

Time vs. Truth: Age-Distortion Tradeoffs and Strategies for Distributed Inference

Présentée le 9 octobre 2023

Faculté informatique et communications
Laboratoire de théorie de l'information
Programme doctoral en informatique et communications

pour l'obtention du grade de Docteur ès Sciences

par

Yunus INAN

Acceptée sur proposition du jury

Prof. R. Urbanke, président du jury
Prof. E. Telatar, directeur de thèse
Prof. A. Tchamkerten, rapporteur
Prof. R. Yates, rapporteur
Prof. P. Thiran, rapporteur

Time vs. Truth: Age–Distortion Tradeoffs and Strategies for Distributed Inference

Yunus İnan

EPFL - Ecole Polytechnique Fédérale de Lausanne

Thesis No. 9868 (July 2023)

Thesis presented to the faculty of computer and communication sciences for
obtaining the degree of Docteur ès Sciences

Accepted by the jury:

Prof. Emre Telatar

Thesis director

Prof. Aslan Tchamkerten

Expert

Prof. Patrick Thiran

Expert

Prof. Roy Yates

Expert

Prof. Rüdiger Urbanke

President of the jury

Ecole Polytechnique Fédérale de Lausanne, 2023

*“Our time here is brief, our risk enormous.
Don’t waste the one or increase the other, if
you please.”*

— Stephen King, *The Dark Tower*

*“The best way to find out if you can trust
somebody is to trust them.”*

— Ernest Hemingway

Abstract

In 1948, Claude Shannon laid the foundations of information theory, which grew out of a study to find the ultimate limits of source compression, and of reliable communication. Since then, information theory has proved itself not only as a quest to find these limits but also as a toolbox which provides a new machinery and new perspectives on problems in various fields. Shannon’s original description of the communication problem omitted the semantic aspects. However, modern communication systems necessitate the consideration of semantics, such as fidelity and freshness of data. Shannon did study a problem related to the fidelity of data — known as rate–distortion theory, which can be seen as an attempt to incorporate semantics in a weak sense. Yet, freshness has not been widely studied until 2011, when Kaul, Yates and Gruetesser introduced a new metric for its assessment, called *age of information* (AoI).

Since 2011, AoI has become a widely studied notion as data freshness becomes increasingly important. But at the same time, not all data is equally important. If, in an attempt to reduce staleness our system drops important pieces of data, the remedy may be worse than the disease. Aligned with this observation, in Part I, we study a discrete-time model where each packet has a cost of not being sent — this cost might depend on the packet content. We study the tradeoff between the age and the cost where the sender is confined to packet-based strategies. We show that the optimal tradeoff can be attained with finite-memory strategies and we devise an efficient policy iteration algorithm to find these optimal strategies. We further study a related problem where the transmitted packets are subject to erasures and show that the optimal policies for our problem are also optimal for this new setup. Allowing coding across packets significantly extends the packet-based strategies and we show that when the packet payloads are small, the performance can be improved by coding. Furthermore, we study a related problem where *some* of the equally important packets must be sent in order to control the output rate. ‘Which packet to send, and when?’ is the relevant question of this problem and we show that if the packet arrival process is memoryless, a simple class of strategies attain the optimal tradeoff. The same class of strategies also solve the analogous continuous-time problem, where packets arrive as a Poisson process.

In Part II, we study two distributed hypothesis testing problems: (i) with a centralized architecture and (ii) with a fully decentralized architecture. In the centralized problem, we consider peripheral nodes that send quantized data to the fusion center in a memoryless fashion. The expected number of bits sent by each node under the null hypothesis is kept limited. We characterize the optimal decay rate of the misdetection probability provided that false alarms are rare, and study the tradeoff between the communication rate and maximal decay rate of misdetection probability. We use the information theory toolbox and resort to rate–distortion methods to provide upper bounds to the tradeoff curve and we also show that at high rates lattice quantization achieves near-optimal performance. We characterize the tradeoff for the case where nodes are allowed to record and quantize a fixed number of samples as well. Moreover, under sum-rate constraints, we show that an upper bound to the tradeoff curve is obtained with a waterfilling solution. In the decentralized problem, we study a locally-Bayesian scheme where at every time instant, each node chooses to receive information from one of its neighbors at random. We show that under this sparser communication scheme, the agents learn the truth eventually and the asymptotic convergence rate remains the same as the baseline algorithms. We also derive large deviation estimates of the log-belief ratios for a special case where each agent replaces its belief with that of the chosen neighbor.

Keywords: Age of information, average age, rate–distortion theory, age–distortion tradeoff, causal policies, Markov decision process, distributed inference, hypothesis testing, social learning

Résumé

En 1948, Claude Shannon a posé les fondements de la théorie de l’information, qui est née d’une étude de la recherche des limites ultimes de la compression des sources et de la communication fiable. Dès lors, la théorie de l’information s’est révélée non seulement comme une quête visant à trouver ces limites, mais également comme une boîte à outils fournissant de nouvelles méthodes et perspectives pour résoudre des problèmes dans divers domaines. La description originale de Shannon du problème de communication a omis les aspects sémantiques. Cependant, les systèmes de communication modernes nécessitent la prise en compte de la sémantique, telle que la fidélité et la fraîcheur des données. Shannon a étudié un problème lié à la fidélité des données, connu sous le nom de théorie du débit–distorsion, qui peut être considéré comme une tentative d’intégrer la sémantique dans un sens faible. Cependant, la fraîcheur n’a été largement étudiée que depuis 2011, lorsque Kaul, Yates et Gruetesser ont introduit une nouvelle mesure pour son évaluation, appelée l’âge de l’information (AoI).

Depuis 2011, l’AoI est devenu un concept largement étudié, car la fraîcheur des données devient de plus en plus importante. Mais toutes les données ne sont pas également importantes. Si, dans le but de réduire le risque d’obsolescence, notre système supprime des éléments importants de données, le remède peut être pire que le mal. En accord avec cette observation, dans la première partie, nous étudions un modèle en temps discret où chaque paquet a un coût s’il n’est pas envoyé — ce coût peut dépendre du contenu du paquet. Nous étudions le compromis entre l’âge et le coût lorsque l’émetteur est limité à des stratégies basées sur les paquets. Nous montrons que le compromis optimal peut être atteint avec des stratégies à mémoire finie, et nous concevons un algorithme d’itération de politique efficace pour trouver ces stratégies optimales. Nous étudions également un problème connexe où les paquets transmis sont soumis à des effacements et montrons que les stratégies optimales pour notre problème sont également optimales pour cette nouvelle configuration. L’utilisation du codage sur les paquets étend considérablement les stratégies basées sur les paquets. Nous montrons que lorsque les charges transportées des paquets sont petites, les performances peuvent être améliorées par le codage. De plus, nous étudions un problème connexe où *certain*s paquets tout aussi

importants doivent être envoyés pour contrôler le débit de sortie. ‘Quel paquet envoyer, et quand?’ est la question pertinente de ce problème, et nous montrons que si le processus d’arrivée des paquets est sans mémoire, une classe simple de stratégies atteint le compromis optimal. La même classe de stratégies résout également le problème en temps continu, où les paquets arrivent selon un processus de Poisson.

Dans la deuxième partie, nous étudions deux problèmes d’inférence distribuée: (i) avec une architecture centralisée et (ii) avec une architecture entièrement décentralisée. Dans le problème centralisé, nous considérons des nœuds périphériques qui envoient des données quantifiées au centre de fusion sans utilisation de mémoire. Le nombre espéré de bits envoyés par chaque nœud sous l’hypothèse nulle est limité. Nous caractérisons le taux de décroissance optimal de la probabilité de non-détection conditionné par le fait que les fausses alarmes soient rares, et étudions le compromis entre le taux de communication et le taux maximal de décroissance de la probabilité de non-détection. Nous utilisons la boîte à outils de la théorie de l’information et recourons aux méthodes de débit–distorsion pour fournir des bornes supérieures à la courbe de compromis. Nous montrons également qu’à des débits élevés, la quantification en treillis atteint des performances proches de l’optimal. Nous caractérisons le compromis dans le cas où les nœuds sont autorisés à enregistrer et à quantifier un nombre fixe d’échantillons également. De plus, sous des contraintes de débit total, nous montrons qu’une borne supérieure à la courbe de compromis est obtenue avec une solution de waterfilling. Dans le problème décentralisé, nous étudions un schéma localement bayésien où à chaque instant, chaque nœud choisit de recevoir des informations d’un de ses voisins de manière aléatoire. Nous montrons que dans ce schéma de communication limité, les agents finissent par apprendre la vérité et le taux de convergence asymptotique reste le même que pour les algorithmes standard. Nous obtenons également des estimations de grande déviation des ratios de croyances logarithmiques pour un cas particulier où chaque agent remplace sa croyance par celle du voisin choisi.

Mots-clés: Âge de l’information, âge moyen, théorie du débit–distorsion, compromis âge–distorsion, politiques causales, processus de décision markovien, inférence distribuée, tests d’hypothèse, apprentissage social.

Acknowledgements

I would like to express my gratitude to all those who have supported me throughout my PhD journey. Without their unwavering encouragement and assistance, this thesis would not have been possible.

First and foremost, I am deeply indebted to my advisor, Emre, for his exceptional mentorship, and insightful guidance. Whenever I met him, even if only for a few minutes and no matter the topic we discussed, I was able to learn something new. Moreover, knowing him has redefined the meaning of ‘humble’ for me, as he is the humblest person I have ever met. His expertise, dedication, and willingness to engage in countless discussions have not only shaped my research but have also played a pivotal role in my personal growth. I am truly fortunate to have had the privilege of working under his supervision.

The second part of this thesis is the result of a fruitful collaboration with my colleague and friend, Mert, and Prof. Ali H. Sayed. I thank Mert for being an excellent collaborator, and Prof. Sayed for his guidance and valuable feedback. During his visit to Lausanne, Prof. Visa Koivunen also generously offered his collaboration, and it was a pleasure to work with him.

I extend my appreciation to the members of my thesis committee, Prof. Aslan Tchamkerten, Prof. Patrick Thiran and Prof. Roy Yates, for their thorough evaluation and for their constructive feedback that have immensely contributed to the refinement of this thesis. I also thank Prof. Rüdiger Urbanke for presiding over the committee.

Being a member of the Information Processing Group (IPG) has been an integral part of my academic journey. I thank the IPG academic staff, both past and present — Bixio, Rüdiger, Michael, Olivier, Nicolas, and Yanina for providing a conducive academic environment. I shall also emphasize that every administrative procedure within IPG is seamlessly managed by Muriel and France, and I am grateful for their efforts in making our lives easier. Whenever technical difficulties arise, Damir is the IPG member who selflessly takes charge, and I also thank him for our delightful coffee talks at Giacometti. My fellow IPG members — Amedeo, Erixhen, Kirill, Dina, Thomas, Andreas, Farzad, Bora and Marco — have been invaluable companions and I am sincerely thankful for the time we enjoyed together. I thank my fellow labmates separately: Pierre, for his help in proof-checking the French translation of the

abstract of this thesis, and my officemate Reka, for both his academic collaboration and for his cooperation in the competitions we have participated in.

My life in Lausanne would not be as joyful without the amazing people I have met. First, I thank Ignacio, one of the most extraordinary people I have ever met and the best flatmate I have ever had, for the fun moments and for countless philosophical discussions we have engaged in. I am thankful to Mert, Baran, Deblina, İlker, Esra, Bahar and Cem for the camaraderie that we shared. A special mention goes to Deniz, a childhood friend, who encouraged me to visit EPFL as an exchange student, and I might not have embarked on this journey without his encouragement.

Last but certainly not least, I want to express my sincere thanks to Bahar for her affection, and to my family — Sirma, Sabri, and Emre — for their enduring patience, understanding, and encouragement during the various highs and lows of my journey. Your support has been my greatest strength, and I am deeply grateful for it.

Contents

Abstract	iii
Résumé	v
Acknowledgements	vii
Contents	ix
List of Figures	xii
1 Introduction	1
1.1 Shannon’s Communication Model	1
1.2 Lossless Source Coding	3
1.2.1 Fixed-to-Variable Length Coding	4
1.2.2 Variable-to-Fixed Length Coding	7
1.3 Channel Coding	9
1.4 Lossy Source Coding — Rate–Distortion Theory	13
1.4.1 Scalar Quantization	15
1.4.2 Vector Quantization	17
1.5 Hypothesis Testing	18
1.5.1 Error Exponents in Hypothesis Testing	20
1.6 Conclusion — Towards Application-Oriented Communication .	22
1.7 Outline of the Thesis and Main Contributions	23
1.7.1 Part I: Age–Distortion Tradeoffs	23
1.7.2 Part II: Strategies for Distributed Inference	25
I Age–Distortion Tradeoffs	27
2 Introduction to Age of Information	29
2.1 The Timely Communication Problem	29
2.2 Age of Information: Moving from Delay to Freshness	30
2.3 Related Work	33

3	Optimal Policies for Age and Distortion	35
3.1	Motivation – Are All Packets of Same Importance?	35
3.2	Notation	37
3.3	Problem Definition	37
3.4	The Age-Distortion Tradeoff	39
3.4.1	Markov Decision Problem Formulation as a Lower Bound	39
3.4.2	Policy Iteration with a Truncated State Space	42
3.4.3	The Exact Buffer Size for an Optimal Policy	44
3.4.4	An Efficient Algorithm to Find the (Δ_e, D) Region	46
3.5	Relation to an Erasure Channel with Feedback	52
3.6	When Timestamps Become Significant	53
3.6.1	Buffer-Ignorant Strategies	54
3.6.2	Revealing Partial Buffer Content	57
3.6.3	About Other Possible Coding Strategies	58
3.7	Discussion	59
3.8	Appendix	60
3.8.1	Proofs of Theorems 3.1 and 3.2	60
3.8.2	Proof of Lemma 3.1	62
3.8.3	Proof of Property 3.1	63
3.8.4	Proof of Theorem 3.3	64
3.8.5	Policy Update Improvement	66
3.8.6	Closed-Form Expressions for the Simple Strategies	67
3.8.7	Proof of Corollary 3.2	70
4	Age-Optimal Causal Labeling of Memoryless Processes	73
4.1	Problem Definition	74
4.2	Finite-State Approximation	77
4.3	Extension To Poisson Processes	82
4.4	Discussion	84
4.5	Appendix	85
4.5.1	Extension of Example 2 to Rational r :	85
4.5.2	Proof of Lemma 4.3	85
4.5.3	Proof of Lemma 4.4	86
4.5.4	Proof of Lemma 4.5	87
5	Part I Conclusion	89
 II Strategies for Distributed Inference		 91
6	Introduction to Distributed Inference	93
6.1	Distributed Inference	93
6.2	Related Work	96
6.2.1	Information Theoretic Approaches	96
6.2.2	Signal Processing Approaches	97

7	Distributed Hypothesis Testing Under Memoryless Quantization	99
7.1	Notation	100
7.2	Problem Formulation	100
7.2.1	Memoryless Quantization and the Communication Constraint	102
7.2.2	Performance Criteria under Memoryless Quantization	103
7.3	Best Performance Under Memoryless Quantization	104
7.3.1	Boundary of the Achievable Region	104
7.3.2	Upper Bound on the Boundary of the Achievable Region	105
7.3.3	Calculating the Upper Bound θ_U	106
7.3.4	Asymptotic Behavior of $R_U(\delta)$	110
7.4	High-Rate Regime and Vector Quantization	112
7.4.1	High-Rate Lattice Quantization	112
7.4.2	Best Performance under Vector Quantization	116
7.5	Multiple-Node Case	118
7.6	Discussion	123
7.7	Appendix	125
7.7.1	Proof of Theorem 1	125
7.7.2	Proof of Lemma 7.1	128
7.7.3	Proof of Lemma 7.2	128
7.7.4	Boundedness of $C_U(r)$	129
8	Introduction to Social Learning	131
8.1	Notation	132
8.2	Bayesian Framework in m -ary Hypothesis Testing	132
8.3	Related Work	133
9	Social Learning Under Randomized Collaborations	135
9.1	Motivation	135
9.2	Problem Formulation	136
9.2.1	A Randomized Diffusion Algorithm	136
9.3	Analysis of the Algorithm	137
9.4	A Special Case: Replacement	142
9.5	Numerical Results	143
9.6	Discussion	145
10	Part II Conclusion	147
	Bibliography	149
	Curriculum Vitae	165

List of Figures

11	Graphical representation of a canonical communication system.	2
12	Tree representation of a prefix-free dictionary	8
13	Graphical representation of a BEC	10
14	The rate–distortion curve $R(D) = 1 - h_b(D)$ for $d(x, y) = \mathbb{1}\{x \neq y\}$	15
15	A scalar quantizer example	16
16	A vector quantizer example	18
21	An example evolution of Δ_t	32
31	A scenario where a network device is throttled.	35
32	Same scenario as in Figure 31 depicted at a speaking time.	36
33	The single-source model	36
34	Comparison of the strategies for $\mathcal{V} = \{1, 20\}$ and $\Pr(V = 1) = 0.7$	52
35	Comparison of the strategies for $\mathcal{V} = \{1, 20\}$ and $\Pr(V = 1) = 0.8$	52
36	Buffer-ignorant strategies and improvement with Tunstall coding	58
41	Graphical representation of the states as (m, n) pairs	80
61	A communication-constrained distributed hypothesis testing setup	95
71	The centralized distributed inference setup	101
72	Bounds for R_U curve for the case where $X \sim \mathcal{N}(0, 1)$ under \mathcal{H}_0 and $X \sim \mathcal{N}(\mu, 1)$ under \mathcal{H}_1 for $\mu = \sqrt{20}$	111
73	Uniform quantizer of radius r	113
74	High-rate behaviors of the lower bound \underline{R}_U , and the curve $R_U^{(L)}$	115
75	Upper and lower bounds under k -dimensional vector quantization, for $k = 1, 2, 4$	118
76	The optimal sum-rate curve for a 3-node instance of the problem	123
91	(a) The network with $K = 10$ nodes and 12 edges (b) Sample paths of $\frac{1}{i}\lambda_{1,i}$ and $\frac{1}{i}\lambda_{6,i}$	143
92	Large deviation estimates of the replacement algorithm	144

Introduction

1

This chapter is intended to provide a brief review of information theory, and some of the main problems it addresses. For readers familiar with information-theoretic concepts, it may serve as a quick reminder; and for those who are unaccustomed to information theory, it could be a short introduction. We will first introduce the canonical communication model. Then, we will adopt an engineering perspective to study the related problems one by one, and consequently encounter information-theoretic quantities as fundamental limits related to these problems. The problems in this chapter are not treated with full rigor and the proofs of relevant theorems/results are not included — except Theorem 1.10 — in order to keep the chapter at an introductory level. For more comprehensive treatments, the reader may refer to the classical reference books in information theory, e.g. [1–4], which have been valuable sources for the author to write this chapter. At the end of this chapter, the reader will be equipped with (or reminded of) the necessary tools and concepts to follow the main chapters of this thesis.

1.1 Shannon’s Communication Model

After successful developments in modulation techniques in communication, such as PCM (Pulse Coded Modulation), a shift towards digital communication was inevitable. In 1948, Claude Shannon formulated a general theory of communication on this basis in his seminal work, *A Mathematical Theory of Communication*. In his words, *the fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point* [5]. A general block diagram for a communication scheme can be given as in Figure 11.




Figure 11 – Graphical representation of a canonical communication system.

We first focus on the source and channel components of Figure 11. The information source in Figure 11 could be a bitstream, data generated from a sensor, or even a sequence of moves in a chess game. It consequently emits a message to send. Later, we will model the source as a stochastic process, and the message to be reproduced at the destination will be a realization of this process. For instance, if the source is a finite bitstream, ‘00110001’ might be the message to be conveyed. The channel is a transmission medium that could be subject to atmospheric noise, scattering, interference from another sources, or in general any phenomenon that injects noise in the system. Although very unlikely to encounter in real-life, one can also model the channel as a noiseless medium, i.e., a simple passthrough.

Appropriate choice of the encoder-decoder pair is the key to designing well-performing communication systems which battle the noise injected by the channel. A significant portion of the communication and coding theories revolves around finding efficient and high fidelity encoder-decoder designs. Simply put, the encoder might represent any kind of processing on the data produced by the source such as compression or adding redundancy to data. Even popular data reduction techniques as Principle Components Analysis (PCA), and neural networks may serve as encoder examples. The decoder is a device whose task is to match the channel output — a possibly corrupted version of the encoder’s output — with an approximate representation of the original message generated by the source.

Now that we have an understanding of how a generic communication scheme is represented, it is natural to ask questions such as ‘How much information does the source carry?’, or ‘How much information per second can we send reliably over a channel?’. To answer these questions, we must have a standard unit of information, similar to the other standard units as kilograms for mass, or seconds for time. Aligned with the development of digital circuits back in the time when Shannon was writing his work, a reasonable unit for measuring information was the number of flip-flops in a circuit. Flip-flops are devices with two stable positions, and hence if a circuit has N flip-flops, it can represent 2^N possible configurations in its memory. Consequently, the configuration of the circuit can be represented with a number i from 0 up to $2^N - 1$, where the state of each flip-flop is a digit in the binary representation of i . Then, it is convenient to say that this circuit contains N binary digits of information. The term *bit*, coined by John Tukey, is an abbreviation of binary digit and is still being used for a measure of information. Now, the questions in the begin-

ning of this paragraph can be restated as (i) ‘What is the number of bits per symbol required to represent a particular information source?’, and (ii) ‘How many bits per second can we send reliably over a given channel?’. *We will refer to these questions as question (i) and (ii) throughout this chapter.* These questions are actually two of the main problems in information theory, and were settled by Shannon in his 1948 work [5] for fairly general models. However for many setups such as the broadcast channel model [1], these fundamental questions still remain open.

1.2 Lossless Source Coding

In the previous section, we have presented a general communication scheme and a unit for measure of information (*bit*). In an attempt to answer the fundamental questions of information theory, we must assume mathematical models for the source and the channel. In this section, we will study a simple source model and find the required number of bits to represent the source. In our setting, we assume that the source is a discrete-time stochastic process X_1, X_2, \dots where each X_i is a random variable taking values in a finite set \mathcal{X} . We further assume that each X_i has the same distribution and independent from all X_j , $j \neq i$. Such stochastic processes are called independent and identically distributed (i.i.d.) processes and will often be encountered in this thesis. We will also assume that one source symbol X_i is generated per time unit — the reader may assume the time unit taken as one second if they feel more comfortable, in this case the source generates one symbol per second.

As a thought experiment, we attempt to make a first estimate for the number of bits required to represent the source. Assume that each X_i is a coin flip with probability of heads being equal to p . If $p = 0$ or $p = 1$, there is no randomness and the destination already knows what the source will say beforehand. Thus we can represent this source with 0 bits per coin toss. On the other extreme, if $p = 1/2$, the coin is purely random and we can guess that we need to allocate 1 bit per coin toss — this guess turns out to be true as we will see in the sequel. If $p = 0.1$, we may believe that most of the time we will observe tails and hopefully we can improve upon 1 bits per coin toss. It is easy to see that it is impossible to improve if we encode each coin toss separately. However, if we encode the coin tosses in groups of two, we indeed get an improvement with a smart encoder choice. Representing heads with H and tails with T , consider the following encoding function c :

$$c(x_1, x_2) = \begin{cases} 0, & x_1x_2 = TT \\ 10, & x_1x_2 = HT \\ 110, & x_1x_2 = TH \\ 111, & x_1x_2 = HH \end{cases} \quad (1.1)$$

where x_1 and x_2 are the outcomes of the first and second coin tosses respectively. Using this source code, let us examine the number of bits allocated per

time unit. At the end of time $2t$, the average number of bits that are output by the encoder is given by

$$\bar{N}_{2t} = \frac{\sum_{i=1}^t c(X_{2i-1}, X_{2i})}{2t} \quad (1.2)$$

and is a random variable due to the source being random. By the law of large numbers, \bar{N}_{2t} converges to $E[c(X_1, X_2)]/2$, where $E[\cdot]$ is the expectation operator, thus approaches to 0.645 bits per unit time, which is an improvement over 1 bit per unit time. One expects that the more random the source behaves, the harder it becomes to represent and we need to allocate more bits per time unit. In the next section, we will see that Shannon's entropy H , a measure of randomness and information, is the fundamental limit for the source coding problem in the sense that the number of bits per unit time cannot be lower than the entropy for a large class of source codes. Luckily, again as we shall mention in the next section, there are sufficiently complex coding techniques to approach this fundamental limit.

1.2.1 Fixed-to-Variable Length Coding

After the intuitive approach to the source coding problem given in the previous section, we now adopt a more formal approach. Let us first focus on the representation of a single source symbol X . A source code c for X is a mapping from \mathcal{X} to $\{0, 1\}^*$, where $\{0, 1\}^*$ denotes the set of all finite-length binary sequences formed by 0's and 1's. Since the output length might depend on the realization of X , these codes are also called *fixed-to-variable length codes*. Obviously, if we want to have a lossless representation of the source, we shall not map different letters in \mathcal{X} to the same binary sequence. Hence, the following class of source codes are of the highest interest.

Definition 1.1 (Injective code). *A code is called injective (non-singular) if for any $x, x' \in \mathcal{X}$, $x \neq x'$ implies $c(x) \neq c(x')$.*

Now assume that we encode all source symbols with the same injective code c . We will denote the resulting source code as c^* . More precisely, c^* is a mapping from \mathcal{X}^* to $\{0, 1\}^*$ with

$$c^*(x_1, x_2, \dots, x_n) = c(x_1)c(x_2) \dots c(x_n) \quad (1.3)$$

where the right-hand side is the concatenation of $c(x_1), c(x_2), \dots, c(x_n)$. With the same reasoning as above, we may want c^* to be injective. Note that c being injective does not imply c^* being injective, thus we need

Definition 1.2 (Uniquely-decodable code). *A code c is called uniquely-decodable if c^* is injective.*

Now is the time when Shannon's entropy, as a fundamental limit in the source coding problem, comes into play. The following theorem gives lower

and upper bounds to the expected length of any uniquely-decodable code in terms of the entropy. Recall that the source symbols take values in a finite set \mathcal{X} . Let $p(x) := \Pr(X = x)$ be the probability that X equals x for $x \in \mathcal{X}$. Also for a source code c , let $l_c(x)$ be the number of bits output for the representation of the realization x . The expected length of c is consequently given by $E[l_c(X)] = \sum_{x \in \mathcal{X}} p(x) l_c(x)$.

Theorem 1.1. *For any uniquely-decodable code c ,*

$$E[l_c(X)] \geq H(X) \quad (1.4)$$

where $H(X) := \sum_{x \in \mathcal{X}} p(x) \log_2 \frac{1}{p(x)}$ is the discrete entropy of the random variable X . Furthermore, there exists a uniquely-decodable code with an explicit construction and such that

$$E[l_c(X)] \leq H(X) + 1. \quad (1.5)$$

One may notice the 1-bit gap as a result of equations (1.4) and (1.5). Recall the coin toss example at the end of the previous section where we have improved by grouping the source symbols. The same methodology will reduce the gap here as well. With an aim to generalize Theorem 1.1 to multiple letters, it may be useful to give definitions of joint entropy and conditional entropy for an ensemble of random variables.

Definition 1.3 (Joint entropy and conditional entropy). *Let X_1, \dots, X_n be an ensemble of discrete random variables taking values in sets $\mathcal{X}_1, \dots, \mathcal{X}_n$ respectively. Let $p(x_1, \dots, x_n) := \Pr(X_1 = x_1, \dots, X_n = x_n)$ be the joint distribution and $p(x_1|x_2) := \Pr(X_1 = x_1|X_2 = x_2)$ be the conditional distributions. Then,*

- The joint entropy of X_1, \dots, X_n is defined as

$$H(X_1, \dots, X_n) := \sum_{x_1, \dots, x_n \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n} p(x_1, \dots, x_n) \log_2 \frac{1}{p(x_1, \dots, x_n)}. \quad (1.6)$$

- The conditional entropy of X_1 given X_2 is defined as

$$H(X_1|X_2) := \sum_{x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2} p(x_1, x_2) \log_2 \frac{1}{p(x_1|x_2)}. \quad (1.7)$$

The following properties also prove useful in the sequel.

Property 1.1. *[Properties of H]*

- $H(X) \geq 0$.
- (Conditioning reduces entropy) $H(X|Y) \leq H(X)$.

- $H(X_1, \dots, X_n) = \sum_{i=1}^n H(X_i | X_1, \dots, X_{i-1}) \leq \sum_{i=1}^n H(X_i)$, with the right-hand side being equality if and only if X_1, \dots, X_n are independent.
- $H(X) \leq \log_2 |\mathcal{X}|$ with the right-hand side being equality if and only if X is uniformly distributed over \mathcal{X} .

One can also view $H(X)$ as a function of the distribution $p(x)$ as it does not depend on the values X takes. In line with this view, we can denote the entropy as $H(p)$. The following property is also important in information theoretic problems:

Property 1.2 (Concavity of H). $H(p)$ is concave in p .

In the coin toss example that we have previously discussed, for a biased coin toss with probability of heads $p = 0.1$, we improved upon 1 bit per unit time and obtained that in the long term, number of bits allocated per unit time converges to 0.645. The improvement relied on grouping of multiple tosses which was a purely intuitive method. We will now show that by grouping a large number of symbols, one can design uniquely-decodable codes with expected bits per source symbol approaching the entropy. From Theorem 1.1 we know that if we can simultaneously encode n symbols, there exists a uniquely-decodable code c with an expected length

$$E[l_c(X_1, \dots, X_n)] \leq H(X_1, \dots, X_n) + 1. \quad (1.8)$$

When X_i 's are i.i.d., Property 1.1 tells that $H(X_1, \dots, X_n) = \sum_{i=1}^n H(X_i) = nH(X_1)$. Thus,

Theorem 1.2. For an i.i.d. sequence X_1, \dots, X_n , there exists a uniquely-decodable code c with

$$\frac{1}{n} E[l_c(X_1, \dots, X_n)] \leq H(X_1) + \frac{1}{n}. \quad (1.9)$$

Consequently, for sufficiently complex uniquely-decodable codes where n is large, one can approach the entropy, which is a lower bound on the expected number of bits per symbol (or unit time).

Observe that the above theorem is valid for uniquely-decodable codes. One may ask how does the theorem change if the requirement that c being uniquely-decodable is relaxed to c being injective. Since with uniquely-decodable c 's we can approach the entropy arbitrarily close, one might guess that with injective c 's, we can beat this fundamental limit. It turns out to be true, one can indeed design injective codes with a lower number of expected bits per unit time, but unfortunately not that much.

Theorem 1.3. [Source Coding Theorem for Injective Codes, [6, 7]] Let X be a discrete random variable (here, \mathcal{X} need not be finite). Then for any injective code c , we have

$$H(X) - \ln(H(X) + 1) - 1 \leq E[l_c(X)] \quad (1.10)$$

and there exists an injective code c such that

$$E[l_c(X)] \leq H(X). \quad (1.11)$$

Consequently, if one encodes length- n blocks of i.i.d. X_i 's, there exists an injective code c with

$$H(X) - \frac{1}{n} \ln(nH(X) + 1) - \frac{1}{n} \leq \frac{1}{n} E[l_c(X_1, \dots, X_n)] \leq H(X). \quad (1.12)$$

Theorem 1.3 indeed shows that one can beat the entropy limit with injective codes. However, as the number of symbols to be encoded grows, the expected number of bits per symbol will be sandwiched between the right and left-hand sides, and hence will be close to H . Another important observation is that the source symbol X in Theorem 1.3 need not take values in a finite set. Such requirement will be important in **Part II** of this thesis, where we study a distributed statistical inference problem subject to communication constraints. In that problem, aligned with the observations of this section, we will impose that the entropy of the statistics sent by a remote node shall not exceed a prescribed limit R . This way, since there exists injective codes satisfying (1.11), we will ensure that the average number of bits sent per unit time cannot exceed R as well.

1.2.2 Variable-to-Fixed Length Coding

In the previous section, the main emphasis was on the *fixed-to-variable* length codes. However, if the encoder's output is strictly limited to N bits per use, one has to come up with a different machinery. As opposed to the previous section, now we must seek an efficient way to (i) choose the number of source symbols to encode and (ii) find a code to represent the selected number of source symbols. Such codes are called *variable-to-fixed* length codes.

For the design and implementation of *variable-to-fixed* codes, one may find an additional device, called *parser*, particularly useful. The parser creates segments from the input stream. Assume for simplicity that $\mathcal{X} = \{0, 1\}$. The following may serve as an example of a binary input parser.

$$\underbrace{0}_{w_1} | \underbrace{110}_{w_2} | \underbrace{10}_{w_3} | \underbrace{0}_{w_4} | \underbrace{1110}_{w_5} | 0 | 11110 | 0 | 0 | 11110 | 0 | 10 | 0 | 1 \dots \quad (1.13)$$

As the above bits are fed into the parser, it decides when to form a word w_i according to a set of rules. Note that each w_i is a concatenation of multiple source symbols. In the above example, the parser ends a word whenever it encounters a 0. More generally, the set of rules the parser obeys are determined by a *dictionary*, which is a subset of \mathcal{X}^* (in the above example, a subset of $\{0, 1\}^*$).

Note that in the above example $X_1, X_2, \dots = w_1, w_2, \dots$, i.e., no source symbols are skipped. Such dictionary choices will be of our interest, which we

will call as valid dictionaries. More formally, a dictionary is valid if every input sequence has a prefix in this dictionary. An example of an invalid dictionary is $\{0\}$. If the parser employs this dictionary, it cannot parse the source stream once a 1 is observed.

Another desired property of a parser could be that it must be able to decide at time t whether it should or should not end the current word it is parsing. In the language of probability theory, the time to end a parsing must be a *stopping time*. Note that for such dictionaries, no word should be a prefix of another. Hence, these dictionaries are called prefix-free dictionaries. The dictionary of the parser in the example (1.13) is also prefix-free. Prefix-free dictionaries can be represented with trees where each word is represented as a leaf of the tree. For the above example, the constituent tree is given in Figure 12.




Figure 12 – The tree representing the prefix-free dictionary based on the rule ‘stop when a 0 is encountered’.

Recall that the encoder’s output was strictly limited to N bits per use. Therefore, we must truncate the above tree to have exactly 2^N leaves. This way, we are able to encode a variable length of input symbols into exactly N bits. So, for the task of variable-to-fixed length coding, one must choose an appropriate tree with 2^N leaves that meets the design criteria.

One design criteria could be to minimize the number of bits sent per unit time over a long-term horizon. Observe that the time it takes to parse a word has now become a random variable. Let T_i be the time taken to parse w_i . After the i th parsing is finished, the number of bits sent per unit time is given by

$$\frac{iN}{T_1 + \dots + T_i}. \quad (1.14)$$

If the dictionary is prefix-free, then T_i ’s are i.i.d.. Therefore, as i tends to infinity, the quantity in (1.14) converges to $N/E[T]$ by the law of large numbers. As a consequence, for long-term horizon, it is reasonable to minimize $N/E[T]$ or equivalently for a given N , to maximize $E[T]$. Luckily, this is an easy task and can be solved iteratively with a simple greedy algorithm devised by Tunstall.

Theorem 1.4 (Tunstall Algorithm).

- Start with the root as intermediate node and all level 1 nodes as leaves.
- If number of leaves is equal to the desired dictionary size stop (if desired).
- Otherwise, pick the highest probability leaf, make it an intermediate node and grow 2 leaves on it. Go to step 2.

In **Part I** of this thesis, we will use variable-to-fixed coding as a means to improve the performance of uncoded strategies in a timely communication setting. We will use the fact that maximizing a similar quantity to $E[T]$ improves the performance of our scheme and we will consequently boost the performance of our original strategies with Tunstall coding.

At this point, we have introduced some important concepts and results to study one of the fundamental questions of information theory, namely question (i): ‘What is the minimum number of required bits to represent a source?’. We have seen that the entropy H , as a measure of randomness and information, turns out to be a fundamental entity in numerous source coding problems. The source coding concepts and tools introduced in this chapter, e.g., injective codes and Tunstall coding, will be of use in the main parts of the thesis. For a more comprehensive introduction to the lossless source coding, the reader may refer to [1, 2], and [4] from a method-of-types perspective. We conclude our introduction to source coding and will study question (ii) in the next chapter.

1.3 Channel Coding

Let us recall the two main questions we have touched upon: (i) ‘What is the minimum number of required bits per unit time to represent a source?’ and (ii) ‘What is the maximum number of bits we can send reliably over a channel per use (or per unit time)?’. We have already studied the first question by assuming a mathematical model for the source. Now, we shall do the same for the channel. Simply put, a channel is a transmission medium that distorts the input by injecting a noise disturbance. Therefore, given a particular input, the output of the channel should be random. Based on this observation, more precisely, a channel is a probability transition matrix (or probability kernel) with given input and output alphabets. As we have done in the previous section, we will mainly discuss discrete memoryless channels (DMC) with finite input and output alphabets.

Definition 1.4 (DMC). *Given an input alphabet \mathcal{X} and an output alphabet \mathcal{Y} , a DMC is characterized with the conditional probabilities*

$$p(y|x) := \Pr(Y = y|X = x), \quad x \in \mathcal{X}, y \in \mathcal{Y}, \quad (1.15)$$

where X and Y are the input and the output of the DMC, respectively. Through multiple uses, if there is no feedback to the sender, the n th output of a DMC

will only depend on the n th input. Thus its conditional distribution for multiple uses can be factorized as

$$p(y_1, \dots, y_n | x_1, \dots, x_n) = \prod_{i=1}^n p(y_i | x_i). \quad (1.16)$$

A well-known example is the Binary Erasure Channel (BEC), whose input X takes values in the binary alphabet $\{0, 1\}$ and whose output Y takes values in the alphabet $\{0, 1, ?\}$. The transition probabilities of a BEC is given by $p(?|0) = p(?|1) = \epsilon$ and $p(0|0) = p(1|1) = 1 - \epsilon$. In words, the channel erases the input with probability ϵ (which is also called the erasure probability). The symbol ‘?’ denotes the erasures in the output. It is sometimes useful and visually appealing to represent channels in a graphical form — more specifically as a bipartite graph. The graphical representation of a BEC is given in Figure 13.




Figure 13 – The graphical representation of a BEC. Directed edges indicate the possible transitions and the edge labels are the probabilities of the corresponding transition.

In the channel coding task, we intend to design a coding scheme which operates in an opposite manner to source codes, i.e., instead of compressing the input, we intend to *add redundancy* to battle the noise. As an example, again think of a bitstream. If the channel is a BEC, a possible input-output pair could be as follows:

$$\begin{array}{l} 0010100010010101011010101 \\ 001?100?100????10??0?0101 \end{array} \quad (1.17)$$

Therefore, the receiver cannot reproduce the exact input sequence. A simple way to combat this noise is to repeat each input symbol twice. If the channel erases approximately the same portion of the input as above, the repeated input sequence and the constituent output might become

$$\begin{array}{l} 00|00|11|00|11|00|00|00|11|00|00|11|00|00|11|00|11|00|11|11|00|11|00|11|00|11 \\ 00|?0|1?|0?|11|0?|0?|00|1?|00|0?|??|00|?1|0?|1?|0?|11|?1|0?|1?|?0|1?|00|11 \end{array} \quad (1.18)$$

where the repeated bits are separated with vertical bars. Note that the receiver now can correctly infer every bit except one, the underlined one, as

both repetitions are erased. Hence, one might think that if the transmitter repeats every bit sufficiently large number of times, the input sequence may be correctly transmitted over the erasure channel. However, in this case, the number of input bits transmitted per channel use becomes small and this results in a very inefficient code. To measure this efficiency, we define the rate of the code as

$$R := \frac{\text{number of bits sent}}{\text{number of channel uses}}. \quad (1.19)$$

In the channel coding problem, we aim to find the possible values of R that allow the reliable transmission of a message. To this end, we will now present a formal setup.

We have previously stated that the main task in communication is to approximately reproduce a message at a remote point. For the channel coding problem, we abstract the source and consider the problem of sending a message among $M = 2^{nR}$ possible messages to a remote point through n uses of a channel — the message can be a direct realization of the source, or its source-coded version in bits (Section 1.2). Similar to source codes, we now define a channel code.

Definition 1.5 (Channel code). *A channel code $C(n, R)$ is a mapping from $\{1, \dots, M\}$, $M = 2^{nR}$ to \mathcal{X}^n . The image of C is called the codebook and each element in the image is called a codeword.*

The encoder then encodes the source based on the channel code. For instance in (1.17) and (1.18), the encoder repeats each input bit twice according to a $C(50, 1/2)$ repetition code. If a message m is to be transmitted, for ease of notation, we say that the encoded version of m is $\text{Enc}_n(m)$, where Enc_n stands for the encoder that bases its encoding to the code $C(n, R)$. Since the decoder's task is to approximately construct the message at a remote point, it must map the channel output to the original message set. Hence, it constitutes a mapping $\text{Dec}_n : \mathcal{Y}^n \rightarrow \{1, \dots, M\}$. The decoding may be based on different rules such as maximum likelihood decoding. Since the output of the decoder is an estimate of the original message, we denote the output of the decoder as \hat{m} .

Given an encoder-decoder pair, a reasonable performance measure is the probability of decoding a message incorrectly. More precisely, we consider the probability of error averaged over all possible messages:

$$P_{e,n} := \frac{1}{M} \sum_{m=1}^M \Pr(\hat{m} \neq m|m). \quad (1.20)$$

We would like $P_{e,n}$ to vanish as n tends to infinity. Hence, we will obtain a reliable communication system at large blocklengths. Considering once more the repetition example in (1.17), we can see that the more we repeat, $P_{e,n}$ is more likely to decrease but R decreases as well. Thus it seems that not all rates are achievable. We define the achievable rates as follows:

Definition 1.6 (Achievable R). R is achievable if there exists a sequence of $(\text{Enc}_n, \text{Dec}_n)$ such that $P_{e,n} \rightarrow 0$, and Enc_n bases its encoding to a $C(n, R)$ code and $M_n = \lceil 2^{nR} \rceil$.

If R is achievable, there is a channel code, and an encoder-decoder pair such that the average error probability shall be arbitrarily small at sufficiently large blocklengths. It turns out that another fundamental quantity, called the *channel capacity* and denoted by C , plays an important role in the characterization of achievable R . The channel capacity is defined as

$$C = \sup_{p(x)} I(X; Y), \quad (1.21)$$

where

$$I(X; Y) := \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} p(x)p(y|x) \log_2 \frac{p(y|x)}{p(y)} = E \left[\log_2 \frac{p(Y|X)}{p(Y)} \right] \quad (1.22)$$

is the mutual information between the random variables X and Y and $p(y) := \sum_{x \in \mathcal{X}} p(y|x)p(x)$ is the marginal distribution of the channel output Y . Now, we are in position to state Shannon's celebrated channel coding theorem.

Theorem 1.5. R is achievable if $R < C$ and not achievable if $R > C$.

In the above theorem, to show that any $R < C$ is achievable, Shannon considers a random ensemble of (Enc, Dec) pairs of rate $R - \epsilon$ for an arbitrarily small $\epsilon > 0$ and shows that at least one of them must drive $P_{e,n}$ to 0, consequently approaching the channel capacity. Hence, as opposed to the source coding theorem, there is no recipe to find the capacity-achieving codes and this has been a long-lasting quest in both information and coding theories. Moreover, a channel code shall have low encoding and decoding complexities to be applicable to real-life communication systems. Polar codes, introduced by Arikan [8], are the first class of channel codes with low encoding and decoding complexities that provably achieve the capacity.

Another important result regarding the channel capacity is that it remains the same even if the encoder has access to all past channel outputs, i.e., has perfect feedback. More specifically, when the encoder has access to Y_1, \dots, Y_{n-1} prior to sending X_n . Still, this additional information is not useful according to the our achievability definition in Definition 1.6.

Theorem 1.6. *Feedback does not increase the capacity.*

As a consequence of Theorems 1.5 and 1.6, we have answered question (ii): Given a discrete-memoryless channel (DMC) $p(y|x)$, one can reliably send at most C bits per channel use. Of course, a finer analysis reveals much more than Theorem 1.5. For instance, one can study the tradeoff between $(R, P_{e,n}, n)$, which gives tighter bounds for finite-length analyses [9]; or one can study the

exponential rate of decrease $E(R)$ such that $\liminf_{n \rightarrow \infty} -\frac{1}{n} \log P_{e,n} \geq E(R)$, also called as the error exponent or the reliability function of the channel [2].

For the BEC, the channel capacity turns out to be $C = 1 - \epsilon$, and if perfect feedback is available, a simple transmission scheme attains the capacity: Send a bit repeatedly until it is received without erasures. Note that in this scheme, the bits are not coded and hence this scheme is reminiscent of packet-based transmission schemes of today's computer networks, e.g., TCP uses acknowledgment messages (ACK) as a feedback mechanism to ensure that a packet is reliably received at the destination. Furthermore, according to this scheme, the expected time it takes to send a packet is $\frac{1}{1-\epsilon}$ time units. Consequently, if the source generates packets with a rate more than $1 - \epsilon$ packets per time unit, then this scheme will require an infinite memory — as the rate of successful transmissions is less than the packet generation rate and the packets that are not yet transmitted will have to be stored somewhere. The transmitter then has no choice but to drop some of the packets in order to operate at finite memory. Which packets to drop is a highly non-trivial question and might depend on the application. For instance if the receiver values fresh (recently generated) packets more, the sender may only attempt to send the freshest packet, such as a last-come-first-served (LCFS) queue. However, if some packets contain more important information than others, the sender might opt to send these important classes of packets. One of the main tasks in **Part I** will be to study the optimal method to choose packets in a resemblant scheme. How to measure the freshness will be explained in **Part I**. A machinery to assess importance, however, is one of the main topics in information theory and we shall now review this topic in the next section.

1.4 Lossy Source Coding — Rate–Distortion Theory

In Section 1.2, we found that a discrete memoryless source cannot be represented with less than H number of bits per unit time for long-term horizon. However, we may not have enough storage to store this source. An example could be an 1024×1024 image, with each pixel taking 256 values (8 bits, or 1 byte). If all existing 1024×1024 images are sampled from a uniform distribution, storing a single image would require 1 gigabyte (GB) — the entropy would be equal to 1 GB/image. Hence the images are encoded with some loss, e.g., JPEG compression, and then stored in our devices; thus allocating less memory than it would without compression. For the image compression task, the loss is usually adjusted such that the compressed image remains appealing to the eye. This example highlights a possible tradeoff. If a source is to be represented with R bits/symbol, where $R < H$, we should sacrifice some of the information provided by the source for the sake of saving memory. The sacrifice should be done such that the reconstructed version of the source obeys

a *fidelity criterion*.

Let us describe the problem formally. As usual, we assume a discrete i.i.d. source X_1, \dots, X_n . Similar to the channel coding problem, we consider a fixed-to-fixed length code. But this time, the mapping will be in the reverse direction: The encoder Enc obeys the codebook to map n source symbols into $\{1, \dots, 2^{nR}\}$. In a dual manner, the decoder Dec maps from $\{1, \dots, 2^{nR}\}$ to \mathcal{Y}^n , where \mathcal{Y} is the reconstruction alphabet. Consequently we obtain Y_1, \dots, Y_n at the decoder's output. We would like the reconstruction to satisfy the following *fidelity criterion*:

$$\frac{1}{n}E[d(X^n, Y^n)] \leq D, \quad (1.23)$$

where $d(x^n, y^n) = \sum_{i=1}^n d(x_i, y_i)$ is an additive distortion metric and D is a prescribed threshold. An example when $\mathcal{X} = \mathcal{Y}$ could be $d(x, y) = \mathbf{1}\{x \neq y\}$. In this case, the system will equally penalize any misrepresentation of x . If both X and Y take real values — the setting can be extended to more general alphabets — one can set for instance $d(x, y) = (x - y)^2$.

All (R, D) pairs are very unlikely to be achieved, and indeed this is the case. When R is low, the reconstruction would be subject to more distortion and conversely, at high R the reconstruction would be close to an exact replica. Similar to the channel coding case, we can define the achievable (R, D) pairs.

Definition 1.7. *An (R, D) pair is achievable if there exists a sequence of $(\text{Enc}_n, \text{Dec}_n)$ such that the output of Enc_n and the input of Dec_n take $[2^{nR}]$ possible values, and $\limsup_{n \rightarrow \infty} \frac{1}{n}E[d(X^n, Y^n)] \leq D$.*

Now, we can state the rate–distortion theorem:

Theorem 1.7. *Define the rate–distortion function*

$$R(D) = \inf_{p(y|x)} I(X; Y) \quad \text{s.t.} \quad E[d(X, Y)] \leq D. \quad (1.24)$$

All pairs strictly above the $R(D)$ curve are achievable, i.e., any (R, D) with $R > R(D)$, whereas the pairs underneath are not.

To give a concrete example, assume $\mathcal{X} = \mathcal{Y} = \{0, 1\}$, $\Pr(X = 1) = \Pr(X = 0) = 0.5$ and $d(x, y) = \mathbf{1}\{x \neq y\}$. It can be shown that

$$R(D) = 1 - h_b(D), \quad 0 \leq D \leq 1/2 \quad (1.25)$$

where $h_b(p) := p \log_2 \frac{1}{p} + (1 - p) \log_2 \frac{1}{(1-p)}$ is the binary entropy function. The curve is depicted below in Figure 14. The red shaded region depicts all achievable pairs whereas the unshaded region underneath is not achievable.

Another example could be when $\mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{0, 1, ?\}$ and

$$d(x, y) = \begin{cases} 0, & x = y \\ 1, & y = ? \\ \infty, & x \neq y \neq ? \end{cases} \quad (1.26)$$




Figure 14 – The rate–distortion curve $R(D) = 1 - h_b(D)$ for $d(x, y) = \mathbb{1}\{x \neq y\}$.

Note the similarity between the BEC and the above metric as no bit flips are allowed to attain a finite distortion value. For this metric, the $R(D)$ curve can be shown to be equal to $R(D) = 1 - D$, $0 \leq D \leq 1$. A very simple coding scheme can attain the $R(D)$ curve: erase an input bit with probability D .

However, in general, similar to the channel coding theorem, there is no explicit recipe to find lossy source codes that approach $R(D)$. As we have stated in the beginning of this section, lossy source coding is crucial in many real-life applications such as image and speech processing. Therefore, there is a plethora of work revolving around finding efficient methods to represent and store signals. A famous example is the Lloyd–Max method (very similar to the Expectation–Maximization algorithm) to convert a continuous random variable into a discrete random variable with an aim to minimize the mean-square distortion, i.e., $d(x, y) = (x - y)^2$. Such conversion methods are usually called *quantization*, and the devices that perform this conversion are called *quantizers*. In the next section, we review some prior work on quantizers and show that with simple methods one can closely approach the rate-distortion curve at high rates. The results will provide insight for **Part II** of this thesis, as the analysis carried out there will guide us to a very similar setting to the rate–distortion theory, and a study of high-rate quantizers will be inevitable.

1.4.1 Scalar Quantization

In this section we are concerned about quantization of a single real-valued random variable X , i.e., X takes values in \mathbb{R} (the set of real numbers). We will view a quantization procedure as a simple function:

Definition 1.8 (Simple function, [10]). *A function on \mathbb{R} that takes finitely*

many values is called a simple function. More precisely, let $\alpha_1, \dots, \alpha_n$ be the distinct values of a simple function f , then any such f is represented as

$$f(l) = \sum_{k=1}^n \alpha_k \mathbb{1}\{l \in B_k\} \quad (1.27)$$

where $\mathbb{1}\{l \in B_k\}$ denotes the indicator function of B_k and $B_1, \dots, B_n \in \mathcal{B}(\mathbb{R})$ form a partition of \mathbb{R} .

Then, for a simple function f , $f(X)$ is the *quantized* version of X . Since f takes a single random variable as its argument, we call it as a *scalar quantizer*. An example scalar quantizer is given in Figure 15:




Figure 15 – An example of a scalar quantizer $f(x)$, defined as in (1.28).

Usually quantizers are non-decreasing, i.e., higher valued inputs are mapped to higher valued outputs. This makes sense, though it is not necessary. The quantizer in the Figure 15 is

$$f(x) = -6\mathbb{1}\{x < -6 + \frac{1}{2}\} + \sum_{k=-5}^5 k\mathbb{1}\{x \in [k - \frac{1}{2}, k + \frac{1}{2})\} + 6\mathbb{1}\{x \geq 6 - \frac{1}{2}\}. \quad (1.28)$$

Since simple functions take finite values, they have to be ‘clipped’ after some threshold. However, we are usually allowed to take limits over the set of simple functions, and consequently we can also work with infinite-valued quantizers such as the following uniform quantizer:

$$f_r(x) = \sum_{k=-\infty}^{\infty} 2rk\mathbb{1}\{x \in [2rk - r, 2rk + r)\}. \quad (1.29)$$

We shall refer to the above quantizer as a uniform quantizer of radius r . It can be shown that if r is sufficiently small, which implies R (bits/symbol) should be

high, then one can approach the boundary of the rate–distortion curve under mild conditions. We now assume that X is a continuous random variable, i.e., its probability density function $p(x)$ exists. For simplicity, we take the mean-square distortion $d(x, y) = (x - y)^2$ although the results below can be generalized to a class of loss functions, e.g., $|x - y|^\alpha$. Under such assumptions, there exist well-known lower and upper bounds to the rate–distortion function given by

$$h(X) - \frac{1}{2} \log(2\pi eD) \leq R(D) \leq \frac{1}{2} \log((\text{Var}(X)/D)^+), \quad (1.30)$$

where

$$h(X) := - \int_{x \in \mathbb{R}} p(x) \log p(x) dx \quad (1.31)$$

is the differential entropy of X . Under mild conditions, e.g., the one given in [11], one can show that there exists a sequence of uniform quantizers f_r where $r \rightarrow 0$, such that

$$\lim_{r \rightarrow 0} H(f_r(X)) + \log(2r) = h(X). \quad (1.32)$$

For a uniform quantizer, the squared error is smaller than r^2 for an interval, i.e., $D \leq r^2$. Hence, the above equation implies that

$$H(f_r(X)) \lesssim h(X) - \frac{1}{2} \log(4D) \quad (1.33)$$

and

$$H(f_r(X)) \lesssim R(D) - \frac{1}{2} \log \frac{\pi e}{2}. \quad (1.34)$$

Consequently, with lossless encoding of $f_r(X)$, as we have seen in Section 1.2, one can approach the rate–distortion bound up to $\frac{1}{2} \log \frac{\pi e}{2} \approx 0.725$ nats — note that we have taken the natural logarithm in the definition of $h(X)$ instead of a base-2 logarithm, thus the information is quantified in natural digits (nats) instead of bits. If one imposes additional conditions on the regularity of $p(x)$, the bound can be improved to $\frac{1}{2} \log \frac{\pi e}{6} \approx 0.176$ nats [12]. In **Part II**, when studying a tradeoff between the communication rate and the learning rate in a distributed inference setup, the above statements will be useful to show that we can approach the optimal tradeoff within a small gap.

The 0.176-nat gap turns out to be a weakness of scalar quantization procedures. If one allows simultaneous quantization of multiple symbols, the gap can be improved. In the next section, shortly, we will discuss such procedures.

1.4.2 Vector Quantization

Vector quantization is the quantization of d symbols X_1, \dots, X_d . Again, as in the scalar quantization case, we view the quantization procedures as simple functions, but this time they form a partition of \mathbb{R}^d . An example of a vector quantizer for $d = 2$ is given in Figure 16 where the quantization regions are




Figure 16 – A vector quantizer example. The red circles represent the quantization values while the blue lines set the boundaries between the quantization regions. Note that any (x_1, x_2) is quantized to the nearest red circle according to the Euclidian distance. The regions formed as such are also called *Voronoi regions*.

separated with lines and the quantized values of the constituent regions are shown with red circles.

Recall the 0.176-nat gap in the scalar quantization case. This turns out to be a result of the covering inefficiency of the 1-dimensional lattice. If, for example, one covers the 2-dimensional space with hexagons, the gap decreases. However, it is a very difficult task to find optimal shapes tuned to the specific distortion function, and whose tessellation entirely covers the d -dimensional space.

1.5 Hypothesis Testing

At first sight, hypothesis testing does not appear to be a main topic in information theory regarding the communication problem perspective. However, some main problems such as channel coding can be viewed as instances of hypothesis tests. In this section, we elaborate on binary hypothesis testing. The setting can be described as follows: Suppose we have observed a random variable X — usually real-valued but can also take values in more general spaces, e.g., a complete and separable metric space. We would like to learn about the true state-of-nature based on our observation. To this end, we have the following assumptions:

- Under the null hypothesis, denoted by \mathcal{H}_0 , X comes from a distribution P .
- Under the alternative hypothesis, denoted by \mathcal{H}_1 , X comes from a distribution Q .

We assume that either \mathcal{H}_0 or \mathcal{H}_1 is the true state-of-nature. In fact, the null and alternative hypotheses help us to understand the true state of a phenomenon. Some examples could be

- (i) \mathcal{H}_0 : a communication signal is not present, \mathcal{H}_1 : a communication signal is present,
- (ii) \mathcal{H}_0 : No risk of fire, \mathcal{H}_1 : Fire risk, etc.

If we work with the example (i), a high value of X^2 may hint to a signal being present. According to our real-valued observation X , we would like to partition the real axis into two sets: \mathcal{D}_0 and \mathcal{D}_1 . If X falls into \mathcal{D}_0 , we decide \mathcal{H}_0 and vice versa. Consequently, two possible types of errors can arise:

- Type-I error (also called false positive, or false alarm): \mathcal{H}_1 is decided when the true hypothesis is \mathcal{H}_0 .
- Type-II error (also called false negative, or misdetection): \mathcal{H}_0 is decided when the true hypothesis is \mathcal{H}_1 .

Consequently, to assess the performance of an hypothesis test, there are two natural criteria:

- Type-I error probability: $\alpha := P(X \in \mathcal{D}_1)$.
- Type-II error probability: $\beta := Q(X \in \mathcal{D}_0)$.

It seems that there exists a tradeoff between α and β . Consider the extreme cases. If \mathcal{H}_0 is always chosen, $\alpha = 0$ and $\beta = 1$. Similarly, if \mathcal{H}_1 is always decided, then $\alpha = 1$ and $\beta = 0$. In fact, there is a tradeoff between α and β ; and the optimal tradeoff is given by the famous Neyman–Pearson lemma. For mathematical convenience we assume that P is absolutely continuous with respect to Q . That is, if $Q(B) = 0$, then $P(B) = 0$ for any B chosen from the Borel set of \mathbb{R} , denoted as $\mathcal{B}(\mathbb{R})$. This ensures the existence of the Radon–Nikodym derivative $\frac{dP}{dQ}$ as a random variable such that

$$E_P[\mathbb{1}\{B\}] = E_Q\left[\frac{dP}{dQ}\mathbb{1}\{B\}\right], \quad B \in \mathcal{B}(\mathbb{R}) \quad (1.35)$$

and where $E_P[\cdot]$ and $E_Q[\cdot]$ denote the expectations taken under P and Q respectively.

Theorem 1.8 (Neyman–Pearson). *Let $\frac{dP}{dQ}$ be the Radon–Nikodym derivative of P with respect to Q . Then the optimal tradeoff is achieved with decision regions in the following form:*

$$\mathcal{D}_0 = \left\{ \frac{dP}{dQ} \geq \eta \right\}, \quad \eta \geq 0. \quad (1.36)$$

In statistics $\frac{dP}{dQ}$ is also called as the likelihood ratio, and (1.36) is called a likelihood ratio test.

Since the Neyman–Pearson lemma states that likelihood ratio tests are optimal, we will confine ourselves to this class of tests in the sequel.

1.5.1 Error Exponents in Hypothesis Testing

Now suppose that instead of a single observation X , we have a collection of i.i.d. observations X_1, \dots, X_n . To highlight the dependency of the errors to n , we denote the type-I and type-II errors as α_n and β_n , respectively. One would expect the performance of the test to improve thanks to having more observations, and both α_n and β_n tend to 0 as $n \rightarrow \infty$. This turns out to be true for likelihood ratio tests. However, one might want to study the speed at which α and β tend to 0. Here, as we will do in the **Part II** of the thesis, we shall adopt the view that \mathcal{H}_1 is a high-risk event, e.g. a possible fire risk, and thus needs to be detected with very low error probability. In fact, we want β_n to decrease exponentially, i.e., $\beta_n \approx \exp(-\theta n)$, and we seek the best θ possible while ensuring that α_n vanishes. Aligned with such criteria, we define an achievable θ below:

Definition 1.9. θ is achievable if there exists a sequence of hypothesis tests such that $\lim_{n \rightarrow \infty} \alpha_n = 0$ and $\liminf_{n \rightarrow \infty} \frac{1}{n} \log \frac{1}{\beta_n} \geq \theta$.

The best possible θ turns out to be an information-theoretic quantity, called the Kullback–Leibler (KL) divergence between P and Q :

$$D(P||Q) := E_P \left[\log \left(\frac{dP}{dQ} \right) \right] \geq 0, \quad (1.37)$$

where $E_P[\cdot]$ denotes the expectation taken under P . The KL divergence can be viewed as a measure of discrepancy between P and Q . As suggested above, it is non-negative and furthermore it is equal to zero if and only if P is equal to Q . Hence, one expects that if $D(P||Q)$ is high, better performance may be achieved as it becomes easier to discern Q from P . This turns out to be true:

Theorem 1.9 (Stein’s Lemma). $\sup\{\theta : \theta \text{ achievable}\} = D(P||Q)$.

The above theorem implies that if a statistician wants β_n to decrease with an exponent θ , the best possible θ is $D(P||Q)$. Hence, it is an ultimate limit on the type-II error exponent. One can also aim to drive both α_n and β_n to 0 exponentially fast. This is also possible and the best error exponent tradeoffs are characterized by entities called Chernoff information. For details, the reader can refer to [1], [4], and to [13] for Chernoff’s original work.

The KL divergence can also be written in its dual form, known as the Donsker–Varadhan representation. Although this variational representation can be derived from duality, it also has connections with hypothesis testing. A special case concerned with the real space can be stated as follows:

Theorem 1.10. [*Donsker–Varadhan Representation*] Let P and Q be distributions on \mathbb{R} and assume that P is absolutely continuous with respect to Q . Then for all bounded and measurable f ,

$$D(P||Q) \geq E_P[f] - \log(E_Q[e^f]), \quad (1.38)$$

and the equality is attained when $f = \log \frac{dP}{dQ}$ (and if not bounded, by taking a limit), i.e., when f is the log-likelihood ratio.

Proof of Theorem 1.10. Let X_1, \dots, X_n be the i.i.d. observations we possess. We design the following hypothesis test:

$$\mathcal{D}_0 = \left\{ \frac{1}{n} \sum_{i=1}^n f(X_i) \geq E_P[f(X)] - \epsilon \right\} \quad (1.39)$$

where f is measurable and bounded and $\epsilon > 0$. Then, the type-I error

$$\alpha_n = P\left(\frac{1}{n} \sum_{i=1}^n f(X_i) < E_P[f(X)] - \epsilon\right) \rightarrow 0 \quad (1.40)$$

by the law of large numbers. To find a lower bound to the type-II error exponent, we upper bound β_n using Markov inequality as

$$\beta_n = Q\left(\frac{1}{n} \sum_{i=1}^n f(X_i) \geq E_P[f(X)] - \epsilon\right) \quad (1.41)$$

$$\leq \frac{E_Q[\exp(\sum_{i=1}^n f(X_i))]}{\exp(n(E_P[f(X)] - \epsilon))} \quad (1.42)$$

$$= \exp(-n(E_P[f(X)] - \epsilon - \log(E_Q[e^{f(X)}]))). \quad (1.43)$$

Hence, $\liminf_{n \rightarrow \infty} -\frac{1}{n} \log \frac{1}{\beta_n} \geq E_P[f(X)] - \epsilon - \log E_Q[e^{f(X)}]$. Stein's lemma (1.9) tells that

$$\begin{aligned} D(P||Q) &\geq \liminf_{n \rightarrow \infty} -\frac{1}{n} \log \frac{1}{\beta_n} \\ &\geq E_P[f(X)] - \epsilon - \log E_Q[e^{f(X)}]. \end{aligned} \quad (1.44)$$

Finally, if one substitutes $f = \frac{dP}{dQ}$ (and if not bounded, takes a limit), it is not difficult to see that the equality holds. \square

In light of the above results, one expects that if any test statistic other than the likelihood ratio is used, the type-II error exponent is likely to decrease. This is indeed observed in many schemes where the likelihood ratios need to be digitally represented, i.e., need to be quantized, or if a remote node has to compress the likelihood ratios to a prescribed rate, say R bits per symbol. In **Part II** of this thesis, we will study a communication-constrained hypothesis testing scheme where remote nodes send quantized versions of the likelihood ratios. With similar definitions and techniques used in this section, we will attempt to find the best achievable type-II error exponent subject to an asymmetric communication constraint.

1.6 Conclusion — Towards Application-Oriented Communication

Hopefully, the reader is introduced to (or reminded of) some of the main problems in information theory and now carries a toolbox of information-theoretic methods. As we have mentioned multiple times, the information theory toolbox will be of help when studying the problems encountered in this thesis, especially in **Part II**. Although we have tried to give a brief overview of the main topics in information theory, the reader should keep in mind that this is only a partial introduction and does not capture the whole beauty of information theory.

We end this section with a brief discussion on application oriented communication, also called semantic communication. Recall that in Section 1.1, we have described the communication problem in Shannon's own words. Those words are followed by

...Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem.

One may therefore guess that the early information-theoretic works omitted the application for which the communication system is designed. However, for different tasks, different schemes and even different capacity notions are needed. A good example is in control theory, where an unstable stochastic process is to be tracked with finite error at all times. The classical notion of capacity is unfortunately not suitable for this task — recall that a rate R was achievable as long as the probability of error vanishes. Yet, even though the block errors are very rare, they eventually accumulate and result in an unbounded error for the task of tracking the unstable process. Thus, a suitable definition for what is achievable was needed and consequently *anytime capacity* was introduced [14]. Another example can be found in the context of timely communication. For timely communication tasks, channel delay is a widely used metric to assess the performance and many works have sought to minimize the channel delay. Although this is a perfectly reasonable engineering approach, ensuring a low channel delay is not equivalent to keeping the receiver as fresh as possible. For many tasks, the receiver may prefer to receive the *freshest data possible*, instead of other data that it has not received yet. For such scenarios, thus, a new metric was needed: Age of Information, which we will study shortly in **Part I**.

1.7 Outline of the Thesis and Main Contributions

This thesis is composed of two parts: **Part I, Age–Distortion Tradeoffs** and **Part II, Strategies for Distributed Inference**. In **Part I**, we study the tradeoff between the freshness and importance of data; where the freshness is quantified with *age of information* and the importance is quantified with a distortion metric reminiscent of that in the rate–distortion theory. **Part II** includes two studies on *distributed inference*; the first one featuring a centralized hierarchy and the second one featuring a fully-decentralized scheme. A more detailed outline is as follows:

1.7.1 Part I: Age–Distortion Tradeoffs

Age of Information (AoI) is a receiver-centric metric to assess the freshness of data. The fundamental entity in many AoI problems is the instantaneous age

$$\Delta_t = t - R_t, \quad (1.45)$$

where R_t is the timestamp of the *freshest* data that the receiver possesses. Normally Δ_t is a stochastic entity as the data arrives to the receiver at random times. Therefore, we will take the expectation of the long-term average of Δ_t ,

$$\Delta = E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \Delta_\tau \right], \quad (1.46)$$

and call it the *average age*. In both main chapters of **Part I**, namely **Chapters 3 and 4**, we study schemes where the data arrival rate is greater than the allowed output rate. Therefore pieces of the data have to be dropped in order to control the output rate. However, all data might not have the same importance. Consequently, the system is penalized for dropping important pieces of data and the distortion metric we propose will be equivalent to

$$D = E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t v(\tau) \mathbb{1}\{\tau^{\text{th}} \text{ packet is not sent}\} \right], \quad (1.47)$$

where $v(t)$ is the importance value of the t^{th} packet and takes values in a finite set \mathcal{V} . If the maximum allowed output rate of the system is

$$R = E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \{\# \text{ of packets sent until } t\} \right], \quad (1.48)$$

one might want to study the attainable (Δ, D, R) tuples, and this is indeed what we will do in **Chapters 3 and 4**. Since it could be of formidable complexity to study all three parameters at once, we study the (Δ, D) tradeoff in **Chapter 3** and the (R, Δ) tradeoff in **Chapter 4**.

Chapter 3: Optimal Policies for Age and Distortion

We consider a discrete-time model where a device (the sender) observes packets at each time instant $t = 1, 2, \dots$ and stores them in its memory. The sender is only allowed to speak at certain times determined by an external scheduler — this restriction fixes the output rate R and leaves us with the study of the (Δ, D) pairs. At a *speaking time*, the sender is allowed to send a single packet from its memory, and the order of the packets shall be preserved. The main results regarding the packet-based strategies described above are as follows:

- We show that among a broad class of strategies (square-integrable), the optimal (Δ, D) tradeoff is attained within finite memory, and we characterize the necessary and sufficient buffer size $K^*(v, R, \eta)$, where η is a weight parameter.
- We devise an efficient policy iteration algorithm that relies on appropriate data structures and problem-specific simplifications. The modified algorithm has $O(N \log N)$ time complexity instead of $O(N^3)$, where $N = |\mathcal{V}|^K$ is the number of possible configurations based on packet importance values in a size- K buffer.
- The optimal policies of the current formulation are also optimal for a problem where packets are subject to erasures and perfect feedback is available.

Later, we allow coding across packets and show that the uncoded strategies can be significantly improved with coding. To this end, we study a simple class of strategies called *buffer ignorant* where the sender bases its decisions on the buffer length, and not on the content. We improve the performance of the buffer ignorant strategies with a *variable-to-fixed* coding scheme, i.e., Tunstall coding.

Chapter 4: Age-Optimal Causal Labeling of Memoryless Processes

In **Chapter 4**, we will study a similar scheme to that in **Chapter 3** except that *equally important* packets arrive intermittently and the arrival process is memoryless. Note that such restriction eliminates D from our formulation and lets us study the (R, Δ) tradeoff. We do not require the packets to arrive in-order, yet it turns out that optimal strategies will preserve the order. Furthermore, a packet can be stored in the buffer indefinitely and can be discarded or forwarded at any time. The main results of this chapter are as follows:

- Although the class of admissible strategies is broad according to the aforementioned formulation, we show that the optimal strategies turn out to be simple and described as ‘wait T , label next’.

- The discrete-time setting can be extended to its continuous-time version where packets arrive as a Poisson process, and ‘wait T , label next’ strategies remain optimal.

1.7.2 Part II: Strategies for Distributed Inference

In **Part II**, we study two different instances of distributed hypothesis testing. In the first problem, studied in **Chapter 7**, peripheral nodes send their observation to a fusion center which makes the ultimate decision. In the second problem, studied in **Chapter 9**, we consider a fully-flat architecture where each node communicates with its neighbors via random selections.

Chapter 7: A Fundamental Limit of Distributed Hypothesis Testing under Memoryless Quantization

We study a scheme where m peripheral nodes make observations and send information to a fusion center at each time instant t . In practice, the nodes may not have sophisticated hardware or sufficient memory, thus we assume that the information sent at time t only depends on the observation made at time t . Furthermore, the *expected* number of bits sent by node i can be at most R_i . The fusion center aims to detect a high-risk event, which is associated with the alternative hypothesis (\mathcal{H}_1) of a binary hypothesis testing problem. Hence, we aim to find the best possible type-II error exponent $\theta^*(R_1, \dots, R_m)$. First, for the single-node case

- we characterize $\theta^*(R_1)$,
- we obtain an upper bound to $\theta^*(R_1)$ via rate–distortion methods and consequently characterize an unachievable region,
- we show that with simple lattice-quantization, the upper bound can be approached within $\frac{1}{2} \log_2(\frac{\pi e}{2}) \approx 1.047$ bits, and
- we provide the upper bound $\theta_k(kR_1)$ for the k -dimensional vector quantization case.

Then, we extend the single-node results to characterize $\theta^*(R_1, \dots, R_m)$ when data is independent across the nodes. If there is a sum-rate constraint, i.e.,

$$\sum_{i=1}^m R_i \leq R, \quad (1.49)$$

we show that an upper bound can be obtained via a waterfilling solution.

Chapter 8: Social Learning under Randomized Collaborations

In **Chapter 8**, we study a network of nodes/agents where each agent communicates with its neighbors. The agents collaboratively seek to find the true state-of-nature (θ°) by sending their beliefs over the possible states. We will formulate the problem as a m -ary hypothesis testing problem in a Bayesian framework. Because of the decentralized nature of the setup, however, fully Bayesian updates are intractable. Hence, we will consider *locally*-Bayesian updates at each agent — this approach is also called *social learning* as initial works on locally-Bayesian updates studied opinion formations in social networks. We will study a sparser and randomized social learning scheme where agents randomly poll a single neighbor for information exchange. The main results of this chapter are as follows:

- We show that all agents learn the truth, i.e., the true hypothesis θ° , eventually and at the same asymptotic learning rate compared to the baseline algorithm.
- Under a special case where agents *replace* their beliefs with a randomly chosen neighbors', we characterize the large deviation estimates which only depend on the marginal distributions at agents.

Part I

Age–Distortion Tradeoffs

Introduction to Age of Information

2

2.1 The Timely Communication Problem

Timeliness is a necessity for many modern communication systems. With emerging Internet-of-Things (IoT) applications and Machine-to-Machine (M2M) communication, stale data have highly undesirable effects; think, for example, of sensor output for autonomous vehicles (this particular example will also be elaborated on Part II), position of an airplane, coolant temperature in a power plant, etc. Therefore, it is crucial to ensure that the destination is kept aware of the status changes as soon as possible in such communication systems — imagine what would happen if data containing an alarming status is backlogged in the network.

Recall Shannon’s ‘non-semantic’ communication problem, where the communication system to be designed is application-independent. Clearly, such formulation does not capture the timeliness aspect of a communication system — the channel capacity being approached at very large blocklengths is a good example that the timeliness is not taken into account since the message is delivered to the destination very late. Consequently, one might need to consider other metrics than the block error probability, and even new channel models that ‘distort’ time for an appropriate formulation of a timely communication problem. This is the point where queueing-theory comes into play: The data input to a channel is received at the destination with some random delay. In the next section, we will explore a simple queueing model that captures the timeliness aspect of communication and we will emphasize the necessity of an appropriate metric to assess the timeliness or freshness. This metric will turn out to be the age of information (AoI), introduced by Kaul, Yates, and Grueteser [15].

2.2 Age of Information: Moving from Delay to Freshness

As a first attempt to formulate a timely communication problem, let us consider a simple first-in-first-out (FIFO) queue; where packets arrive according to a random arrival process, i.e., a point process. Once a packet arrives to the queue, (i) if the queue is empty, it is served. The service time is a random variable, and we assume that the service times are independent and identically distributed for each packet. If (ii) the queue is not empty, the recently arrived packet will wait until all previous packets are served. In the parlance of queueing theory, such queues are called $G/G/1$ queues. The first entry denotes the ‘general’ arrival process of the packets, the second entry denotes the ‘general’ service time, and the third entry indicates the number of servers in the queue, i.e., there is only one server. Note that in this scheme a packet will see (i) a service delay due to the random service time and (ii) a queueing delay due to packets present in the queue. Hence, as we desired, the queueing-theoretic models capture the ‘distortion’ of time and could be suitable candidates for an appropriate timely communication model.

Usually analyses of $G/G/1$ queues are difficult and do not yield closed-form solutions. Therefore, we consider a widely used subset of single-server queues, namely $M/M/1$ queues. Here M denotes the *memorylessness* property of both the arrival process and the service time. A random variable S is memoryless if

$$\Pr(S \geq s + t | S \geq s) = \Pr(S \geq t), \quad \forall t \geq 0. \quad (2.1)$$

In words, if S represents an occurrence time of an event, the fact that ‘this event has not occurred yet’ does not provide additional information about the future occurrence time; hence the random variable S does not have any memory. If S is a continuous random variable, it can be shown that (by taking logarithms of both sides and solving a functional equation) S has an exponential distribution, i.e.,

$$\Pr(S \geq t) = e^{-\lambda t}, \quad \forall t \geq 0 \quad (2.2)$$

for some $\lambda > 0$, and denoted as $\mathcal{E}(\lambda)$. As one may guess, if S is a discrete random variable taking values in $\{1, 2, \dots\}$, S can only have a Geometric distribution, i.e.,

$$\Pr(S > t) = p^t, \quad t \in \{1, 2, \dots\} \quad (2.3)$$

where p is called the failure probability, or the parameter of S . Such a distribution is denoted as $\text{Geom}(p)$, which we will use extensively in the forthcoming chapters.

If one considers the continuous-time communication model, an $M/M/1$ queue model has the property that both (i) interarrival times of packets and (ii) service times are exponentially distributed. As a consequence, the packet arrival process is a Poisson process. In a discrete-time model, the exponen-

tial distributions are replaced with geometric distributions — also denoted as Geo/Geo/. queues.

Let us continue our exploration with an $M/M/1$ queue. Let the rate of the arrival Poisson process be λ , i.e., the interarrival times have $\mathcal{E}(\lambda)$ distribution, and let the service rate be μ , i.e., the service times have $\mathcal{E}(\mu)$ distribution. Furthermore, denote the time elapsed between packet $i - 1$ and packet i (i th interarrival time) as Z_i and denote the service time of packet i as S_i . As we mentioned before, there are two sources of delay: (i) the service delay and (ii) the queueing delay. Obviously, the service delay for packet i is S_i . The queueing delay, denoted as D_i , is a more complex entity and can be found by the following recursive equation [16]:

$$D_i = (D_{i-1} + S_{i-1} - Z_i)^+ := \max\{D_{i-1} + S_{i-1} - Z_i, 0\}, \quad i > 1 \quad (2.4)$$

and $D_1 = 0$. The above relation suggests that unless $E[S] < E[Z]$, or equivalently $\mu > \lambda$, D_i tends to infinity with probability one. Also, the queue will be unbounded. Suppose as a system designer, we have control over the packet arrival rate λ . A possible metric we consider could be the service (or channel) delay. By the law of large numbers, the long-term average of the service delays

$$\frac{1}{n} \sum_{i=1}^n S_i \rightarrow \frac{1}{\mu} \quad (2.5)$$

almost surely, which we do not have any control over. Another possible metric we can consider could be the long-term average of the queueing delays. Equation (2.4) suggests that D_i 's are correlated, hence we need another tool than the law of large numbers. Luckily, Little's law [16] comes to help. An application of Little's law will yield

$$\frac{1}{n} \sum_{i=1}^n (D_i + S_i) \rightarrow \frac{1}{\mu - \lambda} \quad (2.6)$$

and therefore,

$$\frac{1}{n} \sum_{i=1}^n D_i \rightarrow \frac{1}{\mu - \lambda} - \frac{1}{\mu} = \frac{\lambda}{\mu(\mu - \lambda)}. \quad (2.7)$$

We have control over the above quantity. To minimize the long-term average delay, one sets a very small λ . One might notice that such choice of λ is not a smart one. A very small λ tells that there is almost no packet arrival and hence almost no communication takes place. It seems that a sight from individual packets' perspectives does not guide us towards a meaningful performance assessment. We have to change our perspective to the *receiver side*.

Assume that the receiver values *fresh* packets more than the stale ones. This implies that if at time t , the receiver obtains a packet generated at time $s < t$, its penalty might be a non-increasing function of s . I.e., the receiver penalizes the delay as in (2.7). Furthermore, the receiver may also want to penalize the packet interarrival times, as longer interarrival times keep the

receiver less updated. Note that this second penalty is not taken into account in (2.7), and will guide us towards a meaningful entity. For now, we will give an informal definition of this entity. The detailed definitions suitable for our system model will be given in **Chapter 3**. Suppose the packets contain timestamps, i.e., their arrival times to the sender. Let R_t be the timestamp of the most recent packet that the receiver possesses. Assume $R_0 = 0$. Now consider the following process:

$$\Delta_t = t - R_t. \quad (2.8)$$

For simplicity, define the total waiting time

$$W_i := D_i + S_i. \quad (2.9)$$

If we plot the evolution of (2.8), we will obtain a sawtooth-like graph as follows:




Figure 21 – An example evolution of Δ_t . The packet interarrival times Z_i 's and waiting times W_i 's are indicated.

Δ_t is called the *instantaneous age of information (AoI)*. Since it is a random entity, one can study its moments, e.g., $E[\Delta_t]$, $E[\Delta_t^2]$, or their long term behaviors, e.g., $\limsup_{t \rightarrow \infty} E[\Delta_t]$, $\limsup_{t \rightarrow \infty} E[\Delta_t^2]$, e.g., [17]. Furthermore,

- Average Peak AoI: $\Delta_P := \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n E[Z_i + W_i]$
- Average AoI: $\Delta := \limsup_{n \rightarrow \infty} \frac{1}{t} \int_0^t E[\Delta_\tau] d\tau$

are well-studied metrics as well. In **Chapters 3 and 4**, we will optimize a metric similar to the average AoI, which we will also denote as Δ .

We end this section with an important observation. First of all, the $M/M/1$ model is ergodic, hence time averages converge. More precisely,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \Delta_\tau d\tau \quad (2.10)$$

exists with unit probability, and is equal to Δ ; which turns out to be [15]

$$\Delta = \frac{1}{\mu} \left(1 + \frac{1}{\rho} + \frac{\rho^2}{1-\rho} \right), \quad (2.11)$$

where $\rho := \frac{\lambda}{\mu} < 1$ is called the utilization rate. Hence, as we discussed in this section, if the system designer has control over the packet generation rate λ (or equivalently over ρ), the average age is minimized at $\rho \approx 0.53$. Note that this result tells that the age is not minimized when (i) $\lambda \rightarrow 0$, as opposed to minimizing the average queueing delay (2.7) or when (ii) $\lambda \rightarrow \mu$, as opposed to maximizing the throughput. It penalizes both the interarrival times and the latency.

2.3 Related Work

As previously mentioned, freshness of data is recently recognized as a semantic aspect of communication. Initial work by Kaul et al. introduced the AoI as a metric to quantify this aspect [15,18,19]. Following Kaul et al., there have been many studies adopting AoI as a freshness metric. The first strand involved calculation of AoI in simple schemes, e.g. M/M/1 queues [15]. Subsequent extensions involve more general queues [20–24], multiple source streams [25–30], various queue management techniques such as the Last-Come-First-Served (LCFS) protocol [31], and models that allow packet discards [17,32–35] and deadlines [36]. A partial list of studies that seek to compute or to minimize the age under energy or link constraints is [37–52]; and some studies on timely lossless source coding are [53–55]. For a comprehensive survey over the AoI literature, see [56]; and for a tutorial, see [57].

Although the works cited above mostly assume error-free transmissions, some others take into account that packets get lost or erased while passing through the network. In [58], the AoI is studied in a model where transmissions are error-prone. One may notice that a simple method to combat erasures is to send the packet repeatedly — recall the simple feedback example scheme mentioned in **Chapter 1**. A more complicated method could send coded packets, which are then to be conveyed through the erasure channel. A list of works concerning coded transmissions with feedback is [59–66]. An example of a study assuming no feedback is [67], where the authors find the optimal coding strategy. In Section 3.5, we will discuss that the optimal strategy for our discrete-time model is also optimal for a model where communications take place over an erasure channel with feedback.

AoI has also found place in stochastic control literature. For instance, in [68], a tradeoff between the information staleness and performance of a Linear Quadratic Regulator (LQR) is illustrated. In [69], freshness is taken as basis for an algorithm devised for distributed tracking of a linear system. These works are aligned in the sense that using the fresh data to track and control a system might improve the performance. This is because of the Markovian nature of the

processes that are tracked — the next state depends only on the freshest data. However, this may not always be the case as the freshest may not be the most important. This observation is in line with some further studies. For instance in [70], a problem of generating timely updates in a remote estimation setting has been proposed. The authors have investigated the Mean-square-optimal and AoI-optimal strategies for the estimation of a Wiener process through a queue and concluded that they are different; consequently demonstrating a tradeoff between freshness and importance. In [71], the authors generalized the settings to include an Ornstein–Uhlenbeck process. In [49], a tradeoff between timeliness and distortion is shown for the case of estimation through a Gaussian channel in a power constrained setting. There also has been several works on integrating the notion of different data importance and timeliness, e.g., by introducing non-linear cost to stale data [72, 73], by considering separate data streams of different priorities [74, 75], or by modeling the distortion as a decreasing function of the service time [76].

We will see in **Chapter 3** that our setup contains flavors from the above approaches, yet it features novel aspects. For instance, the resource constraint is imposed by an external scheduler, giving the sender turns to speak. Hence, as opposed to the works [70, 71], the sender cannot decide when to send; but it rather decides what to send (In **Chapter 4**, we will also study a problem in an attempt to answer both questions simultaneously). Furthermore, we adopt the view that data is formed into packets of different importance levels, e.g., packets containing abnormal levels of coolant temperature in a nuclear plant could be classified as important. Consequently, the distortion metric we propose depends on whether the packets are received or not, and the accumulated importance levels of the missed data constitutes our distortion metric.

Optimal Policies for Age and Distortion¹

3

3.1 Motivation – Are All Packets of Same Importance?

Consider the following scenario, where data from multiple sources are streamed and multiplexed at a device:




Figure 31 – A scenario where the device is throttled. Source i (Src_i) streams packets at rate R_i , $i \in \{1, 2\}$. The output rate $R_{\text{out}} < R_1 + R_2$, hence the device should drop some packets. R_{out} is controlled by an external scheduler which assigns time slots for the device to forward packets.

Further assume that the output rate is constrained such that not all packets can be forwarded through the bottleneck. Consequently, some of the packets have to be dropped. Suppose the device is assigned time slots to speak by an external scheduler, which we will refer as speaking times. Between the speaking times, the device collects packets in the buffer, and at a speaking time it is allowed to send one packet from its buffer. This way, the output rate is kept under control. If both sources contain almost identical information, the device could discard the stale packets. However, if Source 2 streams valuable

¹The content of this chapter is based on [77, 78].

information, such as our previous example of coolant temperature in a nuclear plant, then the device might choose to send these important packets from Source 2, as illustrated in the figure below.




Figure 32 – Same scenario as in Figure 31 depicted at a speaking time. The timestamps of the previously arrived packets are denoted as S_i (not to be confused with service times of the previous section) with $S_i \leq S_{i+1}$. At this speaking time, the device chooses to send the packet from Source 1 with corresponding timestamp S_4 .

In Figure 32, the device did not forward the freshest packet with timestamp S_6 but rather chose the packet with timestamp S_4 , generated by Source 2. One may ask what could be the best strategy that the device can execute. Moving towards our problem definition, we will make a simplification to the above model. We will consider a single source by merging the different source streams.




Figure 33 – The single-source model subject to the studies in this chapter.

Note that the device is renamed as ‘Sender’ in the figure above. This is to prepare for our main problem definition, which we will provide in the next section. At this point, it may be useful to make an important remark. In our example, there were two sources, and the second source streamed more valuable packets. What if there were more than two sources? How would we assign importances to multiple streams? To this end, as we mentioned in **Chapter 1**, we will assign *fidelity criteria* to strategies, according to which the unsent packets will incur some penalty. Naturally, we will assign higher penalties to the packets of higher importances. In the end, the long-term average of the penalties incurred by the unsent packets will constitute our *fidelity criterion*. Now, we can continue with the rigorous problem definition. But first, it may be beneficial to clarify some notation that we will use in the sequel.

3.2 Notation

Random variables are denoted with uppercase letters (e.g., X); and vectors are denoted with boldface letters (e.g., \mathbf{b}). Sets are denoted with script-style letters (e.g., \mathcal{V}). $l(\mathbf{b})$ is the length of a vector \mathbf{b} , and b_i is its i^{th} element. For vectors \mathbf{b} and \mathbf{b}' , $\mathbf{b}\|\mathbf{b}' := [b_1, b_2, \dots, b'_1, b'_2, \dots]$ is the concatenation of the two vectors. $\mathbf{b}_{\geq i} := [b_i, \dots, b_l]$ is segment of \mathbf{b} from its i^{th} element until the end; and $\mathbf{b}_i^j := [b_i, \dots, b_j]$ is the segment between its i^{th} and j^{th} elements, $\mathbf{b}^i := \mathbf{b}_1^i$. \mathbf{b}' is a suffix of \mathbf{b} if there exists an $i > 1$ such that $\mathbf{b}' = \mathbf{b}_{\geq i}$. If $\mathbf{b}' = \mathbf{b}_{\geq i}$ is suffix of \mathbf{b} , then $\mathbf{b} \setminus \mathbf{b}' = \mathbf{b}^{i-1}$. For $a, b \in \mathbb{R}$, $a \wedge b := \min\{a, b\}$, and $a \vee b := \max\{a, b\}$.

3.3 Problem Definition

In this section, we provide a detailed description of our discrete-time model in terms of the data to be conveyed, the sender-receiver pair with their respective communication protocol, and the channel in between.

Similar to our previous explorations of Section 2.2 and 3.1, we assume that the data is formed into packets, and at each time instant t , a new packet arrives to the sender. The packet payloads originate from a set of finite elements \mathcal{X} , and the probability of a payload taking a particular value is time-invariant and independent of the past. Consequently, the data is an independent and identically distributed (i.i.d.) process $\{X_t\}_{t \in \mathbb{N}}$. The sender observes X_t at time t and keeps X_t in its buffer.

The communication protocol is as follows: The sender is allowed to speak at times T_1, T_2, \dots . The process $\{T_i\}_{i \in \mathbb{N}}$ is independent of the process $\{X_t\}_{t \in \mathbb{N}}$, and has the property that the interspeaking times $Z_i := T_i - T_{i-1}$ are i.i.d.. Moreover, we assume that Z_i 's are strictly positive and square integrable, i.e., $\Pr(Z_i > 0) = 1$ and $E[Z_i^2] \leq \infty$. An example of such a random variable could be a geometric random variable with $\Pr(Z_i = t) = p(1-p)^{t-1}$ for $t \geq 1$. The speaking process $\{T_i\}_{i \in \mathbb{N}}$ is inspired by link layer multiple access protocols where the sender is assigned time slots to speak by an external controller. When the sender is given a turn to speak, i.e., at each T_i , it selects a packet from its buffer with timestamp $S_i \leq T_i$ and forwards X_{S_i} . Once X_{S_i} is forwarded, we restrict the sender to not send a packet with timestamp less than S_i at the subsequent speaking times T_{i+1}, T_{i+2}, \dots . Note that such restriction results in $S_i < S_{i+1}$, and preserves the order of the packets — which is desirable if the receiver has small or no buffer. The increasing sequence $\{S_i\}_{i \geq 0} =: \mathbf{S}$ is henceforth referred as the ‘selection process’.

Transmissions between the sender and the receiver are noiseless and zero-delay. Hence, by time t , the receiver has observed X_{S_i} for every i such that $T_i < t$. We also suppose that the packets are formed to contain their timestamps in the header, i.e., the packet containing X_{S_i} also contains the information that it was generated at time S_i in its header. Consequently, at time t , the receiver

is able to reconstruct the data as $Y_j(t) = X_j$ if X_j is among its observation up to time t ; otherwise it sets $Y_j(t) = ?$, i.e., represents with an erasure symbol.

At this point, we have described our model. Now, we introduce the appropriate distortion (or fidelity criterion) and freshness metrics to study their tradeoff. Specifically, given $d : \mathcal{X} \times \mathcal{X} \cup \{?\} \rightarrow \mathbb{R}_{\geq 0}$, with

$$d(x, x) = 0 \text{ and } d(x, ?) =: v(x), \quad (3.1)$$

and given a selection procedure \mathbf{S} , define

$$D_t^{(\mathbf{S})} := \frac{1}{t} \sum_{i=1}^t d(X_i, Y_i(t)) \text{ and } D^{(\mathbf{S})} := E \left[\limsup_{t \rightarrow \infty} D_t^{(\mathbf{S})} \right]. \quad (3.2)$$

With an analogy to rate-distortion theory, observe that $D_t^{(\mathbf{S})}$ quantifies average distortion between the source and its reconstruction. Similar to what we have described in Section 1.4, $D^{(\mathbf{S})}$ is the expected long-term average distortion.

Timeliness of information will be assessed using a metric similar to the average age Δ introduced in the previous chapter. Namely, with $i(t) := \sup\{i \geq 0 : T_i < t\}$, $T_0 = S_0 = 0$; define for all $t > 0$,

$$\Delta_t^{(\mathbf{S})} := t - S_{i(t)}, \text{ and } \Delta^{(\mathbf{S})} := E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \Delta_\tau^{(\mathbf{S})} \right]. \quad (3.3)$$

As mentioned before, $\Delta_t^{(\mathbf{S})}$ is usually referred as the instantaneous age; and similar to $D^{(\mathbf{S})}$, $\Delta^{(\mathbf{S})}$ is the expectation of the long-term average age.

Note that $Y_i(t)$ can be either equal to X_i or to '?'. Therefore, specifying $d(x, x)$ and $d(x, ?)$ — which is readily determined by $v(x)$ — is sufficient to evaluate $D^{(\mathbf{S})}$. As a consequence, the sender may base its selection S_i on $\mathbf{V}^{T_i} := [V_1, \dots, V_{T_i}]$, where $V_t := v(X_t)$. Therefore, the selection S_i is a mapping from \mathcal{V}^{T_i} to $\{S_{i-1}+1, \dots, T_i\}$ with $\mathcal{V} := \{v(x) : x \in \mathcal{X}\}$. Intuitively, V_i represents an importance score for the packet i ; high V_i is interpreted as the content having high importance and not sending it incurs a high penalty — this interpretation is also consistent with a model where some arrivals are prioritized. Observe that the structure of the problem stays the same if all elements of \mathcal{V} are multiplied by a positive constant. If \mathcal{V} does not contain 0, then without loss of generality one can assume that the minimum element in \mathcal{V} is 1 and it is an ordered set as $1 = v_1 < v_2 < \dots < v_{|\mathcal{V}|} := v_{\max} < \infty$; otherwise $v_1 = 0$.

Now that we have the full description of the setting, we aim to characterize the achievable region of $(\Delta^{(\mathbf{S})}, D^{(\mathbf{S})})$ pairs. We attempt to characterize this region in the sequel and conclude this section with a few remarks.

- (i) The model we propose is reminiscent of a remote estimation problem of a discrete-time stochastic process through a discrete-time queue. However, we require that the sender sends a packet exactly at speaking times, which is equivalent to force the sender to send as soon as the queue is idle in a discrete-time queueing setting. In [79] and [70], it is shown that the

optimal policies need not be of this type. This makes our problem different and allows us to make the relaxation that S_i need not be stopping times (with respect to the natural filtration of the process $\{X_i\}_{i \in \mathbb{N}}$).

- (ii) If $0 \in \mathcal{V}$, there are multiple interpretations. $V_t = 0$ can be interpreted as either the data is totally trivial (need not be reconstructed), or interpreted as the source having not generated data at time t . The second interpretation allows us to model a source which generates data sporadically. Now there is the question of allowing X_t to be sent or not. Our model allows sending of X_t , i.e., in the second interpretation, informs the receiver that there has not been any data generated by the source, and Δ_t decreases accordingly. The reduction of Δ_t can be avoided by appropriate reformulation — to be discussed in Remark 3.1.

3.4 The Age-Distortion Tradeoff

3.4.1 Markov Decision Problem Formulation as a Lower Bound

To study the age-distortion tradeoff, we study the family of weighted costs $\eta\Delta^{(\mathbf{S})} + D^{(\mathbf{S})}$ for $\eta > 0$. It is known that, e.g., [80], the boundary of the achievable $(\Delta^{(\mathbf{S})}, D^{(\mathbf{S})})$ region can be characterized by studying this family. We seek to obtain a tractable lower bound for $\eta\Delta^{(\mathbf{S})} + D^{(\mathbf{S})}$, and then we further optimize this lower bound by choosing the best \mathbf{S} . First, we derive a simpler expression for $\Delta^{(\mathbf{S})}$. Observe that

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \Delta_{\tau}^{(\mathbf{S})} = \limsup_{i \rightarrow \infty} \frac{\sum_{j=1}^i Q_j^{(\mathbf{S})}}{\sum_{j=1}^i Z_j} \quad (3.4)$$

where

$$Q_j^{(\mathbf{S})} := (T_j - S_j)Z_{j+1} + \frac{Z_{j+1}(Z_{j+1} + 1)}{2}. \quad (3.5)$$

Since $\lim_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i \frac{Z_{j+1}(Z_{j+1}+1)}{2} = \frac{1}{2}E[Z_1(Z_1 + 1)] =: \nu_Z$ and $\lim_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i Z_j = E[Z_1] =: \mu_Z$ with probability 1 by the law of large numbers, we obtain

$$\Delta^{(\mathbf{S})} = E \left[\frac{1}{\mu_Z} \limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i (T_j - S_j)Z_{j+1} + \frac{\nu_Z}{\mu_Z} \right]. \quad (3.6)$$

Note that $\Delta^{(\mathbf{S})}$ cannot be smaller than ν_Z/μ_Z . We subtract ν_Z/μ_Z to obtain the excess age, given by

$$\Delta_e^{(\mathbf{S})} := E \left[\frac{1}{\mu_Z} \limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i (T_j - S_j)Z_{j+1} \right] \quad (3.7)$$

and determine the feasible $(\Delta_e^{(\mathbf{S})}, D^{(\mathbf{S})})$ pairs. With the same reasoning as above, we study the family $\eta\Delta_e^{(\mathbf{S})} + D^{(\mathbf{S})}$. When the selection process \mathbf{S} satisfies

a certain square-integrability condition given in the theorem below, we can find alternative expressions for $\Delta_e^{(\mathbf{S})}$ and $D^{(\mathbf{S})}$.

Theorem 3.1. *Let \mathcal{S}_2 be the set of selection processes \mathbf{S} with $\sup_i E[(T_i - S_i)^2] < \infty$. Then for any $\mathbf{S} \in \mathcal{S}_2$,*

$$\Delta_e^{(\mathbf{S})} = E \left[\limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i (T_j - S_j) \right], \quad (3.8)$$

$$D^{(\mathbf{S})} = \frac{1}{\mu_Z} E \left[\limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_{j-1}, S_j) \right], \quad (3.9)$$

where

$$D(\mathbf{V}^{T_j}, S_{j-1}, S_j) := \sum_{j'=S_{j-1}+1}^{S_j-1} V_{j'} \quad (3.10)$$

is the penalty incurred by skipping the portion $[V_{S_{j-1}+1}, \dots, V_{S_j-1}]$ of \mathbf{V}^{T_j} .

Proof. See Appendix 3.8.1. \square

At this point, we would like to eliminate some of the non-optimal selection processes. More specifically, we show that if a packet with importance value v_{\min} and with timestamp less than T_i is selected at time T_i , then one can find another selection process which performs at least as well.

Lemma 3.1. *Consider a selection process \mathbf{S} with $S_{i_0} < T_{i_0}$ and $V_{S_{i_0}} = v_{\min}$ for some i_0 . Then one can find another process $\tilde{\mathbf{S}}$ with $V_{\tilde{S}_{i_0}} > v_{\min}$ and such that $\Delta_e^{(\tilde{\mathbf{S}})} \leq \Delta_e^{(\mathbf{S})}$ and $D^{(\tilde{\mathbf{S}})} \leq D^{(\mathbf{S})}$.*

Proof. See Appendix 3.8.2. \square

Lemma 3.1 helps restrict the search space for possibly optimal selection processes. The processes we study in the sequel will not select a minimum-importance packet if it has not arrived exactly at the speaking time. Let us denote the class of such selection processes as \mathcal{S}'_2 .

A search for an optimal strategy based on the expressions in Theorem 3.1 seems to be a complex task. We aim to obtain a further lower bound for $\eta \Delta_e^{(\mathbf{S})} + D^{(\mathbf{S})}$ which turns out to be an entity that is amenable for analysis. This lower bound is given in

Theorem 3.2. *Define*

$$J_i(\eta)^{(S_1^{T_i})} := E \left[\sum_{j=1}^i \frac{1}{\mu_Z} D(\mathbf{V}^{T_j}, S_{j-1}, S_j) + \eta(T_j - S_j) \right] \quad (3.11)$$

and

$$J(\eta)^{(\mathbf{S})} := \limsup_{i \rightarrow \infty} \frac{1}{i} J_i(\eta)^{(S_1^{T_i})}. \quad (3.12)$$

Then, for any $\mathbf{S} \in \mathcal{S}'_2$, $J(\eta)^{(\mathbf{S})} \leq D^{(\mathbf{S})} + \eta \Delta_e^{(\mathbf{S})}$.

Proof. See Appendix 3.8.1. □

A straightforward implication is

$$J^*(\eta) := \inf_{\mathbf{S} \in \mathcal{S}'_2} J(\eta)^{(\mathbf{S})} \leq \inf_{\mathbf{S} \in \mathcal{S}'_2} \left(D^{(\mathbf{S})} + \eta \Delta_e^{(\mathbf{S})} \right). \quad (3.13)$$

As we shall see later in Section 3.4.3, it turns out that the inequality (3.13) is indeed an equality. The key reason to introduce $J(\eta)^{(\mathbf{S})}$ is that its infimization can be formulated as a Markov Decision Process (MDP), which we do next. This requires identifying the states and the actions of the MDP and verifying that (i) the distribution of the next state and (ii) the one-step cost depend only on the current state-action pair. We claim that the buffer content at the i^{th} speaking time

$$\mathbf{B}_i := \mathbf{V}_{S_{i-1}+1}^{T_i}, \quad (3.14)$$

and the number of packets to be dropped $A_i := S_i - S_{i-1}$ constitutes this state-action pair. To see this, note that given \mathbf{B}_i and A_i , (i) the next buffer content is independent of the past; and (ii) the one-step cost

$$\frac{1}{\mu_Z} D(\mathbf{V}^{T_i}, S_{i-1}, S_i) + \eta(T_i - S_i) \quad (3.15)$$

is a function only of \mathbf{B}_i and A_i . This is because the above is equal to

$$\frac{1}{\mu_Z} \sum_{k=1}^{s-1} b_k + \eta(l(\mathbf{b}) - s); \quad (3.16)$$

with $\mathbf{b} = \mathbf{B}_i$ and $s = A_i$. Since (i) and (ii) are satisfied, the problem is indeed an MDP whose state at instant i is the buffer content \mathbf{B}_i ; and whose action at instant i is A_i . The assumption $S_0 = 0$ allows us to choose \mathbf{B}_0 as an empty buffer.

The formulation above is an infinite-horizon average-cost MDP [81] with states $\mathbf{b} \in \mathcal{V}^* := \cup_{k=1}^{\infty} \mathcal{V}^k$; and the set of possible actions for a state \mathbf{b} is given by $s \in \{1, \dots, l(\mathbf{b})\}$, where $l(\mathbf{b})$ is the length of the buffer \mathbf{b} . Consequently, the sender chooses the packet with timestamp $S_i = T_i + s - l(\mathbf{b})$ at time T_i . Observe that setting $s = l(\mathbf{b})$ corresponds to the selection of the freshest packet whereas setting $s = 1$ corresponds to the selection of the oldest packet in the buffer. Furthermore, since we are interested in the selection procedures in \mathcal{S}'_2 , i.e., a minimum importance packet is only chosen if it has arrived exactly at the speaking time, $s \neq l(\mathbf{b})$ only if $b_s > v_{\min}$.

In general, the optimal policies of an MDP need not be stationary. A policy is stationary if the current action is a function of the current state. We will show in Section 3.4.3 that for our problem, the optimal policy is indeed stationary and deterministic. When we consider a stationary and deterministic selection process \mathbf{S} , we explicitly write the argument (\mathbf{b}) . Let us recall

Definition 3.1 (Unichain policy, [81]). *If a stationary policy $s(\mathbf{b})$ induces a Markov chain with a single recurrent class and a possibly empty set of transient states, it is called unichain.*

In an average-cost dynamic programming setting, we evaluate a unichain policy $s(\mathbf{b})$, $\mathbf{b} \in \mathcal{V}^*$ by solving the linear system with unknowns $h(\mathbf{b})$, $\mathbf{b} \in \mathcal{V}^*$ and λ ; given by [81, Chapter 4]

$$h(\mathbf{b}) + \lambda = \frac{1}{\mu_Z} \sum_{k=1}^{s(\mathbf{b})-1} b_k + \eta(l(\mathbf{b}) - s(\mathbf{b})) + E[h(\mathbf{b}_{\geq s(\mathbf{b})+1} \| \mathbf{V}^Z)], \quad (3.17)$$

where Z is the next interspeaking time, \mathbf{V}^Z is a vector of i.i.d. V 's of length Z , $h(\mathbf{b})$ is called the relative value of state \mathbf{b} , and λ is the average cost induced by this policy. Note that the right-hand side of (3.17) is equal to the sum of distortion cost of skipped packets, weighted age cost, and the expected relative value of the next state. Since (3.17) determines $h(\mathbf{b})$ up to an additive constant, we take a reference state as one of the $\mathbf{b} \in \mathcal{V}^*$ and we set $h(\mathbf{b}) = 0$. We also note that for a unichain policy, the linear system given by (3.17) has a unique solution [81].

Remark 3.1. *To cover the case where $v = 0$ is interpreted as the source having not generated any data and Δ_t should not decrease upon the sending of $v = 0$; one can proceed as follows: The set of possible actions for a state \mathbf{b} is extended to $\{0, 1, \dots, l(\mathbf{b})\}$, and $s > 0$ only if $b_s > 0$. That is, the sender is allowed to choose only the packets with $v > 0$. Also note that for the all-zero buffer, the only possible action is to set $s = 0$, i.e., nothing has arrived since the last selection and hence there is nothing to send. Observe that in this case Δ_t does not drop.*

Although we have characterized a lower bound based on a MDP formulation, the formulated problem has a countably infinite state space, \mathcal{V}^* . It is known that for this class of problems, analysis of optimal policies become formidably complex in general. Moreover, it is not certain that a stationary policy attains the infimum in (3.13). The complexity of this problem leads us to consider a finite-state modification of the problem; and the next section is devoted for this modified version.

3.4.2 Policy Iteration with a Truncated State Space

We now consider a finite-state version of the problem where the sender forgets the packets that have arrived more than K time slots ago. That is, the current buffer content at time T_i becomes

$$\mathbf{B}_i := \mathbf{V}_{(S_{i-1}+1)\vee(T_i-K+1)}^{T_i}. \quad (3.18)$$

Consequently, the buffer length is limited to at most K , and the state space becomes finite. Denote this finite state space by $\mathcal{V}^{\leq K} := \cup_{l \leq K} \mathcal{V}^l$. One may notice that by restricting the state space, we might not attain the infimal value $J^*(\eta)$. However, as we shall see later in 3.4.3, the optimal policy of the original *infinite* state space problem will base its decisions only on a bounded buffer. Thus, we do not lose optimality provided that K is large enough.

A slight modification of (3.17) is enough to obtain the linear system whose solution yields the relative values and the average cost. First, observe that if the buffer state \mathbf{b} has length more than K , i.e., if $l(\mathbf{b}) > K$, then $h(\mathbf{b})$ can be replaced with

$$\frac{1}{\mu_Z} \sum_{k=1}^{l(\mathbf{b})-K} b_k + h(\mathbf{b}_{\geq l(\mathbf{b})-K+1}) \quad (3.19)$$

as the first $l(\mathbf{b}) - K$ terms will be forgotten in the truncated problem. Let $p_i := \Pr(Z = i)$ and $q_i := \Pr(Z \geq i)$. Then the linear system of equations to evaluate a unichain stationary policy $s(\mathbf{b})$, $\mathbf{b} \in \mathcal{V}^{\leq K}$ becomes

$$h(\mathbf{b}) + \lambda = \frac{1}{\mu_Z} \sum_{k=1}^{s(\mathbf{b})-1} b_k + \eta(l(\mathbf{b}) - s(\mathbf{b})) + \frac{1}{\mu_Z} \sum_{k=s(\mathbf{b})+1}^l b_k q_{K+k-l} + \frac{E[V]}{\mu_Z} E[(Z - K)^+]$$

$$+ \sum_{k=1}^{K-l+s(\mathbf{b})} p_k E[h(\mathbf{b}_{\geq s(\mathbf{b})+1} \| \mathbf{V}^k)] + \sum_{k=K-l+s(\mathbf{b})+1}^{K-1} p_k E[h(\mathbf{b}_{\geq k-(K-l)+1} \| \mathbf{V}^k)]$$

$$+ q_K E[h(\mathbf{V}^K)] \quad (3.20)$$

$$=: C_h(\mathbf{b}, s(\mathbf{b})), \quad (3.21)$$

with $h(v_{\min}) = 0$. This choice of the reference state also sets $h(\mathbf{b}) = 0$ for $\mathbf{b} \in \mathcal{V}$. Namely, for buffers that contain only one packet, the relative value will be equal to zero.

At this point, we have obtained the policy evaluation method for our finite-state problem. A first attempt could be to find the optimal policies numerically. We consider the well-known policy iteration algorithm [81]. A brief description is given in Algorithm 1.

Algorithm 1: Policy iteration

- 1 Start with the stationary policy $s^{(0)}(\mathbf{b}) = l(\mathbf{b})$.
- 2 Evaluate $s^{(i)}(\mathbf{b})$ according to (3.20) to find $h^{(i)}(\mathbf{b})$, $\mathbf{b} \in \mathcal{V}^{\leq K}$ and $\lambda^{(i)}$.
- 3 For all $\mathbf{b} \in \mathcal{V}^{\leq K}$, set

$$s^{(i+1)}(\mathbf{b}) = \arg \min_{\substack{s \in \{1, \dots, l(\mathbf{b})\} \\ b_s \neq v_{\min} \text{ if } s < l(\mathbf{b})}} C_{h^{(i)}}(\mathbf{b}, s)$$

- 4 If $s^{(i+1)}(\mathbf{b}) = s^{(i)}(\mathbf{b})$ for all $\mathbf{b} \in \mathcal{V}^{\leq K}$, terminate. Else go to step 2.
-

Although the number of states is now finite, it is not clear that the policy iteration algorithm yields a unichain policy. We shall show this in the following lemma.

Lemma 3.2. *If the buffer size is limited to K , the policy iteration terminates with an optimal unichain policy in \mathcal{S}_2^l .*

Proof. Consider all truncated policies, e.g., non-stationary, history dependent but can only choose the most recent K packets in the buffer. Since $\{V_i\}_{i \geq 0}$ is an i.i.d. process, any policy eventually reaches a state consisting of only v_{\min} 's and must choose the most recent packet. Hence, the buffer must be eventually renewed for any policy in \mathcal{S}'_2 and as a result, there must be a single recurrent class. Therefore, there exists an optimal stationary and deterministic strategy that is unichain and this policy can be found with the policy iteration algorithm [81]. \square

Remark 3.2. *Intuitively, step 3 of the above algorithm modifies $s^{(i)}(\mathbf{b})$ in the following way: Consider two processes starting at the state \mathbf{b} . The first one is iterated with respect to $s^{(i)}$, whereas the second one is iterated with a different action $\tilde{s}(\mathbf{b})$ at the first step and with $s^{(i)}$ subsequently. Now consider the expected accumulated costs of these two processes until they reach the same state. If the second process has a smaller expected accumulated cost, changing all $s^{(i)}(\mathbf{b})$ to $\tilde{s}(\mathbf{b}) = s^{(i+1)}(\mathbf{b})$ results in a better policy; otherwise try another $\tilde{s}(\mathbf{b})$.*

3.4.3 The Exact Buffer Size for an Optimal Policy

Truncating the state space restricts the actions that may be taken. Therefore, in general, the infimum in (3.13) may not be attained with a truncated buffer. As we have said in the previous section, it turns out that this is not the case for our problem and the infimum is indeed attained with a finite buffer size. In this section we quantify this buffer size.

First, consider the policy $s(\mathbf{b}) = l(\mathbf{b})$ for all $\mathbf{b} \in \mathcal{V}^{\leq K}$, i.e., always send the most recent packet in the buffer. One can observe that this policy induces a Markov chain with only $|\mathcal{V}|$ states regardless of K . We shall now show that this policy is optimal for η above some threshold η_{\max} .

Lemma 3.3. *For $\eta \geq \eta_{\max} := \frac{1}{\mu_Z}(v_{\max} - v_{\min})$ and for any $M \geq 1$, the optimal policy among $\mathcal{V}^{\leq M}$ is $s(\mathbf{b}) = l(\mathbf{b})$; which can be implemented with a buffer size of 1. In words, if age-penalty weight η is large enough, the optimal strategy always chooses the freshest packet.*

Proof. We show that the policy $s(\mathbf{b}) = l(\mathbf{b})$ remains unchanged under policy iteration. Start the policy iteration with $s^{(0)}(\mathbf{b}) = l(\mathbf{b})$. Recalling Remark 3.2, we will show that perturbing the policy at the initial step cannot decrease the expected accumulated cost until the original and perturbed processes coincide. Assume the perturbed action is $\tilde{s}(\mathbf{b}) = l(\mathbf{b}) - k$ for a $k > 0$. Notice that the two processes will coincide immediately at the next step and the difference of the two accumulated costs will be $k\eta - \frac{1}{\mu_Z}(b_{l(\mathbf{b})-k} - b_{l(\mathbf{b})}) \geq \eta - \frac{1}{\mu_Z}(v_{\max} - v_{\min}) \geq 0$. Hence the perturbed policy incurs a higher cost and the policy remains unchanged. \square

Considering the truncated state space $\mathcal{V}^{\leq K}$, we give some properties of optimal policies.

Property 3.1. For an optimal stationary policy $s^*(\mathbf{b})$, and optimal relative values $h^*(\mathbf{b})$, the following hold: For any $\mathbf{b}, \mathbf{b}' \in \mathcal{V}^{\leq K}$,

(i) For any state $\mathbf{b} \parallel \mathbf{b}'$, either $s^*(\mathbf{b} \parallel \mathbf{b}') = l(\mathbf{b}) + s^*(\mathbf{b}')$ or $s^*(\mathbf{b} \parallel \mathbf{b}') \leq l(\mathbf{b})$.

(ii) $h^*(\mathbf{b}') \leq h^*(\mathbf{b} \parallel \mathbf{b}') \leq \frac{1}{\mu_Z}(b_1 + \dots + b_{l(\mathbf{b})}) + h^*(\mathbf{b}')$.

Proof. See Appendix 3.8.3. □

To get a sense of how far can the optimal policy go back in time, i.e., to measure how large $T_i - S_i$ can be, it may be informative to consider the following extreme case, which yields a lower bound on the maximal possible value of $T_i - S_i$ (Recall that our ultimate aim is to show that $T_i - S_i$ shall be bounded).

Lemma 3.4. For the state $\mathbf{b} = [v_{\max}, \underbrace{v_{\min}, \dots, v_{\min}}_{L-1}]$,

$$s^*(\mathbf{b}) = \begin{cases} 1, & \eta \leq \frac{1}{\mu_Z} \frac{(v_{\max} - v_{\min})}{L-1} \\ L, & \eta > \frac{1}{\mu_Z} \frac{(v_{\max} - v_{\min})}{L-1} \end{cases}. \quad (3.22)$$

Proof. Since we work with policies in \mathcal{S}'_2 , the two possible actions for this state are either choosing the v_{\max} at the beginning or choosing the v_{\min} at the end. Referring to Remark 3.2, suppose at iteration i we have $s^{(i)}(\mathbf{b}) = l(\mathbf{b})$ and $\tilde{s}(\mathbf{b}) = 1$; and we aim to find the difference of accumulated costs until the original and the perturbed processes coincide. Observe that these processes coincide immediately after the first step and the difference will be $\eta(L-1) + \frac{1}{\mu_Z}(v_{\min} - v_{\max})$. Then, $s^{(i+1)}(\mathbf{b}) = 1$ if $\eta(L-1) \leq \frac{1}{\mu_Z}(v_{\max} - v_{\min})$; otherwise $s^{(i+1)}(\mathbf{b}) = l(\mathbf{b})$. Note that the difference does not depend on i and hence the statement for $s^{(i+1)}(\mathbf{b})$ is also true for $s^*(\mathbf{b})$. □

The above lemma therefore gives a necessary buffer size for a possibly optimal policy as it tells that at $\eta = \frac{1}{\mu_Z} \frac{(v_{\max} - v_{\min})}{L-1}$, the first packet in the buffer given in Lemma 3.4 is chosen by the optimal policy. Hence, attaining the optimal policy requires a buffer size of at least $\lceil \frac{1}{\mu_Z} \frac{(v_{\max} - v_{\min})}{\eta} \rceil$. Observe that this does not imply that the optimal policy is reached within this particular finite buffer size. Nevertheless, we can prove that this is indeed the case.

Theorem 3.3. For $M \geq K(\eta) := \lceil \frac{1}{\mu_Z} \frac{(v_{\max} - v_{\min})}{\eta} \rceil$, the optimal policy among $\mathcal{V}^{\leq M}$ is attained by a policy with buffer size $K(\eta)$. Furthermore, if $b_{s^*(\mathbf{b})} = v_i$, then $l(\mathbf{b}) - s^*(\mathbf{b}) < K_i(\eta) := \lceil \frac{1}{\mu_Z} \frac{(v_i - v_{\min})}{\eta} \rceil$ for all $1 \leq i \leq |\mathcal{V}|$.

Proof. See Appendix 3.8.4. □

Theorem 3.3 implies that when the policy iteration terminates, the policy it outputs not only solves the Bellman equation for the state space $\mathcal{V}^{\leq M}$ for every $M \geq K(\eta)$, but it also solves the Bellman equation for the countable state space \mathcal{V}^* . Since $h^*(\mathbf{b})$ is finite and $s^*(\mathbf{b})$ is attained for every \mathbf{b} , a

straightforward extension of Proposition 2.1 in [81, Chapter 4] concludes that $s^*(\mathbf{b})$ is indeed the optimal policy that attains $J^*(\eta)$. Let $\mathbf{S}^* = \{s^*(\mathbf{B}_i)\}_{i \geq 0}$ be the random sequence of the actions taken by the stationary and deterministic policy $s^*(\mathbf{b})$. Recall the inequality

$$J^*(\eta) \leq D^{(\mathbf{S}^*)} + \eta \Delta_e^{(\mathbf{S}^*)} \quad (3.23)$$

given in Theorem 3.2. Furthermore, observe that the buffer state process $\{\mathbf{B}_i\}$ controlled by \mathbf{S}^* is a renewal process — this follows from Lemma 3.2. Hence, by the renewal reward theorem [16, Theorem 3.6.1] we have

$$\begin{aligned} \Delta_e^{(\mathbf{S}^*)} &= E \left[\limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i (T_j - S_j) \right] = E \left[\lim_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i (T_j - S_j) \right] \\ &= \lim_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i E[T_j - S_j] \end{aligned} \quad (3.24)$$

and similarly

$$D^{(\mathbf{S}^*)} = \frac{1}{\mu_Z} \lim_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i E[D(\mathbf{V}^{T_j}, S_{j-1}, S_j)]. \quad (3.25)$$

The above shows that the inequality (3.23) is indeed an equality. Therefore, the optimal policies for the MDP give the tangent lines to the exact boundary curve of the achievable (Δ_e, D) region, where η is the slope of a tangent. Consequently, by varying η , this curve can be found. We end this section with the following corollary that summarizes the above results.

Corollary 3.1. *The optimal policy among untruncated state space policies is attained with a buffer size $K(\eta)$ and can be found with the policy iteration algorithm run over the state space $\mathcal{V}^{\leq K(\eta)}$, which returns $J^*(\eta)$. Furthermore, the least upper bound to the family of straight lines $D + \eta \Delta_e = J^*(\eta)$, $\eta > 0$ gives the boundary of the achievable (Δ_e, D) region.*

3.4.4 An Efficient Algorithm to Find the (Δ_e, D) Region

Although one can run the generic policy iteration algorithm to find the tangent lines to the achievable (Δ_e, D) region, this turns out to be highly inefficient. In this section, we provide an efficient modification of the policy iteration algorithm. The main idea is to exploit the following property. We omit its proof as it is a straightforward extension of Property 1(i).

Property 3.2. *Consider the policy iteration algorithm (Algorithm 1). For any \mathbf{b} , the policy $s^{(i)}(\mathbf{b})$ is either equal to $s^{(i)}(\mathbf{b}_{\geq 2}) + 1$, or to 1. Furthermore, if $s^{(i)}(\mathbf{b}) = s^{(i)}(\mathbf{b}_{\geq 2}) + 1$, then $h^{(i)}(\mathbf{b}) = b_1/\mu_Z + h^{(i)}(\mathbf{b}_{\geq 2})$.*

Property 3.2 implies that if $s^{(i)}(\mathbf{b}) \neq 1$, then there must exist a \mathbf{b}' , which is a suffix of \mathbf{b} , and with $s^{(i)}(\mathbf{b}') = 1$. Consequently,

$$h^{(i)}(\mathbf{b}) = h^{(i)}(\mathbf{b}') + \sum_{b \in \mathbf{b} \setminus \mathbf{b}'} b/\mu_Z. \quad (3.26)$$

The above observation leads to an improvement in the policy evaluation stage of the algorithm as we shall see shortly. Denote the set of states \mathbf{b}' with $s^{(i)}(\mathbf{b}') = 1$ as \mathcal{B}_1 . In light of (3.26), we see that since the relative values for the other states $\mathbf{b} \notin \mathcal{B}_1$ can be determined based on the states in \mathcal{B}_1 , it is sufficient for the linear system in the policy evaluation step to include the states in \mathcal{B}_1 . We observed empirically that $|\mathcal{B}_1|$ is much smaller compared $|\mathcal{V}^{\leq K}|$. As solving a linear system with n variables has $O(n^3)$ complexity, reducing the set of variables to \mathcal{B}_1 results in a significant improvement.

If the algorithm is modified as suggested above, the policy evaluation step gives $h^{(i)}(\mathbf{b}')$, $\mathbf{b}' \in \mathcal{B}_1$, and the average cost $\lambda^{(i)}$. To find other $h^{(i)}(\mathbf{b})$'s, we refer to equation (3.26), which suggests that the appropriate data structure to represent the states is a $|\mathcal{V}|$ -ary tree structure, denoted as \mathcal{T} , where a state \mathbf{b} has children $\{b \parallel \mathbf{b}\}_{b \in \mathcal{V}}$. That is, $\mathbf{b}_{\geq 2}$ is the parent of \mathbf{b} , denoted by $\text{parent}(\mathbf{b})$, and every suffix of \mathbf{b} is its ancestor. Then the final statement of Property 2 translates into the recursion $h^{(i)}(\mathbf{b}) = b_1/\mu_Z + h^{(i)}(\text{parent}(\mathbf{b}))$ for $\mathbf{b} \notin \mathcal{B}_1$.

Along with Property 3.2, Theorem 3.3 also provides simplifications for the search of an optimal policy. Namely, for a state \mathbf{b} , and for any iteration j , $s^{(j)}(\mathbf{b}) \neq s$ for an $s < l(\mathbf{b}) - K_{b_s}(\eta)$, where $K_{b_s}(\eta) = K_i(\eta)$ if $b_s = v_i$. Using the above facts, we are ready to provide a more efficient version of the policy iteration algorithm, Algorithm 2, which is tuned for our problem. The for loops over the tree \mathcal{T} (lines 9 and 12) are in breadth-first manner. Recall that $h^{(i)}(b) = 0$ for $b \in \mathcal{V}$ and we set $h^{(i)}(\delta) = 0$ where δ denotes the empty string.

Although Algorithm 2 is much more efficient compared to the generic policy iteration, one needs to evaluate $C_{h^{(i)}}(\mathbf{b}, s)$'s — defined in (3.21) — for both policy evaluation and update stages. Observe that calculating one of these quantities takes exactly $|\mathcal{V}^{\leq K}| = O(|\mathcal{V}^K|)$ iterations. We can allocate some memory to store these quantities and reduce the time complexity. First, let $\text{parent}_{\mathcal{B}_1}(\mathbf{b})$ be the longest ancestor of \mathbf{b} that is in \mathcal{B}_1 . Then, $\text{parent}_{\mathcal{B}_1}(\mathbf{b}) = \text{parent}_{\mathcal{B}_1}(\text{parent}(\mathbf{b}))$ and $\text{parent}_{\mathcal{B}_1}(\mathbf{b}') = \mathbf{b}'$ for $\mathbf{b}' \in \mathcal{B}_1$. Moreover, let the cost of an edge between \mathbf{b} and its parent be b_1/μ_Z and let $\text{cost}^{(i)}(\mathbf{b})$ denote the cost of going from \mathbf{b} to $\text{parent}_{\mathcal{B}_1}(\mathbf{b})$. Obviously, $\text{cost}^{(i)}(\mathbf{b}') = 0$ for $\mathbf{b}' \in \mathcal{B}_1 \cup \{\delta\}$.

Now, observe that the linear system of equations in (3.20) can be written as

$$\begin{aligned} h^{(i)}(\mathbf{b}') + \lambda &= \eta(l(\mathbf{b}') - 1) + E[h^{(i)}(\text{parent}(\mathbf{b}') \parallel \mathbf{V}^Z)] \\ &= \eta(l(\mathbf{b}') - 1) + E \left[\text{cost}^{(i)}(\text{parent}(\mathbf{b}') \parallel \mathbf{V}^Z) + h(\text{parent}_{\mathcal{B}_1}(\text{parent}(\mathbf{b}') \parallel \mathbf{V}^Z)) \right]. \end{aligned} \quad (3.27)$$

The policy iteration algorithm starts with $s^{(0)}(\mathbf{b}) = l(\mathbf{b})$. Therefore, $h^{(0)}(\mathbf{b}) = \text{cost}^{(0)}(\mathbf{b}) = \sum_{i < l(\mathbf{b})} b_i$ and one can also set $\text{parent}_{\mathcal{B}_1}(\mathbf{b}) = v_{\min}$. We will ini-

Algorithm 2: Efficient Policy Iteration v1(η)

Input: η
Output: $J^*(\eta)$

- 1 **Initialize**
- 2 $K \leftarrow K(\eta);$
- 3 $s^{(0)}(\mathbf{b}) \leftarrow l(\mathbf{b});$
- 4 $\mathcal{B}_1 \leftarrow \{v_{\min}\};$
- 5 Set δ as the root of \mathcal{T} and add the children $\{b\|\mathbf{b}\}_{b \in \mathcal{V}}$ for every
 buffer $\mathbf{b} \in \mathcal{T}$ such that $l(\mathbf{b}) < K;$
- 6 $i \leftarrow 0;$
- 7 **repeat**
- 8 /* Policy evaluation */
 Find $\lambda^{(i)}, h^{(i)}(\mathbf{b}'), \mathbf{b}' \in \mathcal{B}_1$ by solving (3.20);
- 9 **for** $\mathbf{b} \in \mathcal{T} \setminus \mathcal{B}_1$ *with* $l(\mathbf{b}) > 1$ **do**
- 10 $h^{(i)}(\mathbf{b}) \leftarrow h^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
- 11 /* Policy update */
 $\mathcal{B}_1 \leftarrow \{v_{\min}\};$
- 12 **for** $\mathbf{b} \in \mathcal{T}$ *with* $l(\mathbf{b}) > 1$ **do**
- 13 **if** $l(\mathbf{b}) < K_{b_1}(\eta)$ *and*
 $C_{h^{(i)}}(\mathbf{b}, 1) < C_{h^{(i)}}(\text{parent}(\mathbf{b}), s^{(i+1)}(\text{parent}(\mathbf{b}))) + b_1/\mu_Z$ **then**
- 14 $s^{(i+1)}(\mathbf{b}) \leftarrow 1;$
- 15 Add \mathbf{b} to $\mathcal{B}_1;$
- 16 **else**
- 17 $s^{(i+1)}(\mathbf{b}) \leftarrow s^{(i+1)}(\text{parent}(\mathbf{b})) + 1;$
- 18 $i \leftarrow i + 1.$
- 19 **until** $s^{(i+1)} = s^{(i)};$
- 20 **return** $\lambda^{(i)}$

tialize the procedure accordingly so that the policy evaluation step makes use of these quantities at the first iteration. We aim to update $\text{cost}^{(i)}(\mathbf{b})$ and $\text{parent}_{\mathcal{B}_1}(\mathbf{b})$ in the policy update step. To this end, we need the temporary variable

$$\text{temp}^{(i)}(\mathbf{b}) = \min_{s \leq l(\mathbf{b})} C_{h^{(i)}}(\mathbf{b}, s). \quad (3.28)$$

Then the condition for the policy update becomes

$$C_{h^{(i)}}(\mathbf{b}, 1) < \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z \quad (3.29)$$

and $\text{temp}^{(i)}(\mathbf{b})$ will be updated accordingly. Note that for $b \in \mathcal{V}$, $\text{temp}^{(i)}(b) = \lambda^{(i)}$. The modified version is given in Algorithm 3.

Recall that the number of elements in tree \mathcal{T} is $O(|\mathcal{V}|^K)$. The complexity of a single iteration in Algorithm 3 is found as follows:

Algorithm 3: Efficient Policy Iteration v2(η)

Input: η
Output: $J^*(\eta)$

- 1 **Initialize**
- 2 $K \leftarrow K(\eta);$
- 3 $s^{(0)}(\mathbf{b}) \leftarrow l(\mathbf{b});$
- 4 $\mathcal{B}_1 \leftarrow \{v_{\min}\};$
- 5 Set δ as the root of \mathcal{T} and add the children $\{b \parallel \mathbf{b}\}_{b \in \mathcal{V}}$ for every
 buffer $\mathbf{b} \in \mathcal{T}$ such that $l(\mathbf{b}) < K$;
- 6 For all \mathbf{b} , $\text{cost}(\mathbf{b}) \leftarrow \sum_{i < l(\mathbf{b})} b_i$ and $\text{parent}_{\mathcal{B}_1}^{(0)}(\mathbf{b}) \leftarrow v_{\min};$
- 7 $i \leftarrow 0;$
- 8 **repeat**
- 9 /* Policy evaluation */
- 10 Find $\lambda^{(i)}, h^{(i)}(\mathbf{b}'), \mathbf{b}' \in \mathcal{B}_1$ by solving (3.27);
- 11 **for** $\mathbf{b} \in \mathcal{T} \setminus \mathcal{B}_1$ **with** $l(\mathbf{b}) > 1$ **do**
 $h^{(i)}(\mathbf{b}) \leftarrow h^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z$ **if** $\mathbf{b} \notin \mathcal{B}_1$;
- 12 /* Policy update */
- 13 $\mathcal{B}_1 \leftarrow \{v_{\min}\};$
- 14 **for** $b \in \{v_{\min}, \dots, v_{\max}\}$ **do**
 $\text{temp}^{(i)}(b) \leftarrow \lambda^{(i)}.$
- 15 **for** $\mathbf{b} \in \mathcal{T}$ **with** $l(\mathbf{b}) > 1$ **do**
 if $l(\mathbf{b}) < K_{b_1}(\eta)$ **and** $C_{h^{(i)}}(\mathbf{b}, 1) < \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z$
 then
 $s^{(i+1)}(\mathbf{b}) \leftarrow 1;$
 Add \mathbf{b} to \mathcal{B}_1 ;
 $\text{temp}^{(i)}(\mathbf{b}) \leftarrow C_{h^{(i)}}(\mathbf{b}, 1);$
 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow 0;$
 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \mathbf{b};$
 else
 $s^{(i+1)}(\mathbf{b}) = s^{(i+1)}(\text{parent}(\mathbf{b})) + 1;$
 $\text{temp}^{(i)}(\mathbf{b}^{(i)}) \leftarrow \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow \text{cost}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \text{parent}_{\mathcal{B}_1}^{(i)}(\text{parent}(\mathbf{b}));$
- 16 **if** $l(\mathbf{b}) < K_{b_1}(\eta)$ **and** $C_{h^{(i)}}(\mathbf{b}, 1) < \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z$
 then
 $s^{(i+1)}(\mathbf{b}) \leftarrow 1;$
 Add \mathbf{b} to \mathcal{B}_1 ;
 $\text{temp}^{(i)}(\mathbf{b}) \leftarrow C_{h^{(i)}}(\mathbf{b}, 1);$
 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow 0;$
 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \mathbf{b};$
 else
 $s^{(i+1)}(\mathbf{b}) = s^{(i+1)}(\text{parent}(\mathbf{b})) + 1;$
 $\text{temp}^{(i)}(\mathbf{b}^{(i)}) \leftarrow \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow \text{cost}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \text{parent}_{\mathcal{B}_1}^{(i)}(\text{parent}(\mathbf{b}));$
- 17 $s^{(i+1)}(\mathbf{b}) \leftarrow 1;$
- 18 Add \mathbf{b} to \mathcal{B}_1 ;
- 19 $\text{temp}^{(i)}(\mathbf{b}) \leftarrow C_{h^{(i)}}(\mathbf{b}, 1);$
- 20 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow 0;$
- 21 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \mathbf{b};$
- 22 **else**
- 23 $s^{(i+1)}(\mathbf{b}) = s^{(i+1)}(\text{parent}(\mathbf{b})) + 1;$
- 24 $\text{temp}^{(i)}(\mathbf{b}^{(i)}) \leftarrow \text{temp}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
- 25 $\text{cost}^{(i)}(\mathbf{b}) \leftarrow \text{cost}^{(i)}(\text{parent}(\mathbf{b})) + b_1/\mu_Z;$
- 26 $\text{parent}_{\mathcal{B}_1}^{(i)}(\mathbf{b}) \leftarrow \text{parent}_{\mathcal{B}_1}^{(i)}(\text{parent}(\mathbf{b}));$
- 27 $i \leftarrow i + 1.$
- 28 **until** $\mathbf{s}^{(i+1)} = \mathbf{s}^{(i)};$
- 29 **return** $\lambda^{(i)}$

- (i) In the policy evaluation, the equation system (3.27) is constructed in $O(|\mathcal{B}_1||\mathcal{V}|^K)$ time, and solved in $O(|\mathcal{B}_1|^3)$.
- (ii) The policy update runs over all states. Furthermore, for each state \mathbf{b} , calculation of $C_{h^{(i)}}(\mathbf{b}, 1)$ also requires iterations over all states. Hence, it has $O(|\mathcal{V}|^{2K})$ time complexity.

As we stated before, $|\mathcal{B}_1|$ is usually very small compared to $|\mathcal{V}|^K$. The bottleneck then seems to be the policy update stage, which requires $O(|\mathcal{V}|^{2K})$ steps. However, the policy update stage can be further improved such that the complexity decreases to $O(K|\mathcal{V}|^K)$. This modification is given in Appendix 3.8.5.

Now that we have an efficient algorithm yielding $J^*(\eta)$, and with help of Corollary 1, we should be able to find the boundary curve by varying η . One may notice that initializing the trees and policies for every η can be avoided with a minor modification. The idea is as follows: Choose a decreasing sequence $\eta_1 > \eta_2 > \dots > \eta_n$ with $\eta_1 = \eta_{\max} = \frac{1}{\mu_Z}(v_{\max} - v_{\min})$. Note that $K(\eta_1) = 1$. Run the algorithm in the order of η_m 's and when $K(\eta_{m+1}) > K(\eta_m)$, append new states to the tree \mathcal{T} . At $(m+1)^{\text{th}}$ run, one can also start with the optimal policy found for η_m , i.e., a warm start. If η_1, \dots, η_n are chosen densely, the boundary curve can be well-approximated.

However, finding the boundary region everywhere would be optimistic. Theorem 3.3 implies that the necessary buffer size scales with $\frac{1}{\eta}$. This suggests that even though the algorithm gives the almost exact curve, it is impractical to do so. To overcome this difficulty, one may rely on approximate dynamic programming algorithms; or resort to Monte Carlo estimations for the policy evaluation [81].

Note that any straight line $D + \eta\Delta_e = J^*(\eta)$ in the (Δ_e, D) plane is a lower bound to the feasible region. Hence, any family of straight lines obtained in such manner gives a lower bound in general — and if we were able to run the algorithm for every $\eta > 0$, this would give the exact boundary curve. As before, let $\eta_1 > \eta_2 > \dots > \eta_n$ be a densely chosen sequence for which the algorithm is run. Let $\Delta_e^{(n-1,n)}$ be the abscissa of the point where the last two lines $D + \eta_{n-1}\Delta_e = J^*(\eta_{n-1})$ and $D + \eta_n\Delta_e = J^*(\eta_n)$ intersect. Then, one can see that the supremum of the straight lines obtained for η_1, \dots, η_n approximately gives the tradeoff curve for $\Delta_e \leq \Delta_e^{(n-1,n)}$, while it gives a lower bound for $\Delta_e > \Delta_e^{(n-1,n)}$. This is because all intersection points that lie on the supremum, and the line segments connecting them are achievable. This straight-line converse bound is referred as ‘PI converse’ in the numerical examples, which we will provide shortly.

We end this section with some numerical examples. We have calculated the optimal policies with Algorithm 3, and unfortunately we have not observed any simple structure for optimal policies for $|\mathcal{V}| = 2$. We also evaluated some simple policies described below and compared their performances with the

family of straight lines generated by Algorithm 3, referred as ‘PI’ in Figures 34 and 35. These simple policies are:

- (S1) Send the oldest important data within a maximum buffer size K .
- (S2) Send the newest important data within a maximum buffer size K .
- (S3) Send the newest important data that has arrived more than K slots ago. If there is no such data, send the oldest important one. Note that under such strategies, the instantaneous age is kept above K most of the time.

Each strategy above induces a finite-state Markov chain. Moreover, when Z is geometrically distributed, all the Markov chains induced by these strategies have closed-form stationary distributions. Δ_e and D pertaining to these strategies will accordingly have closed-form expressions. We provide these expressions in Appendix 3.8.6.

To compare these strategies, we also give a simple converse bound and observe their approach towards this bound for large Δ_e .

Lemma 3.5. *Suppose $V = v_i$ with probability α_i . Let j^* be the maximum index such that $\sum_{i=j^*}^{|\mathcal{V}|} \alpha_i \geq \frac{1}{\mu_Z}$. Then for any $\Delta_e^{(S)}$, $D^{(S)} \geq D_{\min} = \sum_{i=1}^{j^*-1} \alpha_i v_i + \left(\sum_{i=j^*}^{|\mathcal{V}|} \alpha_i - \frac{1}{\mu_Z}\right)v_{j^*}$.*

Proof Sketch. The sender can send at most $\frac{1}{\mu_Z}$ fraction of the data. We then optimize over its selection of data in a greedy manner to obtain the result. \square

In the first numerical example, provided in Figure 34, $\Pr(Z = 1) = 0.2$ and $\mathcal{V} = \{1, 20\}$ with $\Pr(V = 1) = 0.7$. The blue family of straight lines correspond to the lines $D + \eta\Delta_e = J^*(\eta)$, obtained for different η values. We could calculate until $\eta = \frac{v_{\max} - v_{\min}}{17\mu_Z}$, which indicates that we used a maximum buffer size of 17. The region lying under this family of lines is unachievable, and the supremum of this family gives the boundary of the feasible region until the red solid line, which is the straight-line converse bound described above. The curves corresponding to strategies S1, S2 and S3 are red dashed, green dashed and cyan dotted curves, and plotted for $K \leq 20$. The simple lower bound $D_{\min} = 2.7$ is drawn as the solid black line. One can see that S2 nearly coincides with PI. Note that we observe an asymptotic behavior as the sender will never be able to allocate all of its resources to send all of the important packets.

The second numerical example differs from the first one with $\Pr(Z = 1) = 0.3$ and $\Pr(V = 1) = 0.8$. The simple strategies calculated for $K \leq 40$ together with the policy iteration results obtained until a buffer size of 17 are plotted in Figure 35. Here, $D_{\min} = 0.7$ could be achieved in finite-age as the sender can send all the important packets while keeping $\sup_i E[(S_i - T_i)^2] < \infty$.




Figure 34 – Comparison of the strategies for $\mathcal{V} = \{1, 20\}$ and $\Pr(V = 1) = 0.7$. Z is taken as a Geometric random variable with success probability 0.2. (S2 almost coincides with PI)




Figure 35 – Comparison of the strategies for $\mathcal{V} = \{1, 20\}$ and $\Pr(V = 1) = 0.8$. Z is taken as a Geometric random variable with success probability 0.3.

3.5 Relation to an Erasure Channel with Feedback

We have observed that the model discussed primarily in this work is similar to transmitting a stream of packets over an erasure channel with feedback. Recall that the setting for transmission over discrete memoryless channels with feedback requires the feedback for the data transmitted at time t to be revealed

just after time t . For an erasure channel, the knowledge of an erasure event indicator at time t , i.e., $\mathbb{1}\{X_t \text{ is erased}\}$ is sufficient for a perfect feedback.

We note that the main difference between our model and a feedback erasure channel stems from the restriction that the feedback for data t is revealed *just after* time t . If we assume that the sender knows about a possible erasure event *just before* time t ; it may send the data t , or it may keep data t in its buffer for a later transmission. This relaxation exactly gives the model we described with interspeaking times distributed according to a geometric distribution, i.e., $\Pr(Z_i = z) = (1 - p)p^{z-1}$ where p is the erasure probability for a discrete memoryless erasure channel. Also note that one can model some erasure channels with memory by varying the distribution of Z .

Consider a modification of our model which requires that the constituent feedback is revealed just after the transmission. Now, we formalize the modified setting. Just before time t , the sender commits to a packet with timestamp $C_t \leq t$. Then, at time t , the committed packet X_{C_t} is transmitted through the erasure channel. If the packet is erased, the sender commits to packet with timestamp C_{t+1} just before time $t + 1$ — which is not necessarily equal to C_t — and the procedure is repeated until the committed data is sent. If the committed data X_{C_t} is sent successfully, then $S_{i(t)+1} = C_t$ where $i(t) = \sup\{i \geq 0 : T_i < t\}$; and the past of $S_{i(t)+1}$ contributes to the distortion and cannot be modified later. The age and distortion metrics are defined similar to the ones in Section 3.3.

Observe that for any sequence of commitments $\{C_t\}_{t>0} =: \mathbf{C}$, there must exist a selection procedure \mathbf{S} in the original problem such that

$$D^{(\mathbf{C})} + \eta\Delta^{(\mathbf{C})} \geq D^{(\mathbf{S})} + \eta\Delta^{(\mathbf{S})}. \quad (3.30)$$

This is because in the original problem, selections are made with more information — the erasure event is known beforehand. Conversely, for any *stationary* policy \mathbf{S} in the original problem, there exists a \mathbf{C} in the latter problem where the sender commits to $C_t = t + s(\mathbf{B}_t) - l(\mathbf{B}_t)$ where $\mathbf{B}_t := V_{S_{i(t)+1}}^t$. As a consequence, the sender with no erasure information beforehand can also attain the infimal value $J^*(\eta)$, as we know from Corollary 3.1 that the optimal policy for the original problem is stationary. Together with the inequality (3.30), we conclude that for both problems the tradeoff curve is the same. In brief, whether erasures are revealed just before or after transmissions do not change the tradeoff between the age and distortion.

3.6 When Timestamps Become Significant

We have shown that the optimal value for an $\eta > 0$ is attained with a bounded buffer policy, say of K . In the model so far, packets contain the timestamps as part of their headers. Consequently, there was no need to send additional information to the receiver to tell it which packet among the K packets in the buffer is being sent. If the packets do not have headers, this additional

information must be included. If K is much smaller compared to $|\mathcal{X}|$, this additional information is insignificant. In this section, we treat the case of headerless packets when K is comparable to $|\mathcal{X}|$. We take binary $\mathcal{X} = \{0, 1\}$ and $v(0) = 1$, $v(1) = v$. We study the setting described in Section 3.3 but with the difference that the sender is allowed to send N bits at each speaking time. We assume that Z is distributed geometrically with success probability p , i.e. $\Pr(Z = 1) = p$.

Consider the optimal policy to attain $J^*(\eta)$ in (3.13), which is of bounded buffer size $K(\eta)$. Recall that at each speaking time, the sender is able to send one packet. If X is binary, then without any coding, the optimal policy is feasible only if $N \geq 1 + \lceil \log(K(\eta)) \rceil$; otherwise it is not able to describe the timestamps of selected data, e.g., for a state \mathbf{b} with $l(\mathbf{b}) = K(\eta)$, the timestamps must be of length $\lceil \log(K(\eta)) \rceil$ and the remaining one bit corresponds to the data. However, it seems unreasonable to use almost all of the N bits for timestamp description. Can one come up with methods that do not require explicit timing information to be sent and allocate more bits to describe the data itself? The rest of this section elaborates on some possible tradeoffs aligned with this perspective.

3.6.1 Buffer-Ignorant Strategies

Think of the following strategy: Always send the most recent $N = \log(K(\eta))$ bits. In this case, it is easy to see $\Delta_e = 0$ and

$$D = \mu_V(1-p)^N = \mu_V K(\eta)^{\log(1-p)} \simeq \mu_V \left(\frac{v-1}{\eta} \right)^{\log(1-p)}, \quad (3.31)$$

where $\mu_V := E[V] = (1-q) + vq$, and $q = \Pr(V = v)$. We know that for the optimal policies described in Section 3.4, D_{\min} given in Lemma 3.5 yields a lower bound for distortion. For a binary X , and thus V , one can obtain

$$D_{\min} = \begin{cases} 1-p, & p \geq q \\ \mu_V - pv & p < q \end{cases}. \quad (3.32)$$

If $\eta \geq \frac{1}{v-1} \left(\frac{\mu_V}{D_{\min}} \right)^{\frac{1}{\log(1-p)}}$ then $D \leq D_{\min}$, implying that the perfect timing information strategies are beaten by the timing ignorant strategy described as sending the most recent N bits. In other words, one does better by sending N bits of most recent data instead of sending one bit together with its timestamp.

The above arguments motivate the following question: What are the limits of these timing ignorant strategies? Note that both the sender and receiver know the speaking times T_i . Suppose for a moment that the receiver knows the selection times S_i as well. With this assumption, the receiver has the perfect knowledge of the buffer length at time T_i , which is $T_i - S_{i-1}$. Hence, if the sender bases its strategies solely on its buffer size, ignoring the buffer content, it does not have to include any timing information and is able to use all its

N -bit budget for sending data. An example could be as follows: Suppose $N = 3$. Then the sender could send 1, 3, 5th bits whenever the buffer size is 5. Since the receiver knows the buffer size and the sender's strategy, it will know upon its reception of 3 bits that they correspond to 1, 3, 5th bits. Although we have previously coined the term 'timing-ignorant' for such strategies, a more suitable term could be 'buffer-ignorant'; as the sender ignores what is in its buffer.

Remark 3.3. *One may notice that this procedure is a simple online compression algorithm for the binary source $\{X_i\}$, with the restriction that the sender can only send N bits at each speaking time. Together with the unsophisticated receiver that only constructs the N bits it receives at each speaking time, this scheme operates at $\Delta_e = 0$ and $D = \mu_V(1 - p)^N$.*

Let us study the 'buffer-ignorant' strategies further. Since the transmitter does not use the buffer content and has to choose N bits among them, a simple choice could be to send size- N bit contiguous chunks. Adopting the terminology from Section 3.4, such strategies correspond to sending $X_{S_i-N+1}^{S_i}$ at time T_i from a buffer of size $L_i := T_i - S_{i-1}$. Nothing from the past X^{S_i} can be sent after time T_i and therefore with similar arguments we have previously done, one can formulate the problem of finding optimal selection times as a MDP. The corresponding MDP will have the buffer lengths as its states, which implies that we encounter another countable state-space problem with its state-space being \mathbb{Z}^+ . When the buffer length is l , and the selection index is s , the one-step cost can be written as

$$g(l, s) = \mu_V p(s - N)^+ + \eta(l - s). \quad (3.33)$$

The corresponding Bellman equation is given by

$$h(l) + \lambda = \min_{s \leq l} \left\{ \mu_V p(s - N)^+ + \eta(l - s) + E[h(l - s + Z)] \right\} \quad (3.34)$$

for $l > 1$ and with $h(1) = 0$, where the buffer of length one is chosen as the reference state. It is not difficult to see that this choice also implies $h(l) = 0$ for $l \leq N$, as the sender can immediately empty the buffer for such states.

Similar to Property 3.2, we must have either $s^*(l) = N$ or $s^*(l) = s^*(l - 1) + 1$ for the optimal policy. Therefore as l increases, the optimal policy tends to leave more bits at the end. Although this observation suggests that the optimal policy may be attained with an unbounded buffer, one can show that this is not the case. Similar coupling arguments as we did in the proof of Theorem 3.3 lead to the conclusion that the optimal policy cannot leave more than $\frac{N\mu_V}{\eta}$ bits at the end. Thus, a simple policy iteration algorithm run for a sufficiently large state-space also solves the Bellman equation for the infinite-state problem.

The optimal policies may not be simple-to-describe. However, when Z is geometrically distributed, the numerical simulations indicate that single-threshold policies are optimal. These policies are characterized as

$$s(l) = \min\{\max\{l - \tau, N\}, l\} \quad (3.35)$$

for some $\tau \geq 0$. In other words, the sender always keeps τ unsent bits in the buffer if possible. We do not have an analytical proof for this result, but it is not unreasonable to believe that these simple policies are optimal because of the memorylessness property of Z . Recall that the states of Markov chains incurred by such strategies are described with the buffer length l . Denote the stationary probabilities as π_l . In this special case, the stationary probabilities (and consequently the average age and distortion) incurred by such strategies have closed-form expressions.

Corollary 3.2. *The (Δ_e, D) curve attained by single-threshold strategies has a parametric description that is available in closed-form. Let $\bar{p} := 1 - p$, $S_j^{(0)} := (1 + jp)$ and $S_j^{(n)} := \sum_{k=0}^j S_k^{(n-1)}$ for $n \geq 1$. Also let $S_j^{(n)} = 0$ for $j < 0$. For $0 \leq j \leq \tau - 1$, the stationary probabilities are given by*

$$\pi_{\tau-j} = \pi_{\tau+1} \frac{1 + \sum_{k=0}^{\lceil \tau/N \rceil} (-1)^{k+1} S_{j-kN}^{(k)} p^k \bar{p}^{(k+1)(N-1)}}{\bar{p}^{j+1}} \quad (3.36)$$

with

$$\pi_{\tau+1} = \left[\sum_{j=0}^{\tau-1} \frac{1 + \sum_{k=0}^{\lceil \tau/N \rceil} (-1)^{k+1} S_{j-kN}^{(k)} p^k \bar{p}^{(k+1)(N-1)}}{\bar{p}^{j+1}} + \frac{1}{p} \right]^{-1}. \quad (3.37)$$

and $\pi_{\tau+1+j} = \bar{p}^j \pi_{\tau+1}$ for $j > 0$. The (Δ_e, D) curve therefore has a parametric description $(\Delta_e(\tau), D(\tau))$ given by

$$\Delta_e(\tau) = \sum_{j=1}^{\tau-1} j \pi_{N+j} + \frac{\tau \pi_{\tau+1} \bar{p}^{N-1}}{p}. \quad (3.38)$$

and

$$D(\tau) = \frac{\mu_V \pi_{\tau+1} \bar{p}^N}{p^2}. \quad (3.39)$$

Remark 3.4. *If the bits sent are equal in terms of their importance, i.e., $v_1 = v_2$, then buffer-ignorant strategies are optimal among the strategies without timestamp coding. This is because strategies as such already assume that 1 and 0 are of equal importance and there is no need to indicate which bit is more important. Consequently, for the binary erasure channel, i.e., $N = 1$; the only optimal buffer-ignorant strategy is to send the last bit if the bits are equally important.*

Until now, we have only considered integer N . However, the erasure channel between the sender and receiver could admit K inputs, where K is not necessarily a power of 2. This would imply a non-integer $N = \log K$. Studying such N requires strategies with some knowledge of the buffer content and possibly requires some coding. We will elaborate on how to handle these cases in the next section.

3.6.2 Revealing Partial Buffer Content

The previous section was devoted to buffer-ignorant strategies. Now, we allow some knowledge of the buffer content to improve buffer-ignorant strategies and also to cover the case of a non-integer N .

Let us first show the improvement by coding over the buffer-ignorant strategies for an integer N . Take the single-threshold strategy $s(l)$ described in (3.35) together with the threshold τ . Consider the state $l > \tau + N$, where the dictated action is to keep τ bits for future and send N of the remaining bits. Then $l - \tau - N$ bits will never be sent and hence the distortion penalty will be $\mu_V(l - \tau - N)$. We aim to show that the distortion penalty can be decreased with some coding while preserving or decreasing the age penalty.

To that end, consider an alternative way of describing the buffer-ignorant strategy above: send the first N bits of the sequence $x_{l-\tau}, x_{l-\tau-1}, \dots, x_1$ regardless of the content. It is therefore reasonable to think that a parser with a dictionary of size 2^N which sends the identity of the first parsed word in the same sequence could result in an improvement. One could resort to some variable-to-fixed length source coding techniques, such as Tunstall coding. As we have already stated in **Chapter 1**, Tunstall coding is known to maximize the expected number of bits parsed among prefix-free and variable-to-fixed length dictionaries. Let $E[L_{\text{Tun}}]$ be the expected number of bits parsed. With Tunstall coding, the expected number of unsent bits will be $l - \tau - E[L_{\text{Tun}}]$ and since $E[L_{\text{Tun}}] \geq N$, i.e., the expected number of compressed bits is greater than N , the distortion cost decreases. Thus, Tunstall algorithm improves the buffer-ignorant strategies. Also note that for non-integer N , Tunstall algorithm can be used to determine parsing methods.

We end this section by presenting some numerical results that illustrate the improvement with buffer-ignorant strategies and Tunstall coding. We use the same source and interspeaking time distribution as in Figure 36, i.e., $\mathcal{V} = \{1, 20\}$ with $\Pr(V = 1) = 0.7$ and $\Pr(Z = 1) = 0.2$. Recall that we were able to find the optimal policies with Algorithm 3 up to a buffer size of 17. This suggests that the sender describes the timing information with $\lceil \log 17 \rceil = 5$ bits and with an additional bit to describe the content, which implies that the left end of the PI converse line segment can be attained with sending $N = 6$ bits — note that this converse bound is valid only for strategies of Section 3.4, and not for buffer-ignorant strategies. As we suggested, the timing information can be sacrificed to allocate the whole budget for the data description in an attempt to improve the performance. This improvement is evident even for $N = 3$, where




Figure 36 – For the same setting in Figure 34, the curves pertaining to optimal buffer-ignorant strategies (BI, colored in red) together with their improved versions with Tunstall coding (BIT, colored in cyan) are plotted for the same source in Figure 34. The curves corresponding to $N = 3$ are dashed, and the curves corresponding to $N = 6$ are dotted.

the optimal buffer-ignorant strategy (BI, $N = 3$) and its improved version with Tunstall coding (BIT, $N = 3$) perform better than the optimal PI curve as seen in Figure 36. For $N = 6$, there is drastical improvement and one can approach zero distortion with finite age — see BI and BIT, $N = 6$ in Figure 36.

3.6.3 About Other Possible Coding Strategies

The buffer-ignorant strategies discussed in the previous section only depend on the buffer size and we have shown that these strategies can be improved by revealing some buffer content. One can argue that these coding strategies might be far from optimal, as they use partial knowledge. The sender could base its strategies on all past speaking times and the past buffer content. Calculating the tightest (Δ_e, D) curve corresponding to this broad class of strategies seems to be formidably complex. Even with the sole knowledge of the current buffer content, the problem becomes difficult. We illustrate this with an example. Suppose the current buffer content is $\mathbf{b} = [v, v, 1, 1, v, v, v, 1, 1, 1, v, v, v, 1]$ and the sender could only send $N = 3$ bits. The single-threshold strategy with $\tau = 4$ will choose the index $s(\mathbf{b}) = 10$ and send $(1, 1, 1)$, which does not contain any important data. If the index $s = 13$ is chosen, (v, v, v) will be sent and this might be a better strategy; however, the description of $s = 13$ must be somehow included in the $N = 3$ bits and all three important data might not be sent. Encoding both the data and their indices into a fixed number of bits

complicates the possible actions and thus the problem appears to be very hard. We also observe an interesting tradeoff in this situation. In our MDP, sacrificing perfect state information results in a smaller policy space with policies of lower penalties, i.e., if k bits are allocated for timestamp description, $N - k$ information bits can be sent and when k decreases, the number of possible actions also decreases but their one step costs can possibly be smaller.

3.7 Discussion

In this chapter, we have studied a discrete-time model where the sender is only allowed to speak at time slots assigned by an external scheduler. In the absence of a distortion measure, it is clear that the optimal strategy is to send the freshest packet in the buffer at each speaking time as this will minimize the age. However, if this freshest packet has low importance, it may be beneficial to send a packet of higher importance instead, sacrificing freshness for lowering distortion. Hence, it is immediate that a tradeoff between the age and the distortion exists. It turns out that the optimal tradeoff can be attained with bounded buffer policies, and these policies can be found with numerical methods. Unfortunately, they turn out to be not simple-to-describe.

We observed that the usual policy iteration methods were inefficient for our specific problem, and we devised an algorithm based on appropriate data structures and problem-specific simplifications. The new algorithm performs significantly better — the time complexity decreases from $O(|\mathcal{V}|^{3K})$ to $O(K|\mathcal{V}|^K)$, where K is the exact buffer size needed for attaining the optimal tradeoff. However, at the high-age regime, K is large and in turn, the optimal tradeoff cannot be computed. One could try to find simple-to-describe policies that are not too far from the optimal tradeoff at this regime.

The main results of this chapter also apply when the process of importance levels $\{V_i\}_{i \geq 0}$ is an ergodic Markov chain — one can verify the conditions (i) and (ii) with the same state and action spaces defined in Section 3.4.1. Consequently, the necessary buffer size will be the same and one is able to find the optimal curve as in the i.i.d. case.

The problem we formulated in the first place turns out to be closely related to a problem of transmitting packets over an erasure channel with perfect feedback. The difference is that in our setting, erasure events are revealed before transmissions. Nevertheless, we have shown that for both problems, the optimal tradeoff is the same.

Until Section 3.6, we assumed that the timing information is contained in the header of a packet, which is much smaller in size compared to the payload. In Section 3.6, we studied the case where the packets need not contain perfect timing information, and consist of at most N bits. As a consequence, if one decides to sacrifice age for lowering distortion, some additional information must be included in the N bits in order to tell the receiver to which time the information bits pertain. Therefore, if the sender decides not to send the

freshest data, it not only sacrifices age but may also decrease the amount of information bits to be sent. We have studied some simple policies, called buffer-ignorant, where the sender ignores what is in the buffer and allocates all its N -bit budget for the information bits. When the timing information dominates the payload, buffer-ignorant policies improve drastically over the optimal policies found in Section 3.4, which include the timing information. Later, we have shown that buffer-ignorant policies can be further improved by revealing some buffer content and using variable-to-fixed length coding. However, it seems very challenging to find the optimal tradeoff when any strategy that sends N bits at a time is allowed.

The reader may ask what if the sender is allowed to choose the speaking times as well. As expected, this relaxation will result in a larger admissible policy space. Moreover, an additional parameter should be introduced, the packet rate R . Consequently, in the relaxed problem, one has to study the (Δ, D, R) tradeoff. This problem might be too complex to solve. However, if the packets are of equal importance, the D component will be constant and we will be left with the (Δ, R) tradeoff. The study of this tradeoff is in fact the topic of the next chapter.

3.8 Appendix

3.8.1 Proofs of Theorems 3.1 and 3.2

Define $W_j := T_j - S_j$ for the rest of the proof and denote ‘almost surely’ by *a.s.*.

Proof of Theorem 3.1

We first prove a convergence result in Lemma 3.6 below, from which the equation (3.8) follows as a corollary.

Lemma 3.6. $\frac{1}{i} \sum_{j=1}^i W_j (Z_{j+1} - \mu_Z) \rightarrow 0$ *a.s.* if $\sup_j E[W_j^2] < \infty$.

Proof. We use the result that if $\sum_i b_i/i$ converges, then $\frac{1}{i} \sum_{j \leq i} b_j \rightarrow 0$ — known as Kronecker’s Lemma [82]. Therefore, it is sufficient to show $\sum_i W_i (Z_{i+1} - \mu_Z)/i$ converges *a.s.* We now show that $M_n := \sum_{i=2}^n W_{i-1} (Z_i - \mu_Z)/(i-1)$, $M_1 := 0$ is a martingale with respect to the filtration $\mathcal{F}_n := \sigma(Z_1, \dots, Z_n, \mathbf{X}_1^{T_n})$.

Observe that $E[M_n | \mathcal{F}_{n-1}] = M_{n-1} + E[W_{n-1} (Z_n - \mu_Z) | \mathcal{F}_{n-1}]/(n-1) = M_{n-1}$ as W_{n-1} is \mathcal{F}_{n-1} -measurable and Z_n is independent of \mathcal{F}_{n-1} with $E[Z_n] = \mu_Z$. Since M_n consists of uncorrelated increments, one can write

$$E[M_n^2] = \sum_{i=2}^n \frac{E[W_{i-1}^2] \text{Var}(Z)}{(i-1)^2}. \quad (3.40)$$

Note that we assumed $E[Z^2] < \infty$, hence $\text{Var}(Z) < \infty$. Moreover, since $\sup_j E[W_j^2] < \infty$, $\sum_i \frac{E[W_i^2]}{i^2} < \infty$. As a consequence, $\sup_n E[M_n^2] < \infty$ and the *a.s.* convergence of M_n follows from the martingale convergence theorem. \square

Now we prove (3.9) and complete the proof of Theorem 3.1. Define

$$Y_j = \begin{cases} X_j, & j \in \mathcal{S} \\ ?, & \text{else.} \end{cases} \quad (3.41)$$

Given t , define $i = i(t) := \sup\{j \geq 0 : T_j \leq t\}$. Observe that for $j \leq i$, $Y_j(t) = Y_j$. Thus,

$$D_t^{(\mathcal{S})} = \frac{1}{t} \sum_{j \leq S_i} d(X_j, Y_j) + \frac{1}{t} \sum_{j=S_i+1}^t d(X_j, Y_j(t)). \quad (3.42)$$

Upper bound $D_t^{(\mathcal{S})}$ as

$$\begin{aligned} D_t^{(\mathcal{S})} &\leq \frac{1}{t} \sum_{j \leq S_i} d(X_j, Y_j) + \frac{1}{t} (T_{i+1} - S_i) v_{\max} \\ &\leq \frac{1}{T_i} \sum_{j \leq S_i} d(X_j, Y_j) + \frac{1}{t} (Z_{i+1} + W_i) v_{\max}. \end{aligned} \quad (3.43)$$

Since $\sup_j E[W_j^2] < \infty$, W_i is *a.s.* finite for all i and hence $Z_{i+1} + W_i$ is *a.s.* finite. Thus $\frac{1}{t} (Z_{i+1} + W_i) v_{\max} \rightarrow 0$ *a.s.* Then, we obtain

$$\limsup_{t \rightarrow \infty} D_t^{(\mathcal{S})} \leq \limsup_{i \rightarrow \infty} \frac{1}{T_i} \sum_{j \leq S_i} d(X_j, Y_j). \quad (3.44)$$

Now, we lower bound $D_t^{(\mathcal{S})}$ as

$$D_t^{(\mathcal{S})} \geq \frac{1}{t} \sum_{j \leq S_i} d(X_j, Y_j) \geq \frac{1}{T_{i+1}} \sum_{j \leq S_i} d(X_j, Y_j) = \frac{1}{T_i + Z_{i+1}} \sum_{j \leq S_i} d(X_j, Y_j) \quad (3.45)$$

and take \limsup on both sides to obtain

$$\limsup_{t \rightarrow \infty} D_t^{(\mathcal{S})} \geq \limsup_{i \rightarrow \infty} \frac{1}{T_i + Z_{i+1}} \sum_{j \leq S_i} d(X_j, Y_j). \quad (3.46)$$

Finally, observe that $\frac{T_i}{i} \rightarrow \mu_Z$ and $\frac{Z_{i+1}}{i} \rightarrow 0$ *a.s.* Hence,

$$\limsup_{t \rightarrow \infty} D_t^{(\mathcal{S})} = \frac{1}{\mu_Z} \limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j \leq S_i} d(X_j, Y_j) = \frac{1}{\mu_Z} \limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_j, S_{j-1}), \quad (3.47)$$

which ends the proof of Theorem 3.1.

Proof of Theorem 3.2

Since $\sup_j E[W_j^2] < \infty$, it follows that $\sup_i E[(\frac{1}{i} \sum_{j=1}^i W_j)^2] < \infty$. Thus, the family $(\frac{1}{i} \sum_{j=1}^i W_j)_{i \in \mathbb{N}}$ is uniformly integrable [82, Chapter 13]. One can then use the reverse Fatou's lemma for uniformly integrable families [83] to obtain

$$\Delta_e^{(\mathbf{S})} = E \left[\limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i W_j \right] \geq \limsup_{i \rightarrow \infty} E \left[\frac{1}{i} \sum_{j=1}^i W_j \right]. \quad (3.48)$$

A similar reasoning for $D^{(\mathbf{S})}$ follows from the fact that each $D(\mathbf{V}^{T_j}, S_j, S_{j-1})$ is smaller than $v_{\max}(W_{j-1} + Z_j)$. Therefore, $\{D(\mathbf{V}^{T_i}, S_i, S_{i-1})\}_{i \geq 0}$ is a uniformly integrable family, and one proceeds in a similar way as above to obtain

$$D^{(\mathbf{S})} \geq \limsup_{i \rightarrow \infty} \frac{1}{\mu_Z} E \left[\frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_j, S_{j-1}) \right]. \quad (3.49)$$

As $\limsup_n a_n + \limsup_n b_n \geq \limsup_n (a_n + b_n)$, the inequality $J(\eta)^{(\mathbf{S})} \leq D^{(\mathbf{S})} + \eta \Delta_e^{(\mathbf{S})}$ holds. This completes the proof.

3.8.2 Proof of Lemma 3.1

We construct $\{\tilde{S}_i\}$ iteratively as follows:

$$\tilde{S}_i = \begin{cases} S_i, & i < i_0 \\ \min\{s > S_{i_0} : V_s > v_{\min}\} \wedge T_{i_0}, & i = i_0 \\ \tilde{S}_{i-1} + 1, & S_i \leq \tilde{S}_{i-1}, i > i_0 \\ S_i, & S_i > \tilde{S}_{i-1}, i > i_0 \end{cases}. \quad (3.50)$$

Verbally, at time i_0 , $\tilde{\mathbf{S}}$ takes the next packet whose importance value is more than v_{\min} (if no such packet, selects the freshest one) and then selects consecutive packets irrespective of their importance values while it cannot choose anything that \mathbf{S} chooses. If $\tilde{\mathbf{S}}$ can choose the packet that \mathbf{S} chooses at some time instant after i_0 , it does so forever. Since $\tilde{S}_j \geq S_j$ for all j , we have for all i

$$\frac{1}{i} \sum_{j=1}^i (T_j - \tilde{S}_j) \leq \frac{1}{i} \sum_{j=1}^i (T_j - S_j), \quad (3.51)$$

and consequently $\Delta_e^{(\tilde{\mathbf{S}})} \leq \Delta_e^{(\mathbf{S})}$. For $i > i_0$, if $S_i = \tilde{S}_i$, then

$$\frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, \tilde{S}_{j-1}, \tilde{S}_j) \leq \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_{j-1}, S_j). \quad (3.52)$$

This is because $V_s \leq V_{\tilde{s}}$ for every $s \in \{S_1, \dots, S_i\}$ and $\tilde{s} \in \{\tilde{S}_1, \dots, \tilde{S}_i\}$. If $S_i < \tilde{S}_i$ at $i > i_0$, then $\tilde{\mathbf{S}}$ must have skipped at most $\tilde{S}_{i_0} - S_{i_0}$ packets with

minimum importance. Then we have

$$\begin{aligned} \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, \tilde{S}_{j-1}, \tilde{S}_j) &\leq \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_{j-1}, S_j) + \frac{1}{i} v_{\min}(\tilde{S}_{i_0} - S_{i_0}) \\ &\leq \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_{j-1}, S_j) + \frac{1}{i} v_{\min}(T_{i_0} - S_{i_0}). \end{aligned} \quad (3.53)$$

Since $T_{i_0} - S_{i_0}$ is almost surely finite, $\frac{1}{i}(T_{i_0} - S_{i_0}) \rightarrow 0$. Therefore

$$\limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, \tilde{S}_{j-1}, \tilde{S}_j) \leq \limsup_{i \rightarrow \infty} \frac{1}{i} \sum_{j=1}^i D(\mathbf{V}^{T_j}, S_{j-1}, S_j) \quad (3.54)$$

and consequently $D(\tilde{\mathcal{S}}) \leq D(\mathcal{S})$.

3.8.3 Proof of Property 3.1

In the following proofs, $g(\mathbf{b}, s)$ refers to the one step costs of the MDP and $E_{\mathbf{b},s}[h(\mathbf{B}')]]$ refers to the expected relative value of the next state accessed under the action s , i.e., $g(\mathbf{b}, s) = \sum_{k=1}^{s-1} b_k + \eta(l(\mathbf{b}) - s)$ and $E_{\mathbf{b},s}[h(\mathbf{B}')] = E[h(\mathbf{b}_{\geq s+1} \| \mathbf{V}^Z)]$.

- (i) Let $\mathbf{q} := \mathbf{b} \| \mathbf{b}'$. Suppose $s^*(\mathbf{q}) \neq l(\mathbf{b}) + s^*(\mathbf{b}')$ and $s^*(\mathbf{q}) > l(\mathbf{b})$. Then by the optimality condition, $s^*(\mathbf{q}) = \arg \min_{s \leq l(\mathbf{q})} g(\mathbf{q}, s) + E_{\mathbf{q},s}[h(\mathbf{B}')] = l(\mathbf{b}) + \arg \min_{s \leq l(\mathbf{b}')} g(\mathbf{b}', s) + E_{\mathbf{b}',s}[h(\mathbf{B}')] = l(\mathbf{b}) + s^*(\mathbf{b}')$, which contradicts the statement. Therefore, if $s^*(\mathbf{q}) > l(\mathbf{b})$, then $s^*(\mathbf{q}) = l(\mathbf{b}) + s^*(\mathbf{b}')$.
- (ii) We use \mathbf{u} instead of \mathbf{b}' for notational convenience and let $\mathbf{q} := \mathbf{b} \| \mathbf{u}$. The second upper bound is easy to show as the policy is restricted to $s > l(\mathbf{b})$. To prove $h^*(\mathbf{u}) \leq h^*(\mathbf{q})$, we will find a lower bound for $h^*(\mathbf{q}) - h^*(\mathbf{u})$ and show that it is non-negative. Observe that

$$\begin{aligned} h^*(\mathbf{q}) - h^*(\mathbf{u}) &= h^*(\mathbf{q}) - \min_s E_{\mathbf{u},s}[g(s, \mathbf{u}) + h^*(\mathbf{U}') - \lambda^*] \\ &\geq h^*(\mathbf{q}) - E_{\mathbf{u},s}[g(s, \mathbf{u}) + h^*(\mathbf{U}') - \lambda^*] \\ &= E_{\mathbf{q},s^*(\mathbf{q})}[g(s, \mathbf{q}) + h^*(\mathbf{Q}')] - E_{\mathbf{u},s}[g(s, \mathbf{u}) + h^*(\mathbf{U}')] \end{aligned} \quad (3.55)$$

for any s . Therefore, changing the actions that govern the Markov chain $\{\mathbf{U}_i\}$ leads to further lower bounds. Now, we will use a similar coupling idea as done in the proof of Lemma 3.1. Consider the two Markov chains above where the first one starts from \mathbf{q} and the other from \mathbf{u} . Assume the two processes are coupled with having the same future arrivals for the consequent buffer states. Moreover, the former chain is controlled with its respective optimal policy \mathbf{s}^* , whereas the latter is controlled as follows: If possible, choose the data chosen by the first process; otherwise choose the oldest data possible. Denote this policy as $\tilde{\mathbf{s}}$ and let $\{\mathbf{Q}_i\}$

be the Markov chain pertaining to the former process. Then these two processes will follow the paths

$$\begin{aligned} \mathbf{q} &= \mathbf{Q}_0 \xrightarrow{s^*(\mathbf{Q}_0)} \mathbf{Q}_1 \xrightarrow{s^*(\mathbf{Q}_1)} \mathbf{Q}_2 \rightarrow \dots \rightarrow \mathbf{Q}_\tau \\ \mathbf{u} &= \mathbf{U}_0 \xrightarrow{\tilde{s}(\mathbf{U}_0)} \mathbf{U}_1 \xrightarrow{\tilde{s}(\mathbf{U}_1)} \mathbf{U}_2 \rightarrow \dots \rightarrow \mathbf{U}_\tau \end{aligned} \quad (3.56)$$

and eventually end in the same state $\mathbf{Q}_\tau = \mathbf{U}_\tau$ after a random time τ . This is because all possible policies induce the same recurrent class as mentioned in Lemma 3.2. Replacing $h^*(\mathbf{Q}')$, $h^*(\mathbf{U}')$ in (3.55) several times, we obtain

$$h^*(\mathbf{q}) - h^*(\mathbf{u}) \geq E \left[\sum_{i=0}^{\tau} g(\mathbf{Q}_i, s^*(\mathbf{Q}_i)) - g(\mathbf{U}_i, \tilde{s}(\mathbf{U}_i)) \right]. \quad (3.57)$$

Observe that the right-hand side is equal to the expectation of the difference of accumulated costs incurred by the two processes until they reach the same state. Note that \mathbf{U}_i is a suffix of \mathbf{Q}_i for all i . Therefore, the age penalty of the former process will be greater. Furthermore, the latter process chooses every data that can be chosen by the former one, except the ones in the first portion \mathbf{b} , whose miss do not contribute to the distortion penalty of the latter process. Since both accumulated age and distortion penalties of the former process $\{\mathbf{Q}_i\}$ cannot be smaller than those of $\{\mathbf{U}_i\}$, the expectation in (3.57) is non-negative. Hence the proof is complete.

3.8.4 Proof of Theorem 3.3

For simplicity, we prove the case with $|\mathcal{V}| = 3$ and we set $v_{\min} = 1$ without loss of generality. The same proof technique may be extended to larger $|\mathcal{V}|$, e.g., with induction.

Assume the policy iteration algorithm is executed with a sufficiently large buffer size M . We first show that packets of importance v_2 are never chosen by the optimal policy if choosing them incurs an age penalty greater than $K_2 := \lceil \frac{v_2-1}{\eta\mu_Z} \rceil$, i.e., they must not be generated more than K_2 time slots ago. We use the same coupling idea of two controlled Markov chains with different policies, as we have done in the proofs of Lemma 3.1 and Property 3.1(ii).

Take a state $\mathbf{b} \in \mathcal{V}^{\leq M}$ with length greater than K_2 , with $b_1 = v_2$ (this equality is without loss of generality) and suppose $s^*(\mathbf{b}) = 1$. We will show that there is a better strategy than $s^*(\mathbf{b}) = 1$, which contradicts the optimality of s^* . Consider now $\tilde{s}(\mathbf{b}) = \min\{k > 1 : b_k > 1\}$, i.e, the index of first non-1 data (whose importance is greater than 1) in the remaining buffer. Optimality of s^* requires the following inequality to hold (the argument \mathbf{b} of $\tilde{s}(\mathbf{b})$ is omitted), otherwise the policy iteration would not have converged.

$$E \left[\eta(\tilde{s} - 1) + h^*(b_2^l \| \mathbf{V}^Z) - h^*(b_{\tilde{s}+1}^l \| \mathbf{V}^Z) \right] - \frac{1}{\mu_Z} v_2 \leq 0 \quad (3.58)$$

Now, consider two coupled processes $\{\mathbf{Q}_i\}$, $\{\mathbf{U}_i\}$ with $\mathbf{Q}_0 := (b_2^l \| \mathbf{V}^Z)$, $\mathbf{U}_0 := (b_{\tilde{s}+1}^l \| \mathbf{V}^Z)$ where the actions of $\{\mathbf{U}_i\}$ are modified. The modification will be the same as in the proof of Property 3.1(ii): Choose the packet chosen by the first process if possible; otherwise choose the oldest possible. As shown in the proof of Property 1(ii), altering the policy increases the expected relative value. Consequently, we obtain a lower bound to the expectation in (3.58). Again, as every policy has the same recurrent class, the two processes coincide with probability one. When they coincide, one of the following occurs:

- (i) $\{\mathbf{Q}_i\}$ misses a non-1 data that is taken by $\{\mathbf{U}_i\}$. Then the accumulated cost is greater than $\frac{1}{\mu_Z} v_2$.
- (ii) $\{\mathbf{U}_i\}$ takes a packet of importance 1 at the end. Then the accumulated penalty incurred by age will be greater than $\eta l(\mathbf{b})$ (using the fact that $\{\mathbf{U}_i\}$ remains a suffix of $\{\mathbf{Q}_i\}$ for all i) and the accumulated cost from distortion will be greater than $\frac{1}{\mu_Z}$ (since $\{\mathbf{U}_i\}$ takes an extra packet of importance 1). The total cost will be greater than $\eta l(\mathbf{b}) + \frac{1}{\mu_Z} \geq \eta K_2(\eta) + \frac{1}{\mu_Z} \geq \frac{1}{\mu_Z} v_2$.

Therefore, we conclude that $\frac{1}{\mu_Z} v_2$ is a lower bound to the expectation in (3.58). Hence the left-hand side of (3.58) can never be negative and s^* cannot be optimal — if it is equal to zero, then s^* and \tilde{s} are indifferent, so one can drive the process with \tilde{s} .

Now we proceed in a similar fashion to prove the case for v_3 . Consider a state \mathbf{b} with $l(\mathbf{b}) > K_3 := \lceil \frac{v_3-1}{\eta\mu_Z} \rceil$, $b_1 = v_3$ and suppose $s^*(\mathbf{b}) = 1$. Choose $\tilde{s}(\mathbf{b}) = \min\{k > 1 : b_k > 1, l(\mathbf{b}) - k < K_2 \text{ if } b_k = v_2\}$, i.e., take the first non-1 data but with the constraint that if it has importance v_2 , it must be generated within the most recent K_2 time slots.

Define $\{\mathbf{Q}_i\}$ and $\{\mathbf{U}_i\}$ in a similar fashion. The modification done to the actions on $\{\mathbf{U}_i\}$ will have a minor difference compared to the previous case. Again, if \mathbf{U}_i cannot choose a data that is chosen by \mathbf{Q}_i , it chooses the first non-1 data but with the following extra condition: If its importance is v_2 , take it if has been generated less than K_2 time slots ago; otherwise skip it and do the same for the next non-1 data. This modification ensures that \mathbf{U}_i will choose every possible data that may be chosen by \mathbf{Q}_i . When the two processes coincide, one of the following occurs:

- (i) $\{\mathbf{Q}_i\}$ misses v_3 that is taken by $\{\mathbf{U}_i\}$. Then the accumulated cost is greater than $\frac{1}{\mu_Z} v_3$.
- (ii) $\{\mathbf{Q}_i\}$ misses v_2 that is taken by $\{\mathbf{U}_i\}$. Then the accumulated cost incurred by age will be greater than $\eta(l(\mathbf{b}) - K_2)$ (as $\{\mathbf{U}_i\}$ has remained a suffix of $\{\mathbf{Q}_i\}$ and has taken this v_2 , whose index must be greater than $l(\mathbf{b}) - K_2$). The accumulated cost from distortion will be greater than $\frac{1}{\mu_Z} v_2$. Consequently, the total cost will be greater than $\eta(l(\mathbf{b}) - K_2) + \frac{1}{\mu_Z} v_2 \geq \frac{v_3-1}{\mu_Z} - \frac{v_2-1}{\mu_Z} + \frac{1}{\mu_Z} v_2 = \frac{1}{\mu_Z} v_3$.

- (iii) $\{U_i\}$ takes a 1 at the end. Then the accumulated cost incurred by age will be greater than $\eta l(\mathbf{b})$ (again, using the fact that $\{U_i\}$ remains a suffix of $\{Q_i\}$) and the accumulated cost from distortion will be greater than $\frac{1}{\mu_Z}$ (since $\{U_i\}$ takes an extra 1). The total cost will be greater than $\eta l(\mathbf{b}) + \frac{1}{\mu_Z} \geq \frac{1}{\mu_Z} v_3$.

Similar to the case of v_2 , the expectation (3.58) can never be negative and hence s^* cannot be optimal.

All in all, we have shown that starting the policy iteration algorithm with a buffer size $M > K_3$, the algorithm terminates with an optimal policy that does not use more than a buffer size of K_3 . This implies that the solution of the Bellman equation lies within buffers of size at most K_3 .

3.8.5 Policy Update Improvement

As discussed before, the policy update stage of Algorithm 3 can be improved to run in $O(K|\mathcal{V}|^K)$ time. Our aim is to improve the calculation of $C_{h^{(i)}}(\mathbf{b}, 1)$ at step 16 of Algorithm 3. Recall that $p_z := \Pr(Z = z)$, and $q_z := \Pr(Z \geq z)$. We omit the argument of $l(\mathbf{b})$ for brevity. Let us rewrite $C_{h^{(i)}}(\mathbf{b}, 1)$ by first conditioning on $Z = z$ as

$$\begin{aligned} C_{h^{(i)}}(\mathbf{b}, 1) &= \eta(l-1) + \sum_{z=1}^{\infty} p_z E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^z)] \\ &= \eta(l-1) + \sum_{z=1}^{K-l} p_z E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^z)] \\ &\quad + p_{K-l+1} E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^{K-l+1})] \end{aligned} \quad (3.59)$$

$$+ \sum_{z=K-l+2}^{\infty} p_z E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^z)] \quad (3.60)$$

and note that for $z > K-l+1$, $h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^z) = b_2/\mu_Z + h^{(i)}(\mathbf{b}_{\geq 3} \| \mathbf{V}^z)$ since the length of $\mathbf{b}_{\geq 2} \| \mathbf{V}^z$ exceeds K . In line with this result, we rewrite the summation of (3.59) and (3.60) as

$$\begin{aligned} \kappa^{(i)}(\mathbf{b}_{\geq 2}, K) &:= p_{K-l+1} E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^{K-l+1})] \\ &\quad + q_{K-l+2} b_2/\mu_Z + \sum_{z=K-l+2}^{\infty} p_z E[h^{(i)}(\mathbf{b}_{\geq 3} \| \mathbf{V}^z)] \\ &= p_{K-l(\mathbf{b}_{\geq 2})} E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^{K-l(\mathbf{b}_{\geq 2})})] \\ &\quad + q_{K-l(\mathbf{b}_{\geq 2})+1} b_2/\mu_Z + \sum_{z=K-l(\mathbf{b}_{\geq 2})+1}^{\infty} p_z E[h^{(i)}(\mathbf{b}_{\geq 3} \| \mathbf{V}^z)] \end{aligned} \quad (3.61)$$

The key observation here is that the last term is equal to $\kappa^{(i)}(\mathbf{b}_{\geq 3}, K) = \kappa^{(i)}(\text{parent}(\mathbf{b}_{\geq 2}), K)$. Therefore, we have the recursive relation

$$\kappa^{(i)}(\mathbf{b}, K) := p_{K-l(\mathbf{b})} E[h^{(i)}(\mathbf{b} \| \mathbf{V}^{K-l(\mathbf{b})})] + q_{K-l(\mathbf{b})+1} b_1/\mu_Z + \kappa^{(i)}(\text{parent}(\mathbf{b}), K) \quad (3.62)$$

with the initial condition

$$\kappa^{(i)}(\delta, K) := q_K E[h^{(i)}(\mathbf{V}^K)] + \frac{E[V]}{\mu_Z} E[(Z - K)^+]. \quad (3.63)$$

The recursive relation (3.62) can be implemented exactly the same as the other updates that take place in lines 17-26 of Algorithm 3. Finally, we have

$$C_{h^{(i)}}(\mathbf{b}, 1) = \eta(l - 1) + \sum_{z=1}^{K-l} p_z E[h^{(i)}(\mathbf{b}_{\geq 2} \| \mathbf{V}^z)] + \kappa^{(i)}(\text{parent}(\mathbf{b}), K). \quad (3.64)$$

Knowing $\kappa^{(i)}(\text{parent}(\mathbf{b}), K)$, both $C_{h^{(i)}}(\mathbf{b}, 1)$ and $\kappa^{(i)}(\mathbf{b}, K)$ can be calculated in $O(|\mathcal{V}|^{K-l})$ steps. Consequently, the total amount of calculations done for all length- l states is $O(|\mathcal{V}|^K)$, and consequently for the depth- K tree \mathcal{T} , the amount of calculations done is $O(K|\mathcal{V}|^K)$.

3.8.6 Closed-Form Expressions for the Simple Strategies

S1 — Send the oldest important data among the most recent K

For a buffer \mathbf{b} , let s be the index of the first important data among the K most recently arrived packets. Let $a = (K \wedge l(\mathbf{b})) - s + 1$, i.e., the number of fresher items plus 1; and if there is no important data among the K most recent packets, set $a = 0$. Let A_i be the value of a at time instant i . It can be verified that the process $\{A_i\}$ is a Markov chain with state space $\{0, 1, \dots, K\}$. Let $p_{a,a'}$ denote the transition probability from state a to a' ; and π_a denote the stationary probability of state a . Recall that Z is a geometric random variable and denote $p := \Pr(Z = 1)$ and $q := \Pr(V = v_2)$. Let $\bar{x} := 1 - x$.

Now, let us calculate $p_{a,a'}$. Observe that when we are at state a , the first element in the buffer is selected and the remaining $a - 1$ elements are untouched and need not be known. Conditioned on the next speaking time being $a' - a < z \leq K - a$, the probability that the next state being a' is equal to $\bar{q}^{z-(a'-a)-1}q$. This is because the first $z - (a' - a) - 1$ elements must be v_1 and the next should be v_2 . For $z > K - a$, since we only check the most recent K data in the buffer, the first $K - a'$ must be v_1 and the next one should be v_2 . Thus we obtain the probability as $\bar{q}^{K-a'}q$. Averaging over z , we have

$$p_{a,a'} = \sum_{z=a'-a+1}^{\infty} \bar{p}^{z-1} p \bar{q}^{(z-a'+a-1) \wedge (K-a')} q. \quad (3.65)$$

With a similar reasoning, we calculate $p_{a',a}$ as

$$\begin{aligned} p_{a',a} &= \sum_{z=1}^{\infty} \bar{p}^{z-1} p \bar{q}^{(z+a'-a-1) \wedge (K-a)} q \\ &= \sum_{z=a'-a+1}^{\infty} \bar{p}^{z-a'+a-1} p \bar{q}^{(z-1) \wedge (K-a)} q = \left(\frac{\bar{q}}{\bar{p}} \right)^{a'-a} p_{a,a'}. \end{aligned} \quad (3.66)$$

It is easy to see that $p_{0,a} = p_{1,a}$ as the buffer is emptied in both cases; and $p_{a,0} = p_{a,1} \frac{\bar{q}}{q}$ as the only difference between the transition from state a to state 0, or to state 1 is the most recent element being v_1 or v_2 .

We now claim that the chain is reversible. The claim is easily verified by noting that the distribution $(\pi_a : a = 0, \dots, K)$ described by

$$\pi_0 = (\bar{q}/q)\pi_1, \quad \pi_a = (\bar{q}/\bar{p})^{K-a}\pi_K, \quad a = 1, \dots, K-1$$

with

$$\pi_K = \left(1 - \frac{\bar{q}}{\bar{p}}\right) / \left(1 - \frac{p}{q} \left(\frac{\bar{q}}{\bar{p}}\right)^K\right) \quad (3.67)$$

satisfies the detailed balance equations $\pi_a p_{a,a'} = \pi_{a'} p_{a',a}$.

The average excess age is then calculated straightforwardly:

$$\Delta_e^{(S1)}(K) = \sum_{k=1}^K (k-1)\pi_k = \frac{K-1}{\left(1 - \frac{p}{q} \left(\frac{\bar{q}}{\bar{p}}\right)^K\right)} - \frac{\frac{\bar{q}}{\bar{p}} \left(1 - \left(\frac{\bar{q}}{\bar{p}}\right)^K\right)}{\left(1 - \frac{p}{q} \left(\frac{\bar{q}}{\bar{p}}\right)^K\right) \left(1 - \frac{\bar{q}}{\bar{p}}\right)}. \quad (3.68)$$

For $q = p$, one takes $\frac{\bar{q}}{\bar{p}} \rightarrow 1$ in (3.68) to obtain $\Delta_e^{(S1)} = \frac{(K-1)K}{2(K+\frac{p}{p})}$.

The average distortion is calculated as follows: Unimportant packets are sent π_0 fraction of the speaking times, and the remaining time is allocated to the transmission of important packets. Hence, unimportant packets are sent $p\pi_0$ fraction of the time and important packets are sent $p(1 - \pi_0)$ fraction of the time. Consequently,

$$D^{(S1)}(K) = (\bar{q} - p\pi_0)v_1 + (q - p(1 - \pi_0))v_2. \quad (3.69)$$

S2 — Send the newest important data among the most recent K

Compared to the other strategies, the analysis will be relatively simpler. For a buffer \mathbf{b} , let s be the index of the newest important data among the K most recent and let the state be $a = l - s + 1$. If there is no important data among the K most recent, set $a = 0$. Observe that regardless of the value of a , the packets (if any) that remain the buffer after transmission are of minimum importance and will be ignored at the next speaking time. Hence, at the next speaking time, the next state a' will not depend on the current state a , and we have an i.i.d. process. Let π_a be the probability of the next state being equal to a . Conditioned on the next speaking time z , this probability is equal to $q\bar{q}^{a-1} \mathbb{1}\{a \leq z\}$. Hence, for $0 < a \leq K$

$$\pi_a = \sum_{z \geq a} q\bar{q}^{a-1} \bar{p}^{z-1} p = q(\bar{p}\bar{q})^{a-1}, \quad (3.70)$$

and $\pi_0 = 1 - \sum_{a=1}^K \pi_a = \frac{\bar{q}p + q(\bar{p}\bar{q})^K}{1 - \bar{p}\bar{q}}$. The age and distortion are calculated similar to (3.68) and (3.69).

S3 — Send the newest important data that has arrived more than K slots ago. If there is no such data, send the oldest important one

We set the state associated to a buffer \mathbf{b} as follows: If \mathbf{b} does not contain an important data arrived more than K time slots ago, set $a = (K \wedge l(\mathbf{b})) - s + 1$ as the state; and if there is no important data in the buffer, set $a = 0$. Note that this is exactly the same as in S1. If \mathbf{b} does contain an important data that has arrived more than K time slots ago, set $a = K + 1$. Hence the state space will be $\{0, 1, \dots, K + 1\}$. Similar to the analysis of S1, $\{A_i\}$ will be a Markov chain. We want to calculate the transition probabilities $p_{K+1,a}$ for $0 < a < K + 1$. Since at state $K + 1$, an important data has arrived more than K slots ago, \mathbf{b} is of the form $\mathbf{b} = [v_2, v_1, \dots, v_1, \underbrace{\dots}_K]$, where the last K data need not be known. Conditioned on the speaking time z , the probability of ending up in state a is then should be $\bar{q}^{z+K-a}q$. Consequently,

$$p_{K+1,a} = \sum_{z=1}^{\infty} \bar{p}^{z-1} p \bar{q}^{K-a+z} q = \frac{pq\bar{q}^{K-a+1}}{1 - \bar{p}\bar{q}}. \quad (3.71)$$

Now, we aim to find $p_{a,K+1}$ for $0 < a < K + 1$. Since \mathbf{b} is of the form $\mathbf{b} = [\dots, v_2, \underbrace{\dots}_{a-1}]$, and the last $a - 1$ data need not be known, the next speaking time should be greater than $K - a + 1$. Moreover, there must be at least one important data among the first $z - (K - a + 1)$. Then, we obtain

$$p_{a,K+1} = \sum_{z>K-a+1} \bar{p}^{z-1} p (1 - \bar{q}^{z-(K-a+1)}) = \frac{q\bar{p}^{K-a+1}}{1 - \bar{p}\bar{q}}. \quad (3.72)$$

Calculation of $p_{a,a'}$, $0 < a \leq a' < K + 1$ is similar to the one in S1. Since $\mathbf{b} = [\dots, v_2, \underbrace{\dots}_{a-1}]$, the next speaking time should be greater than $a' - a$, and the first $z - (a' - a + 1)$ data must be unimportant while the $z - (a' - a)$ th data must be important. Thus,

$$p_{a,a'} = \sum_{z>a'-a} \bar{p}^{z-1} p \bar{q}^{z-a'+a} q = \frac{pq\bar{p}^{a'-a}}{1 - \bar{p}\bar{q}}. \quad (3.73)$$

A similar analysis reveals that for $0 < a' \leq a < K + 1$,

$$p_{a,a'} = \sum_{z=1}^{\infty} \bar{p}^{z-1} p \bar{q}^{z-a'+a-1} q = \frac{pq\bar{q}^{a-a'}}{1 - \bar{p}\bar{q}}. \quad (3.74)$$

Finally, as we discussed in the analysis of S1, $p_{0,a} = p_{1,a}$ and $p_{a,0} = p_{a,1} \frac{\bar{q}}{q}$. Hence, we have found all the transition probabilities. We will verify that $\{A_i\}$ is a reversible Markov chain. The detailed balance equations for the state pairs $(a, K + 1)$ with $a > 0$ require the stationary distribution π to satisfy

$$\pi_{K+1} p_{K+1,a} = \pi_a p_{a,K+1} \quad (3.75)$$

and thus we find $\pi_a = p\left(\frac{\bar{q}}{\bar{p}}\right)^{K+1-a} \pi_{K+1}$ for $a > 0$. Detailed balance equations for the state pair $(0, 1)$ further yield $\pi_0 = (\bar{q}/q)\pi_1 = \frac{p\bar{q}}{q}\left(\frac{\bar{q}}{\bar{p}}\right)^K \pi_{K+1}$. One can easily verify that this choice satisfies not only (3.75) but all the detailed balance equations. Thus, we have verified the reversibility of the chain and found its stationary distribution.

Let us start the chain with the stationary distribution and let $\Delta_i := l(\mathbf{B}_i) - s(\mathbf{B}_i)$ be the instantaneous excess age at the i^{th} speaking time. If $0 \leq A_i \leq K$, then $\Delta_i = (A_i - 1) \vee 0$. When $A_i = K + 1$, however, Δ_i is random. It turns out that the distribution of $\Delta_i - (K - 1)$ conditioned on $A_i = K + 1$ is geometrically distributed. To see this, observe

$$\begin{aligned}
& \Pr(\Delta_i = K - 1 + z' | A_i = K + 1) \\
&= \frac{\Pr(\Delta_i = K - 1 + z', A_i = K + 1)}{\pi_{K+1}} \\
&= \frac{1}{\pi_{K+1}} \sum_{a=1}^{K+1} \pi_a \sum_{z \geq K-a+1+z'} \bar{p}^{z-1} p q \bar{q}^{z'-1} + \frac{\pi_0}{\pi_{K+1}} \sum_{z \geq K+z'} \bar{p}^{z-1} p q \bar{q}^{z'-1} \\
&= \frac{1}{\pi_{K+1}} \sum_{a=1}^{K+1} \pi_a \bar{p}^{K-a+z'} q \bar{q}^{z'-1} + \frac{\pi_0}{\pi_{K+1}} \bar{p}^{K+z'-1} q \bar{q}^{z'-1} \\
&= \sum_{a=1}^{K+1} p^{\mathbb{1}\{a \leq K\}} \left(\frac{\bar{q}}{\bar{p}}\right)^{K+1-a} \bar{p}^{K-a+z'} q \bar{q}^{z'-1} + p \frac{\bar{q}}{q} \left(\frac{\bar{q}}{\bar{p}}\right)^K \bar{p}^{K+z'-1} q \bar{q}^{z'-1} \\
&= (\bar{p}\bar{q})^{z'-1} (1 - \bar{p}\bar{q}).
\end{aligned} \tag{3.76}$$

Then, the average excess age is found as

$$\Delta_e^{(S3)}(K) = \sum_{k=1}^K (k-1)\pi_k + \pi_{K+1} \left(\frac{1}{1-\bar{p}\bar{q}} + K - 1 \right). \tag{3.77}$$

The average distortion is calculated the same as in (3.69).

3.8.7 Proof of Corollary 3.2

Denote the stationary probabilities of each state l as π_l and $p_i := \Pr(Z = i) = p\bar{p}^{i-1}$ with $\bar{p} := 1 - p$. Assume $\tau > N$ for the moment. The derivation for $\tau \leq N$ will easily follow.

Let $s_i^j := \sum_{l=i}^j \pi_l$ and $s_i := s_i^\infty$. The stationary probabilities are the solution

to the linear system

$$\begin{aligned}
\pi_1 &= p_1 s_1^N \\
\pi_2 &= p_2 s_1^N + p_1 \pi_{N+1} \\
&\dots \\
\pi_\tau &= p_\tau s_1^N + p_{\tau-1} \pi_{N+1} + \dots + p_1 \pi_{N+\tau-1} \\
\pi_{\tau+1} &= p_{\tau+1} s_1^N + p_\tau \pi_{N+1} + \dots + p_1 s_{N+\tau} \\
&\dots \\
\pi_{\tau+j} &= p_{\tau+j} s_1^N + p_{\tau+j-1} \pi_{N+1} + \dots + p_j s_{N+\tau} \\
&\dots
\end{aligned} \tag{3.78}$$

Observe that

$$\pi_{j+1} = \bar{p} \pi_j + p \pi_{N+j}, \quad 1 \leq j \leq \tau - 1 \tag{3.79}$$

and

$$\pi_{j+1} = \bar{p} \pi_j, \quad j \geq \tau + 1. \tag{3.80}$$

Summing up the equalities in (3.79), with indices up to $j + 1$, we obtain

$$s_2^{j+1} = \bar{p} s_1^j + p s_{N+1}^{N+j} \tag{3.81}$$

and hence

$$\pi_{j+1} + p s_2^j = \bar{p} \pi_1 + p s_{N+1}^{N+j}. \tag{3.82}$$

Note that the first equation in (3.78) implies $\bar{p} \pi_1 = p s_2^N$ and thus we get

$$\pi_{j+1} + p s_2^j = p s_2^N + p s_{N+1}^{N+j} \tag{3.83}$$

and

$$\bar{p} \pi_{j+1} = p s_{j+2}^{N+j}, \quad 1 \leq j \leq \tau - 1, \tag{3.84}$$

which implies for $j = \tau - 1$

$$\pi_\tau = p \sum_{k=1}^{N+\tau-1} \pi_{\tau+k} = \frac{\pi_{\tau+1} (1 - \bar{p}^{N-1})}{\bar{p}} \tag{3.85}$$

where the last equality follows from (3.80). Now repeated application of (3.79) gives

$$\pi_{\tau-j} = \pi_{\tau+1} \frac{1 - (1 + jp) \bar{p}^{N-1}}{\bar{p}^{j+1}}, \quad 0 \leq j \leq N - 1. \tag{3.86}$$

Our aim is now to find all stationary probabilities in terms of $\pi_{\tau+1}$. With the above, we are able to find $\pi_{\tau-j}$, $0 \leq j \leq N - 1$ in terms of $\pi_{\tau+1}$. For $j > N - 1$, we try to observe a pattern. First, try to calculate $\pi_{\tau-N}$ by (3.79), which gives

$$\pi_{\tau-N} = \pi_{\tau+1} \frac{1 - (1 + Np) \bar{p}^{N-1} + p \bar{p}^{2(N-1)}}{\bar{p}^{N+1}}. \tag{3.87}$$

Once more, repeated application of (3.79) gives

$$\pi_{\tau-N-j} = \pi_{\tau+1} \frac{1 - (1 + (N + j)p)\bar{p}^{N-1}}{\bar{p}^{N+j+1}} + \pi_{\tau+1} \frac{p(\sum_{k=0}^j(1 + kp))\bar{p}^{2(N-1)}}{\bar{p}^{N+j+1}} \quad (3.88)$$

for $0 \leq j \leq N - 1$. Doing the same procedure, we observe the following pattern: Let $S_j^{(0)} := (1 + jp)$ and $S_j^{(n)} := \sum_{k=0}^j S_k^{(n-1)}$ for $n \geq 1$. Also let $S_j^{(n)} = 0$ for $j < 0$. Then,

$$\pi_{\tau-j} = \pi_{\tau+1} \frac{1 + \sum_{k=0}^{\lceil \tau/N \rceil} (-1)^{k+1} S_{j-kN}^{(k)} p^k \bar{p}^{(k+1)(N-1)}}{\bar{p}^{j+1}}. \quad (3.89)$$

Finally, since the probabilities sum up to one, we have

$$\begin{aligned} 1 &= s_1^\tau + s_{\tau+1} = s_1^\tau + \frac{\pi_{\tau+1}}{p} \\ &= \pi_{\tau+1} \sum_{j=0}^{\tau-1} \frac{1 + \sum_{k=0}^{\lceil \tau/N \rceil} (-1)^{k+1} S_{j-kN}^{(k)} p^k \bar{p}^{(k+1)(N-1)}}{\bar{p}^{j+1}} + \frac{\pi_{\tau+1}}{p} \end{aligned} \quad (3.90)$$

and therefore

$$\pi_{\tau+1} = \left[\sum_{j=0}^{\tau-1} \frac{1 + \sum_{k=0}^{\lceil \tau/N \rceil} (-1)^{k+1} S_{j-kN}^{(k)} p^k \bar{p}^{(k+1)(N-1)}}{\bar{p}^{j+1}} + \frac{1}{p} \right]^{-1}. \quad (3.91)$$

After calculating all π_j s, it is straightforward to obtain expressions for Δ_e and D as

$$\Delta_e = \sum_{j=1}^{\tau-1} j \pi_{N+j} + \tau \sum_{j=\tau}^{\infty} \pi_{N+j} = \sum_{j=1}^{\tau-1} j \pi_{N+j} + \frac{\tau \pi_{\tau+1} \bar{p}^{N-1}}{p}. \quad (3.92)$$

and

$$D = \mu_V \sum_{j=1}^{\infty} j \pi_{\tau+N+j} = \frac{\mu_V \pi_{\tau+1} \bar{p}^N}{p^2}. \quad (3.93)$$

Age-Optimal Causal Labeling of Memoryless Processes¹

4

As in the previous chapter, think of a component in the network, which receives data at a higher rate than it is allowed to send. In this case, one may ask (i) what and (ii) when to send with an aim to optimize AoI? Studying the first question may need classification of data according to their importance, which we have already done. The second question is addressed for example in [50, 51, 79], and in [38–40, 84–86] when the throughput is limited due to energy constraints. All of these studies assume *a priori* knowledge of the network dynamics, e.g., the data is known to be conveyed through a single-server queue [70, 71], or controlled by an external device as in the previous chapter. In this chapter, we will assume that the sender is oblivious of the network dynamics except the limit on the output data rate.

We focus on arrivals modeled as a memoryless point process with an arrival rate greater than the limited output rate. The sender thus needs to filter out some of the arrivals. We call this filtering operation a ‘labeling procedure’, where an arrival is passed through the network as soon as it is labeled. We consider causal labelings, i.e., a data can be labeled only after it has arrived. These labeling procedures relate to both questions (i) and (ii) above. As the sender is oblivious of the network dynamics, we study the tradeoff between the rate and the age of the labeled process, which are to be defined in Section 4.1. The reader should keep in mind that in this chapter the packets are treated with equal priority, as opposed to **Chapter 3**.

The outline of this chapter is as follows: In Section 4.1, we provide the problem definition in a discrete-time setting. The rate-age tradeoff is related to an appropriate Markov Decision Problem (MDP) formulation. In Section 4.2, we study a finite-state approximation which allows to characterize the optimal labeling procedures for the original one. The optimal procedures turn

¹The content of this chapter is based on [52].

out to be rather simple: Wait T time slots and label the next arrival, where T is tuned to match the output rate. In Section 4.3, we extend the results to a continuous-time model, where the arrivals are modeled as a Poisson process.

4.1 Problem Definition

The problem definition will be very similar to the one in **Chapter 3**, and *the notation will be the same*. For convenience, we will still provide a self-contained problem formulation for the causal labeling problem. Consider a discrete-time memoryless point process, where the interarrival times Z_1, Z_2, \dots are independent and identically distributed (i.i.d.) with Geometric distribution — recall that as a result of our previous explorations in Section 2.2, we have stated that in discrete-time, the memoryless arrival process should have Geometric interarrival times. Suppose their success probability is p , which is also equal to the rate of the arrival process. The arrivals can also be modeled as i.i.d. Bernoulli random variables X_t with success probability p , with the natural filtration $\{\mathcal{F}_t\}$ where $\mathcal{F}_t := \sigma(X_1, \dots, X_t)$. For convenience assume $X_0 = 1$. We define a causal labeling on the point process as a sequence of (possibly random) labeling functions adapted to the filtration $\{\mathcal{F}_t\}$ as $S_t : \{0, 1\}^t \rightarrow \{1, \dots, t\} \cup \{?\}$, with the following restriction that if $S_t \neq ?$, then (i) X_{S_t} must be 1 and (ii) $S_t \neq S_\tau$ for all $\tau < t$. Observe that as opposed to the previous chapter, we do not impose order-preserving strategies. Arrival X_{S_t} is labeled at time t , and multiple labelings of a single point are not allowed. We assume $S_0 = 0$.

Given a procedure $\{S_t\} := \mathbf{S}$, we define the rate R as the expected long-term average of the number of labelings. More precisely,

$$R^{(\mathbf{S})} := E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \mathbf{1}\{S_\tau \neq ?\} \right]. \quad (4.1)$$

Define the most recent labeling at time t as $M_t := \max\{S_\tau : 0 \leq \tau < t, S_\tau \neq ?\}$ with $M_0 = 0$. The instantaneous age $\Delta_t^{(\mathbf{S})}$ and average age $\Delta^{(\mathbf{S})}$ are then defined as

$$\Delta_t^{(\mathbf{S})} := t - M_t, \quad \Delta^{(\mathbf{S})} := E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \Delta_\tau^{(\mathbf{S})} \right]. \quad (4.2)$$

Sometimes we omit the superscript (\mathbf{S}) from the above expressions for brevity. At this point, it might be useful to give the following examples of different labeling procedures to demonstrate what they resemble in practice.

Example 4.1. *Label every point upon arrival with probability α . This labeling procedure constitutes a renewal process whose interarrival times W are Geometric random variables with parameter αp . Hence, the rate and age of this renewal process are $R = \alpha p$ and $\Delta = \frac{E[W(W+1)]}{2E[W]} = \frac{1}{\alpha p} = \frac{1}{R}$.*

Example 4.2. Label every k^{th} point upon arrival. Likewise, this procedure yields a renewal process whose interarrival times are sum of k Geometric random variables with success probability p . Hence, $R = \frac{p}{k}$, and $\Delta = \frac{k+1}{2p} = \frac{1+1/k}{2} \frac{1}{R}$, which is strictly smaller than the age resulting from the labeling procedure in Example 1 for $k > 1$. This labeling procedure can be extended to cover the rates $R = pr$ for rational r and the resulting age can be shown to be smaller than the one in Example 1 as well. For the extension, see Appendix 4.5.1.

We are interested in finding the achievable region of possible (R, Δ) pairs with causal labeling procedures. More precisely, we are interested in finding the boundary curve of such pairs. Given the large class of possible labelings, this search seems difficult at first sight. However, we can eliminate some of the labeling procedures to make the search tractable.

Definition 4.1 (Strictly Increasing Procedures). A labeling procedure \mathbf{S} is strictly increasing if its subsequence $\{S_t : S_t \neq ?\}$ is strictly increasing.

Lemma 4.1. For any labeling procedure \mathbf{S} , there exists a strictly increasing modification $\tilde{\mathbf{S}}$ such that $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$ and $\Delta^{(\tilde{\mathbf{S}})} = \Delta^{(\mathbf{S})}$.

Proof. Recall that M_t is the most recent labeling at time t , with $M_0 = 0$. Take $\tilde{S}_t = S_t$ if $S_t = M_{t+1}$; otherwise $\tilde{S}_t = ?$. Then, $\Delta^{(\tilde{\mathbf{S}})} = \Delta^{(\mathbf{S})}$ and $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$, where $R^{(\mathbf{S})} = \sum_{\tau=1}^t \mathbb{1}\{S_\tau \neq ?\}$. Thus, $\tilde{R} \leq R$ and $\tilde{\Delta} = \Delta$. \square

Lemma 4.1 can be interpreted as follows: At time t , a strictly increasing modification of \mathbf{S} will consider arrivals after the most recent labeling M_t , and the (R, Δ) pair pertaining to this modification will be closer to the boundary curve we are trying to find. Hence we may focus on strictly increasing labelings that omit arrivals before and including M_t . Observe that the strictly increasing labeling procedures ensure that packets arrive in-order, just as in chapter 3.

We further restrict the space of labeling procedures we are interested in by introducing the lemma below.

Lemma 4.2. For any $t_1 < t_2 \leq \tau$ such that $X_{t_1} = X_{t_2} = 1$, and for every strictly increasing \mathbf{S} with $S_\tau = t_1$, there exists a strictly increasing modification $\tilde{\mathbf{S}}$ with $\tilde{S}_\tau = t_2$ and $\Delta^{(\tilde{\mathbf{S}})} \leq \Delta^{(\mathbf{S})}$, $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$. Consequently, $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$ and $\Delta^{(\tilde{\mathbf{S}})} \leq \Delta^{(\mathbf{S})}$.

Proof. Define $\tilde{\mathbf{S}}$ such that $\tilde{S}_t = S_t$ for $t < \tau$, $\tilde{S}_\tau = t_2$ and for all $t > \tau$,

$$\tilde{S}_t = \begin{cases} S_t, & S_t > t_2 \\ ?, & \text{else} \end{cases}. \quad (4.3)$$

In words, whenever an arrival later than t_2 is labeled by \mathbf{S} , it is also labeled by $\tilde{\mathbf{S}}$. Observe $\tilde{M}_t \geq M_t$ and $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$, thus $\Delta^{(\tilde{\mathbf{S}})} \leq \Delta^{(\mathbf{S})}$ and $R^{(\tilde{\mathbf{S}})} \leq R^{(\mathbf{S})}$. \square

Set the index of the freshest arrival as $I_t := \max\{\tau \leq t : X_\tau = 1\}$. Observe that Lemmas 4.1 and 4.2 imply

Corollary 4.1. *For any \mathcal{S} , there exists a strictly increasing modification $\tilde{\mathcal{S}}$ such that $\tilde{S}_t = I_t$ or $\tilde{S}_t = ?$ and $R(\tilde{\mathcal{S}}) \leq R(\mathcal{S})$, $\Delta(\tilde{\mathcal{S}}) \leq \Delta(\mathcal{S})$.*

Corollary 4.1 tells that one should examine the procedures that label only the freshest arrival. Let \mathcal{S}_F be the space of such labeling procedures. The following theorem gives a lower bound to $CR + \Delta$, $C > 0$ for a specific class (square-integrable, as in **Chapter 3**) of labeling functions and hence gives a lower bound to the boundary curve of feasible (R, Δ) region.

Theorem 4.1. *For $\{S_t\} \in \mathcal{S}_F$ such that $\sup_t E[\Delta_t^2] < \infty$,*

$$CR + \Delta \geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t E[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau]. \quad (4.4)$$

Proof. The machinery is very similar to the proof of Theorem 3.2. Since $\frac{1}{t} \sum_{\tau=1}^t \mathbb{1}\{S_t \neq ?\}$ is bounded for all t , we directly apply Reverse Fatou's lemma to the first term on the left-hand side and obtain

$$R \geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \Pr\{S_t \neq ?\}.$$

Since $\sup_t E[\Delta_t^2]$ is finite, $\sup_t E\left[\left(\frac{1}{t} \sum_{\tau=1}^t \Delta_t\right)^2\right]$ is also finite and constitutes a uniformly integrable family; allowing the use of Reverse Fatou's lemma [83]. Thus, $\Delta \geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t E[\Delta_\tau]$. Lastly, we observe

$$\begin{aligned} CR + \Delta &\geq \limsup_{t \rightarrow \infty} \frac{C}{t} \sum_{\tau=1}^t \Pr\{S_t \neq ?\} + \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t E[\Delta_\tau] \\ &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t E[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau]. \end{aligned}$$

□

The expression on the right-hand side of (4.4) contains a summation whose τ^{th} term is \mathcal{F}_τ -measurable. This tells that for any $\{S_t\}$, this expression is the average reward (cost, in our case) of a Markov Reward Process with state space $\mathcal{B} := \{0, 1\}^*$ and the problem of choosing an appropriate labeling procedure $\{S_t\}$ is a Markov Decision Problem (MDP), which is formulated as finding the infimal limsup average cost

$$\begin{aligned} \lambda^* &:= \inf_{\{S_t\} \in \mathcal{S}_F} J^{\{S_t\}}(C), \text{ where} \\ J^{\{S_t\}}(C) &:= \limsup_t \frac{1}{t} \sum_{\tau=1}^t E[C\mathbb{1}\{S_\tau \neq ?\} + \Delta_\tau]. \end{aligned}$$

In the current problem formulation, and given the definition of the labeling procedures at the beginning of this manuscript, the states and actions seem to be complicated. However, Corollary 4.1 tells that it is sufficient to consider only two actions: (i) label the freshest arrival or (ii) wait, which we denote by l and w respectively. Since all arrivals before I_t are ignored at time t , we can reduce the state space to binary strings \mathbf{B}_t of length $t - M_t$, with its $(I_t - M_t)^{\text{th}}$ element being 1 and its other elements being 0, e.g., [000100].

Remark 4.1. *At this point, we have not imposed that the labeling strategies depend only on the current state \mathbf{B}_t . The procedures can depend on the whole past (also called history-dependent). Thus, one may argue that by reducing the state space, the history may not be recovered. However, this is not true as one is able to construct X_1, \dots, X_t from $\mathbf{B}_1, \dots, \mathbf{B}_t$.*

Note that \mathbf{B}_t 's are binary strings which contain at most a single 1. We can represent such a string by a pair (m, n) of non-negative integers, where m is the number of elements from the beginning of the string until and including the 1, and n is the remaining number of zeros. For instance, the buffer content [000100] becomes $(4, 2)$; and [000] becomes $(0, 3)$. One can also view n as the instantaneous age of the freshest arrival, and m as the time difference between the most recent labeling and the freshest arrival plus one.

Once more, as in **Chapter 3**, we encounter a *countable-state average-cost MDP* [81]. In the next section, we give a finite-state approximation to the problem with an aim to use the methods for finite-state problems; and as we will see, the finite-state formulation fortunately allows us to characterize the optimal strategies for the countable-state model as well.

4.2 Finite-State Approximation

Assume ‘phantom’ arrivals are generated whenever $m + n = L$, with no sending cost, i.e., do not contribute to R . With a similar proof as in Lemma 4.2, one can show that if the length of the buffer reaches L , procedures that label the phantom arrival will have a smaller Δ compared to the procedures that wait instead. Hence, the boundary curve of pairs pertaining to such procedures will lie under the original (R, Δ) curve. This truncated problem is finite state. Furthermore, since the state $(0, L)$ is recurrent under any policy — because the Geometric distribution has infinite tail and thus the buffer length reaches L infinitely often — the problem is unichain, i.e., every policy induces a Markov Chain with a single recurrent class [81]. Define

$$\lambda_L^* := \inf_{\mathcal{S} \in \mathcal{S}_{F,L}} J^{(\mathcal{S})}(C)$$

where $\mathcal{S}_{F,L}$ is the set of labeling procedures that only label the freshest arrival and always label the phantom arrival. Observe that λ_L^* is non-decreasing with L . Moreover, the sequence $\{\lambda_L^*\}$ has a limit. To see this, take the strategy

‘label every point upon arrival’ in the untruncated problem. The average cost corresponding to this strategy will be $\frac{C}{E[Z]} + \frac{E[Z(Z+1)]}{2E[Z]} = Cp + 1/p$. Then, $\lambda_L^* \leq Cp + 1/p$ for all L . Since $\{\lambda_L^*\}$ is bounded from above and is non-decreasing, it has a limit which we denote by λ_∞^* .

At this moment, the problem has become finite state and unichain; and it is known that there exists an optimal stationary policy for such problems. One may therefore focus on stationary policies and their evaluation methods. A stationary policy $s : \mathbb{N} \times \mathbb{N} \rightarrow \{l, w\}$ in our problem is evaluated by solving for λ, \mathbf{h} in the system of linear equations below. [81]

$$h(m, n) + \lambda = \begin{cases} m + n + C \\ + ph(n + 1, 0), & s(m, n) = l, m + n < L, m \geq 1 \\ + qh(0, n + 1) \\ m + n \\ + ph(m + n + 1, 0), & s(m, n) = w, m + n < L \\ + qh(m, n + 1) \\ L + ph(1, 0) \\ + qh(0, 1), & m + n = L \end{cases} \quad (4.5)$$

where $q := 1 - p$. We choose the state $(1, 0)$, i.e., the buffer content [1], as the reference state and set $h(1, 0) = 0$. λ gives the average cost of the unichain stationary labeling policy and $h(m, n)$ is called relative value of the state (m, n) [81].

Remark 4.2. $h(m, n) + \lambda$ is equal to the one-step cost plus the expected relative value of the next state depending on the action. E.g., if $s(m, n) = l$, then the one-step cost is $m + n + C$ and the next state will be $(n + 1, 0)$ with probability p and $(0, n + 1)$ with probability q . Thus for any policy, the state transitions are inferred from (4.5).

Let us recall the policy updates — given in Algorithm 1. Choose a stationary policy $s(m, n)$, evaluate it by solving (4.5) and obtain the relative values $h(m, n), m + n \leq L$ together with the average cost λ . Given the relative values, take a state $(m_0, n_0), m_0 + n_0 < L, m_0 \geq 1$ and consider the policy

$$s'(m_0, n_0) = \begin{cases} l, & C + ph(n_0 + 1, 0) + qh(0, n_0 + 1) \\ & \leq ph(m_0 + n_0 + 1, 0) + qh(m_0, 1) \\ w, & \text{else} \end{cases} \quad (4.6)$$

and $s(m, n) = s'(m, n)$ for all other states. Now, solve (4.5) with respect to $s'(m, n)$ to obtain λ' . It is known that $\lambda' \leq \lambda$ [81], i.e., the policy s is updated to a better policy s' . In fact, if one does the above procedure not only for (m_0, n_0) , but also for every possible state, and then solves (4.5), the procedure is the well-known policy iteration algorithm given in Algorithm 1 of chapter 3.

We aim to characterize the optimal strategies by using policy updates as a tool. Instead of direct application of the generic policy iteration algorithm, we consider the procedure described above where at k^{th} step we choose a single state (m_k, n_k) , update $s(m_k, n_k)$ and solve (4.5). Denote the average cost at the end of k^{th} step as $\lambda_L^{(k)}$, denote the updated policy and the relative values as $\mathbf{s}^{(k)}$ and $\mathbf{h}^{(k)}$ respectively.

Start the procedure with the initial policy ‘label every point upon arrival’. That is, $s^{(0)}(m, n) = l$ if $m > 0$ and $n = 0$; otherwise it is equal to w . Note that $s^{(0)}(0, 0) = w$ and although the state $(0, 0)$ will never be encountered according to our formulation, it provides convenience in the description of the procedure. The average cost corresponding to this policy is $\lambda_L^{(0)} \leq Cp + 1/p$.

At k^{th} step, choose $(m_k, n_k) = (k, 0)$. Apply the update rule in (4.6) to obtain

$$s^{(k)}(k, 0) = \begin{cases} l, & C \leq d^{(k-1)}(k, 0) \text{ or } k = L \\ w, & \text{else} \end{cases} \quad (4.7)$$

where

$$d^{(k-1)}(k, 0) := \left(ph^{(k-1)}(k+1, 0) + qh^{(k-1)}(k, 1) \right) - \left(ph^{(k-1)}(1, 0) + qh^{(k-1)}(0, 1) \right). \quad (4.8)$$

In fact, without solving the linear system in (4.5), it is possible to calculate $d^{(k-1)}(k, 0)$. Repeated application of (4.5) with $h^{(k-1)}(1, 0) = 0$ yields

$$d^{(k-1)}(k, 0) = E[G^{(k,0)}] - \lambda_L^{(k-1)} E[T^{(k,0)}] \quad (4.9)$$

where $T^{(k,0)}$ is the time until return to the reference state under policy $\mathbf{s}^{(k-1)}$ if we start from $(k, 0)$ and opt not to label; and $G^{(k,0)}$ is the accumulated cost until the return. $T^{(k,0)}$ has a truncated Geometric distribution, i.e., $T^{(k,0)} = \min\{T, L-k\}$ for a geometrically distributed T with parameter p . Observe that given $T^{(k,0)} = t^{(k,0)}$, the accumulated cost $G^{(k,0)}$ will be equal to $\sum_{\tau=k+1}^{t^{(k,0)}+k} \tau + \mathbb{1}\{t^{(k,0)} < L-k\}C$. Therefore, the expectations above can be calculated straightforwardly and we obtain

$$d^{(k-1)}(k, 0) = \frac{k}{p} + \frac{1}{p^2} - \frac{\epsilon_k}{p^2} - \frac{L\epsilon_k}{p} + (1 - \epsilon_{k+1})C - \frac{\lambda_L^{(k-1)}}{p}(1 - \epsilon_k) \quad (4.10)$$

where $\epsilon_k := q^{L-k}$. The update rule (4.7) is then equivalent to

$$\lambda_L^{(k-1)}(1 - \epsilon_k) \stackrel{l}{\leq} k + \frac{1}{p} - \frac{\epsilon_k}{p} - L\epsilon_k - Cp\epsilon_{k+1}. \quad (4.11)$$

Continue the procedure until $s^{(K_L)}(K_L, 0) \neq s^{(K_L+1)}(K_L+1, 0)$ for the first time for some $K_L \geq 0$. We are interested in large L as we want to make the ‘phantom’ arrivals as rare as possible. Thus, the limiting behavior of K_L is of interest.

Lemma 4.3. $\lim_{L \rightarrow \infty} K_L =: K$ exists and is finite. Furthermore,

$$K = \left\lceil \frac{-1 + \sqrt{1 + 8C + \frac{4q}{p^2}}}{2} - \frac{1}{p} \right\rceil. \quad (4.12)$$

Proof. See Appendix 4.5.2 □

Since K_L 's are integral, Lemma 4.3 implies the existence of an L_1 such that for all $L > L_1$, $K_L = K$. From now on, assume $L > L_1$ so that the termination time of our procedure is K . The next step is to show that for $m > K$ and $n = 0$, the policy will not be updated if L is large enough.

Lemma 4.4. *There exists an L_2 such that for all $L > L_2$, one obtains a worse policy by the modification $\tilde{s}^{(K)}(m, n) = w$ for any (m, n) such that $m > K$ and $n = 0$.*

Proof. See Appendix 4.5.3 □

At this point, we have covered the states given in the region (a) of Figure 41. It only remains to find the optimal actions for states (m, n) with $n > 0$, i.e., regions (b) and (c) in Figure 41. Our procedure did not modify the actions of these states. Hence, $s^{(K)}(m, n) = w$ for $n > 0$. Now we show that the actions of region (b) should remain unchanged.




Figure 41 – States as (m, n) pairs covered until different steps of our analysis. (a) corresponds to the darkest shaded region and includes the states $n = 0$ and $m + n = L$, which are covered with Lemma 4.4. (b) corresponds to the slightly shaded region including the states $n > 0$, $m \leq K$, $m + n < L$, which are covered with Lemma 4.5. The moderately shaded region (c) corresponds to the remaining states and these states turn out to be transient according to the Markov Chain induced by $\mathbf{s}^{(K)}$.

Lemma 4.5. *There exists an L_3 such that for all $L > L_3$, one obtains a worse policy by the modification $\tilde{s}^{(K)}(m, n) = l$ for any (m, n) such that $m \leq K$ and $n > 0$.*

Proof. See Appendix 4.5.4 □

Lemma 4.5 points out an unintuitive result: If an arrival is not labeled upon arrival, it will remain unlabeled. For instance, suppose an arrival occurred at $K - 1$. The optimal procedure skips it and waits for the next arrival even though it occurs very late.

Together with Lemma 4.5 we have covered the regions (a) and (b). Now observe that for the policy $\mathbf{s}^{(K)}$, the states in region (c) are transient since these states cannot be reached from any other state — check the state transitions in light of Remark 4.2 to see that transitions from regions (a) and (b) are to (a) and (b). Therefore, no matter what action is taken at a transient state in region (c), the average cost remains the same. This completes the proof that an optimal strategy is indeed $\mathbf{s}^{(K)}$ as any of its modification results in a higher cost.

Let us summarize what we have shown so far: There exists an $L' = \max\{L_1, L_2, L_3\}$ such that for all $L > L'$, the optimal strategy is

$$s^{(K)}(m, n) = \begin{cases} l, & m > K, n = 0 \\ w, & \text{else} \end{cases} \quad (4.13)$$

to which we shall refer as ‘wait K , label next’ strategy. One can easily calculate λ_L^* for $L > L'$ as

$$\lambda_L^* = \frac{\frac{K(K+1)}{2} + \frac{K}{p} - \left(\frac{L}{p} + \frac{1}{p^2}\right)\epsilon_K + \frac{1}{p^2} + C(1 - \epsilon_{K+1})}{K + \frac{1}{p}(1 - \epsilon_K)}$$

and thus (recall $\epsilon_k = q^{L-k}$)

$$\lambda_\infty^* = \frac{K^2 + (2/p + 1)K + 2/p^2 + 2C}{2(K + 1/p)}.$$

Since $\lambda_L^* \leq \lambda^*$ for all L , $\lambda_\infty^* \leq \lambda^*$. Moreover, λ_∞^* can be achieved via ‘wait K , label next’ strategy in the untruncated problem. Therefore $\lambda_\infty^* \geq \lambda^*$, and thus $\lambda_\infty^* = \lambda^*$.

The ‘wait K , label next’ strategy satisfies the square integrability condition in Theorem 4.1, i.e., $\sup_t E[\Delta_t^2] < \infty$. To see this, observe that at time t , the previous labeled arrival must have arrived at most $K + Z$ time slots ago. Hence, we obtain $E[\Delta_t^2] \leq E[(K + Z)^2] = (K + \frac{1}{p})^2 + \frac{q}{p^2} < \infty$ for all t . Furthermore, this strategy is stationary and constitutes a renewal process, which implies $R = \lim_t \frac{1}{t} \sum_{\tau=1}^t E[C \mathbb{1}\{S_\tau \neq ?\}]$ and $\Delta = \lim_t \frac{1}{t} \sum_{\tau=1}^t E[\Delta(\tau)]$, thus $CR + \Delta = \lambda^*$. This finally proves

Theorem 4.2. *The boundary of the feasible region of (R, Δ) pairs is characterized as the interpolation of the set of pairs $\{(R_k, \Delta_k)\}_{k \in \mathbb{N}}$ where*

$$\Delta_k = \frac{k^2 + (2/p + 1)k + 2/p^2}{2(k + 1/p)}$$

$$R_k = \frac{1}{k + 1/p}.$$

As a final remark, we point out that the pairs $\{(R_k, \Delta_k)\}_{k \in \mathbb{N}}$ are achievable with ‘wait k , label next’ strategies, and the remaining pairs on the boundary curve are achievable with time sharing between at most two of such strategies, e.g., for $0 < \rho < 1$, do ‘wait k , label next’ ρ fraction of the time and do ‘wait $k + 1$, label next’ $1 - \rho$ fraction of the time.

4.3 Extension To Poisson Processes

In this section, we extend the results obtained for the discrete-time problem to a continuous-time problem: The arrivals are modeled as a Poisson process. Let $\mathcal{N}(t)$ be the counting process associated with a stationary Poisson process of intensity ν , with its natural filtration $\{\mathcal{F}_t\}_{t \geq 0}$. Similar to the discrete case, the causal labelings for the continuous model are defined as the collection of functions $\{S_t\}_{t \in \mathbb{R}^+}$, such that every S_t is \mathcal{F}_t -measurable and $S_t : \mathcal{P}^{[0,t]} \rightarrow [0, t] \cup \{?\}$ where $\mathcal{P}^{[0,t]}$ denotes the space of the sample paths of $\mathcal{N}(t)$ on the interval $[0, t]$. Observe that the process that tracks the number of labelings until t constitutes another counting process. Denote this process by $\mathcal{R}(t)$. Note that $\mathcal{R}(t) \leq \mathcal{N}(t)$.

The rate and the average age are defined analogously to the discrete-time case — see (4.1) and (4.2) — namely

$$\begin{aligned} \mathcal{R} &:= E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \mathcal{R}(t) \right] \text{ and} \\ \delta &:= E \left[\limsup_{t \rightarrow \infty} \frac{1}{t} \int_{\tau=0}^t \Delta_\tau d\tau \right]. \end{aligned} \tag{4.14}$$

Recall that the proofs of Lemma 4.1 and Lemma 4.2 involve pathwise coupling arguments and can be extended to continuous-time straightforwardly. Therefore Lemma 4.1, Lemma 4.2 and thus Corollary 4.1 hold for the continuous-time model as well. Once more, this restricts the class of labelings to the set \mathcal{S}_F : The procedures that only label the freshest arrival. In the following theorem, we extend the results of the discrete model and show that ‘wait T , label next’ type of strategies are optimal for the continuous-time problem as well.

Theorem 4.3. *For $\{S_t\}_{t \in \mathbb{R}^+} \in \mathcal{S}_F$ such that $\sup_{t \in \mathbb{R}^+} E[\Delta_t^2] < \infty$, the boundary curve of achievable (\mathcal{R}, δ) pairs lie on the curve*

$$\delta(\mathcal{R}) = \frac{1}{2\mathcal{R}} + \frac{\mathcal{R}}{2\nu^2}, \quad \mathcal{R} \leq \nu \tag{4.15}$$

which are achieved with ‘wait T , label next’ strategies with $T = \frac{1}{\mathcal{R}} - \frac{1}{\nu}$.

Proof. We discretize the time axis by dividing it into small intervals of length h . Then, for any strategy $\{S_t\} \in \mathcal{S}_F$, consider a modification that only makes decisions at times that are multiples of h . Suppose the continuous-time strategy $\{S_t\}$ has made k labelings at times l_1, \dots, l_k on the interval $((n-1)h, nh]$; and

Δ_t drops down to values a_1, \dots, a_k respectively for each labeling. We can infer that the labeled arrivals have occurred at times $t_1 = l_1 - a_1, \dots, t_k = l_k - a_k$. Observe that only the first labeled arrival can belong to a past interval; otherwise it would not be the freshest arrival. This implies that $l_i - t_i \leq h$ for $i > 1$. Suppose the discrete-time modification labels only the k^{th} arrival at time nh and its instantaneous age $\tilde{\Delta}_n$ drops down to $\lfloor \frac{nh - (l_k - a_k)}{h} \rfloor h$. If $k > 1$, $\tilde{\Delta}_n = 0$, hence smaller than Δ_t . If $k = 1$, then $\tilde{\Delta}_n = \lfloor \frac{nh - (l_1 - a_1)}{h} \rfloor h \leq \lfloor \frac{h + a_1}{h} \rfloor h \leq a_1 + h \leq \Delta_t + h$ for all $t \in ((n-1)h, nh]$. Hence, the discrete age can at most be h higher than the continuous age for all t , yielding

$$\frac{1}{t} \int_{\tau=0}^t \Delta_\tau d\tau \geq \frac{1}{(N+1)} \sum_{n=1}^N (\tilde{\Delta}_n - h) \tag{4.16}$$

where $N = \lfloor t/h \rfloor$. Then, taking limsup on both sides, we have

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_{\tau=0}^t \Delta_\tau d\tau \geq \limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \tilde{\Delta}_n - h. \tag{4.17}$$

Also observe that the number of labelings made by the discrete-time modification is always smaller than $\mathcal{R}(t)$ — as it only labels the k^{th} arrival. Proceeding similarly as above, we have

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \mathcal{R}(t) \geq \limsup_{N \rightarrow \infty} \frac{1}{Nh} \sum_{n=1}^N R_n^{(h)} \tag{4.18}$$

where $R_n^{(h)} := \mathbb{1}\{\exists t \in ((n-1)h, nh] : S_t \neq ?\}$. Using Reverse Fatou's lemma once more as in Theorem 1, we obtain

$$C\mathcal{R} + \delta \geq \limsup_{N \rightarrow \infty} \frac{1}{Nh} \sum_{n=1}^N E[\tilde{\Delta}_n h + CR_n^{(h)}] - h. \tag{4.19}$$

The limsup expression above can be lower bounded by the average cost of a very similar problem that we have considered in discrete-time. Recall that the discrete-time modification is able to label from an interval as long as there exists at least one arrival. This implies that the interarrival times of the discrete model are Geometrically distributed with parameter $p = 1 - e^{-\nu h}$, which is equal to the probability that at least one arrival occurs on an interval of length h . Furthermore, the discrete-time modification is able to label from the same interval more than once. Hence, comparing with our discrete-time formulation in Section 4.1, we see that $R_n^{(h)} \geq \mathbb{1}\{S_n \neq ?\}$, where S_n is defined as in Section 4.1. The expression above is then lower bounded by h times the optimal average cost of an MDP with

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N E[\Delta_{n+1} - 1 + \frac{C}{h^2} \mathbb{1}\{S_n \neq ?\}]. \tag{4.20}$$

where Δ_n is defined as in our discrete-time formulation in Section 4.1. The crucial difference is the first term, which is a time shifted version of the original

problem. With similar truncation arguments as in Section 4.2, we can analyze the modified problem and study its optimal stationary policies. However, observe that any stationary and unichain policy of the modified problem can be simulated with our original problem — a time shift of instantaneous ages does not change the long-term average. Hence, the optimal policy will exactly be the optimal policy of the original problem. Now write down the optimal average cost and do the appropriate scaling to obtain

$$\begin{aligned} \delta + C\mathcal{R} &\geq \lim_{h \rightarrow 0} h \frac{K^2 + (2/p + 1)K + 2/p^2 + 2C/h^2}{2(K + 1/p)} - 2h \\ &= \lim_{h \rightarrow 0} \frac{T^2 + 2T/\nu + 2/\nu^2 + Th + 2C}{2(T + 1/\nu)} \end{aligned} \quad (4.21)$$

where $T := hK(h)$. We have written $K(h)$ to emphasize its dependence in h . Substituting C/h^2 and $p = 1 - e^{-\nu h}$ in (4.12), we see that T tends to $\sqrt{2C + \frac{1}{\nu^2}} - \frac{1}{\nu}$ as $h \rightarrow 0$, which is a constant. Hence,

$$\delta + C\mathcal{R} \geq \frac{T^2 + 2T/\nu + 2/\nu^2 + 2C}{2(T + 1/\nu)}. \quad (4.22)$$

Similar to the end of the proof of Lemma 4.2, we argue that ‘wait T and label’ strategy is stationary, satisfies the square integrability condition and attains the average cost above; concluding that the above is an equality. \square

4.4 Discussion

In this chapter, we have characterized the optimal labeling strategy to reduce the output rate for memoryless processes: ‘Wait T , sample next’. The main difficulty of the problem and also the main difference from energy harvesting scenarios is that once a newly arrived packet is not labeled, it is not immediately discarded and kept in the buffer for a possible labeling in the future. E.g., if no new arrivals occur for a long time, the sender may opt to label the packet in its buffer and forward it through the network. For memoryless processes, it turned out that such buffering is unnecessary — an optimal strategy should either label the packet or discard it upon its arrival. However, this may not be true if the packet arrival process has some memory, i.e., the interarrival times are not geometric (or Poisson). For generic interarrival times, the problem becomes complex and neither simple strategies nor closed-form expressions are expected.

4.5 Appendix

4.5.1 Extension of Example 2 to Rational r :

Let $r = \frac{m}{n}$ with $m \leq n$. Write $n = m\alpha + \beta$, where $\alpha \geq 1$ is the quotient and β is the remainder. Note that $\beta < m$. Write r as

$$r = \frac{m}{(\alpha + 1)\beta + \alpha(m - \beta)}. \quad (4.23)$$

Consider the following strategy: (i) Label every $\alpha + 1$ th arrival, do this β times. Then (ii) label every α th arrival, do this $m - \beta$ times. Then repeat (i) and (ii) consecutively. By renewal theory, the average rate will be pr . Also note that the interarrival times are sum of β Geometric random variables with parameter $\frac{p}{\alpha+1}$ plus sum of $m - \beta$ Geometric random variables with parameter $\frac{p}{\alpha}$. With some algebra, one can obtain

$$\Delta = \frac{r(\alpha + 1)}{2R} \left(1 + \frac{\beta r}{m}\right). \quad (4.24)$$

To check if this expression is smaller than $1/R$, observe

$$\begin{aligned} \frac{r(\alpha + 1)}{2R} \left(1 + \frac{\beta r}{m}\right) &= \frac{1}{2R} \frac{m}{n} (\alpha + 1) (1 + \beta/n) \\ &= \frac{1}{2R} \frac{n - \beta + m}{n} (1 + \beta/n) \\ &= \frac{1}{2R} \left(1 - \frac{\beta}{n} + \frac{m}{n}\right) \left(1 + \frac{\beta}{n}\right) \\ &= \frac{1}{2R} \left(1 - \frac{\beta^2}{n^2} + \frac{m}{n} + r \frac{\beta}{n}\right) \\ &\leq \frac{1}{2R} \left(1 + \frac{m + r\beta}{n}\right) \\ &\leq \frac{1}{2R} (1 + 1) \\ &= \frac{1}{R} \end{aligned} \quad (4.25)$$

as $n = m\alpha + \beta$ and $\alpha \geq 1$ and $r \leq 1$.

4.5.2 Proof of Lemma 4.3

K_L is defined as the minimum non-negative integer satisfying $\lambda_L(k) < g_L(k)$, where

$$\lambda_L(k) = \frac{\frac{k(k+1)}{2} + \frac{k}{p} - \left(\frac{L}{p} + \frac{1}{p^2}\right)\epsilon_k + \frac{1}{p^2} + C(1 - \epsilon_{k+1})}{k + \frac{1}{p}(1 - \epsilon_k)} \quad (4.26)$$

is the average cost pertaining to the ‘wait k , label next’ strategy and

$$g_L(k) := \frac{k + 1 + \frac{1}{p} - \frac{\epsilon_{k+1}}{p} - L\epsilon_{k+1} - Cp\epsilon_{k+2}}{(1 - \epsilon_{k+1})}. \quad (4.27)$$

Define $g_L(L-1) = \infty$ for convenience. This ensures that $K_L \leq L-1$. Observe that for any finite $k \geq 0$ and small ϵ , we can find large enough L such that a sufficient condition for the inequality $\lambda_L(k) < g_L(k)$ can be obtained as

$$k^2 + \left(\frac{2}{p} + 1\right)k + \frac{2}{p^2} + 2C < 2\left(k + \frac{1}{p}\right)\left(k + 1 + \frac{1}{p}\right) - \epsilon \quad (4.28)$$

which is equivalent to

$$\epsilon < \left(k + \frac{1}{p}\right)^2 + \left(k + \frac{1}{p}\right) - 2C - \frac{1}{p^2} + \frac{1}{p} =: f_L(k). \quad (4.29)$$

Since K_L is the smallest integer satisfying the above, we must have $\lambda_L(k-1) \geq g_L(k-1)$. Again, for large enough L , an equivalent sufficient condition will be

$$f_L(k-1) \leq -\epsilon. \quad (4.30)$$

$f(k)$ is a quadratic function of k and it always has two distinct real roots. The smaller root is always negative, hence take the larger root

$$\tilde{k} = \frac{-1 + \sqrt{1 + 8C + \frac{4q}{p^2}}}{2} - \frac{1}{p} \quad (4.31)$$

and note that for small enough ϵ both sufficient conditions are satisfied at $k = \lceil \tilde{k} \rceil$. This shows that $K_L \leq \lceil \tilde{k} \rceil$ for large enough L , which implies $\limsup_L K_L \leq \lceil \tilde{k} \rceil$. Furthermore, for $0 < C \leq \frac{1}{p}$, $-1 < \tilde{k} \leq 0$ and thus $\lceil \tilde{k} \rceil = 0$, so $\lim_L K_L = 0$ for such C . For $C > \frac{1}{p}$, with a similar argument, obtain necessary conditions by choosing large enough L and small enough ϵ ; which are

$$f(k) > -\epsilon \text{ and } f(k-1) \leq \epsilon. \quad (4.32)$$

Observe that $f(\lceil \tilde{k} \rceil - 1)$ is negative and bounded away from zero — as $\tilde{k} > 0$ and the smaller root is always negative. Hence the necessary conditions will not be satisfied for $\lceil \tilde{k} \rceil - 1$ and $K_L > \lceil \tilde{k} \rceil - 1$ eventually. This concludes that $\liminf_L K_L \geq \lceil \tilde{k} \rceil$ and therefore $\lim_L K_L = \lceil \tilde{k} \rceil =: K$.

We remark that $K \leq \lfloor Cp \rfloor$. This fact will be important when proving Lemma 4.4 and Lemma 4.5.

4.5.3 Proof of Lemma 4.4

This is equivalent to showing that the update condition

$$\lambda_L^{(K)}(1 - \epsilon_m) \geq m + \frac{1}{p} - \frac{\epsilon_m}{p} - L\epsilon_m - \frac{Cp}{q}\epsilon_m \quad (4.33)$$

is violated for all $K < m < L$. It is sufficient to check if the minimum of the function

$$f(t) = L + \frac{1}{p} + \frac{\log t}{\log \frac{1}{q}} - \lambda_L^{(K)}(1 - t) - \left(L + \frac{1}{p} + \frac{Cp}{q}\right)t \quad (4.34)$$

on $t \in [q^{L-K-1}, q]$ is greater than zero. f is concave, therefore the minimum occurs at boundaries. Observe $f(q) = Lp - Cp - p\lambda_L^{(K)} \geq Lp - Cp - p(Cp + \frac{1}{p})$ therefore greater than zero for $L \geq L_2 := C + (Cp + \frac{1}{p})$. The other boundary corresponds to $m = K + 1$, and from the definition of K as being the termination time of the procedure — check the definition in the beginning of Appendix 4.5.2 — we must have $f(q^{L-K-1}) \geq 0$. Thus, the condition above is violated for all $m > K$.

4.5.4 Proof of Lemma 4.5

The condition in (4.6) implies that if we alter $s^{(K)}(m, n)$ to l , the policy does not improve if

$$C + ph(n + 1, 0) + qh(0, n + 1) \geq ph(m + n + 1, 0) + qh(m, n + 1). \quad (4.35)$$

Some calculation reveals that this condition is equivalent to

$$\begin{aligned} & C + \sum_{\tau=n+1}^K \tau + \sum_{\tau=K \vee n}^{L-1} (\tau + 1)q^{\tau - K \vee n} + C(1 - q^{L - K \vee n}) \\ & - \lambda \left(\sum_{\tau=n+1}^K 1 + \sum_{\tau=K \vee n}^{L-1} q^{\tau - K \vee n} \right) \\ & \geq \sum_{\tau=m+n+1}^K \tau + \sum_{\tau=K \vee (m+n)}^{L-1} (\tau + 1)q^{\tau - K \vee (m+n)} + C(1 - q^{L - K \vee (m+n)}) \\ & - \lambda \left(\sum_{\tau=m+n+1}^K 1 + \sum_{\tau=K \vee (m+n)}^{L-1} q^{\tau - K \vee (m+n)} \right) \end{aligned} \quad (4.36)$$

Now, consider the three cases (for which $m \leq K$ according to the main statement of the Lemma):

(i) $n > K$

Then (4.36) is equivalent to

$$\frac{Cp - m}{q^{L-m-n} - q^{L-n}} + L + \frac{1}{p} + Cp \geq \lambda_L^{(K)}. \quad (4.37)$$

Recall that $Cp \geq K$ — see the end of Appendix 4.5.2 — hence the first term on the left-hand side is positive. We also know that $\lambda_L^{(K)} \leq Cp + 1/p$ as the latter is the average cost of the untruncated problem with ‘label every point upon its arrival’ policy. Hence the inequality always holds.

(ii) $n \leq K, m + n > K$

Observe that the (m, n) pairs lie on a bounded set that does not grow with L . Hence, for large enough L , the condition will be equivalent to

$$C + (K - n)(K + n + 1)/2 - (m + n - K)/p \geq \lambda(K - n) \quad (4.38)$$

the left-hand side is minimized at $m = K$, which yields

$$C + (K - n)(K + n + 1)/2 - n/p \geq \lambda(K - n). \quad (4.39)$$

Since the left-hand side is a concave quadratic function of n , it is minimized at either $n = 0$, or $n = K$. For $n = 0$, the state corresponds to $(K, 0)$ and we know the inequality holds from definition of K . For $n = K$, we have

$$C - K/p \geq 0 \quad (4.40)$$

which we know is true.

(iii) $n \leq K$, $m + n \leq K$. Again, the possible pairs lie in a bounded set that does not grow with L . For large L , rewrite the condition as

$$C \geq \lambda m - m(m + 2n + 1)/2. \quad (4.41)$$

The right-hand side is maximized at $n = 0$, and similar to (ii) the result is immediate from the definition of K .

As K , m , n are integers, there also exists an L_3 such that the conditions in (ii), (iii) are exact for $L > L_3$.

Part I Conclusion

5

In Part I, we explored two closely related problems concerned with data freshness and importance, and the results of **Chapter 3** support the intuition that ‘the freshest data may not be the most important’. More specifically, we studied the tradeoff between the age and the distortion in a discrete-time setting where the output rate was limited and we characterized the optimal tradeoffs by resorting to the theory of Markov decision processes (MDP).

In the first problem (**Chapter 3**), the output rate was limited by an external scheduler. As a consequence, we could study the tradeoff between the age and the distortion by an appropriate MDP formulation of the problem. We showed that we could efficiently find the optimal strategies to attain the tradeoff in low-age regime. However, at high-age regime, the necessary buffer size to attain the tradeoff grows and the algorithm to find the optimal strategies becomes time consuming. Also, if the number of importance levels increase, the time complexity degrades significantly. Therefore, for applications that include a multitude of importance classes and at high-age regime, one may resort to approximation methods for MDPs, e.g., to Monte-Carlo simulations or to various reinforcement learning techniques. We later explored a similar problem where the packets were prone to erasures and showed that the optimal strategies for this new problem are the same. Towards the end of the chapter, we studied the case where the packet payloads were small and reformulated the problem such that the sender was able to send N bits at a speaking time, instead of a complete packet. The packet-based strategies can be significantly outperformed by ad-hoc strategies (such as Buffer ignorant) and with (Tunstall) coding. However, the optimal tradeoff for all strategies that allow coding remains as an important problem and could also be difficult to solve.

What if the sender is allowed to choose the speaking times as well? The second problem (**Chapter 4**), was formulated based on this question. However,

we assumed packets of equal importance, leaving us with the output rate-age, i.e., (R, Δ) , tradeoff. The sender was allowed to send (i) any packet from its past at (ii) any time instant. This implies that if a packet is not sent upon its arrival, it can be sent at any later time, and this broad class of strategies made the problem complex. Luckily, we showed that if the packets arrive according to a memoryless process, simple threshold-based strategies are optimal and a closed form representation of the optimal (R, Δ) curve can be obtained. Furthermore, the optimal strategies only choose a packet upon its arrival and can be characterized as ‘wait T , label next’. We also extended the results to a continuous-time setting where packets arrive as a Poisson process. If the packet arrival process is not memoryless, however, lookback strategies could be among the optimal ones, i.e., a packet could be stored in the buffer and sent at a later time. In this case, the optimal strategies may not be simple-to-describe.

Note that in both problems we studied two parameters at once. Namely, (Δ, D) and (R, D) . Finding the region of achievable (R, Δ, D) triplets under the discrete-time setting studied in this thesis seems challenging, and could be a future study.

Part II

**Strategies for Distributed
Inference**

Introduction to Distributed Inference

6

6.1 Distributed Inference

Decision making is a common task, both for living and non-living entities. An animal may decide whether to hunt or not, a human may decide whether to buy a stock or not; and as an example for an inanimate decision maker consider a router (though programmed by a human-being) that decides to forward a packet through a particular subnet. Even though these examples are binary decision rules, more complicated rules are also frequently encountered, e.g., to put a bid order for the stock at price P . Usually, the (rational) decision makers seek to maximize their utility and take the action accordingly. To this end, they exploit the information available to them through observations. For instance, a stock broker might read some news, and/or make use of the patterns of historical prices to obtain information; and choose the bid price P according to the risk tolerance.

As part of the decision making process, the decision maker usually aims to *infer* the true state-of-nature in order to take accurate actions. More specifically, the decision maker faces a *statistical inference* problem, which usually refers to the task of learning about a phenomenon of interest through observations whose statistics are governed by the unknown state-of-nature. The statistical inference problem is generally studied under two domains depending on the properties of the unknown state. If the unknown state is assumed to be continuous and real-valued entity, it is called an *estimation* problem; and if the state takes values in a finite set, it is called a *detection*, or *hypothesis testing* problem. Recall that we have already described a special case of a hypothesis testing setting in **Chapter 1**, where the decision maker has to choose either the null hypothesis \mathcal{H}_0 or the alternative hypothesis \mathcal{H}_1 . In the sequel, we will focus on the hypothesis testing problem rather than estimation.

Hypothesis testing is a highly non-trivial problem even in the single-agent case. In multi-agent decision making systems the problem gets even more complex. Some challenges of multi-agent decision problems arise from (i) information heterogeneity across agents, and (ii) dependence on the network structure, e.g., connectivity, hierarchies. To elaborate on the first challenge, suppose we, as a decision maker, have access to two different information sources, e.g., the news about a company and its stock price. The information from these sources are highly heterogeneous and how to combine the two observations to make a decision is a very complex task. For the second challenge, think of the following architectures: Peripheral nodes may send information to a central node (to be referred as the fusion center in the sequel), or the network can be fully decentralized without any central controller that aggregates information. For these two network settings, the hypothesis testing problem differs significantly. In the forthcoming chapters, we will study hypothesis testing problems for both settings.

In addition to the present complexity of the problem, the reader should also keep in mind that the processing capabilities of agents could be limited in reality. Moreover, the communication between agents must be limited as well — obviously a communication link cannot use infinite bandwidth. To illustrate this situation, consider a vehicle equipped with a collision avoidance system that relies on vehicular communication. In such a communication scheme, information can be received from other vehicles (vehicle-to-vehicle, V2V) or from other objects such as mobile phones, base stations etc. (vehicle-to-everything, V2X). The collision avoidance system is activated upon detection of a possible collision — this risky state might be associated with the alternative hypothesis \mathcal{H}_1 in a binary hypothesis testing setting. V2V and V2X communication protocols, e.g., IEEE 802.11p [87], limit the data rate. Hence, the devices in proximity are required to compress or quantize the data they possess before sending it to the vehicle. In this thesis we will focus on the communication link constraints and we will study two related problems:

- In **Chapter 7**, we will consider a centralized scheme where peripheral nodes transmit information to a fusion center, and the expected number of bits per unit time is limited under the null hypothesis \mathcal{H}_0 .
- In **Chapter 9**, we will study a fully decentralized scheme where agents are connected in a network topology; and we will study an algorithm where agents randomly poll their neighbors with an attempt to both decrease the communication rate and mitigate collisions that arise from nodes trying to send data simultaneously.

Now, let us make a short excursion to see how communication constraints affect the performance of a hypothesis test, and motivate the next chapter. In accordance with the hypothesis testing setting given in **Chapter 1**, suppose a remote node observes X_t at each time t , and further suppose that X_t 's are independent and identically distributed (i.i.d.) over time. The node has to




Figure 61 – A communication-constrained distributed hypothesis testing setup where the information has to be compressed before being conveyed to the fusion center.

convey information to a central node, which will make a decision after a fixed time horizon T . A visual representation could be as in Figure 61, where the node, and the fusion center is drawn blue and red respectively. Let the log-likelihood ratio (LLR) pertaining to X_t be L_t . Note that the LLR is a real number in general. Hence, due to the communication link in between having finite bandwidth, the exact LLR cannot be sent. For this reason, the node has to compress the LLR before passing it through the link. Observe that if there were no communication constraints, the center would perform an optimal, i.e. Neyman–Pearson, test by possessing the LLRs, and consequently employ the same achievability scheme in the proof of Theorem 1.10, which yields the optimal type-II error exponent $D(P||Q)$. Now, let us recall the data processing inequality.

Theorem 6.1. [Data Processing Inequality] *Let $w(\cdot|\cdot)$ be a probability kernel on \mathbb{R} , i.e., (i) for every Borel-measurable set Y , $w(Y|\cdot)$ is Borel-measurable, and (ii) for every $x \in \mathbb{R}$, $w(\cdot|x)$ is a probability measure in \mathbb{R} . Let $P \circ w$, and $Q \circ w$ denote the output distributions when P and Q are given as inputs respectively. Then,*

$$D(P||Q) \geq D(P \circ w||Q \circ w). \quad (6.1)$$

Note that in the above theorem w can also be a deterministic mapping. Hence,

$$D(P||Q) \geq D(P \circ f||Q \circ f). \quad (6.2)$$

In accordance with (6.2), the center suffers from a performance loss due to not possessing the exact LLRs. One expects that the more stringent the communication constraints are, the less the type-II error exponent will be. Hence, if the communication constraint is limited to R (in some sense), one might want to study a tradeoff function $\theta(R)$ that gives the best possible type-II error exponent with vanishing type-I error probability as in **Chapter 1**. Indeed, this is what we will do in **Chapter 7**.

Before studying the two aforementioned problems, let us proceed with presenting some related work on communication-constrained hypothesis testing from information theoretic and signal processing perspectives.

6.2 Related Work

We will see that our problem formulation in **Chapter 7** contains flavors from both information theoretic and signal processing approaches. We review the related work under these two approaches respectively.

6.2.1 Information Theoretic Approaches

Distributed hypothesis testing under communication constraints is a long-standing problem studied by the information theory community. An early work by Ahlswede and Csiszár [88] underlies most of the subsequent developments. It is therefore instructive to review their problem setup for a better understanding of the subsequent work. Their setup — henceforth referred to as the Ahlswede–Csiszár setup — is as follows. A remote node possesses a sequence X^n , while the decision maker possesses a Y^n . The pair (X^n, Y^n) is i.i.d. with distribution P under the null hypothesis (\mathcal{H}_0) and with distribution Q under the alternative hypothesis (\mathcal{H}_1). The decision maker estimates the true hypothesis by using both Y^n and an nR -bit side information conveyed by the remote node. The communication constraint is “hard” in the sense that X^n is represented with exactly nR bits under both hypotheses. Their aim is to find the fastest exponential decay rate of the type-II error given a prescribed type-I error probability, say $0 < \epsilon < 1$. It turns out that the fastest decay rate does not depend on ϵ , and it is fully characterized for the special case of dependence testing, i.e., when $Q_{XY} = P_X P_Y$ where P_X, P_Y are the marginals of X and Y under \mathcal{H}_0 . The characterization of the optimal decay rate for the general case turns out to be more involved and it is still unknown although some upper and lower bounds exist.

The Ahlswede–Csiszár setup motivated various subsequent works on distributed hypothesis testing. For instance, [89] presents tighter lower bounds on the optimal decay rate for the Ahlswede–Csiszár setup and further extends the formulation to include zero-rate compression (see also [90]), as well as to include the compression of Y^n . The lower bound on the best possible decay rate for the general case is improved in [91]. One may refer to [92] for a comprehensive survey on the literature considering Ahlswede–Csiszár setup and its variants. Subsequent works on communication-constrained hypothesis testing include studies on tradeoffs between type-I and type-II error exponents [93, 94], performance under finite-blocklength regime [95], and under noisy communication [96–99]. Further extensions of this problem include interactive protocols [100–102], privacy constraints [103–108], the additional task of data reconstruction at the receiver [109]. For dependence testing, [110] concludes that binning schemes are optimal; whereas the recent work [111] shows that the performance can be improved with sequential methods for the general case.

The works cited above elaborate on the “hard” communication constraints, as Ahlswede and Csiszár did. A recent strand of works relax the “hard” com-

munication constraints and study the dependence testing problem by limiting the *expected* number of bits sent. A partial list of the studies adopting this perspective is [112–117]. The next chapter’s study is also in line with this perspective. We remark that for the special case of dependence testing, since the X marginals are the same under both hypotheses, the expected number of bits conveyed does not depend on the true hypothesis. We will, on the other hand, focus on the general case. Hence, given a strategy, the expected number of bits sent might differ under the null and alternative hypotheses; introducing an inherent asymmetry to the problem. We will limit the *expected rate under the null hypothesis* \mathcal{H}_0 . This choice aligns with the view that \mathcal{H}_1 is a rare high-risk event and necessary communication must take place to detect this event with high probability. A more detailed discussion on such choice is given in Section 7.2.

Information theoretic approaches could be criticized because they use high-dimensional vector quantization, i.e., the entire block X^n should be observed before being represented with nR bits. A system designed as such may not be desirable for timing-crucial applications, as the decision maker is kept oblivious of the side information until time n . Furthermore, for large n , such a system is not memory-efficient as the remote node records the whole past and it might also be computationally expensive to compress X^n . These observations suggest that low-dimensional quantization could be of interest for low-latency and memory-efficient applications. Such quantization procedures for distributed detection are often studied in the signal processing literature.

6.2.2 Signal Processing Approaches

As mentioned, signal processing approaches are usually centered around low-dimensional quantizer designs. The scalar quantization procedures specialized for the task of binary hypothesis testing aim to keep the dissimilarity between the distributions of the quantizer output under \mathcal{H}_0 and \mathcal{H}_1 as large as possible while representing the output only with R bits. Various methods for evaluating the dissimilarity include calculation of the Kullback–Leibler divergence $D(\cdot||\cdot)$ — the optimal type-II error rate under vanishing type-I error [13] or vice versa — or one may consider the more general Ali–Silvey distances [118] (or equivalently f -divergences [119]) which prove useful for a variety of signal detection problems [120]. Notable early studies on quantization for binary hypothesis testing include [121–123].

Finding the optimal quantizer is in general a daunting task and there is no standard machinery to obtain such quantizers. However, there exists iterative methods to find suboptimal quantizers as in [123], or studies on the high-rate quantization regime [124–126]. Some extremal properties of likelihood-ratio quantizers is given in [127]. Quantizer designs based on privacy and secrecy constraints are studied in [128, 129]. Error resilient designs are studied in [130–132], as well as Byzantine resilient designs in [133]. A recent work on multilevel quantization is [134].

A similar trend to that in the information theoretic studies is also observed in the signal processing literature — the works cited above rely on “hard” communication constraints. Different from the existing signal processing literature, in the next chapter, we will study the fundamental limits under memoryless (scalar) quantization with *expected* rate constraints under \mathcal{H}_0 , and provide impossibility results for the subject case. Namely, we will see that if the *expected* rate under \mathcal{H}_0 is limited to R bits, then the type-II error rate cannot be greater than $\theta^*(R)$ — which we will define in the next chapter — under vanishing type-I error probability.

A Fundamental Limit of Distributed Hypothesis Testing Under Memoryless Quantization¹

7

In this chapter, we will study a canonical distributed hypothesis testing problem under communication constraints. Among many possible ways of restricting communication, we choose to limit the *average* number of bits sent under riskless or ordinary state, which associates with the null hypothesis; and we seek the fundamental limits of a distributed hypothesis testing problem under such assumption. We focus on the case where nodes compress their data with practically-appealing memoryless quantization procedures. More precisely, under such setting, we initially focus on the single-node case and when the average number of bits sent is at most R under the null hypothesis:

- we characterize the optimal decay rate of the type-II error probability under vanishing type-I error probability, given by $\theta^*(R)$, in Theorem 7.1 of Section 7.3.1;
- we obtain an upper bound to $\theta^*(R)$ via rate-distortion methods and consequently characterize an unachievable region in Corollary 7.1 of Section 7.3.2;
- we show that with simple lattice-quantization, the upper bound can be approached within $\frac{1}{2} \log_2(\frac{\pi e}{2}) \approx 1.047$ bits in Theorem 7.4 of Section 7.4.1;
- we provide the upper bound $\theta_k(R)$ for the k -dimensional vector quantization case in Section 7.4.2.

The results for the single-node case are then extended to multiple nodes in Section 7.5, where the problem is formulated under individual communication constraints at nodes, together with a sum-rate constrained formulation.

¹The content of this chapter is based on [135, 136].

7.1 Notation

Random variables are denoted with uppercase letters whereas their realizations are written lowercase, e.g., X_n and x_n . $\mathcal{B}(\mathbb{R})$ denotes the Borel algebra of \mathbb{R} . For probability measures P and Q , $D(P||Q)$ denotes the Kullback–Leibler (KL) divergence and $E_P[\cdot]$, $H_P(\cdot)$, $I_P(\cdot; \cdot)$ denote the expectation, entropy, and mutual information under P respectively. All logarithms are taken with natural base unless explicitly stated.

7.2 Problem Formulation

Consider m peripheral nodes that communicate with a fusion center (Figure 71). At each time instant t , the node i observes data arising from distribution $P^{(i)}$ under the null hypothesis \mathcal{H}_0 , and from distribution $Q^{(i)}$ under the alternative hypothesis \mathcal{H}_1 . We assume that for all i , $P^{(i)}$ is absolutely continuous with respect to $Q^{(i)}$, vice versa. That is, if $Q^{(i)}(B) = 0$, then $P^{(i)}(B) = 0$ for any $B \in \mathcal{B}(\mathbb{R})$; and the same holds when P is swapped with Q . In this case P and Q are also called *equivalent*, and denoted as $P \sim Q$. The data is independent across nodes, and across time under both hypotheses. Moreover, the data is identically distributed across time. Therefore, the joint distribution of the network until time t and under \mathcal{H}_0 can be characterized on rectangles in \mathbb{R}^{tm} as follows:

$$P(B^t) = \prod_{\tau=1}^t \prod_{i=1}^m P^{(i)}([a_\tau^{(i)}, b_\tau^{(i)}]) \quad (7.1)$$

where $B_\tau := [a_\tau^{(1)}, b_\tau^{(1)}] \times \cdots \times [a_\tau^{(m)}, b_\tau^{(m)}]$ and $B^t := B_1 \times \cdots \times B_t$ are rectangles in \mathbb{R}^m and \mathbb{R}^{tm} respectively. By a standard extension theorem, [82, Theorem 1.7], P can be extended uniquely to $\mathcal{B}(\mathbb{R}^{tm})$. Under \mathcal{H}_1 , since the independence assumptions are the same, the joint distribution of the network is given exactly by (7.1), with P 's replaced by Q 's.

A key assumption in our setup is that *each node i is only aware of $P^{(i)}$ and $Q^{(i)}$, and the fusion center does not have any knowledge about the statistics of the data observed at the nodes*. Such assumption distinguishes our study from many information-theoretic approaches. For instance, in the Ahlswede–Csiszár setup, both the remote observer and the decision maker are aware of the joint distribution. By contrast, our oblivious fusion center trusts the nodes blindly and sums the “scores” sent by them. Knowing this behavior of the center, nodes prepare their scores accordingly. An example of a score might be the log-likelihood ratio (LLR) of the data observed at time t , i.e., node i calculates the LLR $L_t^{(i)}$ based on its freshly observed data $X_t^{(i)}$ as

$$L_t^{(i)} := \log \frac{dP}{dQ}(X_t^{(i)}), \quad (7.2)$$

sets the score $S_t^{(i)} = L_t^{(i)}$, and passes it through the communication link. Note that the above LLR is well-defined as a Radon–Nikodym derivative due to the




Figure 71 – A representation of the setup studied in this chapter. The peripheral nodes are drawn as blue circles, and the fusion center (FC) is drawn as the red square. At each time instant t , node i sends its compressed score $f_t^{(i)}(L_t^{(i)})$, which is solely based on the fresh observation $X_t^{(i)}$.

absolute continuity of $P^{(i)}$ with respect to $Q^{(i)}$. Suppose each node behaves similarly, i.e., calculates and sends its LLR. Since the data is independent across nodes and across time, under such a strategy, the fusion center receives the sufficient statistic $\sum_{\tau=1}^t \sum_{i=1}^m L_{\tau}^{(i)}$ and is able to perform an optimal test, i.e., a Neyman–Pearson test. However, $P^{(i)}$ and $Q^{(i)}$ can be continuous in general and it is impossible to (i) calculate the LLR with an arbitrarily high precision and (ii) represent the score losslessly with a finite number of bits. Due to these restrictions, the nodes are required to compress (quantize) the data they receive, and send their scores with a finite number of bits at each time instant. The finite-bit score sent by node i at time t is represented by $S_t^{(i)}$ and the fusion center performs a threshold test based on the average score

$$\bar{S}_t := \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^m S_{\tau}^{(i)}, \quad (7.3)$$

with the estimate being the result of the following test:

$$\hat{\mathcal{H}} = \begin{cases} \mathcal{H}_0, & \bar{S}_t \geq \eta_t \\ \mathcal{H}_1, & \text{else} \end{cases} \quad (7.4)$$

where η_t is a threshold that can depend on t . In addition to the finite-bit constraint, the rate of communication between the nodes and the center may be subject to limitations. As mentioned in the previous chapter, we will study the distributed hypothesis testing problem under the following communication constraint: *The average number of bits per node sent under \mathcal{H}_0 must be kept limited.*

Remark 7.1. *The communication constraint is not symmetric, i.e., there is no constraint under \mathcal{H}_1 . This aligns well with many real-world scenarios when*

\mathcal{H}_1 represents a high-risk situation in which the system is allowed to violate communication constraints in order to identify the risk — responding to an emergency takes priority over communication constraints — recall the collision avoidance example of the previous chapter. This view of \mathcal{H}_1 also implies that the type-II error must be very rare. In fact, in many hypothesis testing problems, it is desired that the type-II error decays exponentially. This is the approach we follow for the rest of this chapter. \square

7.2.1 Memoryless Quantization and the Communication Constraint

For simplicity, we focus on a single node i at the moment and omit the symbol (i) from the superscripts. In this section, we formally define the memoryless quantization procedures that map the LLR L_t to the score S_t , and the communication constraints for such quantization procedures. We recall the definition of a simple function, which we have also seen in Section 1.5.

Definition 7.1 (Simple function, [10]). *A function on \mathbb{R} that takes finitely many values is called a simple function. More precisely, let $\alpha_1, \dots, \alpha_n$ be the distinct values of a simple function f , then any such f is represented as*

$$f(l) = \sum_{k=1}^n \alpha_k \mathbb{1}\{l \in B_k\} \quad (7.5)$$

where $B_1, \dots, B_k \in \mathcal{B}(\mathbb{R})$ form a partition of \mathbb{R} . \square

We let $S_t = f_t(L_t)$ with a simple function f_t . Observe that such procedures are *memoryless* — quantization at time t depends only on the data arriving at time t , and does not depend on past. From Definition 7.1, it is clear that S_t 's are discrete random variables. For example, if f_t is set as in (7.5), then for $1 \leq k \leq n$:

$$P(S_t = \alpha_k) = P(L_t \in B_k) \quad (7.6)$$

and the entropy of S_t under \mathcal{H}_0 is defined as

$$H_P(S_t) := - \sum_{k=1}^n P(S_t = \alpha_k) \log P(S_t = \alpha_k) \quad (7.7)$$

with $0 \log 0 := 0$. Recall that in **Chapter 1**, we have seen that a discrete random variable can be compressed *losslessly* with a binary code whose expected length is ℓ , which is bounded as

$$H_P(S_t) \log_2 e - \log_2(H_P(S_t) \log_2 e + 1) - \log_2 e \leq \ell \leq H_P(S_t) \log_2 e. \quad (7.8)$$

Therefore, the peripheral node can compress its LLR L_t with a simple function f_t , and can represent its score $S_t = f_t(L_t)$ with an average number of bits less than $H_P(S_t) \log_2 e$ under \mathcal{H}_0 . If we impose

$$\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R / \log_2 e, \quad (7.9)$$

all scores until time t can be represented with an expected number of bits less than Rt under \mathcal{H}_0 ; and the average number of bits sent over the communication link is kept limited to at most R bits. Constraints formed as in (7.9) are then suitable candidates for being the communication constraint in our distributed hypothesis testing setting.

Remark 7.2. *The memoryless quantization procedures we consider are practically appealing since the peripheral devices can be designed in a memory-efficient manner. Moreover, the assumption that each node only knows their own P 's and Q 's allows independent design of the peripheral nodes, as opposed to the joint design of all sensors which may be impractical. Note that without independence across the nodes, joint design might be necessary. We assume that the network subject to this study is designed such that the peripheral nodes have a spatial configuration that yields, or at least approximates, independence across nodes. \square*

7.2.2 Performance Criteria under Memoryless Quantization

As mentioned earlier, the fusion center decides over the hypotheses based on the threshold test given in (7.4). Under this test, the type-I and type-II error probabilities are defined respectively as

$$\begin{aligned}\alpha_t &:= P(\bar{S}_t < \eta_t) \\ \beta_t &:= Q(\bar{S}_t \geq \eta_t)\end{aligned}\tag{7.10}$$

For an $\epsilon > 0$, we assume that the fusion center sets the threshold to

$$\eta_t = \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \epsilon.\tag{7.11}$$

Recalling that the fusion center is unaware of the statistics at the nodes, one might argue that this choice of η_t is not valid. However, such adjustment is without loss of generality: If the nodes send the centered version of the scores, i.e., $S_t - E_P[S_t]$, and if the fusion center performs the test based on $\eta_t = -\epsilon$, the performance of the scheme will be equivalent to the scheme where the threshold is chosen as in (7.11). Note that setting $\eta_t = -\epsilon$ does not require any knowledge on P 's and Q 's. Moreover, the discrete entropy $H_P(S_t)$ does not change under any shift and the communication constraints are not violated. It turns out that the choice in (7.11) achieves the optimal curve $\theta^*(R)$ — to be defined in Theorem 7.1 — and we keep this choice for the rest of the chapter.

7.3 Best Performance Under Memoryless Quantization

7.3.1 Boundary of the Achievable Region

In view of Remark 7.1, and also of Section 1.5.1, our aim is to drive the type-II error probability to zero as fast as possible while ensuring the type-I error probability vanishes. In particular, the type-II error probability must decay exponentially. A suitable definition of an achievable region in line with this perspective is given as follows.

Definition 7.2. *Given P and Q , (R, θ) is an achievable pair if there exists a sequence $\{f_t\}$ of simple functions and thresholds $\{\eta_t\}$ such that*

$$(a) \quad \frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R, \text{ for all } t$$

$$(b) \quad \lim_{t \rightarrow \infty} \alpha_t = 0$$

$$(c) \quad \liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t} \geq \theta$$

where $S_t = f_t(L_t)$ is the quantized LLR, and α_t, β_t are the type-I and type-II errors defined in (7.10). \square

Note that the communication constraint imposed in Definition 7.2(a) is in terms of *nats* for notational simplicity. The achievable region is then defined as the set of the achievable pairs (R, θ) . The theorem below characterizes the boundary of this region in two parts.

Theorem 7.1. *Let $\theta^*(R) := \sup\{\theta : (R, \theta) \text{ achievable}\}$ and define*

$$\theta_t(R) := \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau] - \log E_P[e^{S_\tau - L_\tau}] \right) \quad (7.12)$$

where $\mathcal{F}_t(R)$ is the set of all simple real-valued functions f_1, \dots, f_t on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ such that $\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R$. Then, the following statements hold.

(i) Let

$$\theta_1(R) = \sup_{f_1 \in \mathcal{F}_1(R)} E_P[S_1] - \log E_P[e^{S_1 - L_1}]. \quad (7.13)$$

Then $\lim_{t \rightarrow \infty} \theta_t(R)$ equals to the upper concave envelope $\check{\theta}_1(R)$ of $\theta_1(R)$.

(ii) $\theta^*(R) = \lim_{t \rightarrow \infty} \theta_t(R) = \check{\theta}_1(R)$.

Proof. See Appendix 7.7.1. \square

Theorem 7.1 provides the boundary of the achievable region in a variational form that is reminiscent of a single-letter characterization. However, the optimization problem (7.13) has a non-convex domain, which makes $\mathcal{F}_1(R)$ a non-convex set. We will therefore consider a relaxed version of the optimization problem (7.13) in the next section.

7.3.2 Upper Bound on the Boundary of the Achievable Region

In order to relax the problem (7.13), we (i) allow randomized quantization, and (ii) modify the communication constraint to $I_P(S_1; L_1) \leq R$, where $I_P(S_1; L_1)$ is the mutual information between S_1 and L_1 under \mathcal{H}_0 . Note that since $H_P(S_1) \geq I_P(S_1; L_1)$, $H_P(S_1) \leq R$ implies $I_P(S_1; L_1) \leq R$, hence the communication constraints indeed become less stringent. (These can be verified with the Properties 1.1) Moreover, the randomized quantization procedures can be represented as channels $p_{V|U} : \mathcal{U} \times \mathbb{R} \rightarrow \mathbb{R}_+$ where for each u , $p_{V|U}(v, u)$ is a probability mass function on the finite set $\mathcal{U} \subset \mathbb{R}$. We further relax the problem by taking $\mathcal{U} = \mathbb{R}$, hence the possible channels become $p_{V|U} : \mathcal{B}(\mathbb{R}) \times \mathbb{R} \rightarrow \mathbb{R}_+$, where for each u , $p_{V|U}(v, u)$ is a probability measure on \mathbb{R} . Adopting the modifications we have just described, problem (7.13) then becomes

$$\begin{aligned} \theta_U(R) &:= \sup_{p_{V|U}} E_P[V] - \log E_P[\exp(V - U)] \\ &\text{s.t. } I_P(U; V) \leq R \end{aligned} \quad (7.14)$$

where U has the same distribution as the LLR L_1 . Observe that as R increases, the optimization domain is enlarged and thus $\theta_U(R)$ cannot decrease; which shows that $\theta_U(R)$ is non-decreasing. Moreover, $\theta_U(R)$ also captures the behavior at the extremes. Intuitively, if $R \rightarrow \infty$, then V can be set equal to U and θ_U becomes

$$E_P[U] = E_P \left[\log \frac{dP}{dQ} \right] = D(P||Q) \quad (7.15)$$

which is known from Stein's lemma — which we have already seen as Theorem 1.9 — as the optimal type-II error exponent under vanishing type-I error probability. This intuitive argument will be made rigorous in Lemma 7.2. On the other extreme, if $R = 0$, then the best possible choice is to set V equal to a constant v and θ_U becomes

$$v - \log E_P \left[\left(\frac{dP}{dQ} \right)^{-1} \right] - v = 0, \quad (7.16)$$

which is consistent with the fact that the center is not able to infer the true hypothesis when there is no communication.

Another useful characterization of θ_U is given by the following lemma.

Lemma 7.1. *Let*

$$\begin{aligned} \tilde{\theta}_U(R) &:= \sup_{p_{V|U}} E_P[V] - E_P[\exp(V - U)] + 1 \\ &\text{s.t. } I_P(U; V) \leq R. \end{aligned} \quad (7.17)$$

Then, $\theta_U(R) = \tilde{\theta}_U(R)$.

Proof. See Appendix 7.7.2. □

Observe that $-\tilde{\theta}_U$ is given by

$$\begin{aligned} -\tilde{\theta}_U(R) &= \inf_{p_{V|U}} -E_P[V] + E_P[\exp(V - U)] - 1 \\ &\text{s.t. } I_P(U; V) \leq R. \end{aligned} \quad (7.18)$$

We highlight the equivalence between $-\tilde{\theta}_U(R)$ and the distortion-rate function with the distortion function $d(u, v) = -v + e^{v-u} - 1$. Since it is known that this curve is convex, $-\tilde{\theta}_U(R)$ is also convex and consequently, $\tilde{\theta}_U(R)$ is concave. We then make use of the characterization in Lemma 7.1 and conclude that $\theta_U(R)$ is concave as well.

We end this section with the following corollary, which states that θ_U is a concave upper bound to the boundary of the achievable region given by $\theta^*(R)$.

Corollary 7.1. $\theta_U(R) \geq \check{\theta}_1(R) = \theta^*(R)$.

Proof. As (7.14) is a relaxation of (7.13), we know that $\theta_U(R) \geq \theta_1(R)$. In addition, $\theta_U(R)$ is concave; then it must also dominate the concave envelope $\check{\theta}_1(R)$ of $\theta_1(R)$. □

Remark 7.3. *As stated before, we have shown the equivalence of θ_U and $\tilde{\theta}_U$ in Lemma 7.1. Although it might be tempting to work with $\tilde{\theta}_U$, as it is the optimal value of an optimization problem with a linear objective, (7.17) takes a smaller value than (7.14) if a generic $p_{V|U}$ is substituted; thus leading to tighter bounds. This is due to the inequality $\log x \leq x - 1$. In view of this observation, we work with the formulation (7.14) in the sequel.* □

7.3.3 Calculating the Upper Bound θ_U

In this section, and for the rest of the chapter, we assume all the expectations (including the mutual information $I_P(\cdot; \cdot)$) are taken under P , and we omit it from the subscripts for brevity. Applying Jensen's inequality to the objective function in (7.14), we have

$$\begin{aligned} E[V] - \log E[\exp(V - U)] &\leq E[V] - E[V - U] \\ &= D(P||Q). \end{aligned} \quad (7.19)$$

Thus, it is also convenient to study the gap to $D(P||Q)$. The gap function $\delta_U(R) := D(P||Q) - \theta_U(R)$ is then straightforwardly given by

$$\begin{aligned} \delta_U(R) &= \inf_{p_{V|U}} \log E[\exp(V - U)] - E[V - U] \\ &\text{s.t. } I(U; V) \leq R. \end{aligned} \quad (7.20)$$

Note that since θ_U is concave and non-decreasing, $\delta_U(R)$ is convex and non-increasing by definition, and the following lemma explains the limiting behavior as $R \rightarrow \infty$.

Lemma 7.2. $\lim_{R \rightarrow \infty} \delta_U(R) = 0$. Consequently, $\lim_{R \rightarrow \infty} \theta_U(R) = D(P||Q)$.

Proof. See Appendix 7.7.3. □

We highlight that Lemma 7.2 holds even if P and Q do not admit densities. Now, we intend to derive an upper bound for δ_U . Let $Z := V - U$. Then (7.20) is equivalent to

$$\begin{aligned} \delta_U(R) &= \inf_{P_{Z|U}} \log E[\exp(Z)] - E[Z] \\ &\text{s.t. } I(U; U + Z) \leq R. \end{aligned} \quad (7.21)$$

A simple upper bound to $\delta_U(R)$ can be obtained by choosing Z as a Gaussian random variable independent of U . With such choice, we have

$$\begin{aligned} I(U; U + Z) &= h(U + Z) - h(U + Z|U) \\ &= h(U + Z) - h(Z|U) \\ &= h(U + Z) - h(Z) \\ &\leq \frac{1}{2} \log \left(1 + \frac{\text{Var}(U)}{\text{Var}(Z)} \right) \end{aligned} \quad (7.22)$$

where $h(\cdot)$ denotes the differential entropy. Observe that $U + Z$ always admits a probability density; U need not be continuous. However we assume U is square integrable such that $\text{Var}(U)$ exists. Furthermore, for a Gaussian Z

$$\begin{aligned} \log E[\exp(Z)] - E[Z] &= E[Z] + \log(e^{\frac{1}{2} \text{Var}(Z)}) - E[Z] \\ &= \frac{1}{2} \text{Var}(Z). \end{aligned} \quad (7.23)$$

Denoting the variance of Z by σ^2 , observe that the parametric curve

$$R = \frac{1}{2} \log \left(1 + \frac{\text{Var}(U)}{\sigma^2} \right), \quad \delta = \frac{1}{2} \sigma^2 \quad (7.24)$$

lies above $\delta_U(R)$, and equivalently

$$\delta_U(R) \leq \frac{\text{Var}(U)}{e^{2R} - 1} =: g_U(R). \quad (7.25)$$

The bound (7.25) is however not tight at low rates. Observe that as $R \rightarrow 0$, the right-hand side of (7.25) tends to infinity although we know that the gap δ can at most be $D(P||Q)$ — see (7.19). The bound can be strengthened as follows: Since we know δ_U is convex with $\delta_U(0) = D(P||Q)$, and $\delta_U(R) \leq g_U(R)$, any line segment connecting $(0, D(P||Q))$ with the curve $g_U(R)$ lies above $\delta_U(R)$. Among such line segments, the one which is tangent to $g_U(R)$ gives the tightest bound.

After obtaining this simple upper bound, we direct our attention to the calculation of $\delta_U(R)$. Note that the objective function in (7.21) is concave.

This is because $\log E[e^Z]$ is concave, and $E[Z]$ is linear in $p_{Z|U}$. Hence, it is a concave minimization problem, and might a priori require examining all extreme points of the feasible set. However, we now show that the problem can be formulated as a convex minimization, circumventing the combinatorial challenge. First, note that both the objective function and the constraint in (7.21) remain unchanged if we add a constant to Z . Thus, centering Z does not change the feasible region in (7.21). Consequently, we can add the constraint $E[Z] = 0$ to our problem without changing its value, which yields an equivalent formulation of (7.21):

$$\begin{aligned} \delta_U(R) = \inf_{p_{Z|U}} \quad & \log E[\exp(Z)] \\ \text{s.t.} \quad & I(U; U + Z) \leq R \\ & E[Z] = 0 \end{aligned} \quad (7.26)$$

Any infimizer of the above problem also infimizes the optimization problem with the objective function $E[\exp(Z)]$, and the optimal value of the former problem is the logarithm of the optimal value of the latter. Further note that the objective function becomes linear when changed to $E[\exp(Z)]$. The latter problem is formulated as the convex program

$$\begin{aligned} \Delta_U(R) := \inf_{p_{Z|U}} \quad & E[\exp(Z)] \\ \text{s.t.} \quad & I(U; U + Z) \leq R \\ & E[Z] = 0 \end{aligned} \quad (7.27)$$

with $\log \Delta_U(R) = \delta_U(R)$. Observe that $\Delta_U(R)$ is convex, non-decreasing, and is finite at every $R \geq 0$ — check the feasible choice $Z = -U + E[U]$ and observe $\Delta_U(R) \leq E[e^{-U}]e^{E[U]} = e^{D(P||Q)}$ since $E[e^{-U}] = 1$ and $E[U] = D(P||Q)$. Therefore, $\Delta_U(R)$ can be expressed as

$$\Delta_U(R) = \sup_{\lambda > 0} \mathcal{L}(\lambda) - \lambda R \quad (7.28)$$

where

$$\begin{aligned} \mathcal{L}(\lambda) := \inf_{p_{Z|U}} \quad & E[\exp(Z)] + \lambda I(U; U + Z) \\ \text{s.t.} \quad & E[Z] = 0. \end{aligned} \quad (7.29)$$

Each $\lambda > 0$ describes a straight line $\Delta + \lambda R = \mathcal{L}(\lambda)$ in the (R, Δ) plane. $\Delta_U(R)$ is the supremum in the Δ axis of these lines. The generalized inverse of $\Delta_U(R)$, $R_U(\Delta)$, is then the supremum of these lines in the R axis,

$$\begin{aligned} R_U(\Delta) &= \sup_{\lambda > 0} \frac{1}{\lambda} \mathcal{L}(\lambda) - \frac{1}{\lambda} \Delta \\ &= \sup_{\eta > 0} \eta \mathcal{L}\left(\frac{1}{\eta}\right) - \eta \Delta \end{aligned} \quad (7.30)$$

which is identical to the following convex problem for $\Delta > 1$.

$$\begin{aligned} R_U(\Delta) &= \inf_{p_{Z|U}} I(U; U + Z) \\ \text{s.t.} \quad & E[\exp(Z)] \leq \Delta \\ & E[Z] = 0. \end{aligned} \tag{7.31}$$

An important direction is to obtain a closed-form lower bound for R_U , which consequently gives a lower bound for δ_U . Such a lower bound characterizes an unachievable region as δ_U is a lower bound to the boundary curve of the achievable region.

Assumption 7.1. *For the rest of the chapter, we assume that U admits a probability density p_U . Hence the differential entropy $h(U)$ is well-defined (but not necessarily finite). \square*

Note that (7.31) is exactly the same as the rate–distortion formulation except the additional constraint $E[Z] = 0$. This special structure allows us to derive a lower bound based on maximum-entropy principles, which also led Shannon to derive the well-known lower bound for the rate–distortion problem under mean-square distortion [137]. We shall use the same machinery as well. Note that

$$\begin{aligned} I(U; U + Z) &= h(U) - h(U|U + Z) \\ &= h(U) - h(Z|U + Z) \\ &\geq h(U) - h(Z) \end{aligned} \tag{7.32}$$

where the last inequality is due to Property 1.1, “conditioning reduces entropy”. Hence, we obtain

$$\begin{aligned} R_U(\Delta) &\geq \inf_{p_{Z|U}} h(U) - h(Z) \\ \text{s.t.} \quad & E[\exp(Z)] \leq \Delta \\ & E[Z] = 0. \end{aligned} \tag{7.33}$$

Since the new objective function depends only on the marginal of Z , the problem above is equivalent to finding a maximum-entropy distribution p_Z that satisfies the constraints $E[e^Z] \leq \Delta$ and $E[Z] = 0$. The problem can now be formulated as

$$\sup_{p_Z} h(Z) \quad \text{s.t.} \quad E[e^Z] \leq \Delta, \quad E[Z] = 0. \tag{7.34}$$

The entropy maximizing distribution can be found with the methods in [1, Chapter 12] and is given by

$$f(z) = \frac{\beta^\alpha}{\Gamma(\alpha)} \exp(\alpha z - \beta e^z), \quad \alpha, \beta > 0. \tag{7.35}$$

Observe that $f(z)$ is the distribution of the logarithm of a Gamma random variable, i.e., $Z = \log G$ where $G \sim \text{Gamma}(\alpha, \beta)$. The following entities have closed form expressions:

$$\begin{aligned} E[e^Z] &= \frac{\alpha}{\beta} \\ E[Z] &= \psi(\alpha) - \log \beta \\ h(Z) &= \log \Gamma(\alpha) - \alpha\psi(\alpha) + \alpha \end{aligned} \quad (7.36)$$

where $\Gamma(\cdot)$ and $\psi(\cdot)$ are gamma and digamma functions respectively. Note that $\log E[e^Z] - E[Z] = \log \alpha - \psi(\alpha)$ and does not depend on β . Substituting (7.36) into (7.33), we have just proved

Lemma 7.3. *Denote the gamma and digamma functions by $\Gamma(\cdot)$ and $\psi(\cdot)$ respectively. Define the parametric curve*

$$\begin{aligned} \underline{R}_U(\alpha) &= h(U) - \log \Gamma(\alpha) + \alpha\psi(\alpha) - \alpha, \\ \delta(\alpha) &= \log \alpha - \psi(\alpha), \quad \alpha > 0, \end{aligned} \quad (7.37)$$

Then $\underline{R}_U(\delta) \leq R_U(\delta)$. □

In comparison, the parametric curve in (7.24) gives the upper bound

$$\bar{R}_U(\delta) = \frac{1}{2} \log \left(1 + \frac{\text{Var}(U)}{2\delta} \right). \quad (7.38)$$

Both \bar{R}_U and \underline{R}_U are depicted in Figure 72 for a Gaussian U . As discussed before, the upper bound $\bar{R}_U(\delta)$ is not tight at low rates since we know that $R_U(\delta) = 0$ at $\delta = D(P||Q)$, and the convexity of $R_U(\delta)$ enables tightening the upper bound by drawing the tangent line from the point $(D(P||Q), 0)$ to \bar{R}_U . This straight line bound is denoted as SL in Figure 72.

7.3.4 Asymptotic Behavior of $R_U(\delta)$

Although Figure 72 suggests that \underline{R}_U and \bar{R}_U match closely at high rates, it is not evident if they tend to infinity at the same rate. Therefore, the asymptotic behavior of the exact R_U is still unknown. We will characterize this behavior in this section. We first derive another upper bound than \bar{R}_U and show that this new upper bound behaves the same as the lower bound \underline{R}_U asymptotically. Once again, refer to (7.22) and observe for a Gaussian Z with variance v and independent of U ,

$$I(U; U + Z) = h(U + \sqrt{v}\tilde{Z}) - h(\tilde{Z}) - \frac{1}{2} \log v, \quad (7.39)$$

where \tilde{Z} is a standard Gaussian random variable. We obtain an upper bound to $h(U + \sqrt{v}\tilde{Z})$ with a different method. Suppose U has a differentiable probability density p_U . We use De Bruijn's identity [1, Chapter 17], which states

$$\frac{\partial}{\partial v} h(U + \sqrt{v}\tilde{Z}) = \frac{1}{2} J(U + \sqrt{v}\tilde{Z}), \quad (7.40)$$




Figure 72 – Bounds for R_U curve for the case where $X \sim \mathcal{N}(0, 1)$ under \mathcal{H}_0 and $X \sim \mathcal{N}(\mu, 1)$ under \mathcal{H}_1 for $\mu = \sqrt{20}$. U has the same distribution as the LLR $L \sim \mathcal{N}(10, 20)$. The lower bound \underline{R}_U is drawn with blue color and the shaded region underneath is unachievable. The upper bound \bar{R}_U is drawn with red color, and its tightened version is drawn with a dashed line, denoted as SL. The true R_U curve lies between \underline{R}_U and SL.

where

$$J(X) := E \left[\left(\frac{\partial}{\partial x} \log p_X(x) \right)^2 \Big|_{x=X} \right] \quad (7.41)$$

is the Fisher information of a random variable X with differentiable density p_X . We then resort to Taylor's theorem which implies

$$h(U + \sqrt{v}\tilde{Z}) \leq h(U) + \frac{v}{2} \sup_{s \geq 0} J(U + \sqrt{s}\tilde{Z}). \quad (7.42)$$

A well-known convolution inequality for Fisher information states [1, Chapter 17] for random variables X and Y with differentiable densities

$$\frac{1}{J(X+Y)} \geq \frac{1}{J(X)} + \frac{1}{J(Y)}. \quad (7.43)$$

Therefore,

$$\begin{aligned} h(U + \sqrt{v}\tilde{Z}) &\leq h(U) + \frac{v}{2} \sup_{s \geq 0} \frac{J(U)J(\sqrt{s}\tilde{Z})}{J(U) + J(\sqrt{s}\tilde{Z})} \\ &\stackrel{(a)}{=} h(U) + \frac{v}{2} \sup_{s \geq 0} \frac{J(U)}{sJ(U) + 1} \\ &\stackrel{(b)}{=} h(U) + \frac{v}{2} J(U) \end{aligned} \quad (7.44)$$

where (a) follows from $J(\sqrt{s}\tilde{Z}) = \frac{1}{s}$, and (b) follows from the fact that $J(U)$ is always non-negative. Substituting this upper bound into (7.39), we obtain

$$\begin{aligned} I(U; U + Z) &\leq h(U) + \frac{v}{2}J(U) - h(\tilde{Z}) - \frac{1}{2}\log v \\ &= h(U) + \frac{v}{2}J(U) - \frac{1}{2}\log(2\pi ev). \end{aligned} \quad (7.45)$$

Referring to (7.24), we have $v = 2\delta$ and obtain another upper bound to R_U as

$$R_U(\delta) \leq h(U) + \delta J(U) - \frac{1}{2}\log(4\pi e\delta). \quad (7.46)$$

We intend to obtain a matching lower bound using \underline{R}_U . To this end, we use the following inequalities on gamma and digamma functions that are valid for $\alpha > 0$ [138, 5.11(ii)]:

$$\begin{aligned} \log \Gamma(\alpha) &\leq \alpha \log \alpha - \alpha - \frac{1}{2}\log \alpha + \frac{1}{2}\log(2\pi) + \frac{1}{12\alpha} \\ \log \alpha - \frac{1}{2\alpha} - \frac{1}{12\alpha^2} &\leq \psi(\alpha) \leq \log \alpha - \frac{1}{2\alpha}. \end{aligned} \quad (7.47)$$

Using (7.47) we obtain from (7.37)

$$\underline{R}_U(\alpha) \geq h(U) - \frac{1}{2}\log(2\pi e\alpha^{-1}) - \frac{1}{6\alpha}, \quad \delta(\alpha) \geq \frac{1}{2\alpha}. \quad (7.48)$$

Thus,

$$R_U(\delta) \geq \underline{R}_U(\delta) \geq h(U) - \frac{1}{2}\log(4\pi e\delta) - \frac{\delta}{3}. \quad (7.49)$$

Comparing (7.49) with (7.46), one can characterize the high-rate behavior of R_U . We conclude this section with the following theorem that gives the asymptotic behavior.

Theorem 7.2.

$$\lim_{\delta \rightarrow 0} R_U(\delta) - h(U) + \frac{1}{2}\log(4\pi e\delta) = 0. \quad (7.50)$$

□

7.4 High-Rate Regime and Performance under Vector Quantization

7.4.1 High-Rate Lattice Quantization

Until this point, we have investigated the behavior of R_U and characterized its exact asymptotic behavior. However, we have not yet proposed a concrete quantization scheme that attains (R, δ) pairs comparable with R_U . In this section, we will show that with simple quantization schemes, R_U can be closely




Figure 73 – A visualization of $q_r(\cdot)$. The output is set to $2kr$ (drawn as red dots) whenever the input falls into the bin (separated with vertical lines) corresponding to k . Each bin is of radius r .

approached at high rates. More specifically, we study lattice quantization procedures — a detailed reference is [139]. As we focus on scalar (memoryless) quantization in one dimension, the quantization procedures we consider are simply described as

$$q_r(U) := 2r \arg \min_{k \in \mathbb{Z}} |U - 2kr| \quad (7.51)$$

where r is the covering radius. Consequently, $V = q_r(U)$ is a quantized version of U with $|V - U| \leq r$. A visual representation is given in Figure 73.

At this point, we would like to relate the radius r to the gap δ . Referring to (7.20), under lattice quantization $q_U(r)$, the gap is given by

$$\begin{aligned} \delta &= \log E[e^{V-U}] - E[V - U] \\ &\leq E[e^{V-U}] - E[V - U] - 1 \\ &= E[e^Z - Z - 1] \end{aligned} \quad (7.52)$$

and since $|Z| \leq r$ surely, $e^Z - Z - 1 \leq e^r - r - 1$. Consequently,

$$\delta \leq e^r - r - 1, \quad (7.53)$$

which suggests that in the small- r regime, δ behaves quadratically. In fact, if $r \leq D$ for a constant D , then

$$\begin{aligned} \delta &\leq e^r - r - 1 \\ &= \sum_{k=2}^{\infty} \frac{r^k}{k!} \\ &\leq \frac{r^2}{D^2} \sum_{k=2}^{\infty} \frac{D^k}{k!} \\ &= \frac{r^2}{D^2} (e^D - D - 1). \end{aligned} \quad (7.54)$$

The next step is to relate r with $H(V)$, which is an upper bound to the expected length of an optimal lossless code as discussed in (7.8). Under mild regularity conditions on the distribution of U , the asymptotic behavior of $H(V)$ when $r \rightarrow 0$ can be characterized.

Theorem 7.3 ([11]). *Suppose $H([U]) < \infty$. Then,*

$$\lim_{r \rightarrow 0} H(V) + \log(2r) = h(U). \quad (7.55)$$

□

Using the above theorem, we have

$$H(V) \leq h(U) - \frac{1}{2} \log(4\pi e\delta) + \frac{1}{2} \log(\pi e(e^D - D - 1)D^{-2}) + f(r) \quad (7.56)$$

where $f(r)$ is a function such that $\lim_{r \rightarrow 0} f(r) = 0$. Comparing the above with the asymptotic behavior of R_U given in Theorem 7.2, we conclude that with high-rate lattice quantization, one can approach the boundary of the achievable region with at most $\frac{1}{2} \log_2(\pi e(e^D - D - 1)D^{-2})$ bits of difference. As $D \rightarrow 0$, the difference term tends to $\frac{1}{2} \log_2\left(\frac{\pi e}{2}\right) \approx 1.047$ bits. Remembering that we work in the high-rate regime, i.e., we are allowed to send a large number of bits, a 1.047-bit gap from the optimal curve does not seem to be significant. We summarize our results on the high-rate quantization as follows.

Theorem 7.4. *Suppose $H([U])$ is finite. Then with one-dimensional lattice quantization of sufficiently small radius, the lower bound to the optimal curve R_U can be approached within $\frac{1}{2} \log_2\left(\frac{\pi e}{2}\right) \approx 1.047$ bits. \square*

Although Theorem 7.4 quantifies the gap in the limit $r \rightarrow 0$, one may also be interested to find an upper bound on $H(V)$ for strictly positive values of r . To this end, one might need more stringent regularity conditions than those of Theorem 7.3 and work with nicely-behaved distributions. For the moment, consider $\tilde{V} = V + W$, where W is independent of V and uniformly distributed in $[-r, r]$. Observe that the probability density of \tilde{V} is a “quantized” version of the probability density of U . If U has a nicely-behaved distribution and if r is small, then the distribution of \tilde{V} will not be very different from that of U ; which is desirable for the sake of analysis. The family of the aforementioned nicely-behaved distributions are defined as follows.

Definition 7.3 (*v*-regular density, [140]). *Given $v : \mathbb{R} \rightarrow \mathbb{R}$, a continuous and differentiable density function p is called *v*-regular if $\left|\frac{d}{du}p(u)\right| \leq v(u)p(u)$. \square*

In [140, Theorem 8], it has been proved that if U has a *v*-regular density, then

$$H(V) \leq h(U) - \log 2r + 2rC_U(r) \quad (7.57)$$

where $C_U(r)$ is a function of r depending on the density of U and on the function v . Furthermore, if v is Lipschitz-continuous almost everywhere and if $E[v(U)]$ is finite, then $C_U(r)$ can be shown to be bounded for finite r — see Appendix 7.7.4. In particular, if v has Lipschitz constant L ,

$$C_U(r) \leq \sqrt{J(U)} + 2Lr. \quad (7.58)$$

We then obtain the following parametric curve

$$\begin{aligned} R_U^{(L)}(r) &:= h(U) - \log 2r + 2r\sqrt{J(U)} + 4Lr^2 \\ \delta(r) &:= e^r - r - 1 \end{aligned} \quad (7.59)$$

which is an upper bound to the (R, δ) pairs achievable with lattice quantization. Figure 74 illustrates the comparison of the lattice upper bound $R_U^{(L)}$ and the lower bound \underline{R}_U at high rates. The 1.047-bit gap in between is clearly observed.

Remark 7.4. As discussed in [140], the gap $\frac{1}{2} \log_2 \left(\frac{\pi e}{2} \right) \approx 1.047$ is due to the covering inefficiency of the one-dimensional lattice. If we perform a similar analysis under the mean-square distortion, the gap turns out to be exactly the same [140]. This is expected as the gap $\delta = \log(E[e^{V-U}]) - E[V-U]$ — despite not being a distortion function — behaves like $r^2/2$ for small r . Observe that for sufficiently smooth densities, the mean-square error under lattice quantization behaves exactly the same for sufficiently small r . \square




Figure 74 – High-rate behaviors of the lower bound \underline{R}_U , and the curve $R_U^{(L)}$ achievable with one-dimensional lattice quantization. Same hypothesis testing setup in Figure 72 is considered, where U has a Gaussian distribution. It is not difficult to see that a Gaussian distribution is v -regular. The 1.047-bit difference between $R_U^{(L)}$ and \underline{R}_U , mentioned in Theorem 7.4, is visible.

In light of our results in this section, the remote node's strategy in the high-rate regime is apparent. At time t , the node (i) calculates its LLR L_t , (ii) obtains the lattice-quantized score $S_t = q_r(L_t)$, and (iii) sends S_t with an optimal variable-length lossless code designed for P , i.e., for \mathcal{H}_0 . This strategy ensures the approach to the optimal curve within 1.047 bits.

One might ask what is the expected number of bits sent under \mathcal{H}_1 although the code is designed for \mathcal{H}_0 . It is known that if the true distribution of the quantized score S_t is given by $Q(S_t)$, then an optimal lossless code designed for $P(S_t)$ yields the expected number of bits at most

$$H_Q(S_t) \log_2 e + D(Q(S_t)||P(S_t)) \log_2 e \leq H_Q(S_t) \log_2 e + D(Q||P) \log_2 e, \quad (7.60)$$

where the inequality is due to the data processing inequality. Hence, if Q is also absolutely continuous with respect to P , $D(Q||P)$ is finite, and the expected number of bits sent under \mathcal{H}_1 is finite as well.

We end this section by raising the following question: “Is it possible to eliminate the $\frac{1}{2} \log_2 \left(\frac{\pi e}{2} \right)$ gap with more efficient lattice coverings?” For high-dimensions, it is known that covering-efficient lattices exist [140]. Hence, an obvious attempt would be to allow the quantization of multiple samples, i.e., at time tk , the remote node records $L_{(t-1)k+1}, \dots, L_{tk}$ and sends the k -dimensional lattice-quantized version. Although this approach might alleviate the covering inefficiency problem, it is not certain that for such procedures the R_U curve remains the same. We shall study in the next section the behavior of R_U when vector quantization is allowed.

7.4.2 Best Performance under Vector Quantization

This section addresses the problem of quantizing multiple samples instead of one. We continue to study memoryless schemes, that is, at time tk , the k -tuple of LLRs $(L_{(t-1)k+1}, \dots, L_{tk})$ is quantized and sent. We first highlight a key observation in the proof of Theorem 7.1 — given in Appendix 7.7.1. Observe that for a choice of quantization function f , the (optimal) Neyman–Pearson test pertaining to the quantized $S_t = f(L_t)$ yields the type-II error rate $D(P(S)||Q(S))$ and the rate is optimized over possible f 's to obtain $\theta(R)$. Adapting this observation to the vector quantization case, we have the score $S_{t,k} = f(L_{(t-1)k+1}, \dots, L_{tk})$, where $f : \mathbb{R}^k \rightarrow \mathbb{R}$ is a simple function, and we want to optimize $D(P(S)||Q(S))$ under the constraint $I(L_1, \dots, L_k; S) \leq kR$ to obtain an upper bound. Using the Donsker–Varadhan representation of $D(P(S)||Q(S))$ as we did in the proof of Theorem 7.1, we therefore have the upper bound to the best achievable type-II error exponent, analogous to (7.14):

$$\begin{aligned} \tilde{\theta}_{L,k}(R) := \frac{1}{k} \left(\sup_{p_{S|L_1, \dots, L_k}} E_P[S] - \log E_P \left[\exp \left(S - L_{1,k} \right) \right] \right) \\ \text{s.t. } I_P(L_1, \dots, L_k; S) \leq kR. \end{aligned} \quad (7.61)$$

Define $L_{1,k} := L_1 + \dots + L_k$ and observe that $I_P(L_{1,k}; S) \leq I_P(L_1, \dots, L_k; S)$. Hence,

$$\begin{aligned} \theta_{L,k}(R) := \frac{1}{k} \left(\sup_{p_{S|L_1, \dots, L_k}} E_P[S] - \log E_P \left[\exp \left(S - L_{1,k} \right) \right] \right) \\ \text{s.t. } I_P(L_{1,k}; S) \leq kR \end{aligned} \quad (7.62)$$

is an upper bound to $\tilde{\theta}_{L,k}(R)$ as the optimization domain is enlarged (it is not difficult to show that $\theta_{L,k}$ is in fact equal to $\tilde{\theta}_{L,k}$). Since both the objective and constraint functions in (7.62) only depend on $L_{1,k}$, the feasible set can be

reduced to the set of channels from $L_{1,k}$ to S . Hence,

$$\begin{aligned} \theta_{L,k}(R) &= \frac{1}{k} \left(\sup_{p_{S|L_{1,k}}} E_P[S] - \log E_P \left[\exp \left(S - L_{1,k} \right) \right] \right) \\ &\text{s.t. } I_P(L_{1,k}; S) \leq kR. \end{aligned} \quad (7.63)$$

Note the resemblance of (7.63) to (7.14). Consequently, all results for one-dimensional quantization directly translate to the multi-dimensional case and we obtain the following upper bound to the boundary of the achievable region:

$$\begin{aligned} \theta_{U,k}(R) &:= \frac{1}{k} \left(\sup_{p_{V|U}} E_P[V] - \log E_P[\exp(V - U_k)] \right) \\ &\text{s.t. } I_P(U_k; V) \leq kR \\ &= \frac{1}{k} \theta_{U_k}(kR) \end{aligned} \quad (7.64)$$

where U_k is the random variable that has the same distribution as $L_{1,k}$. Following the same steps we have taken for the one-dimensional case, we can also obtain the gap function and the rate-gap curve for the k -dimensional case as

$$\delta_{U,k}(R) = \frac{1}{k} \delta_{U_k}(kR), \quad R_{U,k}(\delta) = \frac{1}{k} R_{U_k}(k\delta). \quad (7.65)$$

The previously obtained upper and lower bounds for the one-dimensional case are therefore valid for k -dimensional case as well:

$$\underline{R}_{U,k}(\delta) := \frac{1}{k} \underline{R}_{U_k}(k\delta) \leq R_{U,k}(\delta) \leq \frac{1}{k} \bar{R}_{U_k}(k\delta) =: \bar{R}_{U,k}(\delta). \quad (7.66)$$

For various k values, the lower bounds $\underline{R}_{U,k}(\delta)$ and upper bounds $\bar{R}_{U,k}(\delta)$ are drawn in Figure 75 for the same scenario in Figure 72.

Observe that $R_{U,k}$'s obey the subadditive relation

$$(k+l)R_{U,k+l}(\delta) \leq kR_{U,k}(\delta) + lR_{U,l}(\delta) \quad (7.67)$$

as the admissible strategies for the quantization of $k+l$ samples include the strategies that quantize k samples and l samples separately. However, note that this does not imply $R_{U,k}(\delta) \leq R_{U,l}(\delta)$ for $k \geq l$. Nevertheless, from a well-known result on subadditive sequences, e.g. [2], we know

$$\lim_{k \rightarrow \infty} R_{U,k}(\delta) = \inf_k R_{U,k}(\delta) \quad (7.68)$$

and using the upper bound (7.38), we obtain for $\delta > 0$,

$$\begin{aligned} \lim_{k \rightarrow \infty} R_{U,k}(\delta) &\leq \lim_{k \rightarrow \infty} \frac{1}{2k} \log \left(1 + \frac{\text{Var}(U_k)}{2k\delta} \right) \\ &= \lim_{k \rightarrow \infty} \frac{1}{2k} \log \left(1 + \frac{\text{Var}(U_1)}{2\delta} \right) \\ &= 0. \end{aligned} \quad (7.69)$$




Figure 75 – Upper and lower bounds under k -dimensional vector quantization, for $k = 1, 2, 4$. Again, the same setup in Figure 72 is considered. The lower bounds $\underline{R}_{U,k}$ are drawn as solid curves and the upper bounds $\bar{R}_{U,k}$ are drawn as dashed curves. Although all upper and lower bounds are pointwise decreasing with k , it is not certain that the true $R_{U,k}$ curves exhibit the same behavior.

This is in contrast with the classical rate–distortion function as it is already defined for $k \rightarrow \infty$.

Although (7.69) shows that the lower bound $R_{U,k}$ tends to zero, this is also true for the true boundary curve. A simple achievability scheme at large k is as follows: Since the remote node records the data until k , it can make its own decision \mathcal{H}_0 or \mathcal{H}_1 and send the one-bit result to the fusion center. The average number of bits sent is then kept arbitrarily small and since the node makes the estimate of the true hypothesis based on an optimal test, the type-II error rate will be close to $D(P||Q)$, which is the best possible decay rate. Although such a design might seem appealing in terms of the performance of type-II error rate, the peripheral node needs to have sufficient computational power as a requirement of this design. Also recall that in the end of Section 7.4.1, we mentioned that the covering efficiency of lattices may improve at high dimensions. However, (7.69) and the strategy we have just described suggest that there is no need for lattice quantization for high dimensions — the node only sends its one-bit decision.

7.5 Multiple-Node Case

All the previous results obtained for the single-node case can be extended to the multiple-node case. This is due to the fact that the data is *independent across nodes*. To make this extension, we provide a modified definition of

achievable pairs. Recall that at time t , node i observes data coming from $P^{(i)}$ under \mathcal{H}_0 , and from $Q^{(i)}$ under \mathcal{H}_1 ; calculates the LLR $L_t^{(i)}$, and compresses it with a simple function $S_t^{(i)} = f_t^{(i)}(L_t^{(i)})$. Furthermore, as discussed before, if

$$\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau^{(i)}) \leq R_i / \log_2 e, \quad (7.70)$$

then the compressed scores can be sent losslessly with an average number of bits less than R_i . After recalling the system dynamics, we provide the modified version of Definition 7.2.

Definition 7.4. Given $\{P^{(i)}\}_{i=1}^m$ and $\{Q^{(i)}\}_{i=1}^m$, $(R_1, \dots, R_m, \theta)$ is an achievable pair if there exists m sequences $\{f_t^{(1)}\}, \dots, \{f_t^{(m)}\}$ of simple functions and a sequence of thresholds $\{\eta_t\}$ such that

- (a) $\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau^{(i)}) \leq R_i$, for all t and for all i
- (b) $\lim_{t \rightarrow \infty} \alpha_t = 0$
- (c) $\liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t} \geq \theta$

where $S_t^{(i)} = f_t^{(i)}(L_t^{(i)})$, and α_t, β_t are the type-I and type-II errors respectively, defined in (7.10). \square

Let $\theta_t(R_1, \dots, R_m) := \sum_{i=1}^m \theta_t^{(i)}(R_i)$, where

$$\theta_t^{(i)}(R_i) := \sup_{\substack{\{f_1^{(i)}, \dots, f_t^{(i)}\} \\ \in \mathcal{F}_t(R_i)}} \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau^{(i)}] - \log E_P[e^{S_\tau^{(i)} - L_\tau^{(i)}}] \right) \quad (7.71)$$

is defined as in Theorem 7.1. Observe that

$$\begin{aligned} \theta_t(R_1, \dots, R_m) &= \sum_{i=1}^m \theta_t^{(i)}(R_i) \\ &= \sum_{i=1}^m \sup_{\substack{\{f_1^{(i)}, \dots, f_t^{(i)}\} \\ \in \mathcal{F}_t(R_i)}} \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau^{(i)}] - \log E_P[e^{S_\tau^{(i)} - L_\tau^{(i)}}] \right) \\ &= \sup_{\substack{\{f_1^{(i)}, \dots, f_t^{(i)}\} \\ \in \mathcal{F}_t(R_i), i \leq m}} \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^m \left(E_P[S_\tau^{(i)}] - \log E_P[e^{S_\tau^{(i)} - L_\tau^{(i)}}] \right) \end{aligned} \quad (7.72)$$

and

$$\begin{aligned} \sum_{i=1}^m \log E_P[e^{S_\tau^{(i)} - L_\tau^{(i)}}] &\stackrel{(a)}{=} \log E_P[e^{\sum_{i=1}^m (S_\tau^{(i)} - L_\tau^{(i)})}] \\ &\stackrel{(b)}{=} \log E_Q[e^{\sum_{i=1}^m S_\tau^{(i)}}] \\ &\stackrel{(c)}{=} \sum_{i=1}^m \log E_Q[e^{S_\tau^{(i)}}] \end{aligned} \quad (7.73)$$

where (a), (c) are due to the independence assumption across nodes and (b) is because $e^{-\sum_{i=1}^m L_\tau^{(i)}} = \prod_{i=1}^m \frac{dP^{(i)}}{dQ^{(i)}} = \frac{dP}{dQ}$ is the Radon–Nikodym derivative of P with respect to Q . Therefore, we have

$$\theta_t(R_1, \dots, R_m) = \sum_{i=1}^m \sup_{\substack{\{f_1^{(i)}, \dots, f_t^{(i)}\} \\ \in \mathcal{F}_t(R_i)}} \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau^{(i)}] - \log E_Q[e^{S_\tau^{(i)}}] \right). \quad (7.74)$$

Comparing (7.74) with (7.12) and following exactly the same steps in the proof of Theorem 7.1, we obtain the analogous version of Theorem 7.1 (ii):

$$\begin{aligned} \theta^*(R_1, \dots, R_m) &:= \sup\{\theta : (R_1, \dots, R_m, \theta) \text{ achievable}\} \\ &= \lim_{t \rightarrow \infty} \theta_t(R_1, \dots, R_m) \\ &= \sum_{i=1}^m \lim_{t \rightarrow \infty} \theta_t^{(i)}(R_i) \\ &= \sum_{i=1}^m \check{\theta}_1^{(i)}(R_i), \end{aligned} \quad (7.75)$$

which characterizes the boundary of the optimal curve for rate constraints (R_1, \dots, R_m) . We also know that $\check{\theta}_1(R)$ is upper bounded by $\theta_U(R) = D(P||Q) - \delta_U(R)$ from Corollary 7.1, hence

Corollary 7.2.

$$\theta^*(R_1, \dots, R_m) \leq \sum_{i=1}^m \left[D(P^{(i)}||Q^{(i)}) - \delta_{U_i}^{(i)}(R_i) \right] \quad (7.76)$$

where U_i has the same distribution as $L_1^{(i)}$, LLR of node i . \square

One might also consider an extension of the problem to sum-rate constraints. Namely, the communication constraint is redefined as $R_1 + \dots + R_m \leq R_{\text{sum}}$. The characterization in (7.75) is readily adapted to sum-rate constraint as

$$\begin{aligned} \theta^*(R_{\text{sum}}) &:= \sup\{\theta : (R_1, \dots, R_m, \theta) \text{ achievable}, R_1 + \dots + R_m = R_{\text{sum}}\} \\ &= \max_{R_1 + \dots + R_m = R_{\text{sum}}} \sum_{i=1}^m \check{\theta}_1^{(i)}(R_i), \end{aligned} \quad (7.77)$$

and

$$\theta^*(R_{\text{sum}}) \leq \max_{R_1 + \dots + R_m = R_{\text{sum}}} \sum_{i=1}^m \theta_{U_i}(R_i). \quad (7.78)$$

The (θ, R) pairs that lie above the curve on right-hand side are unachievable under the sum-rate constraint. Hence, one may be interested in the optimal rate sharing that maximizes the right-hand side with an aim to characterize an unachievable region. We first provide a simple property of a possible optimal allocation: Intuitively, the optimal rate sharing must not exclude the more informative nodes.

Definition 7.5 ([141]). A node j is said to be more informative² than node i , and denoted as $i \prec j$ if there exists a probability transition kernel $w : \mathcal{B}(\mathbb{R}) \times \mathbb{R} \rightarrow \mathbb{R}_+$ such that for all $\mathcal{A} \in \mathcal{B}(\mathbb{R})$

$$\int w(\mathcal{A}, x) dP^{(j)}(x) = P^{(i)}(\mathcal{A}) \quad (7.79)$$

and

$$\int w(\mathcal{A}, x) dQ^{(j)}(x) = Q^{(i)}(\mathcal{A}). \quad (7.80)$$

□

With the above definition, we have

Theorem 7.5. Let $\mathcal{R}^* \subset \mathbb{R}^m$ be the set of optimal allocations in (7.78). If $i \prec j$ for some j , then \mathcal{R}^* contains a (R_1^*, \dots, R_m^*) such that $R_i^* > 0$ only if $R_j^* > 0$.

Proof. Recall the definition of $\theta_{U_i}(R_i)$:

$$\begin{aligned} \theta_{U_j}(R) &= \sup_{\substack{p_{V|U_j} \\ \text{s.t. } I(U_j; V) \leq R}} E_P[V] - \log E_P[\exp(V - U_j)] \\ &\stackrel{(a)}{=} \sup_{\substack{p_{V|U_j} \\ \text{s.t. } I(U_j; V) \leq R}} E_P[V] - \log E_Q[\exp(V)] \\ &\geq \sup_{\substack{p_{V|U_i} \circ w_{U_i|U_j} \\ \text{s.t. } I(U_j; V) \leq R}} E_P[V] - \log E_Q[\exp(V)] \\ &\stackrel{(b)}{\geq} \sup_{\substack{p_{V|U_i} \circ w_{U_i|U_j} \\ \text{s.t. } I(U_i; V) \leq R}} E_P[V] - \log E_Q[\exp(V)] \\ &= \theta_{U_i}(R) \end{aligned} \quad (7.81)$$

where (a) is due to the measure change as U_j is distributed as the logarithm of the Radon–Nikodym derivative $\frac{dP^{(j)}}{dQ^{(j)}}$ and (b) is due to the data processing inequality $I(U_j; V) \leq I(U_i; V)$ for the choice of $p_{V|U_j} = p_{V|U_i} \circ w_{U_i|U_j}$. Suppose $R_i > 0$ and $R_j = 0$. Since $\theta_{U_j}(R)$ pointwise dominates $\theta_{U_i}(R)$, one cannot do worse with the modification $(R_i, 0) \rightarrow (0, R_i)$. □

It might be tempting to think that the optimal allocation assigns $R_i^* = 0$ if $i \prec j$ for some j , as $\theta_{U_j}(R)$ dominates $\theta_{U_i}(R)$ pointwise. This is not true in general. Suppose $P^{(i)}$ and $Q^{(i)}$ are obtained by passing $P^{(j)}$ and $Q^{(j)}$ through an additive Gaussian channel with almost zero noise and suppose R_{sum} is very large. Allocating all the rate to node j will yield an exponent close to $D(P^{(i)}||Q^{(i)})$ whereas an equal rate allocation gives an exponent close to

²Although the definition of more informativeness is different for m -ary hypothesis tests, it is shown in [141] that for $m = 2$, the definition given here is equivalent.

$D(P^{(i)}||Q^{(i)}) + D(P^{(j)}||Q^{(j)})$. Therefore, if the sum-rate constraint is large enough, it is preferred to observe two (almost uncompressed) independent samples instead of one, which surely increases the type-II decay rate.

Since θ_{U_i} 's (or equivalently δ_{U_i} 's) are difficult to calculate in general, one may consider the optimal allocation based on the (R, δ) pairs that lie on the $\underline{R}_{U_i}(\delta)$ curves given by the parametric form (7.37). Observe that such pairs depend on U_i only through shifts of a parametric curve by its differential entropy $h_i := h(U_i)$. Denoting the inverse of $\underline{R}_{U_i}(\delta)$ by $\underline{\delta}_{U_i}(R)$, we have thus the property

$$\underline{\delta}_{U_i}(R) = \underline{\delta}(R - h_i) \quad (7.82)$$

where $\underline{\delta}(R)$ is given by the parametric form

$$\begin{aligned} R(\alpha) &= -\log \Gamma(\alpha) + \alpha\psi(\alpha) - \alpha, \\ \underline{\delta}(\alpha) &= \log \alpha - \psi(\alpha), \quad \alpha > 0. \end{aligned} \quad (7.83)$$

The sum-rate optimization is then formulated as

$$\begin{aligned} \underline{\theta}^*(R_{\text{sum}}) &:= \max_{\sum_{i=1}^m R_i = R_{\text{sum}}} \sum_{i=1}^m [D(P^{(i)}||Q^{(i)}) - \underline{\delta}(R_i - h_i)] \\ &= \sum_{i=1}^m D(P^{(i)}||Q^{(i)}) - \min_{\sum_{i=1}^m R_i = R_{\text{sum}}} \sum_{i=1}^m \underline{\delta}(R_i - h_i), \end{aligned} \quad (7.84)$$

which can be shown to admit a water-filling solution.

Lemma 7.4. *The sum-rate constrained problem (7.84) has a solution given by*

$$R_i^* = (\mu + h_i)^+, \quad (7.85)$$

where $(x)^+$ denotes the positive part of x ; and μ is a constant chosen to satisfy the sum-rate constraint $\sum_{i=1}^m R_i^* = R_{\text{sum}}$.

Proof. First, observe that $\underline{\delta}(R)$ is convex. This is a consequence of its formulation in (7.34). As $\underline{\delta}(R)$ is convex, the Karush–Kuhn–Tucker (KKT) conditions are necessary and sufficient to characterize the solutions. For the sum-constraint $\sum_{i=1}^m R_i = R_{\text{sum}}$, it is known that the KKT conditions are given by [2]

$$\begin{aligned} \underline{\delta}'(R_i - h_i) &= \lambda, & R_i > 0 \\ \underline{\delta}'(-h_i) &> \lambda, & R_i = 0 \end{aligned} \quad (7.86)$$

for some constant λ . Since $\underline{\delta}'(R_i - h_i)$ is non-decreasing due to the convexity of $\underline{\delta}(R)$, the KKT conditions are also equivalent to

$$\begin{aligned} R_i - h_i &= \mu, & R_i > 0 \\ -h_i &> \mu, & R_i = 0 \end{aligned} \quad (7.87)$$

with $\mu = R(\alpha)|_{\alpha=1/\lambda}$, and they characterize the claimed solution in (7.85). \square

Although the sum-rate optimization in (7.84) does not give the exact boundary of the achievable $(R_1, \dots, R_m, \theta)$ pairs, $\underline{\theta}^*(R_{\text{sum}})$ is an upper bound to the boundary. This implies that no $(R_1, \dots, R_m, \theta)$ pair lying above $\underline{\theta}^*(R_{\text{sum}})$ is achievable. A numerical example is illustrated in Figure 76.




Figure 76 – A 3-node instance of the problem. Nodes observe zero-mean Gaussian data under \mathcal{H}_0 , and with a mean vector $[\sqrt{0.2}, \sqrt{0.4}, \sqrt{0.6}]$ under \mathcal{H}_1 . The data has unit variance under both hypotheses. Then, the LLRs also have Gaussian distributions with means $[0.1, 0.2, 0.3]$ and variances $[0.2, 0.4, 0.6]$ under \mathcal{H}_0 . The individual $\underline{\theta}_{U_i}(R) := D(P^{(i)}||Q^{(i)}) - \underline{\delta}_{U_i}(R)$ curves are drawn dashed, whereas the optimal sum-rate curve $\underline{\theta}^*(R)$ is drawn solid. The shaded region is unachievable under sum-rate constraints.

We conclude this section by noting that when the data is not independent across nodes, the question of how to combine the scores is highly non-trivial even if there were no communication constraints. Under communication constraints, the problem for this general case could be of formidable complexity.

7.6 Discussion

In this chapter, we have studied a fundamental limit of a distributed hypothesis testing problem when remote nodes compress their data in a memoryless fashion and the expected number of bits sent under \mathcal{H}_0 should be kept limited to a prescribed quantity R . This asymmetric communication constraint is in line with the view that \mathcal{H}_1 is a rare high-risk event and must be detected with high probability. Thus, nodes are allowed to send a large number of bits under \mathcal{H}_1 . With such a communication constraint, we characterized the maximum attainable type-II error (i.e., mis-detection of \mathcal{H}_1) exponent (Theorem 7.1) for

vanishing type-I error probability and derived a closed-form upper bound to this error exponent (Lemma 7.3).

In the high-rate regime, we show that the upper bound is approached with simple scalar lattice quantization within $\frac{1}{2} \log_2(\pi e/2) \approx 1.047$ bits. This gap is due to the covering inefficiency of the 1-dimensional lattice and due to the fact that the gap to the optimal error rate $D(P||Q)$ behaves quadratically at high rates. Therefore, it is expected that the results for the high-rate regime coincide with the results on the rate–distortion problem for mean-square distortion. This is the reason that the asymptotic behavior of the rate–gap curve (Theorem 7.2) is reminiscent of the Shannon lower bound for the rate–distortion curve under mean-square distortion. It is also because of this quadratic behavior that one can approach the lower bound within $\frac{1}{2} \log_2(\pi e/2)$ bits under scalar lattice quantization.

We have also obtained a simple upper bound for the vector quantization case that can be expressed in terms of its scalar quantization analog (7.64)–(7.65). Hence, the same upper and lower bounds for the scalar case are also valid for the vector quantization case. We have also shown in (7.69) that as the dimension tends to infinity, the rate–gap curve is identically equal to zero for $\delta > 0$. This is consistent with the following simple achievability scheme: The remote node performs its own Neyman–Pearson test and sends its decision with one bit. Hence, the fusion center is informed of the optimal decision, and the average number of bits sent is arbitrarily small.

With an independence assumption across nodes, the results for the single-node problem can be easily extended to the multiple-node problem, and a simple upper bound to the optimal type-II error exponent can be obtained in terms of the sum of individual upper bounds for each remote node. We formulated a sum-rate constrained problem and studied some of its properties (Theorem 7.5 and Lemma 7.4).

As a final remark, we note that the results for the vector quantization case are also applicable to the multiple-node case. This implies that when the dimension tends to infinity, the rate–gap curve will be again equal to zero and is attained with a simple scheme that is similar to the single-node case: Each node performs its optimal test and sends the 1-bit result to the fusion center. The center decides \mathcal{H}_1 if at least one node decides \mathcal{H}_1 . This scheme ensures vanishing type-I error probability and the type-II error exponent is equal to $\sum_i D(P^{(i)}||Q^{(i)})$. Hence the center is able to attain the optimal rate with an arbitrarily low amount of communication. However, this scheme allows each node to dictate a \mathcal{H}_1 decision to the center. This results in a system that is vulnerable to manipulation. By contrast, schemes with scalar quantization, or with low-dimensional vector quantization, give the center the opportunity to detect errors or manipulations and therefore these schemes could be of interest when faulty or malicious nodes are present.

The reader may recall that we have mentioned two types of distributed settings in **Chapter 6**: (i) a centralized one and (ii) decentralized one. Now that we have completed our study of the centralized scheme, we shall proceed

with a study of a decentralized inference setting in the next chapter, called social learning.

7.7 Appendix

7.7.1 Proof of Theorem 1

Proof of (i)

Let $\check{\theta}_1(R)$ be the concave envelope of $\theta_1(R)$. Recall $S_t = f_t(L_t)$ and since all the expectations are taken under P , we omit P from the subscripts. Also recall

$$\theta_t(R) = \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t \left(E[S_\tau] - \log E[e^{S_\tau - L_\tau}] \right).$$

We first show that $\check{\theta}_1(R) \geq \theta_t(R)$ for all t . Let us modify the definition of $\theta_1(R)$ as

$$\begin{aligned} \theta_1(R) = \sup_{f \text{ simple}} E[S_1] \\ \text{s.t. } E[e^{S_1 - L_1}] = 1 \\ H(S_1) \leq R \end{aligned} \quad (7.88)$$

since shifting S_1 does not change the entropy. The supremization is over simple functions. Similarly, $\theta_t(R)$ can be defined as

$$\begin{aligned} \theta_t(R) = \sup_{\substack{f_1, \dots, f_t \\ \text{simple}}} \frac{1}{t} \sum_{i=1}^t E[S_i] \\ \text{s.t. } E[e^{S_i - L_i}] = 1, \forall i \leq t \\ \frac{1}{t} \sum_{i=1}^t H(S_i) \leq R. \end{aligned} \quad (7.89)$$

For any $\{f_i\}$ in the feasible set of (7.89), there exists $\{R_i\}$ with $\frac{1}{t} \sum_{i=1}^t R_i \leq R$ and $H(S_i) \leq R_i$ for all $i \leq t$; and consequently $\frac{1}{t} \sum E[S_i] \leq \frac{1}{t} \sum \theta_1(R_i) \leq \check{\theta}_1(R)$. Thus, $\theta_t(R) \leq \check{\theta}_1(R)$.

It remains to prove the reversed inequality for $t \rightarrow \infty$, i.e., $\theta_t(R) \geq \check{\theta}_1(R) - \epsilon$ for large enough t , given $\epsilon > 0$. Suppose $\theta_1(R)$ is attained in the limit of the sequence of simple functions $\{f_t^*\}$. This implies for all $\epsilon_1 > 0$, there exists a simple function f that maps $L_1 \mapsto S_1$ such that $E[S_1] \geq \theta_1(R) - \epsilon_1$, $E[e^{S_1 - L_1}] = 1$ and $H(S_1) \leq R$. Carathéodory's theorem [142, Section 17] ensures that every point on the concave envelope $\check{\theta}_1(R)$ is achieved by a convex combination of at most two points on $\theta_1(R)$. This implies the existence of functions f, \tilde{f} , and $\lambda \in [0, 1]$, such that $\lambda E[S_1] + (1 - \lambda) E[\tilde{S}_1] \geq \check{\theta}_1(R) - \epsilon_2$ for all $\epsilon_2 > 0$, and $\lambda H(S_1) + (1 - \lambda) H(\tilde{S}_1) \leq R$. Assume $H(\tilde{S}_1) \leq R$ without loss of

generality. Consider the sequence $\{f_t\}$ such that $f_i = f$ for $i \leq \lceil \lambda t \rceil$ and $f_i = \tilde{f}$ otherwise. Observe $\frac{1}{t} \sum H(S_i) \leq R$ and thus $\theta_t(R) \geq \frac{1}{t} \sum E[S_i] \geq \check{\theta}_1(R) - 2\epsilon_2$ for t large enough. The proof of part (i) is complete.

In the achievability proof of part (ii), we have to show that the supremizers of $\theta_t(R)$ have to drive $\alpha_t \rightarrow 0$. For completeness, we provide the proof here. Take $\{f_t\}$ as above and use Chebyshev's inequality to bound the type-I error probability under the threshold test with the threshold η_t chosen as in (7.11):

$$\alpha_t = P\left(\sum_{i=1}^t f_i(L_i) < E_P[g_i(L_i)] - \epsilon\right) \leq \frac{\sum_{i=1}^t \text{Var}(f_i(L_i))}{\epsilon^2 t^2} \quad (7.90)$$

Recall that f_i is defined to be equal either to f or \tilde{f} . As f and \tilde{f} take finitely many values, the variances of $f(U)$, $\tilde{f}(U)$ are bounded for any U . Therefore,

$$\alpha_t \leq \frac{\max\left\{\text{Var}(f(L_1)), \text{Var}(\tilde{f}(L_1))\right\}}{\epsilon^2 t} \rightarrow 0, \quad (7.91)$$

which shows that any sequence in the achievability part of (i) indeed satisfies the property (b) in Definition 1.

Proof of (ii)

We will follow an approach very similar to the one in the proof of Theorem 1.10.

(Achievability) As mentioned, choose $\eta_t = \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \epsilon$ as in (7.11). We upper bound the type-II error β_t as

$$\begin{aligned} Q(\bar{S}_t \geq \eta_t) &= Q\left(\exp(t\bar{S}_t) \geq \exp(t\eta_t)\right) \\ &\stackrel{(a)}{\leq} E_Q\left[\exp\left(\sum_{\tau=1}^t (S_\tau - E_P[S_\tau] + \epsilon)\right)\right] \\ &\stackrel{(b)}{=} \prod_{\tau=1}^t E_Q[e^{S_\tau}] \exp(-E_P[S_\tau] + \epsilon) \end{aligned} \quad (7.92)$$

where (a) follows from Markov inequality and (b) follows from independent processing of LLRs. Therefore,

$$\frac{1}{t} \log \frac{1}{\beta_t} \geq \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau] - \log E_Q[e^{S_\tau}]\right) - \epsilon. \quad (7.93)$$

Optimizing the right-hand side with respect to the choice of f_t 's satisfying the communication constraints we have

$$\frac{1}{t} \log \frac{1}{\beta_t} \geq \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \log E_Q[e^{S_\tau}] - \epsilon. \quad (7.94)$$

Consider the transformation $\tilde{S}_t = \log \frac{P(S_t)}{Q(S_t)}$, i.e., the LLR of S_t . Observe that the mapping $L_t \mapsto \tilde{S}_t$ is a simple function and since S_t is discrete, $H(S_t) \geq H(\tilde{S}_t)$ as the mapping $S_t \mapsto \tilde{S}_t$ is deterministic. Therefore, communication constraints are still satisfied. Furthermore, \tilde{S}_t is a sufficient statistic and the fusion center is able to deploy a Neyman–Pearson test based on \tilde{S}_t 's. It is known from Donsker–Varadhan representation [143] of divergence that

$$D(P||Q) = \sup_{g:\mathbb{R}\rightarrow\mathbb{R}} E_P[g(X)] - \log E_Q[e^{g(X)}] \quad (7.95)$$

where the supremum is over the set of bounded measurable functions on \mathbb{R} , and is attained at $g(X) = \log \frac{dP}{dQ}$, the logarithm of the Radon–Nikodym derivative of P with respect to Q . We combine (7.95) with (7.94) to obtain

$$\frac{1}{t} \log \frac{1}{\beta_t} \geq \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(\mathbb{R})} \frac{1}{t} \sum_{\tau=1}^t D(P_{f_\tau} || Q_{f_\tau}) - \epsilon, \quad (7.96)$$

where $P_{f_\tau} := P \circ f_\tau$ and $Q_{f_\tau} := Q \circ f_\tau$. It only remains to show that the supremizers in (7.96) must drive $\alpha_t \rightarrow 0$, which we have already proved at the end of part (i).

(Converse) Now, following similar steps to Stein's lemma, we apply data processing inequality twice to see that for any sequence of f_t 's:

$$tD(P||Q) \geq \sum_{\tau=1}^t D(P_{f_\tau} || Q_{f_\tau}) \geq d(\alpha_t || 1 - \beta_t) \quad (7.97)$$

where $d(p||q) := p \log \frac{p}{q} + \bar{p} \log \frac{\bar{p}}{\bar{q}}$ is the binary divergence. Hence,

$$\sum_{\tau=1}^t D(P_{f_\tau} || Q_{f_\tau}) \geq -h_e(\alpha_t) - \alpha_t \log(1 - \beta_t) - (1 - \alpha_t) \log(\beta_t), \quad (7.98)$$

with $h_e(p) := -p \log p - (1 - p) \log(1 - p)$. Suppose $\alpha_t \rightarrow 0$ and β_t bounded away from 1. Then, it must be true that

$$\liminf_{t \rightarrow \infty} \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(\mathbb{R})} \frac{1}{t} \sum_{\tau=1}^t D(P_{f_\tau} || Q_{f_\tau}) \geq \liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t}. \quad (7.99)$$

Taking \liminf and $\epsilon \rightarrow 0$ in (7.96), combining with (7.99); we therefore have

$$\liminf_{t \rightarrow \infty} \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(\mathbb{R})} \frac{1}{t} \sum_{\tau=1}^t \sup_{f_\tau} D(P_{f_\tau} || Q_{f_\tau}) = \liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t}. \quad (7.100)$$

In other words, β_t decays with an exponent at least

$$\liminf_{t \rightarrow \infty} \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(\mathbb{R})} \frac{1}{t} \sum_{\tau=1}^t D(P_{f_\tau} || Q_{f_\tau}). \quad (7.101)$$

Recall that the set $\mathcal{F}_t(R)$ also includes the supremal function g_t 's in the Donsker–Varadhan formulation (7.95), thus we have

$$\begin{aligned}
 & \liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t} \\
 &= \liminf_{t \rightarrow \infty} \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \log E_Q[e^{S_\tau}] \\
 &\stackrel{(a)}{=} \liminf_{t \rightarrow \infty} \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \log E_P[e^{S_\tau - L_\tau}] \\
 &= \lim_{t \rightarrow \infty} \theta_t(R)
 \end{aligned} \tag{7.102}$$

where (a) follows from the assumption that $P \sim Q$; hence both $\frac{dP}{dQ}$, $\frac{dQ}{dP}$ exist and $\frac{dQ}{dP} = \left(\frac{dP}{dQ}\right)^{-1}$ almost surely.

7.7.2 Proof of Lemma 7.1

We consider the inequality (7.94) and obtain a more relaxed lower bound for it using $\log x \leq x - 1$ as

$$\frac{1}{t} \log \frac{1}{\beta_t} \geq \frac{1}{t} \sum_{\tau=1}^t \sup_{f_\tau \in \mathcal{F}} E_P[S_\tau] - E_Q[e^{S_\tau}] + 1 - \epsilon. \tag{7.103}$$

It is known that $D(P||Q)$ can also be represented as

$$D(P||Q) = \sup_{g: \mathbb{R} \rightarrow \mathbb{R}} E_P[g(X)] - E_Q[e^{g(X)}] + 1. \tag{7.104}$$

Proceeding similarly to the proof of Theorem 7.1, we obtain $\theta_X(R) = \tilde{\theta}_X(R)$.

7.7.3 Proof of Lemma 7.2

Consider the following quantization of U . Let $\mathcal{I}_n := [-n, n]$ and $v_{n,k} := -n + k2^{-n}$ for $1 \leq k \leq 2^{n+1}n$. If $U \in \mathcal{I}_n$, then set V as the closest $v_{n,k}$ to U . Otherwise, set $V = 0$. Then,

$$E[e^{V-U}] \leq e^{2^{-n}} P(U \in \mathcal{I}_n) + E[e^{-U} \mathbf{1}\{U \notin \mathcal{I}_n\}] \tag{7.105}$$

and

$$E[V - U] \geq -2^{-n} P(U \in \mathcal{I}_n) - E[U \mathbf{1}\{U \notin \mathcal{I}_n\}]. \tag{7.106}$$

Since $E[e^{-U}] = 1$, $E[|U|] < \infty$ and $P(U \in \mathcal{I}_n) \rightarrow 1$, the rightmost terms above tend to zero and for large n ,

$$\log E[e^{V-U}] - E[V - U] \tag{7.107}$$

is close to zero. The proof will be complete if $I(U; V) < \infty$ for finite n . But since V has finite cardinality, this is indeed the case.

7.7.4 Boundedness of $C_U(r)$

We give the definition of $C_U(r)$ suitable for our setting, adopted from [140]. Recall that the output of the lattice quantization takes values in $\{2kr\}_{k \in \mathbb{Z}}$. Let $I_k := [(2k-1)r, (2k+1)r]$. Then $C_U(r) := \sum_k P(I_k) \max_{I_k} v(u)$.

Now, observe

$$E[v(U)] = \sum_k \int_{I_k} p(u)v(u)du. \quad (7.108)$$

Since both p and v are continuous, by the mean value theorem, for each k there exists a c_k such that

$$\int_{I_k} p(u)v(u)du = v(c_k) \int_{I_k} p(u)du. \quad (7.109)$$

Suppose $v(u)$ is Lipschitz with constant L . Since $a_k := \arg \max_{u \in I_k} v(u)$ has distance at most $2r$ to c_k , we have $|a_k - c_k| \leq 2r$ and $v(a_k) \leq 2Lr + v(c_k)$. Then,

$$C_U(r) \leq \sum_k (2Lr + v(c_k))P(I_k) = 2Lr + E[v(U)]. \quad (7.110)$$

Therefore, finiteness of $E[v(U)]$ guarantees the finiteness of C_U . Also observe that if $v(u) = |\frac{d}{du} \log p(u)|$ is Lipschitz, and if $J(U) < \infty$, then $C_U(r)$ is upper bounded as

$$C_U(r) \leq \sqrt{J(U)} + 2Lr \quad (7.111)$$

and is guaranteed to be finite.

For the case when v is not Lipschitz, but is differentiable and its derivative is Lipschitz with constant L , $C_U(r)$ can be upper bounded as

$$C_U(r) \leq E[v(U)] + 2rE[|v'(U)|] + 4Lr^2 \quad (7.112)$$

and if $E[v(U)]$, $E[|v'(U)|]$ are finite, $C_U(r)$ can be bounded from above. Note that similar arguments generalize to higher order derivatives of v .

As an example, suppose U has a density given by $p(u) = K \exp(-|u|^3)$, with K being the appropriate normalization constant. Then since $\frac{d}{du} \log p(u) = 3u^2$, one cannot find a Lipschitz v . Nevertheless, if one sets $v(u) = 3u^2$, $|v'(u)| = 6|u|$ is Lipschitz with constant $L = 6$ and since $E[U^2]$, $E[|U|]$ are finite, $C_U(r)$ is bounded according to (7.112).

Introduction to Social Learning

8

As mentioned in the previous chapter, we will now introduce a decentralized inference framework. As opposed to the hierarchical fusion-center based setup, we will now consider a fully-flat architecture and there will be no center to aggregate information. Our architecture consists of nodes (or agents) that share their beliefs with a subset of nodes at certain times. An example could be a social network, e.g., Twitter, where a user shares their belief about a certain phenomenon with their followers — in graph-theory parlance, neighbors — formed by the private information they possess. For instance a finance expert might have an insider information that the CEO of ABC company runs a Ponzi scheme and might share his updated belief about the ABC stock price (most probably the experts would advise to sell in this situation if they are trustworthy). Although he has valuable information, this particular expert is not the only person who has private information about ABC, or other companies. Hence, agents might have an incentive to share their honest beliefs so as to receive honest beliefs from their companions as well. With the aggregated information (we will see how), agents in the social network might be able to make accurate guesses about more general phenomena — in this example agents' ultimate aim could be to detect the market direction.

The above example is indeed a social learning instance, where social beings collaboratively seek to learn more about the true state-of-nature. Social learning is a paradigm that investigates how opinions are formed over social networks by modeling the interactions between networked agents to learn the true state of a phenomenon. We will describe and study a social network model as a m -ary hypothesis testing problem; where agents will seek to find the true hypothesis among the m possible. At this point, it might be informative to see the Bayesian framework as opposed to the Neyman–Pearson framework we have studied so far. Later, we will emphasize how a Bayesian approach would

be intractable in a social learning problem and *locally* Bayesian learning becomes necessary and helpful.

8.1 Notation

In this and subsequent chapter, we will extensively deal with vectors, matrices and scalars which can be random or not. Hence, for convenience we will use a different notation than used in the previous chapters of this thesis. Random variables are denoted with boldface letters whereas their realizations are denoted with plain letters (e.g., \mathbf{z}_i and x_i). For a collection of random variables, $\sigma(\cdot)$ denotes the smallest σ -algebra pertaining to the collection. Sets and events are denoted with script-style letters (e.g., \mathcal{A}). $|\mathcal{A}|$ denotes the cardinality of set \mathcal{A} . For vectors u and v , $\langle u, v \rangle$ denotes the inner product between u and v ; and $\|u\|_1$, $\|u\|$ denote the ℓ_1 and ℓ_2 norms of u respectively. All the logarithms are assumed to be natural logarithms. For a graph $G = (V, E)$, \mathcal{N}_k denotes the neighbor of vertex k and $\deg(k) \triangleq |\mathcal{N}_k|$ is the degree of vertex k .

8.2 Bayesian Framework in m -ary Hypothesis Testing

Similar to the hypothesis testing problem described in Section 1.5, the statistician aims to learn about a phenomenon based on their observation. This time, there are m possible states of nature, $\theta_1, \dots, \theta_m$, among which there is a true state denoted by θ° . Prior to having any observation, i.e., at time $i = 0$, the statistician has a belief vector $\boldsymbol{\mu}_0$, which is a distribution on $\Theta := \{\theta_1, \dots, \theta_m\}$, where $\boldsymbol{\mu}_0(\theta)$ quantifies the statistician's confidence in the proposition ' $\theta = \theta^\circ$ '.

We assume that at each time instant i , the statistician observes data $\boldsymbol{\xi}_i$. We assume $\boldsymbol{\xi}_i$'s are i.i.d., and under hypothesis θ , $\boldsymbol{\xi}_i$ is sampled from the distribution/likelihood $L(\cdot|\theta)$. After making their observation at time i , the statistician updates their belief vector according to the Bayes rule:

$$\boldsymbol{\mu}_i(\theta) = P(\theta|\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_i) = \frac{\boldsymbol{\mu}_{i-1}(\theta)L(\boldsymbol{\xi}_i|\theta)}{\sum_{\theta} \boldsymbol{\mu}_{i-1}(\theta)L(\boldsymbol{\xi}_i|\theta)}. \quad (8.1)$$

One should expect that the belief vector should eventually assign unit belief to the true hypothesis, i.e., $\boldsymbol{\mu}_i(\theta) \rightarrow \delta(\theta^\circ)$, where $\delta(\cdot)$ is the Kronecker delta function. This is indeed true if the initial belief $\boldsymbol{\mu}_0(\theta^\circ) > 0$; otherwise it is easy to see that $\boldsymbol{\mu}_i(\theta^\circ)$ is stuck at zero for all i . Assume $\boldsymbol{\mu}_0(\theta^\circ) > 0$, and observe that the log-belief ratios

$$\boldsymbol{\lambda}_i(\theta) := \log \frac{\boldsymbol{\mu}_i(\theta^\circ)}{\boldsymbol{\mu}_i(\theta)} \quad (8.2)$$

evolve as a random walk

$$\boldsymbol{\lambda}_i(\theta) = \boldsymbol{\lambda}_{i-1}(\theta) + \log \frac{L(\boldsymbol{\xi}_i|\theta^\circ)}{L(\boldsymbol{\xi}_i|\theta)}. \quad (8.3)$$

Since the true hypothesis is θ° , the expectation of each random increment is

$$E_{\theta^\circ} \left[\log \frac{L(\boldsymbol{\xi}_1|\theta^\circ)}{L(\boldsymbol{\xi}_1|\theta)} \right] = D(L(\cdot|\theta^\circ)||L(\cdot|\theta)). \quad (8.4)$$

Therefore if $D(L(\cdot|\theta^\circ)||L(\cdot|\theta)) > 0$, which is true if $L(\cdot|\theta^\circ) \neq L(\cdot|\theta)$, we conclude from standard results on random walks that $\boldsymbol{\lambda}_i(\theta) \rightarrow \infty$ and hence $\boldsymbol{\mu}_i(\theta) \rightarrow 0$ for all $\theta \neq \theta^\circ$. This consequently shows that $\boldsymbol{\mu}_i(\theta) \rightarrow 1$. We will refer to this property as *truth learning*.

Now suppose that there are K statisticians. Consequently, at each time instant i , a random vector $\boldsymbol{\xi}_i$ of length K is observed. Note that for a fully-Bayesian update as in (8.1), (i) the data $\boldsymbol{\xi}_i$ has to be aggregated at a center (or at least must be known by each agent). Moreover, (ii) the *joint* distribution across agents must be known. Both conditions (i) and (ii) suggest that in a decentralized framework as in the previous section, fully-blown Bayesian updates are intractable. The situation could worsen even more when the number of agents is large. To circumvent the technical difficulties, social learning, as a decentralized inference framework, uses *locally* Bayesian updates. Simply put, in a local Bayesian update, agents first (i) update their belief vectors according to their local data and local likelihood function, e.g., for agent k , $\boldsymbol{\xi}_{k,i}$ and $L(\boldsymbol{\xi}_{k,i}|\theta^\circ)$ (a.k.a. the marginal distribution at k th coordinate). This updated vector is also called the *intermediate belief*. Then, (ii) agents share their updated (intermediate) beliefs with their peers/neighbors. Note that agents might receive multiple belief vectors from their neighbors; therefore they have to aggregate these distinct belief vectors (as well as their own). The two widely-studied methods consist of arithmetic and geometric averaging of beliefs. In both methods, it can be shown that all agents learn the truth almost surely under mild assumptions on network connectivity. However, the speed of learning is an important criterion and there is no certain answer that tells geometric is superior compared to arithmetic or vice versa — although we have some hint that geometric averaging performs better [144–146]. The randomized social learning model proposed in the next chapter will use a geometric averaging method, where each agent randomly polls a neighbor and aggregates the two belief vectors. Before studying this model, let us mention some of the related studies on social learning.

8.3 Related Work

As we have discussed in the previous section, a desirable property of social learning algorithms is that all agents eventually learn the truth, i.e., their beliefs will assign unit probability to the true hypothesis θ° , e.g., [146–149]. Social learning is also studied to understand how opinions are formed over social networks [150–153]. For instance, fully Bayesian strategies to find the global posterior belief is studied in [151]. However, as we have discussed in the previous section, this approach may not be tractable in large networks and

consequently *local* Bayes updates — sometimes called non-Bayesian learning — are studied extensively. Some references along this strand are [146–148, 154–160].

Agents usually incorporate their beliefs with their neighbors' beliefs with consensus [148, 161], diffusion [160, 162–174], or gossip [175, 176] updates. In consensus, agents share their belief vectors with their neighbors whereas in diffusion updates they share their *intermediate beliefs*. In the next chapter, we will study a diffusion-based randomized social learning algorithm — the reader might have noticed that the method we described in the previous section is also diffusion. These two methods are very similar but they require different conditions to ensure truth learning. For instance, in consensus, all agents must have positive self-reliance, i.e., must not discard their own belief during the fusion of neighbors' beliefs, to ensure truth learning. In diffusion, it is enough for one agent to have positive self-reliance. We will also study in the next chapter a special case where agents replace their beliefs with a randomly chosen neighbor, hence no agents will have positive self-reliance with unit probability. Still, we will see that truth learning is ensured in this case under mild conditions.

In the social learning literature, the closest work to what we will study in the next chapter appears to be [176], where the authors have considered symmetric gossip schemes and shown that the agents learn the truth eventually with high probability. Unlike [176], we will consider diffusion [162] algorithms and the communication is not necessarily symmetric. For instance, we assume full-duplex communication at nodes such that at time i agent 1 can receive data from agent 2 while agent 2 receives from agent 3. Furthermore, we will show that truth learning takes place almost surely as opposed to 'with high probability'. Diffusion algorithms with random neighbor selections are included in [177, 178], for optimization purposes rather than social learning framework.

Social Learning Under Randomized Collaborations¹

9

9.1 Motivation

A common assumption in social learning is that agents communicate with *all* of their neighbors at each time instant. While this assumption helps modeling microblogging social media such as Twitter, it falls short when modeling private and personal communication social media platforms such as WhatsApp. In many situations, people exchange beliefs with a subset of their contacts. This might happen because, for example, data may arrive at high rates such that agents might not be able to communicate with all their neighbors between two consecutive arrivals. Furthermore, for designing communication-efficient networked systems, sparse interactions between the devices can be preferred. For instance, consider an agent that attempts to receive data from multiple neighbors. These transmissions are likely to collide and to avoid such issues, each neighbor can be given turns to communicate by the receiver agent, similar to a MAC layer protocol. The above observations motivate us to study the social learning problem when agents update their beliefs based on *only one randomly chosen neighbor* at each time instant. The main results of this chapter are listed below.

- (i) The agents learn the truth eventually with *probability one*.
- (ii) Despite the decreased amount of communication, the asymptotic rate of learning is the same as the standard social learning algorithms where agents interact with all their neighbors.
- (iii) For a special case where the agents replace their beliefs with the neighbor's belief, we will provide a large deviations analysis of log-belief ratios

¹The content of this chapter is based on [179].

that only uses the marginal distributions of the data across the agents, i.e., the result does not depend on any coupling between the agents.

9.2 Problem Formulation

In accordance with the multiple statistician scenario of Section 8.2, consider K agents with peer-to-peer communication on a graph topology. These agents aim to infer the true hypothesis θ° among a finite set of hypotheses Θ . The belief $\boldsymbol{\mu}_{k,i}(\theta)$ quantifies the confidence that agent k has at time i in the proposition “ $\theta = \theta^\circ$ ”. The vector $\boldsymbol{\mu}_{k,i}$ lives in a $|\Theta|$ -dimensional probability simplex. The agents observe partially informative and private observations, i.e., agent k observes $\boldsymbol{\xi}_{k,i}$ at time i , which is distributed according to the likelihood/distribution $L_k(\boldsymbol{\xi}_{k,i}|\theta^\circ)$. Agent k knows its likelihood functions $L_k(\cdot|\theta)$ for all $\theta \in \Theta$. We assume data is identically and independently distributed (i.i.d.) across time; but is not necessarily independent across agents, unlike **Chapter 7**. As opposed to prior work where agent k receives beliefs from all its neighbors at each time instant i , in the study of this chapter, it randomly selects one neighbor at each time instant, independent from the past and receives information from that neighbor.

9.2.1 A Randomized Diffusion Algorithm

In our algorithm, agents update their beliefs with a *local* Bayesian rule to obtain intermediate beliefs as in standard social learning algorithms [146, 148, 149, 160]:

$$\boldsymbol{\psi}_{k,i}(\theta) = \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta)\boldsymbol{\mu}_{k,i-1}(\theta)}{\sum_{\theta'} L_k(\boldsymbol{\xi}_{k,i}|\theta')\boldsymbol{\mu}_{k,i-1}(\theta')}. \quad (\text{Adapt}) \quad (9.1)$$

Then, each agent chooses one of their neighbors and updates its intermediate belief by taking a weighted geometric average with the chosen neighbor’s belief. Specifically, for agent k and $\ell \in \mathcal{N}_k$,

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\boldsymbol{\psi}_{k,i}(\theta)^\alpha \boldsymbol{\psi}_{\ell,i}(\theta)^{\bar{\alpha}}}{\sum_{\theta'} \boldsymbol{\psi}_{k,i}(\theta')^\alpha \boldsymbol{\psi}_{\ell,i}(\theta')^{\bar{\alpha}}}, \text{ with prob. } a_{\ell k} \text{ (Combine)} \quad (9.2)$$

where $\alpha \in [0, 1)$ is a confidence weight and $\bar{\alpha} := 1 - \alpha$. We assume $\boldsymbol{\mu}_{k,0}(\theta) > 0$ for all k and θ . Observe that $\alpha = 0$ corresponds to the case where the agent replaces its intermediate belief with the chosen neighbor’s belief. If we allow $\alpha = 1$, it would correspond to the non-cooperative mode of operation; hence this case is of no interest to our work.

Observe that the matrix $A := [a_{\ell k}]$, called as the combination matrix in social learning, is left-stochastic as $a_{\ell k}$ ’s represent probabilities. We have the following assumption regarding the communication topology and A .

Assumption 9.1. *The network is strongly-connected, which means there is a path with positive probability between every agent pair (k, ℓ) and there is at least one agent k with $a_{kk} > 0$ — analogous to self-reliance property of an agent.*

In contrast to prior work where $a_{\ell k}$ represents the weight agent k assigns to the belief obtained from agent ℓ during convex combination of all neighbors, in this work, this weight represents the probability that agent k chooses and receives information from agent ℓ . This procedure decreases the number of transmissions made at each time slot, i.e., the number of transmissions decreases $\frac{1}{K} \sum_{k=1}^K \deg(k)$ times on average compared to the standard algorithm.

9.3 Analysis of the Algorithm

Recall that θ° is the true hypothesis and let $\theta \in \Theta \setminus \{\theta^\circ\}$ denote an arbitrary hypothesis. We study the evolution of the log-belief ratios:

$$\lambda_{k,i} := \log \frac{\mu_{k,i}(\theta^\circ)}{\mu_{k,i}(\theta)}, \quad (9.3)$$

which can be verified to evolve according to

$$\lambda_{k,i} = \alpha(\mathbf{x}_{k,i} + \lambda_{k,i-1}) + \bar{\alpha}(\mathbf{x}_{\ell,i} + \lambda_{\ell,i-1}), \text{ with prob. } a_{\ell k} \quad (9.4)$$

where

$$\mathbf{x}_{k,i} := \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta^\circ)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)} \quad (9.5)$$

is the log-likelihood ratio (LLR) of the data calculated by agent k at time i . An equivalent way to express (9.4) is

$$\begin{aligned} \lambda_{k,i} &= \langle \lambda_i, \mathbf{w}_0^{(k)} \rangle \\ &= \langle \mathbf{x}_i + \lambda_{i-1}, (\alpha I + \bar{\alpha} \mathbf{A}_i) \mathbf{w}_0^{(k)} \rangle \\ &= \langle \mathbf{x}_i, (\alpha I + \bar{\alpha} \mathbf{A}_i) \mathbf{w}_0^{(k)} \rangle + \langle \lambda_{i-1}, (\alpha I + \bar{\alpha} \mathbf{A}_i) \mathbf{w}_0^{(k)} \rangle \end{aligned} \quad (9.6)$$

where $\mathbf{x}_i := [\mathbf{x}_{1,i} \dots \mathbf{x}_{K,i}]^T$, $\lambda_i := [\lambda_{1,i} \dots \lambda_{K,i}]^T$, $\mathbf{w}_0^{(k)} := [0 \dots 1 \dots 0]^T$ is an all-zero vector with a single 1 at its k^{th} element; and $\mathbf{A}_i := [\mathbf{a}_{\ell k,i}]$ is a random matrix such that $\mathbf{a}_{\ell k,i} = 1$ if node k chooses to communicate with node ℓ at time i . Note that $\sum_{\ell} \mathbf{a}_{\ell k,i} = 1$ surely, i.e., is a left-stochastic matrix, and $E[\mathbf{A}_i] = \mathbf{A}$ for all i . Furthermore, since each node selects its neighbor identically and independently across time, $\mathbf{A}_1, \dots, \mathbf{A}_i$ are i.i.d.

Now, define

$$\mathbf{w}_n^{(k)} := \prod_{j=i-n+1}^i (\alpha I + \bar{\alpha} \mathbf{A}_j) \mathbf{w}_0^{(k)} \quad (9.7)$$

for $1 \leq n \leq i$. Note that $\mathbf{w}_n^{(k)}$ is a random probability vector. We iterate (9.6) to obtain

$$\boldsymbol{\lambda}_{k,i} = \sum_{n=1}^i \langle \mathbf{w}_n^{(k)}, \mathbf{x}_{i-n+1} \rangle. \quad (9.8)$$

Our aim is to show that (with a.s. denoting almost sure convergence)

$$\frac{1}{i} \boldsymbol{\lambda}_{k,i} \xrightarrow{\text{a.s.}} \langle \boldsymbol{\pi}, \mathbf{d} \rangle, \quad (9.9)$$

where \mathbf{d} is the divergence vector with its k^{th} element being the KL divergence $D(L_k(\cdot|\theta^\circ) || L_k(\cdot|\theta))$; and $\boldsymbol{\pi}$ is the Perron vector of A . Recall that $\langle \boldsymbol{\pi}, \mathbf{d} \rangle$, the asymptotic rate of convergence, is the same as the standard algorithm where agents benefit from all neighbors at each time instant [146, 149]. As an initial step to prove (9.9), we first establish the following result.

Lemma 9.1. *For all k ,*

$$\frac{1}{i} \sum_{n=1}^i \mathbf{w}_n^{(k)} \xrightarrow{\text{a.s.}} \boldsymbol{\pi}. \quad (9.10)$$

Proof. The statement does not depend on k , so without loss of generality we take $k = 1$ and omit all the superscripts (k). We first show convergence in probability. From Markov's inequality, we have

$$P\left(\left\|\frac{1}{i} \sum_{n=1}^i \mathbf{w}_n - \boldsymbol{\pi}\right\|^2 > \epsilon\right) \leq \frac{E\left[\left\|\sum_{n=1}^i (\mathbf{w}_n - \boldsymbol{\pi})\right\|^2\right]}{\epsilon i^2}. \quad (9.11)$$

The expected norm on the right-hand side of (9.11) is equal to

$$E\left[\sum_{n=1}^i \|\mathbf{w}_n - \boldsymbol{\pi}\|^2\right] + 2E\left[\sum_{m < n} \langle \mathbf{w}_m - \boldsymbol{\pi}, \mathbf{w}_n - \boldsymbol{\pi} \rangle\right]. \quad (9.12)$$

Define for $m \leq i$, $\mathcal{F}_{m,i} := \sigma(\{\mathbf{A}_j\}_{i-m+1 \leq j \leq i})$ and observe

$$\begin{aligned} E[\langle \mathbf{w}_m - \boldsymbol{\pi}, \mathbf{w}_n - \boldsymbol{\pi} \rangle] &= E\left[E[\langle \mathbf{w}_m - \boldsymbol{\pi}, \mathbf{w}_n - \boldsymbol{\pi} \rangle | \mathcal{F}_{m,i}]\right] \\ &= E[\langle \mathbf{w}_m - \boldsymbol{\pi}, E[\mathbf{w}_n - \boldsymbol{\pi} | \mathcal{F}_{m,i}] \rangle] \end{aligned} \quad (9.13)$$

Furthermore, $\mathbf{w}_n = (\alpha I + \bar{\alpha} \mathbf{A}_{i-n+1}) \dots (\alpha I + \bar{\alpha} \mathbf{A}_{i-m}) \mathbf{w}_m$ and since \mathbf{A}_n 's are i.i.d.,

$$E[\mathbf{w}_n | \mathcal{F}_{m,i}] = (\alpha I + \bar{\alpha} A)^{n-m} \mathbf{w}_m. \quad (9.14)$$

Substituting (9.14) into (9.13), we obtain

$$\begin{aligned} E[\langle \mathbf{w}_m - \boldsymbol{\pi}, \mathbf{w}_n - \boldsymbol{\pi} \rangle] &= E\left[\langle \mathbf{w}_m - \boldsymbol{\pi}, (\alpha I + \bar{\alpha} A)^{n-m} (\mathbf{w}_m - \boldsymbol{\pi}) \rangle\right] \\ &\stackrel{(a)}{\leq} E\left[\|\mathbf{w}_m - \boldsymbol{\pi}\| \|(\alpha I + \bar{\alpha} A)^{n-m} (\mathbf{w}_m - \boldsymbol{\pi})\|\right] \end{aligned} \quad (9.15)$$

where (a) follows from Cauchy-Schwarz inequality. For a strongly-connected A and for any $\alpha \in [0, 1)$, it is known that there exists a $\rho < 1$ such that

$$\|(\alpha I + \bar{\alpha}A)^{n-m}(\mathbf{w}_m - \pi)\| \leq C(\rho)\rho^{n-m}, \quad (9.16)$$

where $C(\rho)$ is a constant that only depends on ρ [180, Chapter 4]. Hence, (9.15) is further upper bounded by

$$E[\|\mathbf{w}_m - \pi\|]C\rho^{n-m}. \quad (9.17)$$

Also note that

$$E[\|\mathbf{w}_m - \pi\|] \leq E[\|\mathbf{w}_m\|] + \|\pi\| \leq 2 \quad (9.18)$$

and

$$E[\|\mathbf{w}_m - \pi\|^2] \leq 2(E[\|\mathbf{w}_m\|^2] + \|\pi\|^2) \leq 4. \quad (9.19)$$

Using (9.18), (9.19) we upper bound (9.12), and consequently upper bound (9.11) as

$$P\left(\left\|\frac{1}{i}\sum_{n=1}^i \mathbf{w}_n - \pi\right\|^2 > \epsilon\right) \leq \frac{4 + 4C\rho/(1 - \rho)}{\epsilon i}. \quad (9.20)$$

This shows that $\frac{1}{i}\sum_{n=1}^i \mathbf{w}_n^{(k)} \rightarrow \pi$ in probability. Also note that for a K -dimensional vector u , $\|u\|_1 \leq \|u\|\sqrt{K}$. Therefore

$$P\left(\left\|\frac{1}{i}\sum_{n=1}^i \mathbf{w}_n - \pi\right\|_1 > \epsilon\right) \leq \frac{4\sqrt{K} + 2\sqrt{K}\rho/(1 - \rho)}{\epsilon i}. \quad (9.21)$$

The last step of the proof follows by a standard trick to obtain the strong law of large numbers [181, Chapter 7]. Let $\mathbf{u}_i := \left\|\sum_{n=1}^i(\mathbf{w}_n - \pi)\right\|_1$. From Borel–Cantelli lemma [82, Chapter 2], the subsequence $\mathbf{u}_{i^2}/i^2 \xrightarrow{\text{a.s.}} 0$ — replace i with i^2 in (9.21) and observe that the right-hand side is summable. Moreover, observe for all $i^2 \leq m \leq (i + 1)^2$:

$$\frac{\mathbf{u}_m}{m} \stackrel{(a)}{\leq} \frac{\mathbf{u}_{i^2} + \sum_{n=i^2}^m \|\mathbf{w}_n - \pi\|_1}{i^2} \stackrel{(b)}{\leq} \frac{\mathbf{u}_{i^2} + 2((i + 1)^2 - i^2)}{i^2}. \quad (9.22)$$

To obtain (a) we upper bounded the numerator with triangle inequality and lower bounded the denominator by using $m \geq i^2$. (b) holds because $\|\mathbf{w}_n - \pi\|_1 \leq 2$ for any w , and $m \leq (i + 1)^2$. Since $\frac{\mathbf{u}_{i^2}}{i^2} \xrightarrow{\text{a.s.}} 0$ so does $\frac{\mathbf{u}_m}{m}$. \square

Taking the inner product of both sides in (9.10) with d , we obtain

Corollary 9.1. $\frac{1}{i}\sum_{n=1}^i \langle \mathbf{w}_n^{(k)}, d \rangle \xrightarrow{\text{a.s.}} \langle \pi, d \rangle$ for all k .

Lemma 1 suggests that the convergence results will not depend on k . Hence, we assume $k = 1$ and omit all the superscripts (k) if not needed. The next step is to show that (9.9) holds under a square-integrability assumption on divergences. More precisely,

Lemma 9.2. *Suppose $E[(\mathbf{x}_{k,i})^2] < \infty$ for all k . Then (9.9) holds.*

Proof. It is sufficient to show that

$$\frac{1}{i} \sum_{n=1}^i \langle \mathbf{w}_n, \mathbf{x}_{i-n+1} - d \rangle \xrightarrow{\text{a.s.}} 0. \quad (9.23)$$

Note that \mathbf{x}_i is independent of \mathbf{w}_j for all j ; and of all \mathbf{x}_j except itself. The same holds for the \mathbf{A}_i 's as well. Therefore, for convenience, let us perform an index change $i - n + 1 \mapsto n$ on $(\mathbf{x}_i, \mathbf{A}_i)$ to have the equivalent statement

$$\frac{1}{i} \sum_{n=1}^i \langle \tilde{\mathbf{w}}_n, \mathbf{x}_n - d \rangle \xrightarrow{\text{a.s.}} 0 \quad (9.24)$$

where $\tilde{\mathbf{w}}_n := (\alpha I + \bar{\alpha} \mathbf{A}_n) \tilde{\mathbf{w}}_{n-1}$ and $\tilde{\mathbf{w}}_0 = [1, \dots, 0]^T$ according to the index change described above. Kronecker's lemma [82, Chapter 12] implies that it is sufficient to check if

$$\mathbf{z}_i := \sum_{n=1}^i \frac{1}{n} \langle \tilde{\mathbf{w}}_n, \mathbf{x}_n - d \rangle \text{ converges a.s.} \quad (9.25)$$

Observe that \mathbf{z}_i is a martingale with respect to the filtration $\{\tilde{\mathcal{F}}_i\}$, where $\tilde{\mathcal{F}}_i := \sigma((\mathbf{A}_n, \mathbf{x}_n)_{n \leq i})$. This is because

$$\begin{aligned} E[\mathbf{z}_{i+1} | \tilde{\mathcal{F}}_i] &= \mathbf{z}_i + \frac{1}{i+1} E[\langle \tilde{\mathbf{w}}_{i+1}, \mathbf{x}_{i+1} - d \rangle | \tilde{\mathcal{F}}_i] \\ &\stackrel{(a)}{=} \mathbf{z}_i + \frac{1}{i+1} \langle E[\tilde{\mathbf{w}}_{i+1} | \tilde{\mathcal{F}}_i], E[\mathbf{x}_{i+1} - d | \tilde{\mathcal{F}}_i] \rangle \\ &\stackrel{(b)}{=} \mathbf{z}_i \end{aligned} \quad (9.26)$$

where (a) from conditional independence of \mathbf{x}_{i+1} and $\tilde{\mathbf{w}}_{i+1}$ and (b) follows from $E[\mathbf{x}_{i+1} - d | \tilde{\mathcal{F}}_i]$ being equal to the all-zero vector. Furthermore, \mathbf{z}_i is a square-integrable martingale as

$$\begin{aligned} \sup_i E[\mathbf{z}_i^2] &= \sum_{n=1}^{\infty} \frac{E[\langle \tilde{\mathbf{w}}_n, \mathbf{x}_n - d \rangle^2]}{n^2} \\ &\stackrel{(a)}{\leq} \sum_{n=1}^{\infty} \frac{E[\|\tilde{\mathbf{w}}_n\|^2 \|\mathbf{x}_n - d\|^2]}{n^2} < \infty \end{aligned} \quad (9.27)$$

where (a) follows from Cauchy-Schwarz inequality. The final expression is bounded since $\|\tilde{\mathbf{w}}_n\|^2 \leq 1$ and $E[\|\mathbf{x}_n\|^2] < \infty$ by our assumption. The above allows the use of the martingale convergence theorem [82, Chapter 11] and therefore \mathbf{z}_i a.s. converges. The proof is complete since we have shown (9.25). \square

Our final aim is to relax the square-integrability condition $E[(\mathbf{x}_{k,i})^2] < \infty$ to $E[\|\mathbf{x}_{k,i}\|] < \infty$. We show a sufficient condition for the latter. Note that there exists a $C' > 0$ such that $|x| \leq C' e^{-x} + x$ for all x . Hence

$$E[\|\mathbf{x}_{k,i}\|] \leq C' E[e^{-\mathbf{x}_{k,i}}] + E[\mathbf{x}_{k,i}] = C' + D(L_k(\cdot | \theta^\circ) \| L_k(\cdot | \theta)). \quad (9.28)$$

Therefore, if all the elements of d are finite, this implies $E[|\mathbf{x}_{k,i}|] < \infty$ for all k .

To extend the result (9.9) under absolute integrability, we use the following lemma.

Lemma 9.3 (Kolmogorov’s Truncation Lemma [82, Chapter 12]). *Consider i.i.d. $\mathbf{x}_1, \mathbf{x}_2, \dots$ where $E[|\mathbf{x}_1|] < \infty$. Let $\mu := E[\mathbf{x}_1]$ be the common mean. Define*

$$\mathbf{y}_i := \begin{cases} \mathbf{x}_i, & |\mathbf{x}_i| \leq i \\ 0, & |\mathbf{x}_i| > i \end{cases} \quad (9.29)$$

Then (i) $\mu_i := E[\mathbf{y}_i] \rightarrow \mu$; (ii) $(\mathbf{y}_i - \mathbf{x}_i) \xrightarrow{\text{a.s.}} 0$; (iii) $\sum i^{-2} E[\mathbf{y}_i^2] < \infty$.

We truncate the vector \mathbf{x}_i elementwise and obtain \mathbf{y}_i , i.e., we relate $\mathbf{x}_{k,i}$ to $\mathbf{y}_{k,i}$ as in (9.29). Let $d_i := E[\mathbf{y}_i]$ and repeat the same steps in the proof of Lemma 9.2 with (\mathbf{y}_i, d_i) instead of (\mathbf{x}_i, d) . Observe that the martingale \mathbf{z}_i is square-integrable because the sum in (9.27) is finite according to (iii) of Lemma 9.3. Then we have

$$\frac{1}{i} \sum_{n=1}^i \langle \tilde{\mathbf{w}}_n, \mathbf{y}_n - d \rangle \xrightarrow{\text{a.s.}} 0. \quad (9.30)$$

Finally, (i) and (ii) in Lemma 9.3 together imply

$$\frac{1}{i} \sum_{n=1}^i \langle \tilde{\mathbf{w}}_n, \mathbf{x}_n - d \rangle \xrightarrow{\text{a.s.}} 0, \quad (9.31)$$

which yields the following conclusion.

Theorem 9.1 (Asymptotic Convergence Rate). *Suppose all elements of d are finite. Then for all k ,*

$$\frac{1}{i} \boldsymbol{\lambda}_{k,i} \xrightarrow{\text{a.s.}} \langle \boldsymbol{\pi}, d \rangle. \quad (9.32)$$

From Theorem 9.1, it is immediate that all agents learn the truth eventually if $\langle \boldsymbol{\pi}, d \rangle > 0$. Since all elements of $\boldsymbol{\pi}$ are positive, this condition — also called global identifiability — holds if at least one element of d is strictly positive.

Corollary 9.2 (Truth Learning). *Suppose all elements of d are finite and at least one element of d is strictly positive. Then for all k ,*

$$\boldsymbol{\mu}_{k,i}(\theta^\circ) \xrightarrow{\text{a.s.}} 1. \quad (9.33)$$

We conclude this section by noting a straightforward extension of our result. Allow each agent to have a different confidence weight α_k . Extensions of Theorem 9.1 and Corollary 9.2 can be obtained as follows: Define the diagonal matrix $J := [j_{\ell k}]$ where $j_{kk} = \alpha_k$; and 0 otherwise. Replace $(\alpha I + \bar{\alpha} \mathbf{A}_n)$ with $(J + \mathbf{A}_n(I - J))$ and $\boldsymbol{\pi}$ with $\tilde{\boldsymbol{\pi}}$, the Perron vector of $(J + A(I - J))$. By following the same steps as above, the extension will be immediate.

9.4 A Special Case: Replacement

In this section we study the special case $\alpha = 0$, where agents replace their beliefs with their neighbors'. At first sight, it is not intuitive whether all agents would learn the truth eventually. This is because the truthful beliefs might be lost upon replacement. However, all the results of Section 9.3 hold for this case as well. Furthermore, when agents replace their beliefs, Theorem 9.1 has a much shorter proof, which we provide below.

Short Proof of Theorem 9.1 for $\alpha = 0$. Again assume $k = 1$. Observe in this case $\tilde{\mathbf{w}}_n$'s — recall the index change $i - n + 1 \mapsto n$ in Lemma 2 — are vectors with their ℓ_n^{th} element being one for some ℓ_n and the others being zero. Now consider a Markov chain governed by the transition kernel A , with the state space $\{1, \dots, K\}$ and observe that the random variable ℓ_n is the state at time n with the initial state being 1. Now, write down $\frac{1}{i}\boldsymbol{\lambda}_i$ as

$$\begin{aligned} \frac{1}{i}\boldsymbol{\lambda}_i &= \frac{1}{i} \sum_{n=1}^i \langle \tilde{\mathbf{w}}_n, \mathbf{x}_n \rangle = \sum_{k=1}^K \frac{1}{i} \sum_{n:\ell_n=k} \mathbf{x}_{k,n} \\ &= \sum_{k=1}^K \frac{\mathbf{m}_i^{(k)}}{i} \frac{1}{\mathbf{m}_i^{(k)}} \sum_{n:\ell_n=k} \mathbf{x}_{k,n} \end{aligned} \quad (9.34)$$

where $\mathbf{m}_i^{(k)}$ is the number of visits to state k in i transitions. Since the Markov chain governed by A is communicating and aperiodic, it is known from a standard result in Renewal theory [182, Chapter 3] that $\frac{\mathbf{m}_i^{(k)}}{i} \xrightarrow{\text{a.s.}} \pi_k$. This also implies $\mathbf{m}_i^{(k)} \xrightarrow{\text{a.s.}} \infty$ and therefore $\frac{1}{\mathbf{m}_i^{(k)}} \sum_{n:\ell_n=k} \mathbf{x}_{k,n} \xrightarrow{\text{a.s.}} d_k$ by the strong law of large numbers. Combining these with (9.34), we see (9.9) holds. \square

Suppose the truth learning terminates at step i for a large i . We now shift our attention to the event that a node makes an error upon termination. More precisely, we are interested in evaluating the probability of this event. To this end, we emphasize an important observation from the proof above: $\tilde{\mathbf{w}}_i$ evolves according to a finite-state Markov chain. Moreover, $\frac{1}{i}\boldsymbol{\lambda}_i$ can be viewed as the average reward of a Markov reward process. Knowing the underlying dependence structure of $\tilde{\mathbf{w}}_i$'s allows us to invoke the known results in large deviations theory. We have a special case of Gärtner–Ellis theorem that implies — see also [183, 184]:

Theorem 9.2 (Theorem 3.1.2 in [185]). *Set $A(t) := [a_{\ell k} E[e^{t\mathbf{x}_\ell}]]$ and let $\Lambda(t)$ be the Perron–Frobenius eigenvalue of $A(t)$. Then for any $\Gamma \subseteq \mathbb{R}$,*

$$-\inf_{s \in \Gamma^\circ} I(s) \leq \liminf_{i \rightarrow \infty} \frac{1}{i} \log P\left(\frac{1}{i}\boldsymbol{\lambda}_i \in \Gamma\right) \quad (9.35)$$

$$\leq \limsup_{i \rightarrow \infty} \frac{1}{i} \log P\left(\frac{1}{i}\boldsymbol{\lambda}_i \in \Gamma\right) \leq -\inf_{s \in \bar{\Gamma}} I(s) \quad (9.36)$$

where $I(s) := \sup_{t \in \mathbb{R}} st - \log \Lambda(t)$ is the Legendre–Fenchel transform of $\log \Lambda(t)$ and $\Gamma^\circ, \bar{\Gamma}$ denote interior and closure of Γ respectively.

Note that Theorem 9.2 only requires the marginals to be known, i.e., $L_k(\cdot|\theta)$; and $E[e^{tx_\ell}]$ to be finite. Therefore, without any knowledge of the joint distribution of the data, the rate function $I(s)$ can be calculated. This allows to approximate the error probabilities of agents given their respective decision rules. For instance, suppose node k decides based on a maximum-a-posteriori rule, i.e., believes $\arg \max_{\theta} \mu_{k,i}(\theta)$. Let $\mathcal{E}_{k,i}$ denote the event that the node k makes an error at time i , whose probability at $i \rightarrow \infty$ can be lower and upper bounded as

$$-\max_{\theta \neq \theta^\circ} \inf_{s < 0} I_\theta(s) \leq \liminf_{i \rightarrow \infty} \frac{1}{i} \log P(\mathcal{E}_{k,i}) \tag{9.37}$$

$$\leq \limsup_{i \rightarrow \infty} \frac{1}{i} \log P(\mathcal{E}_{k,i}) \leq -\max_{\theta \neq \theta^\circ} \inf_{s \leq 0} I_\theta(s) \tag{9.38}$$

where $I_\theta(s)$ is the rate function corresponding to $\lambda_{k,i}(\theta)$. The bounds above do not depend on k .




Figure 91 – (a) The network with $K = 10$ nodes and 12 edges. (b) Sample paths of $\frac{1}{i} \lambda_{1,i}$ and $\frac{1}{i} \lambda_{6,i}$, shown by blue and red curves respectively. The asymptotical convergence rate $\langle \pi, d \rangle$ is drawn as a straight horizontal line at 4.55.

9.5 Numerical Results

In this section we present numerical results based on the simulations performed over the network of $K = 10$ nodes in Figure 91a. We set $\alpha = 0$, i.e., we simulate the replacement algorithm of Section 9.4, and aim to solve a binary hypothesis

testing problem between θ° and θ . For simplicity, we assume that the data is independent across agents — note that the rate function $I_\theta(s)$ is unaffected by such assumption. Moreover, node k observes a Gaussian random variable with unit variance under each hypothesis; with zero mean under θ° , and with mean ν_k under θ . We have set $\nu := [3, 8, 0, 0, 3, 0, 3, 0, 0, 0]$, so for instance $\nu_5 = 3$. Observe that θ° can only be identified by nodes 1,2,5 and 7. The combination matrix A is chosen according to a lazy Metropolis rule [186], namely we set $B := [b_{\ell k}]$ with $b_{\ell k} = \max\{\deg(\ell), \deg(k)\}^{-1}$ for $\ell \neq k$ and $b_{\ell\ell} = 1 - \sum_{k \neq \ell} b_{\ell k}$. Then we set $A = \frac{1}{2}I + \frac{1}{2}B$.

A is symmetric, hence doubly stochastic; which implies $\pi_k = \frac{1}{K}$ for all k . Furthermore, we can straightforwardly calculate $d_k = E\left[\frac{L_k(\xi_k|\theta^\circ)}{L_k(\xi_k|\theta)}\right] = \frac{1}{2}\nu_k^2$; which gives the asymptotic convergence rate as $\langle \pi, d \rangle = 4.55$. Figure 91b shows two sample paths corresponding to $\frac{1}{i}\lambda_{1,i}$ and $\frac{1}{i}\lambda_{6,i}$ for $i \leq 2500$; and is consistent with our theoretical results from Section 9.3 as the paths seem to converge to 4.55.




Figure 92 – Large deviation estimates. I_θ is found numerically and drawn as the black curve; which touches zero exactly at $\langle \pi, d \rangle$. The Monte-Carlo estimates for $\hat{p}_{2,i}(s)$ and $\hat{p}_{6,i}(s)$ are drawn as blue and red curves respectively.

In Figure 92, we have drawn the rate function $I_\theta(s)$ as a black solid line. Note that $I_\theta(s)$ touches zero exactly at $s = 4.55$, indicated with the solid diamond. We have also obtained Monte Carlo estimates of deviations $p_{k,i}(s) := P(\frac{1}{i}\lambda_{k,i} > s)$ for $s > 4.55$ and $P(\frac{1}{i}\lambda_{k,i} < s)$ for $s < 4.55$; $k = 2, 6$, $i = 2500$ to check if they fit in with $I_\theta(s)$. However, the rate function suggests that one should expect $p_{k,i}(s)$ to become exponentially small with i . Hence, standard Monte Carlo method requires that the number of experiments should be exponentially large with i , which is impractical. We therefore resort to an importance sampling method by using a tilted Markov chain and tilted

Gaussians, where the tiltings depend on s . We omit the details of the tilting procedure, and refer the reader to [187, 188] for details. Denote the tilted measure as Q_s . Then our Monte Carlo estimate (for $s < 4.55$) is

$$\hat{p}_{k,i}(s) := \frac{1}{N} \sum_{n=1}^N \frac{dP^{(n)}}{dQ_s^{(n)}} \mathbb{1}\{\lambda_{k,i}^{(n)} < s\} \quad (9.39)$$

where the superscript (n) denotes the n^{th} realization under the tilted measure Q_s . $\frac{dP^{(n)}}{dQ_s^{(n)}}$ is a measure change variable, i.e., Radon–Nikodym derivative, which turns out to be expressed as a product $\prod_{j=1}^i f(x_j^{(n)}, w_j^{(n)})$ for some f . Note that $E_{Q_s}[\hat{p}_{k,i}(s)] = p_{k,i}(s)$. For $s > 4.55$, we replace the indicator function with $\mathbb{1}\{\lambda_{k,i}^{(n)} > s\}$. We performed $N = 60$ experiments to obtain each marker in Figure 92.

9.6 Discussion

Under randomized collaborations, agents still learn the truth, and at the same rate compared to the standard algorithms. Although the asymptotic rate $\lim_{i \rightarrow \infty} \frac{1}{i} \lambda_{k,i}$ does not depend on α , the statistics of $\lambda_{k,i}$ depend on α for finite i . For $\alpha = 0$, we provided a finite-time analysis based on large deviation estimates for finite-state Markov chains. However, for $\alpha > 0$, the corresponding Markov chain has a continuous state space — the states take values in the k -dimensional probability simplex — and a finite-time analysis requires more advanced machinery.

Furthermore, one needs to know the joint distribution across agents in general to make a large deviation analysis of a social learning scheme. This may not be a realistic assumption, if agents only know their marginal likelihoods, who else could know the global likelihood function? Consider the conventional deterministic social learning setting, where $\alpha = 0$ and \mathbf{A} is deterministic. In this case, it is easy to see that $w_i^{(k)}$'s become deterministic and $w_i^{(k)} \rightarrow \pi$ for any agent k . One then faces a standard large-deviation analysis where the rate function is given by the Legendre–Fenchel transform I_θ of

$$\Lambda(t) = \log E[e^{t\langle \pi, \mathbf{x} \rangle}], \quad (9.40)$$

which obviously depends on the joint distribution across agents. If agents perform maximum a posteriori rule, i.e., $s = 0$, the optimal decay rate for rejection of $\theta \neq \theta^\circ$ turns out to be (under some mild conditions on continuity)

$$I_\theta(0) = - \inf_{t \geq 0} \log E[e^{-t\langle \pi, \mathbf{x} \rangle}]. \quad (9.41)$$

One may want to study the worst-case scenario, i.e, assume the worst possible coupling across agents by solving the problem

$$I_\theta^{(\text{worst})}(0) = - \sup_{L(\cdot|\theta^\circ)} \inf_{t \geq 0} \log E[e^{-t\langle \pi, \mathbf{x} \rangle}] \quad (9.42)$$

where $L(\cdot|\theta^\circ)$ is a K -dimensional probability distribution with the constraint that the k th marginal of $L(\cdot|\theta^\circ)$ should be equal to $L_k(\cdot|\theta^\circ)$. This seems a challenging problem and we do not know if $I_\theta^{(\text{worst})}(0)$ is strictly positive. If the nodes, however, are also allowed to randomize their strategies another dimension is added to the problem in (9.42); and it turns out that the worst case exponent can be made positive. In case the reader wonders how this is possible, we have already given the answer in Section 9.4: Agents randomly replace their beliefs.

Part II Conclusion

10

In Part II, we have reviewed two strands of works in distributed inference, adopting information-theoretic and signal processing perspectives. We then studied two distributed inference problems which contain flavors from both perspectives.

The first problem (**Chapter 7**) was formulated under a centralized architecture. At each time instant t , n peripheral nodes were sending quantized data $f^{(i)}(X_t)$, $i \leq n$ to a fusion center and the expected number of bits was kept limited to at most R_i bits per second for each node. For a binary hypothesis testing problem under this setup, and for the single-node case, we found the best possible type-II error exponent $\theta^*(R)$; and obtained a closed form upper bound to $\theta^*(R)$. We later showed that the upper bound can be approached via lattice quantization within a 1.047-bit gap, which is due to the inefficiency of 1-dimensional lattice quantization. We then extended the results to the vector quantization case and also to the multiple-node case when data is independent across nodes. In the multiple-node case, we studied the best type-II error exponent when the sum rate across nodes is kept limited to R and showed that an analogous upper bound can be obtained by waterfilling algorithms.

A possible future work could be to study a slightly reformulated version of the problem. Instead of the expected rate constraint, one may restrict

$$p_t := P\left(\text{len}(f(X_1)) + \cdots + \text{len}(f(X_t)) > Rt\right) \quad (10.1)$$

where $\text{len}(f(X_\tau))$ is the length of the codeword obtained by compressing X_τ (or equivalently L_τ). A modified achievability scheme might impose p_t to vanish as $t \rightarrow \infty$. For this case, we have some evidence that the optimal curve stays the same. However, if one imposes the stronger condition

$$p_t \leq e^{-\gamma t}, \quad (10.2)$$

a new dimension is added to the problem and it might be worth studying.

The second problem (**Chapter 9**) was formulated under a decentralized scheme where nodes/agents were only allowed to communicate with their neighbors. The agents sought to learn the true-state of nature among m possible states and we formulated the problem in a Bayesian framework. Since it was intractable to carry out fully Bayesian updates, we have studied a scheme where nodes update their beliefs locally. More specifically, we studied a social learning scheme where nodes collaborate randomly and we showed that nodes learn the truth eventually with unit probability and at the same rate as the baseline algorithms. For a special case where nodes replaced their beliefs we were able to study the large deviations without knowing the joint distribution across nodes. However, for general communication schemes between the nodes, large deviation analyses seem to require the knowledge of the joint distribution. Thus, for a general scheme, a worst-case problem as in equation 9.42 could be formulated, and if the random replacement strategy belongs to the set of admissible strategies, the worst-case error rate can be shown to be positive. If, however, randomized strategies — and in particular the replacement strategies — are not admissible, the problem given in equation 9.42 could be formidably difficult.

Bibliography

- [1] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, July 2006.
- [2] R. G. Gallager, *Information Theory and Reliable Communication*. USA: John Wiley & Sons, Inc., 1968.
- [3] A. El Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge University Press, 2011.
- [4] I. Csiszár and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011. [Online]. Available: <https://books.google.ch/books?id=2gsLkQlb8JAC>
- [5] C. E. Shannon, “A mathematical theory of communication,” *The Bell System Technical Journal*, vol. 27, pp. 379–423, 1948. [Online]. Available: <http://plan9.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf>
- [6] A. Wyner, “An upper bound on the entropy series,” *Information and Control*, vol. 20, no. 2, pp. 176–181, 1972. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0019995872903658>
- [7] N. Alon and A. Orlitsky, “A lower bound on the expected length of one-to-one codes,” *IEEE Transactions on Information Theory*, vol. 40, no. 5, pp. 1670–1672, 1994.
- [8] E. Arikan, “Channel polarization: a method for constructing capacity-achieving codes for symmetric binary-input memoryless channels,” *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3051–3073, 2009. [Online]. Available: <https://doi.org/10.1109/TIT.2009.2021379>
- [9] Y. Polyanskiy, H. V. Poor, and S. Verdú, “Channel coding rate in the finite blocklength regime,” *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.
- [10] W. Rudin, *Real and Complex Analysis, 3rd Ed.* USA: McGraw-Hill, Inc., 1987.

- [11] I. Csiszár, “Generalized entropy and quantization problems,” *Trans. Sixth Prague Conference on Inform. Theory etc., Academia, Prague*, pp. 159–174, 1971.
- [12] H. Gish and J. Pierce, “Asymptotically efficient quantizing,” *IEEE Transactions on Information Theory*, vol. 14, no. 5, pp. 676–683, 1968.
- [13] H. Chernoff, “A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on the sum of Observations,” *The Annals of Mathematical Statistics*, vol. 23, no. 4, pp. 493 – 507, 1952. [Online]. Available: <https://doi.org/10.1214/aoms/1177729330>
- [14] A. Sahai and S. K. Mitter, “The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link - part I: scalar systems,” *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3369–3395, 2006. [Online]. Available: <https://doi.org/10.1109/TIT.2006.878169>
- [15] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?” in *2012 Proceedings IEEE INFOCOM*, 2012, pp. 2731–2735.
- [16] S. Ross, *Stochastic processes*, ser. Wiley series in probability and statistics: Probability and statistics. Wiley, 1996.
- [17] R. D. Yates, “The age of information in networks: Moments, distributions, and sampling,” *IEEE Transactions on Information Theory*, vol. 66, no. 9, pp. 5712–5728, 2020.
- [18] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, “Minimizing age of information in vehicular networks,” in *2011 8th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, 2011, pp. 350–358.
- [19] S. Kaul, R. Yates, and M. Gruteser, “On piggybacking in vehicular networks,” in *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, 2011, pp. 1–5.
- [20] J. P. Champati, H. Al-Zubaidy, and J. Gross, “Statistical guarantee optimization for age of information for the D/G/1 queue,” in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2018, pp. 130–135.
- [21] A. Soysal and S. Ulukus, “Age of information in G/G/1/1 systems,” in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, 2019, pp. 2022–2027.

- [22] Y. Inoue, H. Masuyama, T. Takine, and T. Tanaka, "A general formula for the stationary distribution of the age of information and its application to single-server queues," *IEEE Transactions on Information Theory*, vol. 65, no. 12, pp. 8305–8324, 2019.
- [23] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *2013 IEEE International Symposium on Information Theory*, 2013, pp. 66–70.
- [24] E. Najm and R. Nasser, "Age of information: The gamma awakening," in *2016 IEEE International Symposium on Information Theory (ISIT)*, 2016, pp. 2574–2578.
- [25] R. D. Yates, J. Zhong, and W. Zhang, "Updates with multiple service classes," in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 1017–1021.
- [26] S. K. Kaul and R. D. Yates, "Age of information: Updates with priority," in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2644–2648.
- [27] R. D. Yates and S. Kaul, "Real-time status updating: Multiple sources," in *2012 IEEE International Symposium on Information Theory Proceedings*, 2012, pp. 2666–2670.
- [28] M. Moltafet, M. Leinonen, and M. Codreanu, "On the age of information in multi-source queueing models," *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 5003–5017, 2020.
- [29] E. Najm and E. Telatar, "Status updates in a multi-stream M/G/1/1 preemptive queue," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2018, pp. 124–129.
- [30] V. Tripathi and S. Moharir, "Age of information in multi-source systems," in *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, 2017, pp. 1–6.
- [31] S. K. Kaul, R. D. Yates, and M. Gruteser, "Status updates through queues," in *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, 2012, pp. 1–6.
- [32] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Transactions on Information Theory*, vol. 62, no. 4, pp. 1897–1910, 2016.
- [33] R. D. Yates, "The age of gossip in networks," in *2021 IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 2984–2989.

- [34] —, “Age of information in a network of preemptive servers,” in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2018, pp. 118–123.
- [35] R. Nasser, I. Issa, and I. Abou-Faycal, “Age distribution in arbitrary preemptive memoryless networks,” in *2022 IEEE International Symposium on Information Theory (ISIT) (ISIT 2022)*, Espoo, Finland, Jun. 2022.
- [36] C. Kam, S. Kompella, G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, “Age of information with a packet deadline,” in *2016 IEEE International Symposium on Information Theory (ISIT)*, 2016, pp. 2564–2568.
- [37] X. Wu, J. Yang, and J. Wu, “Optimal status update for age of information minimization with an energy harvesting source,” *IEEE Transactions on Green Communications and Networking*, vol. 2, no. 1, pp. 193–204, 2018.
- [38] R. D. Yates, “Lazy is timely: Status updates by an energy harvesting source,” in *2015 IEEE International Symposium on Information Theory (ISIT)*, 2015, pp. 3008–3012.
- [39] S. Farazi, A. G. Klein, and D. R. Brown, “Age of information in energy harvesting status update systems: When to preempt in service?” in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2436–2440.
- [40] B. T. Bacinoglu, E. T. Ceran, and E. Uysal-Biyikoglu, “Age of information under energy replenishment constraints,” in *2015 Information Theory and Applications Workshop (ITA)*, 2015, pp. 25–31.
- [41] A. Arafa and S. Ulukus, “Age-minimal transmission in energy harvesting two-hop networks,” in *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, 2017, pp. 1–6.
- [42] —, “Timely updates in energy harvesting two-hop networks: Offline and online policies,” *IEEE Transactions on Wireless Communications*, vol. 18, no. 8, pp. 4017–4030, 2019.
- [43] —, “Age minimization in energy harvesting communications: Energy-controlled delays,” in *2017 51st Asilomar Conference on Signals, Systems, and Computers*, 2017, pp. 1801–1805.
- [44] B. T. Bacinoglu and E. Uysal-Biyikoglu, “Scheduling status updates to minimize age of information with an energy harvesting sensor,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 1122–1126.

- [45] S. Feng and J. Yang, “Minimizing age of information for an energy harvesting source with updating failures,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2431–2435.
- [46] ———, “Optimal status updating for an energy harvesting sensor with a noisy channel,” in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2018, pp. 348–353.
- [47] A. Arafa, J. Yang, and S. Ulukus, “Age-minimal online policies for energy harvesting sensors with random battery recharges,” in *2018 IEEE International Conference on Communications (ICC)*, 2018, pp. 1–6.
- [48] B. T. Bacinoglu, Y. Sun, E. Uysal-Bivikoglu, and V. Mutlu, “Achieving the age-energy tradeoff with a finite-battery energy harvesting source,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 876–880.
- [49] Y. Dong, P. Fan, and K. B. Letaief, “Energy harvesting powered sensing in iot: Timeliness versus distortion,” *IEEE Internet of Things Journal*, vol. 7, no. 11, pp. 10 897–10 911, 2020.
- [50] E. T. Ceran, D. Gündüz, and A. György, “Average age of information with hybrid ARQ under a resource constraint,” in *2018 IEEE Wireless Communications and Networking Conference (WCNC)*, 2018, pp. 1–6.
- [51] B. Wang, S. Feng, and J. Yang, “To skip or to switch? minimizing age of information under link capacity constraint,” in *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2018, pp. 1–5.
- [52] Y. Inan and E. Telatar, “Age-Optimal causal labeling of memoryless processes,” in *2022 IEEE International Symposium on Information Theory (ISIT) (ISIT 2022)*, Espoo, Finland, Jun. 2022.
- [53] J. Zhong and R. D. Yates, “Timeliness in lossless block coding,” in *2016 Data Compression Conference (DCC)*, 2016, pp. 339–348.
- [54] J. Zhong, R. D. Yates, and E. Soljanin, “Timely lossless source coding for randomly arriving symbols,” in *2018 IEEE Information Theory Workshop (ITW)*, 2018, pp. 1–5.
- [55] ———, “Backlog-adaptive compression: Age of information,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 566–570.
- [56] R. D. Yates, Y. Sun, D. R. Brown III, S. K. Kaul, E. Modiano, and S. Ulukus, “Age of information: An introduction and survey,” 2020.

- [57] A. Kosta, N. Pappas, and V. Angelakis, *Age of Information: A New Concept, Metric, and Tool*, 2017.
- [58] K. Chen and L. Huang, “Age-of-information in the presence of error,” in *2016 IEEE International Symposium on Information Theory (ISIT)*, 2016, pp. 2579–2583.
- [59] S. Feng and J. Yang, “Age-optimal transmission of rateless codes in an erasure channel,” in *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 2019, pp. 1–6.
- [60] A. Baknina and S. Ulukus, “Coded status updates in an energy harvesting erasure channel,” in *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, 2018, pp. 1–6.
- [61] S. Feng and J. Yang, “Age of information minimization for an energy harvesting source with updating erasures: Without and with feedback,” *IEEE Transactions on Communications*, vol. 69, no. 8, pp. 5091–5105, 2021.
- [62] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor, “Using erasure feedback for online timely updating with an energy harvesting sensor,” in *2019 IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 607–611.
- [63] P. Parag, A. Taghavi, and J.-F. Chamberland, “On real-time status updates over symbol erasure channels,” in *2017 IEEE Wireless Communications and Networking Conference (WCNC)*, 2017, pp. 1–6.
- [64] R. D. Yates, E. Najm, E. Soljanin, and J. Zhong, “Timely updates over an erasure channel,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 316–320.
- [65] H. Sac, T. Bacinoglu, E. Uysal-Biyikoglu, and G. Durisi, “Age-optimal channel coding blocklength for an M/G/1 queue with HARQ,” in *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2018, pp. 1–5.
- [66] E. Najm, R. Yates, and E. Soljanin, “Status updates through M/G/1/1 queues with HARQ,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 131–135.
- [67] E. Najm, E. Telatar, and R. Nasser, “Optimal age over erasure channels,” *IEEE Transactions on Information Theory*, pp. 5901–5922, 2022.
- [68] T. Soleymani, J. S. Baras, and K. H. Johansson, “Stochastic control with stale information—part i: Fully observable systems,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 4178–4182.

- [69] A. Mitra, J. A. Richards, S. Bagchi, and S. Sundaram, “Finite-time distributed state estimation over time-varying graphs: Exploiting the age-of-information,” in *2019 American Control Conference (ACC)*, 2019, pp. 4006–4011.
- [70] Y. Sun, Y. Polyanskiy, and E. Uysal, “Sampling of the wiener process for remote estimation over a channel with random delay,” *IEEE Transactions on Information Theory*, vol. 66, no. 2, pp. 1118–1135, 2020.
- [71] T. Z. Ornee and Y. Sun, “Sampling and remote estimation for the ornstein-uhlenbeck process through queues: Age of information and beyond,” *IEEE/ACM Transactions on Networking*, vol. 29, no. 5, pp. 1962–1975, 2021.
- [72] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, “Age and value of information: Non-linear age case,” in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 326–330.
- [73] ———, “The cost of delay in status updates and their value: Non-linear ageing,” *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 4905–4918, 2020.
- [74] S. K. Kaul and R. D. Yates, “Age of information: Updates with priority,” in *2018 IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 2644–2648.
- [75] E. Najm, R. Nasser, and E. Telatar, “Content based status updates,” *IEEE Transactions on Information Theory*, vol. 66, no. 6, pp. 3846–3863, 2020.
- [76] M. Bastopcu and S. Ulukus, “Age of information for updates with distortion,” in *2019 IEEE Information Theory Workshop (ITW)*, 2019, pp. 1–5.
- [77] Y. Inan, R. Inovan, and E. Telatar, “Optimal policies for age and distortion in a discrete-time model,” in *2021 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.
- [78] Y. Inan, R. Inovan, and E. Telatar, “Optimal policies for age and distortion in a discrete-time model,” 2022.
- [79] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksall, and N. B. Shroff, “Update or wait: How to keep your data fresh,” *IEEE Transactions on Information Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.
- [80] L. I. Sennott, “Constrained average cost markov decision chains,” *Probability in the Engineering and Informational Sciences*, vol. 7, no. 1, p. 69–83, 1993.

- [81] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. II*, 3rd ed. Athena Scientific, 2007.
- [82] D. Williams, *Probability with Martingales.*, ser. Cambridge mathematical textbooks. Cambridge University Press, 1991.
- [83] S. Resnick, *A Probability Path*, ser. Modern Birkhäuser Classics. Birkhäuser Boston, 2003.
- [84] B. T. Bacinoglu, Y. Sun, E. Uysal, and V. Mutlu, “Optimal status updating with a finite-battery energy harvesting source,” *Journal of Communications and Networks*, vol. 21, no. 3, pp. 280–294, 2019.
- [85] A. Arafa, J. Yang, S. Ulukus, and H. V. Poor, “Age-minimal transmission for energy harvesting sensors with finite batteries: Online policies,” *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 534–556, 2020.
- [86] O. Ozel, “Timely status updating through intermittent sensing and transmission,” in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1788–1793.
- [87] “IEEE standard for information technology—part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications,” *IEEE Std 802.11-2012*, pp. 1–2793, 2012.
- [88] R. Ahlswede and I. Csiszár, “Hypothesis testing with communication constraints,” *IEEE Transactions on Information Theory*, vol. 32, no. 4, pp. 533–542, 1986.
- [89] T. Han, “Hypothesis testing with multiterminal data compression,” *IEEE Transactions on Information Theory*, vol. 33, no. 6, pp. 759–772, 1987.
- [90] H. Shalaby and A. Papamarcou, “Multiterminal detection with zero-rate data compression,” *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 254–267, 1992.
- [91] H. Shimokawa, T. S. Han, and S. Amari, “Error bound of hypothesis testing with data compression,” in *Proc. IEEE International Symposium on Information Theory*, 1994, pp. 114–.
- [92] T. S. Han and S. Amari, “Statistical inference under multiterminal data compression,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2300–2324, 1998.
- [93] T. Han and K. Kobayashi, “Exponential-type error probabilities for multiterminal hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 35, no. 1, pp. 2–14, 1989.

- [94] N. Weinberger and Y. Kochman, "On the reliability function of distributed hypothesis testing under optimal detection," *IEEE Transactions on Information Theory*, vol. 65, no. 8, pp. 4940–4965, 2019.
- [95] S. Watanabe, "Neyman-Pearson test for zero-rate multiterminal hypothesis testing," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 116–120.
- [96] E. Tuncel, "On error exponents in hypothesis testing," *IEEE Transactions on Information Theory*, vol. 51, no. 8, pp. 2945–2950, 2005.
- [97] S. Sreekumar and D. Gündüz, "Distributed hypothesis testing over discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 2044–2066, 2020.
- [98] M. Wigger and R. Timo, "Testing against independence with multiple decision centers," in *2016 International Conference on Signal Processing and Communications (SPCOM)*, 2016, pp. 1–5.
- [99] S. Salehkalaibar and M. Wigger, "Distributed hypothesis testing over noisy broadcast channels," *Information*, vol. 12, no. 7, 2021. [Online]. Available: <https://www.mdpi.com/2078-2489/12/7/268>
- [100] G. Katz, P. Piantanida, and M. Debbah, "Collaborative distributed hypothesis testing," Jan. 2016, *IEEE Trans. on Information Theory* (submitted). [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01436767>
- [101] Y. Xiang and Y.-H. Kim, "Interactive hypothesis testing against independence," in *2013 IEEE International Symposium on Information Theory*, 2013, pp. 2840–2844.
- [102] P. Escamilla, M. Wigger, and A. Zaidi, "Distributed hypothesis testing: Cooperation and concurrent detection," *IEEE Transactions on Information Theory*, vol. 66, no. 12, pp. 7550–7564, 2020.
- [103] M. Mhanna and P. Piantanida, "On secure distributed hypothesis testing," in *2015 IEEE International Symposium on Information Theory (ISIT)*, 2015, pp. 1605–1609.
- [104] S. Sreekumar, D. Gündüz, and A. Cohen, "Distributed hypothesis testing under privacy constraints," in *2018 IEEE Information Theory Workshop (ITW)*, 2018, pp. 1–5.
- [105] A. Gilani, S. Belhadj Amor, S. Salehkalaibar, and V. Y. F. Tan, "Distributed hypothesis testing with privacy constraints," *Entropy*, vol. 21, no. 5, 2019. [Online]. Available: <https://www.mdpi.com/1099-4300/21/5/478>

- [106] J. Liao, L. Sankar, V. Y. F. Tan, and F. du Pin Calmon, "Hypothesis testing under mutual information privacy constraints in the high privacy regime," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 4, pp. 1058–1071, 2018.
- [107] R. Abbasalipour and M. Mirmohseni, "Privacy-aware distributed hypothesis testing in Gray-Wyner network with side information," *arXiv preprint arXiv:2202.02307*, 2022.
- [108] J. Liao, L. Sankar, F. P. Calmon, and V. Y. F. Tan, "Hypothesis testing under maximal leakage privacy constraints," in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 779–783.
- [109] G. Katz, P. Piantanida, and M. Debbah, "Distributed binary detection with lossy data compression," *IEEE Transactions on Information Theory*, vol. 63, no. 8, pp. 5207–5227, 2017.
- [110] M. S. Rahman and A. B. Wagner, "On the optimality of binning for distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 58, no. 10, pp. 6282–6303, 2012.
- [111] S. Watanabe, "On sub-optimality of random binning for distributed hypothesis testing," 2022. [Online]. Available: <https://arxiv.org/abs/2201.13005>
- [112] S. Salehkalaibar and M. Wigger, "Distributed hypothesis testing with variable-length coding," in *Proc. International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2020, pp. 1–5.
- [113] M. Hamad, M. Wigger, and M. Sarkiss, "Cooperative multi-sensor detection under variable-length coding," in *Proc. IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–5.
- [114] —, "Two-hop network with multiple decision centers under expected-rate constraints," in *GLOBECOM 2021: IEEE Global Communications Conference*. Madrid, Spain: IEEE, Dec. 2021, pp. 1–6. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03349810>
- [115] —, "Optimal exponents in cascaded hypothesis testing under expected rate constraints," in *Proc. IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.
- [116] —, "Cooperative multi-sensor detection under variable-length coding," in *ITW 2020: IEEE Information Theory Workshop*, ser. 2020 IEEE Information Theory Workshop (ITW). Riva del Garda, Italy: IEEE, Apr. 2021, pp. 1–5. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03349654>

- [117] M. Hamad, M. Sarkiss, and M. Wigger, “Benefits of rate-sharing for distributed hypothesis testing,” 2022. [Online]. Available: <https://arxiv.org/abs/2202.02282>
- [118] S. M. Ali and S. D. Silvey, “A general class of coefficients of divergence of one distribution from another,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 28, no. 1, pp. 131–142, 1966.
- [119] I. Csiszár and P. Shields, “Information theory and statistics: A tutorial,” *Foundations and Trends® in Communications and Information Theory*, vol. 1, no. 4, pp. 417–528, 2004. [Online]. Available: <http://dx.doi.org/10.1561/0100000004>
- [120] T. Kailath, “The divergence and bhattacharyya distance measures in signal selection,” *IEEE Transactions on Communication Technology*, vol. 15, no. 1, pp. 52–60, 1967.
- [121] S. Kassam, “Optimum quantization for signal detection,” *IEEE Transactions on Communications*, vol. 25, no. 5, pp. 479–484, 1977.
- [122] H. Poor and J. Thomas, “Applications of ali-silvey distance measures in the design generalized quantizers for binary decision systems,” *IEEE Transactions on Communications*, vol. 25, no. 9, pp. 893–900, 1977.
- [123] M. Longo, T. Lookabaugh, and R. Gray, “Quantization for decentralized hypothesis testing under communication constraints,” *IEEE Transactions on Information Theory*, vol. 36, no. 2, pp. 241–255, 1990.
- [124] H. Poor, “Fine quantization in signal detection and estimation,” *IEEE Transactions on Information Theory*, vol. 34, no. 5, pp. 960–972, 1988.
- [125] R. Gupta and A. Hero, “High-rate vector quantization for detection,” *IEEE Transactions on Information Theory*, vol. 49, no. 8, pp. 1951–1969, 2003.
- [126] J. Villard and P. Bianchi, “High-rate vector quantization for the Neyman–Pearson detection of correlated processes,” *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5387–5409, 2011.
- [127] J. Tsitsiklis, “Extremal properties of likelihood-ratio quantizers,” *IEEE Transactions on Communications*, vol. 41, no. 4, pp. 550–558, 1993.
- [128] V. S. S. Nadendla and P. K. Varshney, “Design of binary quantizers for distributed detection under secrecy constraints,” *IEEE Transactions on Signal Processing*, vol. 64, no. 10, pp. 2636–2648, 2016.
- [129] M. Mhanna, P. Duhamel, and P. Piantanida, “Quantization for distributed binary detection under secrecy constraints,” in *2016 IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.

- [130] V. Saligrama, M. Alanyali, and O. Savas, “Distributed detection in sensor networks with packet losses and finite capacity links,” *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4118–4132, 2006.
- [131] B. Liu and B. Chen, “Channel-optimized quantizers for decentralized detection in sensor networks,” *IEEE Transactions on Information Theory*, vol. 52, no. 7, pp. 3349–3358, 2006.
- [132] L. Cao and R. Viswanathan, “Divergence-based soft decision for error resilient decentralized signal detection,” *IEEE Transactions on Signal Processing*, vol. 62, no. 19, pp. 5095–5106, 2014.
- [133] S. Marano, V. Matta, and L. Tong, “Distributed detection in the presence of byzantine attacks,” *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 16–29, 2009.
- [134] G. Gül and M. Baßler, “Scalable multilevel quantization for distributed detection,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5200–5204.
- [135] Y. Inan, M. Kayaalp, A. H. Sayed, and E. Telatar, “A fundamental limit of distributed hypothesis testing under memoryless quantization,” in *2022 IEEE International Conference on Communications (ICC): Communication Theory Symposium (IEEE ICC’22 - CT Symposium)*, Seoul, Korea (South), May 2022.
- [136] —, “A fundamental limit of distributed hypothesis testing under memoryless quantization,” *arXiv:2206.12232*, 2022.
- [137] C. E. Shannon, *Coding Theorems for a Discrete Source With a Fidelity Criterion* *Institute of Radio Engineers, International Convention Record, vol. 7, 1959.*, 1993, pp. 325–350.
- [138] “NIST Digital Library of Mathematical Functions,” <http://dlmf.nist.gov/>, Release 1.1.5 of 2022-03-15, f. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. [Online]. Available: <http://dlmf.nist.gov/>
- [139] J. Conway and N. Sloane, “Fast quantizing and decoding and algorithms for lattice quantizers and codes,” *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 227–232, 1982.
- [140] V. Kostina, “Data compression with low distortion and finite block-length,” in *Proc. Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2015, pp. 1127–1134.
- [141] D. Blackwell, “Comparison of experiments,” *Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102, 1951.

- [142] R. T. Rockafellar, *Convex analysis*, ser. Princeton Mathematical Series. Princeton, N. J.: Princeton University Press, 1970.
- [143] Y. Polyanskiy and Y. Wu, “Lecture notes on information theory,” *Lecture Notes for ECE563 (UIUC) and*, vol. 6, no. 2012-2016, p. 7, 2014.
- [144] M. Kayaalp, Y. Inan, E. Telatar, and A. H. Sayed, “On the arithmetic and geometric fusion of beliefs for distributed inference,” *arXiv:2204.13741*, 2022.
- [145] M. Kayaalp, Y. Inan, V. Koivunen, E. Telatar, and A. H. Sayed, “On the fusion strategies for federated decision making,” *arXiv:2303.06109*, 2023.
- [146] A. Lalitha, T. Javidi, and A. D. Sarwate, “Social learning and distributed hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.
- [147] V. Matta, P. Braca, S. Marano, and A. H. Sayed, “Distributed detection over adaptive networks: Refined asymptotics and the role of connectivity,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 4, pp. 442–460, 2016.
- [148] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, “Non-Bayesian social learning,” *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.
- [149] A. Nedić, A. Olshevsky, and C. A. Uribe, “Fast convergence rates for distributed non-Bayesian learning,” *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.
- [150] C. P. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, 2003.
- [151] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, “Bayesian learning in social networks,” *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.
- [152] V. Krishnamurthy and H. V. Poor, “Social learning and Bayesian games in multiagent signal processing: how do local and global decision makers interact?” *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.
- [153] C. Chamley, A. Scaglione, and L. Li, “Models for the diffusion of beliefs in social networks: An overview,” *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 16–29, 2013.
- [154] V. Bala and S. Goyal, “Learning from neighbours,” *The Review of Economic Studies*, vol. 65, no. 3, pp. 595–621, 1998.

- [155] ———, “Conformism and diversity under social learning,” *Economic Theory*, vol. 17, no. 1, pp. 101–120, 2001.
- [156] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 909–968, 08 2003.
- [157] L. G. Epstein, J. Noor, and A. Sandroni, “Non-Bayesian updating: a theoretical framework,” *Theoretical Economics*, vol. 3, no. 2, pp. 193–229, 2008.
- [158] ———, “Non-Bayesian learning,” *The BE Journal of Theoretical Economics*, vol. 10, no. 1, pp. 1–16, Jan. 2010.
- [159] B. Golub and M. O. Jackson, “Naïve learning in social networks and the wisdom of crowds,” *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–49, February 2010.
- [160] X. Zhao and A. H. Sayed, “Learning over social networks via diffusion adaptation,” in *Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, 2012, pp. 709–713.
- [161] M. H. DeGroot, “Reaching a consensus,” *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [162] A. H. Sayed, “Adaptation, learning, and optimization over networks,” *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, July 2014.
- [163] H. Salami, B. Ying, and A. H. Sayed, “Social learning over weakly connected graphs,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 222–238, 2017.
- [164] V. Shumovskaia, K. Ntemos, S. Vlaski, and A. H. Sayed, “Online graph learning from social interactions,” in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 1263–1267.
- [165] K. Ntemos, V. Bordignon, S. Vlaski, and A. H. Sayed, “Social learning with disparate hypotheses,” in *2022 30th European Signal Processing Conference (EUSIPCO)*, 2022, pp. 2171–2175.
- [166] V. Bordignon, V. Matta, and A. H. Sayed, “Partial information sharing over social learning networks,” *IEEE Transactions on Information Theory*, vol. 69, no. 3, pp. 2033–2058, 2023.
- [167] V. Shumovskaia, M. Kayaalp, M. Cemri, and A. H. Sayed, “Discovering influencers in opinion formation over social graphs,” *IEEE Open Journal of Signal Processing*, vol. 4, pp. 188–207, 2023.

- [168] K. Ntemos, V. Bordignon, S. Vlaski, and A. H. Sayed, “Self-aware social learning over graphs,” *IEEE Transactions on Information Theory*, pp. 1–1, 2023.
- [169] P. Hu, V. Bordignon, S. Vlaski, and A. H. Sayed, “Optimal aggregation strategies for social learning over graphs,” *IEEE Transactions on Information Theory*, pp. 1–1, 2023.
- [170] V. Bordignon, V. Matta, and A. H. Sayed, “Adaptive social learning,” *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.
- [171] M. Kayaalp, V. Bordignon, and A. H. Sayed, “Social opinion formation under communication trends,” *arXiv:2203.02466*, 2023.
- [172] M. Cemri, V. Bordignon, M. Kayaalp, V. Shumovskaia, and A. H. Sayed, “Asynchronous social learning,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [173] M. Cirillo, V. Bordignon, V. Matta, and H. Sayed, “The role of memory in social learning when sharing partial opinions,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [174] V. Shumovskaia, M. Kayaalp, and A. H. Sayed, “Identifying opinion influencers over social networks,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [175] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, “Gossip algorithms for distributed signal processing,” *Proceedings of the IEEE*, vol. 98, no. 11, pp. 1847–1864, 2010.
- [176] S. Shahrampour and A. Jadbabaie, “Exponentially fast parameter estimation in networks using distributed dual averaging,” in *Proc. IEEE Conference on Decision and Control*, 2013, pp. 6196–6201.
- [177] X. Zhao and A. H. Sayed, “Asynchronous adaptation and learning over networks—part ii: Performance analysis,” *IEEE Transactions on Signal Processing*, vol. 63, no. 4, pp. 827–842, 2015.
- [178] —, “Asynchronous adaptation and learning over networks—part iii: Comparison analysis,” *IEEE Transactions on Signal Processing*, vol. 63, no. 4, pp. 843–858, 2015.
- [179] Y. Inan, M. Kayaalp, E. Telatar, and A. Sayed, “Social learning under randomized collaborations,” in *2022 IEEE International Symposium on Information Theory (ISIT) (ISIT 2022)*, Espoo, Finland, Jun. 2022.

- [180] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. American Mathematical Society, 2006.
- [181] G. Grimmett and D. Stirzaker, *Probability and random processes*. Oxford university press, 2001.
- [182] S. Ross, *Stochastic processes*, ser. Wiley series in probability and statistics: Probability and statistics. Wiley, 1996.
- [183] V. Matta, P. Braca, S. Marano, and A. H. Sayed, “Distributed detection over adaptive networks: Refined asymptotics and the role of connectivity,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 4, pp. 442–460, 2016.
- [184] —, “Diffusion-based adaptive distributed detection: Steady-state performance in the slow adaptation regime,” *IEEE Transactions on Information Theory*, vol. 62, no. 8, pp. 4710–4732, 2016.
- [185] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, ser. Stochastic Modelling and Applied Probability. Springer Berlin Heidelberg, 2009.
- [186] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of State Calculations by Fast Computing Machines,” *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, Jun. 1953.
- [187] J. F. Collamore, “Importance sampling techniques for the multidimensional ruin problem for general markov additive sequences of random vectors,” *The Annals of Applied Probability*, vol. 12, no. 1, pp. 382–421, 2002.
- [188] I. Iscoe, P. Ney, and E. Nummelin, “Large deviations of uniformly recurrent markov additive processes,” *Advances in Applied Mathematics*, vol. 6, no. 4, pp. 373–412, 1985.

Curriculum Vitae

Education

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Ph.D. in Computer and Communication Sciences (2018–2023)

Dissertation Title: *“Time vs. Truth: Age–Distortion Tradeoffs and Strategies for Distributed Inference”*

Thesis Advisor: Prof. Emre Telatar

Bilkent University, Turkey

B.Sc. in Electrical and Electronics Engineering (2013–2018)

GPA: 3.99/4.00

Academic Exchange Semester at EPFL, Grade: 6.00/6.00

Research Experience

Research Assistant at EPFL (2018–2023)

Advisor: Prof. Emre Telatar

Research Project Student at Bilkent University (Fall 2017)

Advisor: Prof. Sinan Gezici

Teaching Experience

Teaching Assistant at EPFL

- Information Theory and Coding (2018–2022)
 - Advanced Probability and Applications (2020–2022)
-

Working Experience

Bilkent University/Meteksan Defense, Türkiye

Industrial Design Project; FM Based Indoor Localization (September 2017–June 2018)

Signal Processing Laboratory (LTS4), EPFL/Hôpital Ophtalmique Jules–Gonin

Semester Project and Internship

Software Development for Automatization of Strabismus Tests (March 2017–September 2017)

ASELSAN Defense Industries, Türkiye

Internship on digital system design (July 2016–August 2016)

Meteksan Defense, Türkiye

Internship on Microwave filter design (July 2015–August 2015)

Honors

- 22nd IEEE Statistical Signal Processing Workshop (SSP) Best Student’s Paper Award (2023)
In recognition of “On the Fusion Strategies for Federated Decision Making”.
- Fellowship Ph.D. candidate at EPFL IC doctoral school (2018–2023)
- Ranked 2nd among the 2014 cohort of Bilkent University EE Department (2018)
- Bilkent University Department of EE Academic Excellence Award (2018)
- Bilkent University Comprehensive Scholarship (2013–2018)

Publications and Preprints

- M. Kayaalp, Y. İnan, V. Koivunen, E. Telatar, and A. H. Sayed, “On the Fusion Strategies for Federated Decision Making”, 2023. [Accepted to IEEE Statistical Signal Processing Workshop, 2023].
Preprint: <https://arxiv.org/abs/2303.06109>.
- Y. İnan, R. Inovon, and E. Telatar, “Optimal policies for age and distortion in a discrete-time model,” 2022. [Full version].
Preprint: <https://arxiv.org/abs/2210.12086>.
- Y. İnan, M. Kayaalp, A. H. Sayed, and E. Telatar, “A fundamental limit of distributed hypothesis testing under memoryless quantization,” 2022. [Full version]. Preprint: <https://arxiv.org/abs/2206.12232>.

- M. Kayaalp, Y. İnan, E. Telatar, and A. H. Sayed, “On the arithmetic and geometric fusion of beliefs for distributed inference,” 2022. [Accepted to IEEE Transactions on Automatic Control].
Preprint: <https://arxiv.org/abs/2204.13741>.
- Y. İnan, M. Kayaalp, E. Telatar, and A. H. Sayed, “Social learning under randomized collaborations,” in Proc. IEEE International Symposium on Information Theory (ISIT), Helsinki, 2022.
- Y. İnan and E. Telatar, “Age-optimal causal labeling of memoryless processes,” in Proc. IEEE International Symposium on Information Theory (ISIT), Helsinki, 2022.
- Y. İnan, M. Kayaalp, A. H. Sayed, and E. Telatar, “A fundamental limit of distributed hypothesis testing under memoryless quantization,” in Proc. IEEE International Conference on Communications (ICC), pp. 1–5, Seoul, 2022.
- Y. İnan, R. Inovan, and E. Telatar, “Optimal policies for age and distortion in a discrete-time model,” in Proc. IEEE Information Theory Workshop (ITW), 2021.
- Y. İnan and E. Telatar, “Finite-level quantization procedures for construction and decoding of polar codes,” in Proc. IEEE International Symposium on Information Theory (ISIT), 2020.

Invited Talks

- “Age–Distortion Tradeoffs”, *IISc Bangalore Networks Seminar Series*, March 2023. [Link: <https://www.youtube.com/watch?v=8ryHuwy5p5I>]