, for any policy Π with $k_t < C$ for all $t \in \mathbb{N}$, we have

$$\mathrm{P}\left\{\text{ there exists a } t, \hat{X}^t_{k_t}(\Pi) \neq X^t\right\} \geq p^{(t^*)}_{x_1,x^*} \min_{y\in\mathcal{X}} p_{x^*,y} > 0,$$

a contradiction.

To prove the sufficiency, we show by induction that the simple and well-known binary search [19] yields a policy Π using which we can follow a target when $C \geq \max_{x\in\mathcal{X}} \log |N_x|$. Since the initial location of the target is known a priori, $\mathrm{P}\left\{\hat{X}^1_{k_1} = X_1\right\} = 1$ trivially. For a fixed $t > 0$, assume $\mathrm{P}\left\{\hat{X}^t_{k_t} = X^t\right\} = 1$. It suffices to show that $\mathrm{P}\left\{\hat{X}^{t+1}_{k_{t+1}} = X^{t+1}\right\} = 1$. It is easy to see that, using the binary search procedure, the tracker can pinpoint the target location x_{t+1} with at

most $\lceil \log |N_{x_t}| \rceil$ number of queries the set of sensors to be queried is repeatedly reduced by about half until the target location x_{t+1} is estimated with certainty. Noting C is an integer, we have

$$C \geq \max_{x \in \mathcal{X}} \log |N_x| \geq \log |N_{x_t}| \geq \lceil \log |N_{x_t}| \rceil.$$

Hence, we are able to estimate the value of X_{t+1} with certainty among all possible $|N_{x_t}|$ choices, that is, $\mathrm{P}\left\{ \hat{X}_{k_{t+1}}^{t+1} = X^{t+1} \right\} = 1$. Thus, it suffices to have $C \geq \max_{x \in \mathcal{X}} \log |N_x|$ to follow a target. □

5.6.2 Proof of Theorem 5.4.2

We first prove part (b)—we show that if there exists a *tracking* policy then $C \geq H$. It is to be proved by contraposition. The proof uses the idea of *strong typicality* and a result from large deviation theory.

First, we extend the concept of *strong typicality* [20, Ch.5] to ergodic finite-state Markov chains. Consider a ergodic time-homogeneous Markov chain $\{X_t : t \in \mathbb{N}\}$ with a finite state space $\mathcal{X}$ and one-step transition probabilities $p_{x,y}$, $x, y \in \mathcal{X}$. Denote its stationary distribution by π_x, $x \in \mathcal{X}$. Fixed a $t \in \mathbb{N}$, we define a counting function $N_{x,y} : \mathcal{X}^t \to [t]$ for each transition (x, y) as follows:

$$\mathbf{N}_{x,y}(x^t) = \sum_{k=1}^{t-1} \mathbf{1}_x(x_k)\, \mathbf{1}_y(x_{k+1}),$$

where $\mathbf{1}_x(x_t)$ is the indicator function taking the value of 1 if $x_t = x$ and 0 otherwise. Given a $\delta > 0$, the set $\Delta_{t,\delta}$ defined as follows is called a (t, δ)-*strongly-typical-set*:

$$\Delta_{t,\delta} = \left\{ x^t \in \mathcal{X}^t : \left| \frac{\mathbf{N}_{x,y}(x^t)}{t-1} - \pi_x p_{x,y} \right| < \delta,\ \forall (x, y) \in \mathcal{X}^2 \right\}.$$

The sequences x^t in this set are called the *strongly-typical-sequences*. Strongly-typical-sets and strongly-typical-sequences have the following useful properties.

Lemma 5.6.1. *a) The probability of every strong typical sequence* $x^t \in \Delta_{t,\delta}$ *satisfies*

$$2^{-t(H+c_1\delta)} < p(X^t) < 2^{-t(H-c_1\delta)},$$

where H *is the entropy rate of the Markov chain and the constant* $c_1 > 0$ *a constant.*

b) The probability of the complement of strong typicality sets eventually decreases exponentially fast there exist a $T \in \mathbb{N}$ *and constants* $c_2 > 0$ *and* $c_3 > 0$ *such that*

$$\Pr\left\{X^t \notin \Delta_{t,\delta}\right\} \leq c_2 2^{-c_3 t}, \text{ for all } t > T.$$

The part (a) of the lemma is established by modifying the proof of the similar results for i.i.d. sequences as in [20, Ch.5]. The part (b) is proved by using a result from large deviation theory for finite Markov chains. To avoid interrupting the main idea, we delay the proof of Lemma 5.6.1 to Appendix 5.8.1.

We assume $C < H$ to prove the part (b) of Theorem 5.4.2 by contraposition. Using Lemma 5.6.1, we bound the probability of the "catching-up" event $\{\hat{X}_{k_t}^t = X^t\}$ as follows:

$$\begin{aligned}
\Pr\left\{\hat{X}_{k_t}^t = X^t\right\} &= \sum_{x^t} \Pr\left\{\hat{X}_{k_t}^t = x^t | X^t = x^t\right\} p(x^t) && (58)\\
&= \sum_{x^t \in \Delta_{t,\delta}} \Pr\left\{\hat{X}_{k_t}^t = x^t | X^t = x^t\right\} p(x^t) && (59)\\
&\quad + \sum_{x^t \notin \Delta_{t,\delta}} \Pr\left\{\hat{X}_{k_t}^t = x^t | X^t = x^t\right\} p(x^t) && (60)\\
&\leq \sum_{x^t \in \Delta_{t,\delta}} \Pr\left\{\hat{X}_{k_t}^t = x^t | X^t = x^t\right\} p(x^t) + c_2 2^{-c_3 t} && (61)\\
&< \sum_{x^t \in \Delta_{t,\delta}} \Pr\left\{\hat{X}_{k_t}^t = x^t | X^t = x^t\right\} 2^{-t(H-c_1\delta)} + c_2 2^{-c_3 t} && (62)
\end{aligned}$$

Because there are at most 2^{tC} choices for $\hat{X}^t_{k_t}$, we have

$$\sum_{x^t \in \Delta_{t,\delta}} \Pr\left\{ \hat{X}^t_{k_t} = x^t \middle| X^t = x^t \right\} \leq 2^{tC}.$$

The inequality (62) turns into

$$\Pr\left\{ \hat{X}^t_{k_t} = X^t \right\} < 2^{tC} 2^{-t(H-c_1\delta)} + c_2 2^{-c_3(n-1)} = 2^{-t[(H-C)-c_1\delta]} + c_2 2^{-c_3 t}.$$

Hence, we have

$$\sum_{t \in \mathbb{N}} \Pr\left\{ \hat{X}^t_{[1:t]} = X^t \right\} < \infty,$$

by choosing some $\delta > 0$ such that $H - C - c_1\delta > 0$ holds. By the first Borel-Cantelli lemma [21], we have

$$\Pr\left\{ \hat{X}^t_{k_t} = X^t \text{ i.o.} \right\} = 0,$$

which contradicts our assumption that the *tracking* policy $\mathcal{S}$ "catches up" infinitely often with probability 1.

The part (a) of the theorem — if $C \geq H+1$ then there exists a *tracking* policy — is to be proved by construction. We construct a *tracking* policy $\mathcal{S}$ with the at most C number of queries at each time step. The policy to be constructed is based on the idea of Huffman codes. For this reason, we call it the *Huffman policy*. Under the Huffman policy, the tracker proceeds to query by "traversing" on a growing decision tree built from the Huffman-code trees of individual "transitional" random variables, as we shall elaborate in the following.

Recall that, for a finite state-space random variable X, we can construct the Huffman code for X such that the expected code word length $\bar{l} = \sum_x l(x)p(x)$ is minimized [18, Thm. 5.8.1]. During the construction of the Huffman codes for X, a code tree, which we call the *Huffman-code-tree* is constructed as byproduct. This Huffman-code-tree in fact provides an optimal query policy to "guess" the outcome of the random variable X [18, Sec. 5.7]. Relevant to our case, the expected number

of queries, equal to the expected length of the code words, used to pinpoint the outcome of X is bounded as follows:

$$H(X) \leq \bar{l} < H(X) + 1, \tag{63}$$

where $H(X)$ is the entropy of the random variable X.

Based on Huffman-code-trees, a querying policy is constructed as follows:

- At time step 1, since the initial location x_1 of the target is known to the tracker, the tracker needs to do nothing.

- At time step 2, the tracker constructs the *Huffman-code-tree* T_2 for the random variable Y_2 with probability distribution

  $$\Pr\{Y_2 = y\} = p_{x_1,y},\ y \in N_{x_1}.$$

  The tracker proceeds to query the sensors by traversing T_1 with at most C steps. The tracker

  - either reaches a leaf, at which point the target location x_2 is identified. Denote the number of queries taken by $l(x_2|x_1)$. The tracker construct a new *Huffman-code-tree* T_3, to be used in the following time step, for the new random variable Y_3 with distribution

    $$\Pr\{Y_3 = y\} = p_{x_2,y},\ y \in N_{x_2}.$$

  - or is still in the middle of the tree T_2.

- At time step 3,

  - if in the previous time step, the target location x_2 has already been identified, then the tracker starts to traverse T_3 and query the sensors accordingly, similar to the time step 2.

  - otherwise, the tracker continues on T_1 with at most C steps. If the tracker reaches a leaf of T_1, then it constructs T_3 and continues on T_2 until C queries are used up. In the end of time step 2, the tracker
    * either reaches a leaf of T_3, at which point the target location x_3 is identified, which takes $l(x_3|x_2)$ queries,
    * or is still in the middle of the tree T_3.

- Continuing this process,

- At time step t,
  - If the tracker has identified x_{t-1} in the previous time step, then it start to traverse T_t and query accordingly about x_t
  - Otherwise, the tracker must be in the middle of some previous Huffman-code-tree T_τ, $\tau \leq t-1$. In this case, the tracker continues traverses T_k, $k = \tau, \tau+1.\cdots.t$, sequentially until it either uses up the C queries or reaches a leaf of T_t.

Fig. 9 illustrates the construction of Huffman policy by concatenating Huffman-code-trees T_t, $t = 2, 3, \cdots$, along the time. As we see from the above construction, the tracker may lag behind the current target track temporarily. But, as we show in the following, the tracker can always "catch up" the above policy is a tracking policy. To see this, consider the following sequence of random variables:

$$\{l_t\} = \left\{ \frac{\sum_{k=1}^{t} l(X_{k+1}|X_k)}{n} : t = 1, 2, \cdots \right\}$$

where $l(X_{k+1}|X_k)$ denotes the number of queries used, under the above Huffman policy, to pinpoint the target location X_{k+1} given the target is at location X_k at time step k. We claim that the sequence converges almost surely and its limit $\bar{l}$ is bounded from above by $H+1$, where H is given by (57).

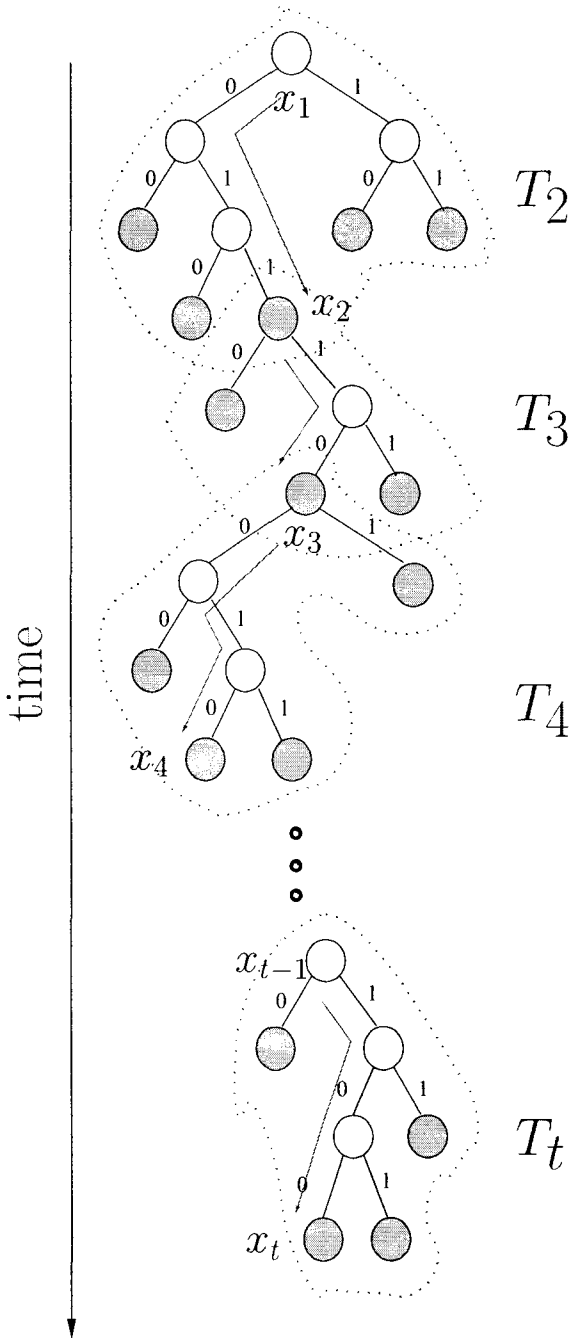

Figure 9. Concatenating Huffman-code-trees to form Huffman policy

Assuming the claim is true, then we have

$$\Pr\{\text{ there exists } t \in \mathbb{N} \text{ such that } l_t < C\} = 1 \tag{64}$$

This can be shown by the following argument based on contradiction. Suppose otherwise, i.e.,

$$\Pr\{l_t > C, \forall t > 0\} = p > 0.$$

Then, since $C \geq H + 1 > \bar{l}$,

$$\Pr\{l_t \geq \bar{l}, \forall t > 0\} \geq \Pr\{l_t \geq C, \forall t > 0\} = p > 0$$

This contradicts to our claim:

$$\Pr\{l_t \to \bar{l}\} = 1.$$

Note that when $l_t < C$ first happens the tracker catches up the current track. Hence, by induction and (64), the tracker catches up the current track infinitely often with probability 1. Therefore, the Huffman policy is a tracking policy.

It remains to show the claim that the sequence l_t

$$\left\{ \frac{\sum_{i=1}^{t} l(X_i|X_{i-1})}{t} : t = 1, 2, \cdots \right\}$$

converges almost surely and its limit $\bar{l}$ is bounded $H + 1$. Consider the following discrete-time random process

$$\left\{ l(X_{k+1}|X_k) : t = 1, 2, \cdots \right\},$$

where $l(y|x)$, $x, y \in \mathcal{X}$ is a time-homogeneous function with $l(y|x) = l(X_{k+1} = y|X_k = x)$, denoting the number of queries used by the Huffman policy to identify the target location y if the target moves from a previous location x. Clearly the function $l(y|x)$ is bounded. By the generalized convergence theorem for bounded functions of discrete-time finite ergodic Markov chains [22], we have

$$\frac{\sum_{k=1}^{t} l(X_{k+1}|X_k)}{t} \xrightarrow{a.s.} \bar{l} = \sum_{x,y} \pi_x p_{x,y} l(y|x),$$

where π_x is the stationary probability of state $x \in \mathcal{X}$. By the property of Huffman code [18, Thm. 5.4.1], the expected code length $\sum_y p_{x,y} l(y|x)$ for any given $x \in \mathbf{X}$ satisfies:

$$-\sum_y p_{x,y} \log p_{x,y} \leq \sum_y p_{x,y} l(y|x) < -\sum_y p_{x,y} \log p_{x,y} + 1.$$

Therefore, we have

$$\sum_i \pi_x [-\sum_y p_{x,y} \log p_{x,y}] \leq \bar{l} < \sum_x \pi_x [-\sum_y p_{x,y} \log p_{x,y} + 1].$$

In other words, the following inequality holds:

$$H \leq \bar{l} < H + 1,$$

where $H = -\sum_{x,y} \pi_x p_{x,y} \log p_{x,y}$ is the entropy rate of the Markov chain measured in bits per time step. Thus, the claim is established. This concludes the proof of both the part (a) of and, hence, the entire theorem. □

5.6.3 Proof of Theorem 5.4.3

Part (b) is proven once again using contradiction and strong typicality. Similarly to the tracking case, we assume that $C < H$. Then, the track estimate $\hat{X}_{k_{t+d}}^{t+d}$ has at most $2^{(t+d)C}$ choices. Consequently, the probability of the event $\left\{\hat{X}_{k_{t+d}=X^t}^{t}\right\}$ is bounded by

$$\Pr\left\{\hat{X}_{k_{t+d}}^{t} = X^t\right\} < 2^{-(t+d)[(H-C)-c_1\delta]} + c_2 2^{-c_3(t+d)}.$$

Therefore, we have

$$\sum_{t\in\mathbb{N}} \Pr\left\{\hat{X}_{k_{t+d}}^{t} = X^t\right\} < \infty.$$

Again, by the first Borel-Cantelli lemma, we have

$$\mathrm{P}\left\{\hat{X}_{k_{t+d}}^{t} = X^t \quad \text{i.o.}\right\} = 0,$$

a contradiction. Thus, $C \geq H$.

We show part (a) using a block version of the catch-up policy described in the proof of Theorem 5.4.2. Consider the sequence of random variables $\{W_n : n \in \mathbb{N}\}$, where $W_n = (X_{d(n-1)+1}, \ldots, X_{dn})$, $d > 0$, that is, each random variable W_n is a segment of length d of the sequence $\{X_t : t \in \mathbb{N}\}$. We call the sequence $\{W_n : n \in \mathbb{N}\}$ a *block Markov chain* taking values in the state space $\mathcal{X}^d$. Assuming $C \geq H + \frac{1}{d}$, and given the initial target location x_0, we skip querying during the first d time steps. For each time step t, from $t = d+1$ to $t = 2d$, we apply the catch-up policy to get $\hat{X}^d$. This is done using the transition probabilities of the Markov chain $\{W_n : n \in \mathbb{N}\}$ to generate Huffman codewords. Thus, at $t = 2d$, we have the estimate $\hat{W}_1 = \left(\hat{X}_1, \hat{X}_2, \ldots, \hat{X}_d\right)$. This procedure is repeated for every "block" of d time steps, hence with at most $C_W = dC$ number of queries for each "block" of d time steps. Moreover, the entropy rate H_W of the Markov chain $\{W_n : n > 0\}$ can be calculated in terms of the entropy rate H of the original Markov chain as

$$\begin{aligned} H_W &= \\ &= \lim_{n\to\infty} \frac{-\log \Pr\left\{X^{nd} = x^{nd}\right\}}{n} \\ &= \lim_{n\to\infty} d \left\{ \frac{-\log \Pr\left\{X^{nd} = x^{nd}\right\}}{nd} \right\} \\ &= -d \sum_{x,y\in\mathcal{X}} \pi_x p_{x,y} \log p_{x,y}, \end{aligned}$$

that is, $H_W = dH$.

By Theorem 5.4.2, if $C_W \geq H_W + 1$, that is, if $C \geq H + \frac{1}{d}$ and $d > 0$, a target is d-trackable. □

5.6.4 Proof of Theorem 5.5.7

The sufficiency part is obvious—with $\lceil \log |\mathcal{X}| \rceil$ number of queries, the tracker can pinpoint the location of the target at any time step.

We prove the converse if a target is universally followable, then $C \geq \log |\mathcal{X}|$ by contradiction. Suppose there exists a universal-following policy Π that

can follow the target with $C < \log|\mathcal{X}|$. Since the policy Π is required to be independent of the transition function $p_{x,y}$, $x, y \in \mathcal{X}$, the policy Π should be invariant with respect to the numerical values of $p_{x,y}$, $x, y \in \mathcal{X}$. Consider that the target moves according to a new transition function with $p'_{x,y} > 0$ for all $x, y \in \mathcal{X}$. It is clear that no policy can follow such a target with $C < \lceil \log|X| \rceil$, contradicting to the assumption that the universal policy Π follows the target. This concludes the proof. □

5.6.5 Proof of Theorem 5.5.8

The part (b) is obvious logically since $C \geq H$ is necessary for policies that are not required to be universal to track the target. In the following, we use $\mathbf{p}$ to denote the transition probability function (matrix) of the Markov chain:

$$(p_{x,y})_{x,y \in \mathcal{X}},$$

where $p_{x,y} = \Pr\{X_t = y | X_{t-1} = x\}$.

We prove part (a) by adding a "learning" components to the Huffman policy, constructed to prove part (a) of Theorem 5.4.2. We call the policy to be constructed the *learning Huffman policy*.

To learn the target motion law, the tracker keeps updating its estimate $\hat{\mathbf{p}}^t$, $t = 1, 2, \cdots$, of the transition probability function $\mathbf{p}$ during the tracking process, as follows. To help the estimation of the transition probability function, the tracker keeps a table, which we call *transition-counting-table*, to record the number of individual transitions $m_{x,y}$, $x, y \in \mathcal{X}$. Denote the transition-counting-table by $\mathbf{m}$. The transition table $\mathbf{m}$ is initialized with $m_{x,y} = 1$ for all $x, y \in \mathcal{X}$. The estimate $\hat{\mathbf{p}}^1$ of the probability function is initialized by normalizing $\mathbf{m}$ as follows:

$$\hat{p}^1_{x,y} = \frac{m_{x,y}}{\sum_{y \in \mathcal{X}} m_{x,y}}, \text{ for all } x, y \in \mathcal{X}.$$

Note that we refer to a particular entry of the transition probability matrix estimate $\hat{\mathbf{p}}^t$ by $\hat{p}^t_{x,y}$ and similarly a particular entry of $\mathbf{m}$ by $m_{x,y}$.

The tracker uses $\hat{\mathbf{p}}^1$ to construct the Huffman-code-tree T_2 according to the probability distribution

$$\big(\hat{p}^1(x_1, y) : y \in \mathcal{X}\big),$$

and then queries according to T_2. The target location x_2 is identified after $l_{\hat{p}^1}(x_2|x_1)$ number of queries. The tracker increases the entry m_{x_1,x_2} of $\mathbf{m}$ by 1 and keep the other entries unchanged, and then updates $\hat{\mathbf{p}}^2$ by normalizing $\mathbf{m}$:

$$\hat{p}^2_{x,y} = \frac{m_{x,y}}{\sum_{y\in\mathcal{X}} m_{x,y}}, \text{ for all } x, y \in \mathcal{X}$$

Then, the tracker constructs the Huffman-code-tree T_3 using the probability distribution

$$\big(\hat{p}^2_{x_2,y} : y \in \mathcal{X}\big),$$

and queries according to T_3.

At a typical time step t, after the target location x_τ, $\tau \leq t$, is pinpointed, the tracker first updates the transition-counting-table $\mathbf{m}$ by increasing the entry $m_{x_{\tau-1},x_\tau}$ by 1 and then updates its estimate $\hat{\mathbf{p}}^\tau$ of the transition probability function by normalizing $\mathbf{m}$:

$$\hat{p}^\tau_{x,y} = \frac{m_{x,y}}{\sum_{y\in\mathcal{X}} m_{x,y}}, \text{ for all } x, y \in \mathcal{X}$$

Then, the tracker constructs the Huffman-code-tree T_τ according to the probability distribution

$$\big(\hat{p}^\tau_{x_\tau,y} : y \in \mathcal{X}\big),$$

and proceeds to query according to T_τ.

It is clear that the learning Huffman policy constructed in the above is a universal policy. We are left to show it is a *tracking* policy. To see this, note

$$\hat{\mathbf{p}}^t \xrightarrow{a.s.} \mathbf{p},$$

because every entry $\hat{p}^t_{x,y}$, $x, y \in \mathcal{X}$, converges almost surely to $p_{x,y}$ by the generalized convergence theorem for Markov chain [22] and the transition probability matrix has only finitely many ($|\mathcal{X}|$) entries.

Consider the sequence of the number of queries used by the policy we constructed above:

$$\left\{ l_{\hat{\mathbf{p}}^t}(X_t|X_{t-1}) \right\}.$$

Applying again the generalized convergence theorem for bounded functions of an ergodic Markov chain, we have

$$\frac{\sum_{k=1}^{t} l_{\hat{\mathbf{p}}^k}(X_{k+1}|X_k)}{t} < \frac{\sum_{k=1}^{t} \left(\hat{H}^k(X_k) + 1\right)}{t} \xrightarrow{a.s.} H + 1,$$

where $\hat{H}^k : x \mapsto \sum_{y \in \mathcal{X}} -\hat{p}^k_{x,y} \log \hat{p}^k_{x,y}$, $x \in \mathcal{X}$, is the empirical transition entropy (function) with respect to the estimated transition probability function $\hat{\mathbf{p}}^t$. Following the same argument in the proof for the part (a) of Theorem 5.4.2, we can see that the learning policy tracks the target if $C \geq H + 1$. This concludes the proof of the theorem. □

5.7 Discussion

In this chapter, we have studied the number of queries required to follow, track, and d-track a target that moves according to a Markov chain for both the cases with the motion law known and unknown a priori to the tracker. Necessary and sufficient conditions have been presented for all cases, as well as corresponding following, tracking, and d-tracking policies. Our results can be applied to the multi-target scenario by considering a larger state space Markov chain by taking as states the vectors of the locations of multiple targets. It is of interest to consider the case where sensors are faulty (i.e., their query responses may be wrong), and where noise is present in the communication between sensors and estimator. In this direction, it would be natural to introduce the notion of distance between sensors (states)

and analyze tracking performance under criteria such as the mean squared error. We conjecture that results related to rate-distortion theory are possible. Another interesting variation is to take sensor responses to be the number of sensors that reply "yes" to a query. Future work also includes investigating the mean number of time steps (in terms of number of queries) involved in the *catch-up* policy before the target track can be estimated, i.e., the mean *lag* time. Although the simplicity of the (learning) Huffman policy is particularly attractive, it is of interest to find the policy that incurs the minimum lag time.

5.8 Appendix

5.8.1 Proof of Lemma 5.6.1

To show part (a), we investigate $-\log p(x^t)$. It can be bounded from the above as follows:

$$\begin{aligned} -\log p(x^t) &= \sum_{k=0}^{t-1} -\log p(x_{k+1}|x_k) && (65)\\ &= \sum_{x,y} N_{x,y}(x^t)(-\log p_{x,y}) && (66)\\ &< t\sum_{x,y}(\pi_x p_{x,y}+\delta)(-\log p_{x,y}) && (67)\\ &= t\Big[\sum_{x,y}(\pi_x p_{ij})(-\log p_{x,y}) + \sum_{x,y}\delta(-\log p_{x,y})\Big] && (68)\\ &= t(H + c_1\delta) && (69) \end{aligned}$$

where $c_1 = \sum_{x,y} -\log p_{x,y}$.

Similarly, we can bound it from the below:

$$t(H - c_1\delta) < -\log p(x^t).$$

Therefore, we have

$$2^{-t(H+c_1\delta)} < p(x^t) < 2^{-t(H-c_1\delta)},$$

which concludes our proof for part (a).

To prove part (b), we use a result from large deviation theory as stated in the following Lemma 5.8.1 [23].

Lemma 5.8.1. *Suppose that $\{X_t\}$ is an ergodic finite state chain with state space $\mathcal{X}$ and let b_t denote its L^1 convergence parameter:*

$$b_t = \sup_x \sup_y |p^t(x,y) - \pi_y|.$$

Then the series $b = \sum_{t>0} b_t$ converges and for any bounded function $F : \mathcal{X} \to \mathbb{R}$ and any $\delta > 0$ we have,

$$\log \Pr\left\{ \frac{\sum_{k=1}^t F(X_k)}{t} - \pi(F) \geq \delta \right\} \leq -\frac{t-1}{2}\left(\frac{\delta}{b\bar{F}} - \frac{3}{t-1}\right)^2,$$

as long as $t \geq 1 + 3b\bar{F}/\delta$, where $\bar{F} = \max_x |F(x)|$ and $\pi(F)$ is the mean of the function F with respect to the stationary distribution π of the Markov chain, that is, $\pi(F) = \sum_{x\in\mathcal{X}} F(x)\pi_x$.

To apply Lemma 5.8.1 and prove part (b), we first construct an ergodic finite-state Markov chain $\{Y_t\}$ from the original Markov chain $\{X_t\}$ by taking $Y_t = (X_{t-1}, X_t)$. It is well known that $\{Y_t\}$ is again an ergodic finite state Markov chain with state space $\mathcal{Y} = \mathcal{X}^2$. It is easy to verify that the stationary distribution $\lambda_{x,y}$ of $\{Y_t\}$ equals to $\pi_x p_{x,y}$, $(x,y) \in \mathcal{X}^2$. Let b_t be the L^1 convergence parameter sequence of $\{Y_t\}$ and $b = \sum_{t>0} b_t$. Let function $F(Y_t)$ be an indication function $\mathbf{1}_{x,y}(Y_t)$. Note that F is bounded and $\bar{F} = \sup_y |F(y)| \leq 1$. Applying Lemma 5.8.1 to function F, we have,

$$\log \Pr\left\{ \frac{\sum_{k=1}^t 1_{ij}(Y_k)}{t} - \lambda_{ij} \geq \delta \right\} \leq -\frac{t-1}{2}\left(\frac{\delta}{b} - \frac{3}{t-1}\right)^2 \leq -\frac{t-1}{2}\left(\frac{\delta}{b} - 3\right)^2,$$

for all $(x,y) \in \mathcal{X}^2$ for t large ($\geq 1 + 3b/\delta$). Writing the above inequality in terms of the counting function $N_{x,y}(X^t)$ and substituting $\lambda_{x,y}$ by $\pi_x p_{x,y}$, we get,

$$\log \Pr\left\{ \frac{N_{x,y}(X^t)}{t} - \pi_x p_{x,y} \geq \delta \right\} \leq -\frac{t-1}{2}\left(\frac{\delta}{b} - 3\right)^2, \tag{70}$$

for t large.

Similarly, applying Lemma 5.8.1 to the function $F' = 1 - F$, we get,

$$\log \Pr\left\{\frac{N_{x,y}(X^t)}{t} - \pi_x p_{x,y} \leq -\delta\right\} \leq -\frac{t-1}{2}\left(\frac{\delta}{b} - 3\right)^2, \tag{71}$$

for t large.

Combining inequality (70) and (71), we have

$$\log \Pr\left\{\left|\frac{N_{ij}(X^t)}{t} - \pi_i p_{ij}\right| \geq \delta\right\} \leq -(t-1)\left(\frac{\delta}{b} - 3\right)^2, \tag{72}$$

for t large. With the above inequality (72), we can now bound the probability of the complement of the strong typical set $\Delta_{t,\delta}$ as follows:

$$\begin{aligned}
\Pr\left\{X^t \notin \Delta_{t,\delta}\right\} &= \Pr\left\{\bigcup_{(x,y)\in\mathcal{X}^2}\left\{\left|\frac{N_{x,y}(X^t)}{t} - \pi_x p_{x,y}\right| \geq \delta\right\}\right\} && (73)\\
&\leq \sum_{(x,y)\in\mathcal{X}^2} \Pr\left\{\left|\frac{N_{x,y}(X^t)}{t} - \pi_x p_{x,y}\right| \geq \delta\right\} && (74)\\
&\leq |\mathcal{X}|^2 2^{-(t-1)(\frac{\delta}{b}-3)^2} && (75)\\
&\leq c_2 2^{-c_3 t} && (76)
\end{aligned}$$

for t large, where the constants $c_2 = |\mathcal{X}|^2 2^{(\frac{\delta}{b}-3)^2}$ and $c_3 = (\frac{\delta}{b} - 3)^2$. This concludes our proof for the part (b) of Lemma 5.6.1 and for the proof of the part (b). □

List of References

[1] S. M. Ulam, *Adventures of a Mathematician.* New York: Scribner, 1976.

[2] A. Rényi, "On a problem of information theory," *MTA Mat. Kut. Int. Kozl.*, vol. 6B, pp. 505–516, 1961.

[3] A. Pelc, "Searching with known probability of error," *Theoretical Computer Science*, vol. 63, pp. 185–202, 1989.

[4] A. Dhagat, P. Gacs, and P. Winkler, "On playing twenty questions with a liar," in *Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Orlando, Florida, Jan. 1992, pp. 16–22.

[5] R. Hill, J. Karim, and E. Berlekamp, "The solution of a problem of Ulam on searching with lies," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, Cambridge, Massachusetts, August 1998.

[6] A. Ambainis, S. A. Bloch, and D. L. Schweizer, "Delayed binary search, or playing twenty questions with a procrastinator," *Algorithmica*, vol. 32, pp. 641–651, 2002.

[7] S. Ganeriwal, R. Kumar, and M. Srivastava, "Timing-sync protocol for sensor networks," in *Proceedings of ACM Conference on Embedded Networked Sensor Systems (SENSYS)*, Los Angeles, California, November 2003.

[8] J. Elson and D. Estrin, "Time synchronization for wireless sensor networks," in *Proceedings of the* 15^{th} *International Parallel and Distributed Processing Symposium*, San Francisco, California, April 2001, p. 186.

[9] R. Iyengar and B. Sikdar, "Scalable and distributed gps free positioning for sensor networks," in *Proceedings of IEEE International Conference on Communications(ICC)*, Anchorage, Alaska, May 2003.

[10] N. Bulusu, J. Heidemann, and D. Estrin, "Gps-less low-cost outdoor localization for very small devices," *IEEE Personal Communications*, vol. 7, no. 5, pp. 28–34, October 2000.

[11] D. Tian and N. Georganas, "Energy efficient routing with guaranteed delivery in wireless sensor networks," in *Proceedings of IEEE Wireless Communications and Networking Conference(WCNC)*, New Orleans, Louisiana, March 2003, pp. 1923–1929.

[12] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly-resilient, energy-efficient multipath routing in wireless sensor networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 11–25, 2001.

[13] J. Carle and D. Simplot-Ryl, "Energy-efficient area monitoring for sensor networks," *IEEE Computer*, vol. 37, no. 2, pp. 40–46, 2004.

[14] J. Aslam, Z. Butler, F. Constantin, V. Crespi, G. Cybenko, and D. Rus, "Tracking a moving object with a binary sensor network," in *Proceedings of ACM Conference on Embedded Networked Sensor Systems (SENSYS)*, Los Angeles, California, November 2003, pp. 150–161.

[15] J. Liu, P. Cheung, L. Guibas, and F. Zhao, "A dual-space approach to tracking and sensor management in wireless sensor networks," in *Proceedings of the* 1^{st} *ACM International Workshop on Wireless Sensor Networks and Applications*, Atlanta, Georgia, April 2002, pp. 131–139.

[16] J. Liu, J. Reich, and F. Zhao, "Collaborative in-network processing for target tracking," *EURASIP JASP: Special Issues on Sensor Networks*, vol. 2003, no. 4, pp. 378 391, March 2003.

[17] R. Evans, V. Krishnamurthy, G. Nair, and L. Sciacca, "Networked sensor management and data rate control for tracking maneuvering targets," *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 1979 1991, 2005.

[18] T. M. Cover and J. A. Thomas, *Elements of information theory*, ser. Wiley series in telecommunications. New York: Wiley, 1991.

[19] D. Knuth, *The Art of Computer Programming, Vol. 3: Sorting and Searching*. Addison-Wesley, 1997.

[20] R. W. Yeung, *A First Course in Information Theory*. Kluwer Academic/Plenum Publishers, 2002.

[21] P. Billingsley, *Probability and Measure*, 3rd ed. John Wiley & Sons, 1995.

[22] S. Meyn and R. Tweedie, *Markov Chains and Stochastic Stability*. London: Springer-Verlag, 1993.

[23] I. Kontoyiannis, L. Lastras-Montaño, and S. Meyn, "Relative entropy and exponential deviation bounds for general markov chains," in *Proceedings of the IEEE International Symposium on Information Theory*, 2005.