DELAY-SENSITIVE COMMUNICATIONS:

CODE-RATES, STRATEGIES, AND DISTRIBUTED CONTROL

A Dissertation

by

PARIMAL PARAG

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

December 2011

Major Subject: Electrical Engineering

DELAY-SENSITIVE COMMUNICATIONS:

CODE-RATES, STRATEGIES, AND DISTRIBUTED CONTROL

A Dissertation

by

PARIMAL PARAG

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Approved by:

Co-Chairs of Committee,	Jean-François Chamberland
	Srinivas Shakkottai
Committee Members,	Costas Georghiades
	Aniruddha Datta
	Thomas Schlumprecht
Head of Department,	Costas Georghiades

December 2011

Major Subject: Electrical Engineering

ABSTRACT

Delay-sensitive Communications:

Code-Rates, Strategies, and Distributed Control. (December 2011)

Parimal Parag, B. Tech., Indian Institute of Technology Madras;

M. Tech., Indian Institute of Technology Madras

Co–Chairs of Advisory Committee: Dr. Jean-François Chamberland
Dr. Srinivas Shakkottai

An ever increasing demand for instant and reliable information on modern communication networks forces codewords to operate in a non-asymptotic regime. To achieve reliability for imperfect channels in this regime, codewords need to be retransmitted from receiver to the transmit buffer, aided by a fast feedback mechanism. Large occupancy of this buffer results in longer communication delays. Therefore, codewords need to be designed carefully to reduce transmit queue-length and thus the delay experienced in this buffer. We first study the consequences of physical layer decisions on the transmit buffer occupancy. We develop an analytical framework to relate physical layer channel to the transmit buffer occupancy. We compute the optimal code-rate for finite-length codewords operating over a correlated channel, under certain communication service guarantees. We show that channel memory has a significant impact on this optimal code-rate.

Next, we study the delay in small ad-hoc networks. In particular, we find out what rates can be supported on a small network, when each flow has a certain end-to-end service guarantee. To this end, service guarantee at each intermediate link is characterized. These results are applied to study the potential benefits of setting

up a network suitable for network coding in multicast. In particular, we quantify the gains of network coding over classic routing for service provisioned multicast communication over butterfly networks. In the wireless setting, we study the trade-off between communications gains achieved by network coding and the cost to set-up a network enabling network coding. In particular, we show existence of scenarios where one should not attempt to create a network suitable for coding.

Insights obtained from these studies are applied to design a distributed rate control algorithm in a large network. This algorithm maximizes sum-utility of all flows, while satisfying per-flow end-to-end service guarantees. We introduce a notion of effective-capacity per communication link that captures the service requirements of flows sharing this link. Each link maintains a price and effective-capacity, and each flow maintains rate and dissatisfaction. Flows and links update their respective variables locally, and we show that their decisions drive the system to an optimal point. We implemented our algorithm on a network-simulator and studied its convergence behavior on few networks of practical interest.

To my siblings and parents

ACKNOWLEDGMENTS

This manuscript is a result of many years of my research in graduate school. There are many who have influenced my research path, directly and indirectly during my graduate years at Texas A&M University.

First and foremost, this thesis is due in large part to my co-advisors Prof. Jean-François Chamberland and Prof. Srinivas Shakkottai. I would like to thank them both for being very patient with me. I have been extremely fortunate to get an opportunity to work with young, enthusiastic, and capable advisors in them. They did their best to help me advance in my career and achieve my potential. If I was found wanting then it was not due to lack of effort and guidance from them, but my own shortcomings. I express sincere gratitude to my thesis committee members; Prof. Costas Georghiades, Prof. Aniruddha Datta, and Prof. Thomas Schlumprecht, for their support and guidance whenever I needed it.

This work could not have been possible without my collaborators, Prof. Kannan Ramachandran, Prof. Krishna Narayanan, Prof. Henry Pfister, Prof. Lei Ying, Dr. Ishai Menache, Prof. Lingjia Liu, Dr. Salim El Rouayheb, Dr. Jia Tang, Wei-Yu Chen, among many others. I also received inputs from several other faculty members in both ECEN and Math departments, notable mentions being Prof. Gautam Natarajan, Prof. Gregory Berkolaiko, Prof. Alex Sprintson, Prof. Natarajan Sivakumar, and Prof. Shuguang Cui. I am indebted to all my teachers in graduate school for their excellent teaching, who got me excited about the beautiful and meaningful research. Some of them deserve special mention, Prof. Marcelo Aguiar, Prof. Dante Deblassie, Prof. William Jonhnson, and Prof. Gautam Natarajan.

Life in graduate school would have been much tougher without the help of ever so helpful staff in our department, thanks to Tammy, Jeanie, Claudia, Paula, Gayle,

Carolyn, Annie, and Linda.

I am glad that Teh-Hsuan, Armin, Engin, Santhosh, and Pankaj tolerated me as their office mate. I received valuable advice from senior graduate students Kapil, Lingjia, Jing, and Wein-Yan. I had extraordinary classmates in Dana, Julia, Daniel, Salim, Ryan, Myung-Jun, Farhan, Samuel and Jia, Sofia, Svetlana, Daniel, Alejandro, Abner, Chris at the ECEN and Math department, respectively. It would be impossible for me to list all my friends who made my stay at Texas A&M a very pleasant one. Nevertheless, I am indebted to all of them and would like to mention some: Keerthi, Khushboo, Richa, Neeharika, Saira, Mariana, Nafisatu, Moritz, Gregory, Aleksander, Rahul, Vishal, Maneesh, Nabeel, Manish, Hao, Devika, Aditya, Maria, and Divya.

Last but not the least, I would not have been able to complete my graduare studies but for my family's unwavering support, love and encouragement. I would like to express my gratitude to my parents Aparna and Prabhat, and my siblings Pragati, Prachand, Prakriti, Pankaj, and Peeyush.

TABLE OF CONTENTS

CHAPTER Page

LIST OF TABLES

LIST OF FIGURES

FIGURE Page

FIGURE Page

CHAPTER I

INTRODUCTION

Cellular networks and the Internet are two celebrated examples of widely successful telecommunication systems. The former offers great flexibility and ubiquitous access, while the latter provides high data rates with negligible delays. The mobility offered by wireless communication comes at the price of fading and severe channel conditions. As a consequence, cellular service providers need to setup a large local infrastructure equipped with highly centralized control. The Internet on the other hand is a collection of several inter-connected networks, with several being added everyday to the world-wide web. Internet users are rational and selfish agents trying to maximize their individual experience. It is hard to have a central entity controlling such a large interconnected network. Therefore, any global objective needs to be achieved in a distributed local fashion.

In both of the above communication networks, the last hop is usually the bottleneck link. Either, we have severely fading wireless channels in cellular networks, or constrained wired link from home to cable head-end in the Internet. Recent advances in technology such as smart phones and tablet personal computers are bringing these two networks closer together to create a smarter, ubiquitous modern communication system. These networks strive to offer a combination of communication, computation, and entertainment on a single hand-held mobile device. All of the above three applications can be very sensitive to delay and jitter. Some of the delay-sensitive applications such as video-conferencing, real-time computing, and streaming videos constitute a large part of communication traffic. To support delay-sensitive applica-

---

The journal model is *IEEE Transactions on Automatic Control.*

tions on modern communication networks one needs to efficiently utilize the available resources.

To efficiently utilize a communication link, one employs transmitter buffers that can store the incoming packets to be sent over the channel at the opportune times. Buffering can improve communication efficiency, but it comes at the price of increasing transmission delay. This illustrates the throughput-delay tradeoff, that we investigate in detail in this work. In particular, we compute the achievable reliable communication rates under statistical guarantees on occupancy of transmit buffer.

Depending on the application in question, there may be variability in the generation of packets to be sent over a communication channel. Furthermore, wireless channels are stochastic in nature, leading to fluctuating service available to the transmitter during every transmission opportunity. Variability in communication links and arrival processes, is often captured by their statistical models. This modeling enables the design of efficient and robust systems. However, statistical models make it difficult to provide deterministic service guarantees. Therefore, throughout this dissertation we will be interested in statistical guarantees on communication.

This dissertation addresses the problem of delay-sensitive communication at multiple levels and from different perspectives. The common theme being provision of statistical service guarantees to each communication flow. We are interested in obtaining the fundamental rates of communication that can be supported for a communication channel with a known statistical characterization. Moreover, we utilize the insights obtained to extend this study to service guaranteed communication over small networks. To provide service guarantees over large networks, one needs to provide a distributed algorithm that utilizes local information to arrive at a global optimum. Below, we introduce three related problems in detail.

A. Bit-Erasure Channel

Given a statistical characterization of a communication link, one can identify a fundamental limit on the rate that can be reliably communicated over this link. This information theoretic bound is known as the Shannon capacity [1] of the channel. Recent advances in coding theory have led us to capacity achieving codes. These codes are quite sophisticated and can be of large lengths, thereby increasing decoding and hence communication delays. For applications intolerant to delays, these coding schemes become impractical. Therefore, a natural question to ask is: what are the bounds on rates of reliable communication with service guarantees?

It is easy to see that no finite delay guarantee can be satisfied for reliable communication over stochastic channels. In fact, there is a non-vanishing probability of reception failure, for any finite-length code. A simple way of achieving reliability over such channels is simple retransmission of packets. If the communication is aided by acknowledgments from the receiver, the transmitter buffer can store the packets and keep retransmitting until successful reception. We focus our attention on statistical service guarantees, such as queue-occupancy of this transmitter buffer.

There are two main questions we wish to address in this context of delay-sensitive communication with finite-length codes. What is the optimal code-rate for a given service guarantee, and what is the impact of channel correlation on this code-rate? It is known that one can achieve reliable communication by operating at a code-rate below Shannon capacity. However, this is an asymptotic result and the answer is not obvious for finite code blocks and specific service guarantees. A low rate code would have low probability of transmission failure and hence less number of retransmissions in an expected sense. However, a low-rate code also sends fewer bits across a communication link during every transmission opportunity. This seems

to suggest an optimal operating point for finite-length codes with fixed code-rates. Furthermore, correlation is not a huge issue in asymptotic results, since one can always interleave transmitted bits such that successive coded bits are independent from each other. However, this does not hold true in a non-asymptotic regime. In fact, some preliminary results suggest that the optimal code-rate would depend on channel memory. We seek to quantify this dependence as well.

Apart from a quantitative answer to the above two questions, one of the major contributions in this work is the careful modeling of the physical and data link layers. This modeling allows us to conduct our analysis utilizing existing mathematical tools. In particular, we employ a Markov modulated bit-erasure channel model for the communication link. In this model, transmitted bits see one of the Markov channel states. In each of these states, the transmitted bit can either be detected successfully or erased completely. The erasure probability is characteristic of each channel state. Further, the mixing time of this Markov chain models the channel correlation. Packet arrival is modeled as a Bernoulli random process, with packet length being geometrically distributed. Incoming packets are stored in large buffer and await transmission. Packets are broken into segments, and encoded into a finite length code and transmitted over the bit-erasure channel with memory. Each segment is retransmitted to the receiver, until a successful reception arrives at the transmitter. We employ random coding for analytical tractability. The mathematical techniques utilized in this section can be utilized to study channels with errors and more sophisticated coding schemes.

At the outset, we would like to admit that there are far more sophisticated ways of trying to achieve reliability with service constraints. For example, one can employ a rateless coding scheme at the transmitter. In this scheme, the transmitter initially sends pure information bits, and then sends additional redundancy bits for

each instance of reception failure. This is akin to changing code-rates based on the partial channel state information sent by the receiver. Given complete channel state information at the transmitter, temporal water filling maximizes system performance. However, we stick to simpler models to keep the analysis tractable and understand the simplest non-trivial problem in this area.

B. Network Coding

In modern communication networks, packets have to traverse through multiple hops before they reach their desired destination. Traditionally, multi-hop communication is modeled as commodity flows. To provide service guarantees on such communication flows is a highly non-trivial task, due to three main reasons. First, end-to-end service guarantees for a flow translate into service guarantees on each link traversed by the flow in a convoluted fashion. Second, service guarantees at each communication link are influenced by the various flows sharing it. Third, buffers at intermediate nodes have a non-linear effect of on the characteristics of the incoming flow.

Before addressing provision of service guarantees to communication flows, we must question the validity of considering multi-hop transmission as a commodity flow. The main reason for keeping the packets from unique source-destination pairs separated arises from the simplicity in implementation of network flows. However, data is a special type of commodity that can be encoded and hence transformed at various intermediate nodes. One can take advantage of the redundancy in the network to send at higher information rate, through the network. It has been shown that sending coded information at intermediate nodes can achieve the cut-set bound for multicast [2]. This idea is popularly referred to as network coding. We wish to address two important questions related to network coding in this work. First,

we wish to quantify the gains of network coding for a given statistical end-to-end service guarantee. Second, we wish to investigate the scenarios where it makes sense to employ network coding.

We compare network coding to classical routing to quantify the gains under a delay-sensitive setting. It is intuitive to see that network coding gains are maximal when communication links are completely utilized. Source burstiness, channel variations, and service guarantees collectively imply that links may not always be fully loaded. Furthermore, classical routing enjoys statistical multiplexing gains at the intermediate nodes, which may offset the gains of network coding. Therefore, it is not clear how much we gain by applying network coding in a communication system subject to Quality of Service (QoS) constraints. We provide quantitative results on the benefits of network coding for a simple butterfly network in the context of delay-sensitive applications.

Network coding works well for multicast and requires a certain redundancy in the network. Further, if one sends coded information between a source-destination pair through multiple routes, then the bottleneck route determines the rate of the coded information. In the case of wireless mesh networks, one can allocate available physical-layer resources to individual nodes in a small area to form a network. Therefore, there is a tradeoff in terms of the amount of resource required to set up a network suitable for network coding and the performance gains offered. Given a budget of physical resources, one can design optimal network topologies for classic-routing and network coding. One must compare the system performance in these two cases, to see when one should attempt to create a network suitable for coding.

This brings us back to the question of end-to-end service guarantees for the flows in a general network. One way to deal with this problem is working with effective bandwidth. Looking at service guarantees on the tail-decay rate of transmitter

buffer occupancy, effective bandwidth identifies the data-rate needed by a source to fulfill its service requirement [3, 4]. In general, the effective bandwidth of a source can be much larger than the average throughput it produces. Three properties of effective bandwidths can be utilized to ensure end-to-end service guarantees. First, the aggregate effective bandwidth of independent sources is additive. Second, so long as the capacity of any link is higher than the aggregate effective-bandwidths of the flows sharing it, the service requirements the corresponding sources are met. Third, for "nice" sources, one can characterize the effective bandwidth of the departure processes from each intermediate node. Thus, if end-to-end service guarantees for a flow are characterized in terms of the tail-decay rate of aggregate buffers of each node it traverses, then an end-to-end service guarantee can be provided to this flow if its effective bandwidth requirement is met at each link. We derive the effective bandwidth of the departure processes and utilize it to find the end-to-end service guarantee for flows on small mesh networks. Classic-routing and network coding cases are analysed with the proposed effective-bandwidth method. Analysis for network coding case gets simplified, since mixing flows at the bottleneck link is equivalent to the flows being oblivious of each other.

C. Distributed Control

We discussed in the previous section the challenges associated with providing end-to-end service guarantees to flows traversing networks. We also discussed effective bandwidth as one of the techniques to overcome these challenges. Such techniques are popularly referred to as stochastic network calculus. Network calculus works well for small network, however it faces two main challenges as the network size grows. First, the service guarantee on the aggregate buffer occupancy cannot be treated on each

link separately. This is due to the correlation between successive link buffers. The second problem is scaling; computing effective bandwidth as function of required service guarantee per flow at each link gets cumbersome. Therefore, we look for alternate local solutions that lead to global per-flow service guarantees.

In this work, we provide a per-flow based service guarantee in terms of an abstract quantity "flow-degradation". Degradation of each flow is guaranteed to be lower than their accepted threshold. Degradation can capture many types of service requirements, such as delay and queue-lengths. We assume that the total flow-degradation is the sum of the degradations seen by the flow at each traversed link. Further, we assume that degradation is inherent to a link and is convex increasing with the total load. Under these assumptions, our global goal is to maximize the sum-utility of all flows such that their per-flow constraints are satisfied.

It turns out that a simple dual-decomposition fails to provide a truly distributed algorithm. This results from the strong coupling of per-flow constraints at each shared link. Specifically, the flow-degradations are sums of the link degradations in routes, and each link-degradation is caused and seen by all the flows sharing the link. To overcome this hurdle, we introduce an "effective-capacity" inherent to every link which acts as a service-limited capacity. In terms of this variable, we are able to decouple the per-flow service constraints. Thus, we achieve a distributed solution that maximizes sum-utility of individual flows while satisfying their per-flow service guarantees. Our findings are supported by numerical studies on few representative networks. We also simulated an implementation on a network simulator.

CHAPTER II

CODE-RATE SELECTION, QUEUEING BEHAVIOR AND THE CORRELATED ERASURE CHANNEL

This chapter considers the relationship between code-rate selection and queueing performance for communication systems subject to time-varying channel conditions. While error-correcting codes offer protection against channel uncertainties, there exists a natural tradeoff between the enhanced protection of low-rate codes and the rate penalty imposed by additional redundancy. In the limiting regime where codewords are asymptotically long, this tradeoff is well-understood and characterized by the Shannon capacity. However, for delay-sensitive communication systems and finite block-lengths, a complete characterization of this tradeoff is not fully developed. This chapter offers a new perspective on the queueing performance of communication systems with finite block-lengths operating over correlated erasure channels. A rigorous framework that links code-rate to overall system performance for random codes is presented. Guidelines for code-rate selection in delay-sensitive systems are identified. These findings are supported by a numerical study.

A. Introduction

The transmission of digital information over noisy channels has become commonplace in modern communication systems. The dependability of contemporary data links is due, partly, to the many successes of information theory and error-control coding [5]. In particular, the reliable transmission of digital information is possible at rates approaching Shannon capacity using asymptotically long codewords [6, 7]. Indeed, many notable communication systems employ long codewords to provide high throughput and low error probabilities [8].

One context where the insights offered by classical information theory do not apply directly is the broad area of delay-constrained communications [9]. Real-time traffic and live interactive sessions are very sensitive to latency. Long codewords are not particularly well-suited for real-time applications because they entail lengthy encoding/decoding delays. Alternative engineering methods, including power control, automatic repeat-request, scheduling and feedback, can be leveraged to establish rapid end-to-end connections [10, 11]. Often, delay considerations force a system to operate well below its Shannon limit [12].

Several articles in information theory are focused on the tradeoff between throughput and delay. Coding performance as a function of delay has been assessed in the information theory literature using the reliability function [6]. This performance criterion identifies the error exponent of a code family as a function of data rate. The notion of reliability function can be extended to variable-length codes in the presence of feedback, leading to the famous Burnashev error exponent [13–15]. While significant, these results remain asymptotic in nature and do not capture the queueing aspect of communication systems. Alternative approaches include effective capacity [16, 17], outage capacity [18, 19], average delay characterizations [20], fluid analysis [21] and heavy-traffic limits [22]. While these contributions provide valuable insights about the design of delay-sensitive systems, many such articles make idealized assumptions about the behavior of coded transmissions. For instance, some authors adopt the notion of instantaneous capacity: individual data blocks are assumed, implicitly or explicitly, to possess enough degrees of freedom to support sophisticated coding schemes and thereby approach Shannon capacity within every time-slot. Perhaps reasonable for long codewords, such assumptions become more of a concern for short data blocks. This is especially problematic for channels with memory, where correlation over time promotes deviations from expected behavior.

For a delay-constrained communication system that utilizes short codewords, two opposite considerations seem to underlie the selection of an error-correcting code. A low-rate code will, in general, result in a small probability of decoding failure; whereas the same system with a high-rate code is more prone to errors. Still, the successful decoding of a codeword associated with a higher-rate code leads to the transmission of a larger number of information bits. In the limit of asymptotically long codewords, it is clear that code-rate should only be slightly below Shannon capacity. However, the optimal operating point for a system with short block length is not so obvious [23]. It may depend on the physical resources available and the service constraints imposed on the system.

Many previous inquiries in the area adopt a higher-layer viewpoint, using simplistic models for the physical layer; or they embrace a channel-coding perspective, intentionally disregarding queueing considerations. Herein, we seek to bridge the gap between these extremes to address an important question. What is the optimal code-rate for a particular implementation? Our approach in obtaining an answer to this question differs from established work in that we strive to provide exact solutions.

To facilitate the type of queueing analysis we wish to carry, we make the following assumptions. The packet arrival process at the transmitter is Bernoulli, with packet length having a geometric distribution at the bit level. The communication medium is a bit-erasure channel with memory. Random codes, with maximum-likelihood decoding, are employed to protect the sent information against erasures. Collectively, these assumptions are sufficient to conduct a rigorous analysis of the probability of block decoding failure at the receiver as well as a complete characterization of the ensuing queueing behavior at the source. Implicit to our system model is the ability to acknowledge the reception of packets through instantaneous feedback. We emphasize that model components are selected with the intent to keep

analysis manageable. The focus is on developing tools and techniques that can be used to bridge communication, coding and queueing. Still, the framework admits several extensions beyond the formulation presented in this chapter; some of these extensions are discussed alongside the main results wherever appropriate.

The remainder of this chapter is organized as follows. The system model is introduced in Section B, with the probability of block decoding failure being computed explicitly in Section 3. Packet arrivals and departures form the main topic of Section C. Altogether, this dictates the queueing behavior of the packetized systems, which is analyzed in Section D. Numerical results are contained in Section E. Finally, new insights, concluding remarks and avenues of future research are discussed in Section F.

B. System Model

We initiate our exposition of the system we wish to study with a description of the underlying communication channel. Bits are sent from a source to a destination over a Gilbert-Elliott erasure channel. The channel can be in one of two states which we denote by integers $\{1, 2\}$. In state 1, every transmitted bit is erased with probability ε_1 independently of other bits. Similarly, in state 2, every bit is lost with probability ε_2. Throughout, we assumes that $\varepsilon_2 \leq \varepsilon_1 \leq 1$. Transitions between channel states occur according to a Markov process. The probability of jumping to state 2 given that the Markov chain is currently in state 1 is denoted by α. The reverse transition probability from state 2 to 1 is written as β. The parameters of this Markov chain can be expressed in the form of a transition probability matrix,

$$\mathbf{P} = \begin{bmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{bmatrix}. \tag{2.1}$$

A graphical interpretation of this communication channel appears in Fig. 1. We note

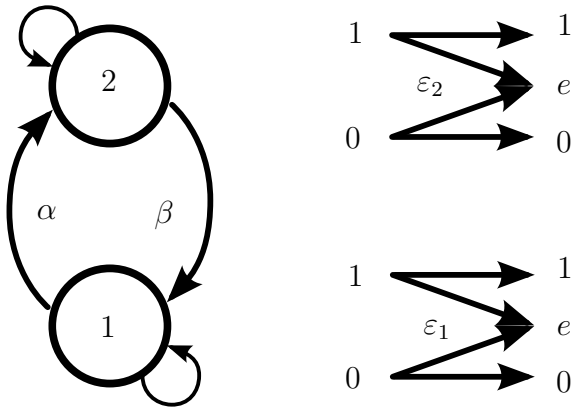

Fig. 1. A Gilbert-Elliott bit erasure channel is employed to model the operation of a communication link with memory. This model captures both the uncertainty associated with transmitting bits over a noisy channel and the correlation over time typical of several communication channels.

that the methodology adopted in this chapter allows a larger number of channel states and can be applied to more intricate physical links. The only fundamental aspect of the Gilbert-Elliott channel is that it represents the simplest non-trivial instance of a finite-state channel with memory, which leads to a more accessible treatment of the problem. Markov models have been employed to capture the behavior of communication channels in the past, and several studies point to methods of selecting parameters to best match the profiles of communication links at the physical layer [24, 25]. In our framework, correlation over time is captured through the transition probability matrix of the channel. At this point, we leave the parameters of our Markov channel in an abstract form, seeking general solutions.

The state of the channel at instant n is a random variable, which we denote by C_n. Using this notation, one can write the progression of the Markov chain over time as $\{C_n : n \in \mathbb{N}\}$. Finding the conditional probability $\Pr(C_{n+1} = d | C_n = c)$, where

$c, d \in \{1, 2\}$, amounts to selecting an entry in $\mathbf{P}$. Likewise, $\Pr(C_{n+N} = d | C_n = c)$ is obtained by locating the corresponding entry in $\mathbf{P}^N$, where

$$\mathbf{P}^N = \frac{1}{\alpha+\beta} \begin{bmatrix} 1 & \alpha \\ 1 & -\beta \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & (1-\alpha-\beta)^N \end{bmatrix} \begin{bmatrix} \beta & \alpha \\ 1 & -1 \end{bmatrix}$$
$$= \begin{bmatrix} \frac{\beta+\alpha(1-\alpha-\beta)^N}{\alpha+\beta} & \frac{\alpha-\alpha(1-\alpha-\beta)^N}{\alpha+\beta} \\ \frac{\beta-\beta(1-\alpha-\beta)^N}{\alpha+\beta} & \frac{\alpha+\beta(1-\alpha-\beta)^N}{\alpha+\beta} \end{bmatrix}.$$

This Markov chain converges to its stationary distribution at an exponential rate that depends on the second largest eigenvalue of $\mathbf{P}$, which is $1-\alpha-\beta$. This quantity can therefore be employed to quantify channel memory.

1. Segments, Block-Length, and Code-Rate

To transmit information over this erasure channel, data packets must first be processed and encoded. In our framework, a packet of length L is sectioned into M data segments, each containing K information bits. Packing loss is treated implicitly as the last data segment of a packet is zero-padded to K bits. Thus, the number of segments within a packet of length L is equal to $M = \lceil L/K \rceil$. Every segment is encoded separately into a codeword of length N, which is subsequently stored in the queue for eventual transmission over the Gilbert-Elliott erasure channel. The transmission of a codeword then requires N consecutive uses of the channel. We assume that decoding failures are handled through immediate retransmission of the missing data. The block-length, N, remains fixed throughout; it is determined by system requirements and the availability of physical resources. On the other hand, the size of a data segment, K (and therefore the code-rate $r = K/N$), is a parameter that should be optimized.

2. Distribution of Erasures

A quantity that is of fundamental importance in our analysis is the probability of decoding failure at the destination. An intermediate step in identifying this probability is to derive expressions for the distributions of the number of erasures, E, within a codeword. This, in turn, depends on the number of visits to each state within N consecutive realizations of the channel. Specifically, we are interested in conditional probabilities of the form

$$\Pr(E = e, C_{N+1} = d | C_1 = c), \tag{2.2}$$

where $c, d \in \{1, 2\}$. The generating functions for these conditional probabilities can be derived based on generalizing the entries of $\mathbf{P}$ to the vector space of real polynomials in x, with

$$\mathbf{P}_x = \begin{bmatrix} (1-\alpha)(1-\varepsilon_1+\varepsilon_1 x) & \alpha(1-\varepsilon_1+\varepsilon_1 x) \\ \beta(1-\varepsilon_2+\varepsilon_2 x) & (1-\beta)(1-\varepsilon_2+\varepsilon_2 x) \end{bmatrix}. \tag{2.3}$$

Proposition B.1. *Let $[\![x^k]\!]$ be the linear functional that maps a polynomial in x to the coefficient of x^k. Then, the conditional probability* $\Pr(E = e, C_{N+1} = d | C_1 = c)$ *is given by*

$$\Pr(E = e, C_{N+1} = d | C_1 = c) = [\![x^e]\!] \left[\mathbf{P}_x^N\right]_{c,d}, \tag{2.4}$$

where $\mathbf{P}_x$ is the matrix defined in (2.3).

Proof. This result can be shown using mathematical induction. Let $E_{i:j}$ denote the number of channel erasures occurring between times i and j, inclusively. By construction, the proposition holds for $N = 1$. As an inductive step, assume that

(2.4) is satisfied for $N = n - 1 > 0$. Then, one can write

$$\Pr\left(E_{1:n} = e, C_{n+1} = c_{n+1} | C_1 = c_1\right) = \sum_{c_n \in \{1,2\}} \Pr(E_{1:n} = e, C_{n+1} = c_{n+1}, C_n = c_n | C_1 = c_1)$$

$$= \sum_{c_n \in \{1,2\}} \sum_{k \in \{0,1\}} \Pr(E_{n:n} = k, C_{n+1} = c_{n+1} | C_n = c_n) \Pr(E_{1:n-1} = e - k, C_n = c_n | C_1 = c_1)$$

$$= \sum_{c_n \in \{1,2\}} \sum_{k \in \{0,1\}} [\![x^k]\!] \left[\mathbf{P}_x\right]_{c_n, c_{n+1}} [\![x^{e-k}]\!] \left[\mathbf{P}_x^{n-1}\right]_{c_1, c_n} = [\![x^e]\!] \left[\mathbf{P}_x^n\right]_{c_1, c_{n+1}}.$$

That is, (2.4) also holds for $N = n$. Since both the basis and the inductive step have been verified, we conclude that the proposition is true for all integers $N > 1$. □

We note that one can employ this method or alternative combinatorial means to obtain closed-form expressions for the desired conditional probabilities [26, 27].

3. Probability of Decoding Failure

At the onset of every transmission attempt, a new code is created to encode K information bits. The code is defined by a random parity-check matrix $\mathbf{H}$ of size $(N-K) \times N$. The entries of $\mathbf{H}$ are selected independently and uniformly over $\{0, 1\}$. This scheme assumes shared randomness between the source and its destination. Maximum likelihood decoding is used at the destination to decode the received messages. Consequently, the probability of decoding failure becomes a function only of the number of erasures contained within a block. Once the value of E is known, one can compute the probability of decoding failure using the following result.

Proposition B.2. *The probability of decoding failure, given e erasures within a codeword of length N, is equal to*

$$P_{\mathrm{f}}(N - K, e) = 1 - \prod_{i=0}^{e-1} \left(1 - 2^{i-(N-K)}\right).$$

Proof. Conditioned on $E = e$, decoding at the destination will succeed if and only if the submatrix of $\mathbf{H}$ formed by choosing the e erased columns has rank e. Furthermore, the probability that a random $p \times e$ matrix over $\mathbb{F}_2$ has rank e is equal to $\prod_{i=0}^{e-1} (1 - 2^{i-p})$. Collecting these two results [28, p. 73], we obtain the probability of a successful transmission, which in turn determines the probability of decoding failure, given $E = e$. □

The unconditioned probability of decoding failure at the destination is equal to

$$P_{\mathrm{f}}(N - K) = \mathbb{E}\left[P_{\mathrm{f}}\left(N - K, E\right)\right],$$

where the distribution of E accounts for all the possible channel realizations within a block. While the probability of decoding failure offers an important measure of performance, it alone does not capture the queueing behavior of the system. Correlation among decoding-failure events may also influence the behavior of the queue at the transmitter. Having introduced a precise model for the physical layer, we turn to the description of the arrival and departure processes of the queue.

C. Arrival and Departure Processes

Packets enter the queue according to a discrete-time Bernoulli process whose clock is synchronized with the codeword transmission cycle. During every codeword transmission attempt, a new packet arrives at the source with probability γ, independently of other time instants. The number of bits in every data packet is random, with packet sizes forming a sequence of independent and identically distributed random variables. The marginal distribution of a packet size is geometric with parameter ρ. In other words, the probability that a packet contains exactly ℓ bits is given by

$$\Pr(L = \ell) = (1 - \rho)^{\ell - 1}\rho \qquad \ell = 1, 2, \ldots$$

where $\rho \in (0,1)$. The arrival process and the packet-length distribution have been selected, partly, to facilitate the analysis we wish to carry. In particular, the memoryless property of the geometric distribution and the independence over time of the Bernoulli process are crucial properties that make for a tractable characterization of queueing behavior. Adopting an intricate arrival process more in tune with a specific application can easily render analysis intractable. This explains why our arrival process conforms to a model commonly found in the queueing literature.

Departures from the queue are governed by the underlying Gilbert-Elliott channel and the selected number of parity bits, $N-K$, of our random code. Probability of decoding failure is monotonically increasing in code-rate, given a fixed block-length N. However, the successful decoding of a high-rate codeword leads to the transmission of a larger number of information bits. As mentioned before, these competing considerations create a natural tradeoff between data content and probability of decoding failure. Accordingly, the code-rate $r = K/N$, or equivalently the number of information bits per data segment K, is a parameter that should be optimized.

Once the code-rate is specified, the number of successfully decoded codewords needed to complete a packet transmission is random with $M = \lceil L/K \rceil$. We note that L being geometric with parameter ρ implies that M is also geometric with parameter

$$\rho_r = \sum_{\ell=1}^{K} (1-\rho)^{\ell-1}\rho = 1-(1-\rho)^K.$$

Thus, the probability that a data packet requires the successful transmission of exactly m codewords becomes

$$\Pr(M=m) = (1-\rho_r)^{m-1}\,\rho_r \qquad m = 1,2,\ldots$$

We emphasize that, in the current setting, the number of coded blocks per data packet M retains the memoryless property.

In our formulation, we assume that K is independent of channel state, which simplifies analysis. When side information is present at the transmitter, one can enhance performance by picking K as a function of the current state. Furthermore, even without explicit state knowledge, it is possible to estimate the channel state through available feedback, i.e., the automatic repeat-request sequence. In the latter scenario, the selection of K as a function of state estimates becomes a partially observable Markov decision process; such problems are often computationally intractable and necessitate careful consideration. Accounting for the presence of partial state information at the transmitter is beyond the scope of this chapter, and we leave this matter as a possible future endeavor. This section completes the description of the communication system under consideration. We proceed below with the characterization of overall performance.

D. Queueing Behavior

Packets are stored in the queue upon generation by the source, and they remain in this buffer until all the corresponding data segments are decoded successfully at the destination. We assume that there are no packet losses and, as such, the transmit buffer has no hard limit. When discussing the size of the queue at the transmitter, two distinct characterizations are possible. The first option is to keep track of the number of packets contained in the queue. The second choice is to track the amount of data awaiting transmission. Although the latter alternative provides a more accurate representation of buffer occupancy in bits, the former option is closely related to the concept of packet delay and it is simpler to analyze. For these reasons, we elect to define the state of the queue as the number of data packets in the queue, as is customary in classical queueing literature [29–31].

Recall that, in the proposed setting, a packet of length L is first subdivided into M data segments. Each segment is encoded separately into a codeword of length N, and the resulting message is subsequently sent over the communication channel. Successful receptions are acknowledged instantaneously through feedback, whereas decoding failures trigger immediate retransmission of the missing block. Upon confirmation of an accurate transfer, a data segment is marked as delivered and transmission of the next data block begins. Though the presence of instantaneous feedback is assumed for mathematical convenience, it may be approximated in practice using high-speed decoders and high-power reverse-link communication.

For the head packet to depart from the queue, the destination must successfully decode the received message and this codeword must be carrying the final segment of information pertaining to this head packet. Specifically, a packet composed of L bits will require the successful reception of $M = \lceil L/K \rceil$ codewords before it is removed from the queue. The length of the queue at the onset of block s is denoted by Q_s. The state of the Gilbert-Elliott channel at this instant is represented by C_{sN+1}. Together, these two quantities form the state of our Markov process, $Y_s = (C_{sN+1}, Q_s)$. We emphasize that this state space is countable, with Y_s belonging to $\{1, 2\} \times \mathbb{N}_0$. Furthermore, the Markov chain underlying the evolution of our system possesses a special structure. It forms an instance of a discrete-time *quasi-birth-death process*. Luckily, there are many established techniques to analyze such mathematical objects [32–34].

Our next step is to examine the transition probabilities of this augmented Markov chain. The probability of jumping from Y_s to Y_{s+1} is given by

$$\begin{aligned} &\Pr(Y_{s+1} = (d, q_{s+1}) | Y_s = (c, q_s)) \\ &= \sum_{e=0}^{N} \Pr\left(Q_{s+1} = q_{s+1} | E = e, Q_s = q_s\right) \Pr(E = e, C_{(s+1)N+1} = d | C_{sN+1} = c). \end{aligned} \tag{2.5}$$

A methodology was introduced in Section 2 to compute conditional probabilities of the form $\Pr(E = e, C_{(s+1)N+1} = d | C_{sN+1} = c)$. Accordingly, it suffices to focus on the other component of each summand, $\Pr\left(Q_{s+1} = q_{s+1} | E = e, Q_s = q_s\right)$, to characterize (2.5).

We first consider conditional events $\{Q_s = q_s\}$ for which $q_s > 0$. In this case, admissible values for Q_{s+1} are given by $\{q_s - 1, q_s, q_s + 1\}$. Two factors can affect the length of the queue: the arrival of a new data packet and the completion of a packet transmission. The latter occurrence will only take place if a codeword is successfully decoded at the destination and the corresponding data block is the last segment of the head packet. Keeping this fact in mind and using independence between arrivals and departures, we get

$$\begin{aligned}
\Pr\left(Q_{s+1} = q_s + 1 | E = e, Q_s = q_s\right) &= \gamma\big(P_{\mathrm{f}}(N-K, e) + (1 - P_{\mathrm{f}}(N-K, e))(1 - \rho_r)\big) \\
\Pr\left(Q_{s+1} = q_s | E = e, Q_s = q_s\right) &= \gamma\left(1 - P_{\mathrm{f}}(N-K, e)\right)\rho_r + (1-\gamma)\big(P_{\mathrm{f}}(N-K, e) \\
&\qquad + (1 - P_{\mathrm{f}}(N-K, e))(1 - \rho_r)\big) \\
\Pr(Q_{s+1} = q_s - 1 | E = e, Q_s = q_s) &= (1-\gamma)\left(1 - P_{\mathrm{f}}(N-K, e)\right)\rho_r.
\end{aligned}$$

When the queue is empty, there are no departures. As such, only two possibilities remain

$$\begin{aligned}
\Pr(Q_{s+1} = 1 | E = e, Q_s = 0) &= \gamma \\
\Pr(Q_{s+1} = 0 | E = e, Q_s = 0) &= 1 - \gamma.
\end{aligned}$$

Assembling these results and using (2.5), we obtain the probability transition matrix of the Markov process $\{Y_s\}$. A graphical representation of possible state transitions appears in Fig. 2.

To proceed with the analysis of our queued system, a compact representation of

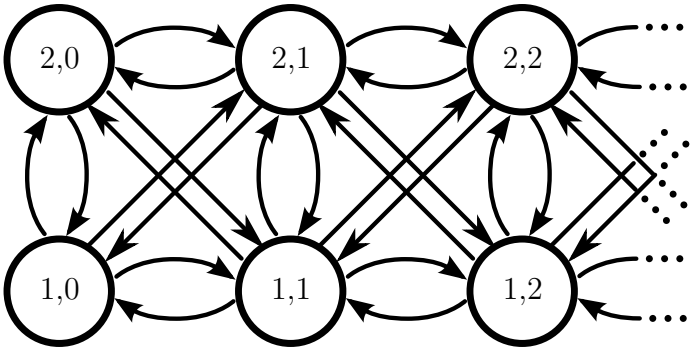

Fig. 2. State space and transition diagram for the aggregate queued process $\{Y_s\}$; self-transitions are intentionally omitted.

the conditional probabilities defined in (2.5) is apropos. For $q \in \mathbb{N}$ and $c, d \in \{1, 2\}$, we introduce the following convenient notation,

$$\mu_{cd} = \Pr(Y_{s+1} = (d, q-1) | Y_s = (c, q))$$
$$\kappa_{cd} = \Pr(Y_{s+1} = (d, q) | Y_s = (c, q))$$
$$\lambda_{cd} = \Pr(Y_{s+1} = (d, q+1) | Y_s = (c, q)).$$

Similarly, when the queue is empty, we use

$$\kappa^0_{cd} = \Pr(Y_{s+1} = (d, 0) | Y_s = (c, 0))$$
$$\lambda^0_{cd} = \Pr(Y_{s+1} = (d, 1) | Y_s = (c, 0)).$$

Collectively, these labels define the 12 transition probabilities associated with a non-empty queue, and the 8 transition probabilities subject to the non-negativity constraint at zero.

At last, we are ready to derive the equilibrium distribution of our aggregate system. We note that this system is stable when mean packet arrival is less than

mean packet departure over a codeword transmission period [35], i.e.,

$$\gamma < \rho_r \mathbb{E}\left[1 - P_{\mathrm{f}}\left(N - K, E\right)\right].$$

Under this stability condition, the Markov chain $\{Y_s\}$ is positive recurrent and possesses a unique stationary distribution [36]. Assuming that the system is stable, let $Y = (C, Q)$ be a random vector with the aforementioned probability distribution,

$$\Pr(Y = (c, q)) = \lim_{s \to \infty} \Pr(Y_s = (c, q)).$$

We employ the semi-infinite vector π to represent the equilibrium distribution of our system, with

$$\pi(2q + i) = \begin{cases} \Pr(C = 1, Q = q) & \text{if } i = 0 \\ \Pr(C = 2, Q = q) & \text{if } i = 1. \end{cases}$$

The states $\{(1, q), (2, q)\}$ are known as the qth level of the chain and $\pi_q = [\pi(2q)\ \pi(2q+1)]$ is the stationary distribution associated with the qth level.

Using this compact notation, one can express the Chapman-Kolmogorov equations for the queued system as $\pi\mathbf{T} = \pi$, where $\mathbf{T}$ denotes the transition probabilities associated with the aggregate Markov chain $\{Y_s\}$. We can represent the transition probability operator $\mathbf{T}$ as a semi-infinite matrix of the form

$$\mathbf{T} = \begin{pmatrix} \mathbf{C}_1 & \mathbf{C}_0 & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \tag{2.6}$$

where the submatrices $\mathbf{C}_1$, $\mathbf{C}_0$, $\mathbf{A}_2$, $\mathbf{A}_1$, and $\mathbf{A}_0$ are 2×2 real matrices. Specifically,

we have

$$\mathbf{A}_0 = \begin{bmatrix} \lambda_{11} & \lambda_{12} \\ \lambda_{21} & \lambda_{22} \end{bmatrix} \qquad \mathbf{A}_1 = \begin{bmatrix} \kappa_{11} & \kappa_{12} \\ \kappa_{21} & \kappa_{22} \end{bmatrix}$$

$$\mathbf{A}_2 = \begin{bmatrix} \mu_{11} & \mu_{12} \\ \mu_{21} & \mu_{22} \end{bmatrix}.$$

When the queue is empty, the relevant submatrices become

$$\mathbf{C}_0 = \begin{bmatrix} \lambda_{11}^0 & \lambda_{12}^0 \\ \lambda_{21}^0 & \lambda_{22}^0 \end{bmatrix} \qquad \mathbf{C}_1 = \begin{bmatrix} \kappa_{11}^0 & \kappa_{12}^0 \\ \kappa_{21}^0 & \kappa_{22}^0 \end{bmatrix}.$$

When $\mathbf{A} = \mathbf{A}_0 + \mathbf{A}_1 + \mathbf{A}_2$ is irreducible, this quasi-birth-death process is recurrent if and only if $(v\mathbf{A}_0\mathbf{1})/(v\mathbf{A}_2\mathbf{1}) \leq 1$, where v is the stationary probability distribution of $\mathbf{A}$ [35]. One possible approach to identify the stationary distribution of the Markov chain $\{Y_s\}$ is to employ spectral representations and ordinary generating functions. This technique is described in Section 1. An alternate numerical means for computing the stationary distribution is the matrix geometric method discussed in Section 2. As we will see, both approaches have their advantages and drawbacks.

1. Transform Method

The first approach we present makes use of generating functions [27, 31]. Let $\mathbf{\Pi}(z)$ be the transform vector defined by

$$\mathbf{\Pi}(z) = \sum_{q=0}^{\infty} z^q \pi_q. \tag{2.7}$$

Theorem D.1. *The invariant distribution of the Markov chain can be derived from the recurrence relation induced by* $\mathbf{T}$*. Finding the stationary distribution of the aug-*

mented Markov chain is equivalent to solving a matrix equation of the form

$$\mathbf{\Pi}(z)\mathbf{D}(z) = \pi_0 \left(\mathbf{D}(z) - \mathbf{D}_0(z)\right), \tag{2.8}$$

where the entries in $\mathbf{D}(z)$ *and* $\mathbf{D}_0(z)$ *are quadratic polynomials*

$$\mathbf{D}(z) = z^2\mathbf{A}_0 + z\left(\mathbf{A}_1 - \mathbf{I}\right) + \mathbf{A}_2 \tag{2.9}$$

$$\mathbf{D}_0(z) = z^2\mathbf{C}_0 + z\left(\mathbf{C}_1 - \mathbf{I}\right). \tag{2.10}$$

The elements of π_0 *can be determined from the requirements imposed by stability and normalization.*

Proof. We begin this demonstration by writing the balance equations governing the Markov chain $\{Y_s\}$. From the Chapman-Kolmogorov equations $\pi\mathbf{T} = \pi$ and the form of $\mathbf{T}$ given in (2.6), we obtain

$$\pi_{q-1}\mathbf{A}_0 + \pi_q\left(\mathbf{A}_1 - \mathbf{I}\right) + \pi_{q+1}\mathbf{A}_2 = 0 \quad q \geq 2 \tag{2.11}$$

$$\pi_0\mathbf{C}_0 + \pi_1\left(\mathbf{A}_1 - \mathbf{I}\right) + \pi_2\mathbf{A}_2 = 0 \tag{2.12}$$

$$\pi_0\left(\mathbf{C}_1 - \mathbf{I}\right) + \pi_1\mathbf{A}_2 = 0. \tag{2.13}$$

Next, we multiply (2.11) by z^{q+1} and sum over all $q \geq 2$ to get

$$\left(\mathbf{\Pi}(z) - \pi_0\right) z^2\mathbf{A}_0 + \left(\mathbf{\Pi}(z) - \pi_1 z - \pi_0\right) z\left(\mathbf{A}_1 - \mathbf{I}\right) + \left(\mathbf{\Pi}(z) - \pi_2 z^2 - \pi_1 z - \pi_0\right)\mathbf{A}_2 = 0.$$

Leveraging boundary conditions (2.12) and (2.13), the equation above reduces to (2.8). □

Using the results of Theorem D.1, one can write

$$\mathbf{\Pi}(z) = \pi_0\left(\mathbf{I} - \mathbf{D}_0(z)\mathbf{D}^{-1}(z)\right)$$

where $\mathbf{D}^{-1}(z)$ is a matrix whose entries are rational functions of z. Note that one

can express the inverse of $\mathbf{D}(z)$ in terms of its adjugate matrix and determinant [37],

$$\mathbf{D}^{-1}(z) = \frac{\operatorname{adj} \mathbf{D}(z)}{\det \mathbf{D}(z)}.$$

Moreover, the entries of $\mathbf{D}_0(z)\mathbf{D}^{-1}(z)/z$ are rational functions where polynomial numerators have at most degree three and the common polynomial denominator $\det \mathbf{D}(z)$ is of degree four. Through careful inspection, we find that $(z-1)$ is a factor common to all numerator and denominator polynomials. After cancellation, the entries of $\mathbf{D}_0(z)\mathbf{D}^{-1}(z)$ can be expressed as quadratic polynomials over a common cubic polynomial. Using the general formula for the roots of cubic polynomials, it is then possible to carry partial fraction expansion for the entries of $\mathbf{D}_0(z)\mathbf{D}^{-1}(z)$ and thereby obtain an expression for $\mathbf{\Pi}(z)$, which is invertible in closed-form. The coefficients of π at level zero are obtained using stability and the fact that

$$\mathbf{\Pi}(1) = \begin{bmatrix} \frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta} \end{bmatrix}. \tag{2.14}$$

Although a closed-form parametric solution for the stationary distribution of this system exists and can be obtained using symbolic equation solvers, it is unfortunately too cumbersome to be included in this chapter. Still, we emphasize that existence of such a solution provides an efficient means to conduct numerical studies. A downside to the approach outlined above lies in the fact that it does not scale well with the number of states in the bit erasure channel, thereby precluding straightforward generalizations to alternate environments. This is due to the difficulty associated with finding the roots of high-degree polynomials. This impediment is addressed in the next section.

2. Matrix Geometric Method

The Markov chain associated with operator (2.6) belongs to the class of random processes with repetitive structures. As such, one can apply standard techniques from the rich literature on matrix analytic methods [33, 38]. The essence of this approach is to take advantage of the symmetric interactions among different levels of the Markov chain. For $q \geq 2$, the recursive structure of our system can be expressed through the formula

$$\pi_{q+1}\mathbf{A}_2 = \pi_q(\mathbf{I} - \mathbf{A}_1) - \pi_{q-1}\mathbf{A}_0.$$

In finding a solution to this matrix equation, it seems that the general form of the embedded Markov structure and, specifically, its block partitioning are far more important than the precise values of each submatrix. The stationary distribution of the queue, in matrix-geometric form, is characterized in the following theorem.

Theorem D.2. *Consider a positive recurrent, irreducible Markov chain on a countable state space with transition probabilities given by* (2.6)*. Let the matrix* $\mathbf{U}$ *be defined such that the* (c, d) *entry is the probability that, starting from state* $(1, c)$*, the Markov chain* $\{Y_s\}$ *first re-enters level one by visiting* $(1, d)$ *and does so without visiting any state at level zero. The substochastic matrix* $\mathbf{U}$ *may be computed as the limit, starting from* $\mathbf{U}_1 = \mathbf{A}_1$*, of the sequence defined by*

$$\mathbf{U}_{j+1} = \mathbf{A}_1 + \mathbf{A}_0 \left(\mathbf{I} - \mathbf{U}_j\right)^{-1} \mathbf{A}_2. \tag{2.15}$$

Let matrix $\tilde{\mathbf{T}}$ *be given by*

$$\tilde{\mathbf{T}} = \begin{bmatrix} \mathbf{C}_1 & \mathbf{C}_0 \\ \mathbf{A}_2 & \mathbf{A}_1 + \mathbf{R}\mathbf{A}_2 \end{bmatrix} \tag{2.16}$$

where $\mathbf{R} = \mathbf{A}_0 \left(\mathbf{I} - \mathbf{U}\right)^{-1}$*. Then,* $\tilde{\mathbf{T}}$ *is a stochastic matrix associated with an irre-*

ducible Markov chain. If we denote the invariant distribution associated with $\tilde{\mathbf{T}}$ *by* $[\tilde{\pi}_0\ \tilde{\pi}_1]$*, then the stationary distribution associated with* $\mathbf{T}$ *can be expressed as*

$$\pi_0 = \frac{\tilde{\pi}_0}{(\tilde{\pi}_0 + \tilde{\pi}_1(\mathbf{I}-\mathbf{R})^{-1})\,\mathbf{1}} \qquad \pi_q = \frac{\tilde{\pi}_1\mathbf{R}^{q-1}}{(\tilde{\pi}_0 + \tilde{\pi}_1(\mathbf{I}-\mathbf{R})^{-1})\,\mathbf{1}} \tag{2.17}$$

where $q \geq 1$.

Proof. For completeness, we provide a succinct outline of a proof to this theorem; our arguments are motivated, partly, by the derivation presented by Latouche and Ramaswami [33]. For quasi-birth-death processes, several authors have reported similar results [32, 34, 39].

The transitions of the Markov chain $\{Y_s\}$, excluding states at level zero, are governed by the substochastic matrix

$$\begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \mathbf{A}_0 & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_2 & \mathbf{A}_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{2.18}$$

We note that the transitions of levels in $\{Y_s\}$, excluding levels zero and one, are dictated by the same semi-infinite matrix (2.18). Exploiting this symmetry and the fact that $\{Y_s\}$ can only jump to neighboring levels, one can use the definition of $\mathbf{U}$ to obtain the following implicit equation,

$$\begin{aligned} \mathbf{U} &= \mathbf{A}_1 + \mathbf{A}_0 \left(\sum_{i=0}^{\infty} \mathbf{U}^i \right) \mathbf{A}_2 \\ &= \mathbf{A}_1 + \mathbf{A}_0 \left(\mathbf{I} - \mathbf{U} \right)^{-1} \mathbf{A}_2. \end{aligned} \tag{2.19}$$

This is equivalent to a quadratic matrix equation and it can be solved efficiently using numerical methods. For instance, multiplying both sides of (2.19) by $\mathbf{R}$, substituting

$\mathbf{R} = \mathbf{A}_0 (\mathbf{I} - \mathbf{U})^{-1}$, and rearranging terms, we obtain

$$\mathbf{A}_0 + \mathbf{R}\mathbf{A}_1 + \mathbf{R}^2\mathbf{A}_2 = \mathbf{R}. \tag{2.20}$$

We proceed to show that the iterative algorithm of (2.15) is one possible method to obtain $\mathbf{U}$. Consider the probability space of paths on $\{1, 2\} \times \mathbb{N}_0$ with the measure induced by Markov chain $\{Y_s\}$. Let $S_j(c, d)$ be the event such that, starting from state $(1, c)$, the Markov chain $\{Y_s\}$ first re-enters level one by visiting $(1, d)$ and with the excursions constrained to lie between levels one and j. It follows from this definition that $S_j(c, d) \subset S_{j+1}(c, d)$ and, therefore, $\lim_{j\to\infty} S_j(c, d) = \cup_{j\in\mathbb{N}} S_j(c, d)$. Utilizing the monotone convergence theorem, we gather that

$$\lim_{j\to\infty} \Pr(S_j(c, d)) = \Pr\left(\bigcup_{j\in\mathbb{N}} S_j(c, d)\right).$$

By construction, $[\mathbf{U}]_{c,d} = \Pr\left(\bigcup_{j\in\mathbb{N}} S_j(c, d)\right)$. To complete the proof, it remains to show that $[\mathbf{U}_j]_{c,d}$, as defined in iteration (2.15), is the probability of event $S_j(c, d)$.

Mathematically, this is equivalent to the statement $[\mathbf{U}_j]_{c,d} = \Pr(S_j(c, d))$ for all $j \in \mathbb{N}$, which we verify using induction. From the definition of $S_j(c, d)$, we immediately obtain $\Pr(S_1(c, d)) = [\mathbf{A}_1]_{c,d}$ and, consequently, $[\mathbf{U}_1]_{c,d} = \Pr(S_1(c, d))$ because $\mathbf{U}_1 = \mathbf{A}_1$. To continue, we assume this proposition holds for all integers less than or equal to j and show this implies that it holds for $j + 1$. First, we note that $S_{j+1}(c, d)$ is the event such that, starting from state $(1, c)$, the Markov chain $\{Y_s\}$ first re-enters level one by visiting $(1, d)$ and with the excursions constrained to lie between levels one and $j + 1$. The elements in $S_{j+1}(c, d)$ can be partitioned into sets according to their number of visits to level two. In particular, the Markov chain $\{Y_s\}$ remains at level one with probability $[\mathbf{A}_1]_{c,d}$. Alternatively, it can immediately transition to level two, revisit this level a number of times while remaining between

levels two and $j+1$, and then jump back down to level one.

Key to the proof is the symmetric nature of the chain: the probability that, starting from state $(2,c)$, the Markov chain $\{Y_s\}$ first re-enters level two by visiting $(2,d)$ and with the excursions constrained to lie between levels two and $j+1$ is equal to $\Pr(S_j(c,d))$. Indeed, there is a natural, probability-preserving bijection between paths in $S_j(c,d)$ and paths from $(2,c)$ that first re-enter level two at $(2,d)$ and remain between levels two and $j+1$. By the Markov property and our inductive hypothesis, we can write the probability that, starting from $(1,c)$, the Markov chains immediately goes up to level two, visits this level exactly $k+1$ times before it first re-enters level one at $(1,d)$ as $\left[\mathbf{A}_0\mathbf{U}_j^k\mathbf{A}_2\right]_{c,d}$. The proposed partition of $S_{j+1}(c,d)$ is a countable union of disjoint events, where each set accounts for a distinct number of visits to level two. It follows from the renewal property of Markov chains and the symmetry of the problem that

$$\begin{aligned}\Pr\left(S_{j+1}(c,d)\right) &= \left[\mathbf{A}_1 + \mathbf{A}_0\sum_{k=0}^{\infty}\mathbf{U}_j^k\mathbf{A}_2\right]_{c,d} \\ &= \left[\mathbf{A}_1 + \mathbf{A}_0\left(\mathbf{I}-\mathbf{U}_j\right)^{-1}\mathbf{A}_2\right]_{c,d} = \left[\mathbf{U}_{j+1}\right]_{c,d}\end{aligned}$$

where the second equality follows from the Neumann expansion and the third equality is an application of definition (2.15). Hence, for every $j\in\mathbb{N}$, we have $\Pr(S_j(c,d)) = [\mathbf{U}_j]_{c,d}$. This establishes that the iterative algorithm of (2.15) converges to $\mathbf{U}$, as desired.

To complete the proof, it remains to show that the candidate distribution specified in Theorem D.2 is indeed the invariant distribution of $\mathbf{T}$. Notice that (2.20) immediately ensures that

$$\pi_{q-1}\mathbf{A}_0 + \pi_q\mathbf{A}_1 + \pi_{q+1}\mathbf{A}_2 = \pi_q \tag{2.21}$$

for $q \geq 2$. Consider the finite matrix $\tilde{\mathbf{T}}$ with non-negative entries introduced in (2.16). We wish to prove that this is a stochastic matrix. Since $\mathbf{T}$ represents a probability transition matrix, we already have $(\mathbf{C}_0 + \mathbf{C}_1)\mathbf{1} = \mathbf{1}$. To establish the second equality, we examine the following progression,

$$\begin{aligned}(\mathbf{A}_2 + \mathbf{A}_1 + \mathbf{R}\mathbf{A}_2 - \mathbf{I})\mathbf{1} &= (\mathbf{R}\mathbf{A}_2 - \mathbf{A}_0)\mathbf{1} \\ &= \left(\mathbf{R}\mathbf{A}_2 + \mathbf{R}\mathbf{A}_1 + \mathbf{R}^2\mathbf{A}_2 - \mathbf{R}\right)\mathbf{1} \\ &= \mathbf{R}(\mathbf{A}_2 + \mathbf{A}_1 + \mathbf{R}\mathbf{A}_2 - \mathbf{I})\mathbf{1}\end{aligned}$$

where the first step relies on the identity $(\mathbf{A}_0 + \mathbf{A}_1 + \mathbf{A}_2)\mathbf{1} = \mathbf{1}$ and the second step follows from (2.20). Since the matrix $\mathbf{I}-\mathbf{R}$ is invertible, one can move all terms to the LHS to see that $(\mathbf{A}_2+\mathbf{A}_1+\mathbf{R}\mathbf{A}_2-\mathbf{I})\mathbf{1} = \mathbf{0}$ and therefore $(\mathbf{A}_2+\mathbf{A}_1+\mathbf{R}\mathbf{A}_2)\mathbf{1} = \mathbf{1}$. That is, (2.16) is a stochastic matrix. This implies that it admits an invariant distribution which satisfies $\tilde{\pi}_0\mathbf{C}_1 + \tilde{\pi}_1\mathbf{A}_2 = \tilde{\pi}_0$ and $\tilde{\pi}_0\mathbf{C}_0 + \tilde{\pi}_1(\mathbf{A}_1 + \mathbf{R}\mathbf{A}_2) = \tilde{\pi}_1$. Then, for the distribution defined in (2.17), we get

$$\pi_0\mathbf{C}_1 + \pi_1\mathbf{A}_2 = \pi_0$$

$$\pi_0\mathbf{C}_0 + \pi_1\mathbf{A}_1 + \pi_2\mathbf{A}_2 = \pi_0\mathbf{C}_0 + \pi_1(\mathbf{A}_1 + \mathbf{R}\mathbf{A}_2) = \pi_1.$$

These equations, together with (2.21), imply that the distribution defined in (2.17) is invariant under $\mathbf{T}$, as desired. □

Corollary D.3. *When the appropriate inverse matrices exist, one can write the first two levels of the stationary distribution π associated with* (2.6) *as*

$$\pi_1 = \begin{bmatrix} \frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta} \end{bmatrix} \left(\mathbf{A}_2(\mathbf{I} - \mathbf{C}_1)^{-1} + (\mathbf{I} - \mathbf{R})^{-1}\right)^{-1}$$

and $\pi_0 = \pi_1\mathbf{A}_2(\mathbf{I} - \mathbf{C}_1)^{-1}$. The remaining levels are obtained through the recursion $\pi_{q+1} = \pi_q\mathbf{R}$, where $q \geq 1$.

Proof. The coefficients of π_1 can be derived from the channel equilibrium condition,

$$\begin{aligned}\begin{bmatrix} \frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta} \end{bmatrix} &= \sum_{q=0}^{\infty} \pi_q = \pi_0 + \sum_{q=1}^{\infty} \pi_q \\ &= \pi_1 \left(\mathbf{A}_2 (\mathbf{I} - \mathbf{C}_1)^{-1} + (\mathbf{I} - \mathbf{R})^{-1} \right).\end{aligned}$$

Given that an inverse exists, one can solve for π_1 in terms of the invariant distribution of the channel. From there, the distribution at other levels are obtained in a straightforward manner. □

In summary, we have presented an algorithmic method to derive the stationary distribution of Y and obtain the stationary distribution of the queue

$$\Pr(Q = q) = \pi_q \mathbf{1}.$$

The matrix $\mathbf{R}$ is closely related to the asymptotic behavior of the complementary cumulative distribution function of the queue.

Corollary D.4. *The decay rate of the complementary cumulative distribution function of the queue satisfies*

$$\lim_{\tau \to \infty} \frac{1}{\tau} \log \Pr(Q \geq \tau) = \log \varrho(\mathbf{R}),$$

where $\varrho(\mathbf{R})$ *is the spectral radius of* $\mathbf{R}$.

Proof. Since $\mathbf{R}$ is a positive matrix, the Perron-Frobenius theorem implies that $\mathbf{R}$ has a unique positive eigenvalue $\lambda = \varrho(\mathbf{R})$ of maximum modulus [37]. Furthermore, this eigenvalue is associated with a positive left-eigenvector v, and a positive right-eigenvector w. It follows that

$$\frac{u\mathbf{R}^j}{\varrho(\mathbf{R})^j} = u \left(\frac{w^T v}{v w^T} + o(1) \right) = \frac{u w^T}{v w^T} \left(v + o(1) \right)$$

for any non-negative, non-zero vector u. The tail probability of the queue is consequently governed by

$$\begin{aligned}\Pr(Q \geq \tau) &= \pi_1 \left(\sum_{q=\tau-1}^{\infty} \mathbf{R}^q \right) \mathbf{1} = \pi_1 \mathbf{R}^{\tau-1} (\mathbf{I} - \mathbf{R})^{-1} \mathbf{1} \\ &= \frac{\pi_1 w^T}{v w^T} \varrho(\mathbf{R})^{\tau-1} \left((v + o(1)) (\mathbf{I} - \mathbf{R})^{-1} \mathbf{1} \right).\end{aligned}$$

Taking the normalized limit of the logarithm completes the proof. □

3. Linking Generating Functions to the Matrix Geometric Method

We presented two approaches to compute the stationary distribution of the aggregate Markov process. Naturally, these methods must be related. In this section, we explore their connection and we link the generating function procured by first principles to the matrix geometric method. First, we note that $\mathbf{U}$ satisfies the implicit equation

$$\mathbf{U} = \mathbf{A}_1 + \mathbf{A}_0 (\mathbf{I} - \mathbf{U})^{-1} \mathbf{A}_2. \tag{2.22}$$

Using the relation $\mathbf{R} = \mathbf{A}_0 (\mathbf{I} - \mathbf{U})^{-1}$ and rearranging terms in (2.22), we get $\mathbf{A}_0 = \mathbf{R} (\mathbf{I} - \mathbf{A}_1) - \mathbf{R}^2 \mathbf{A}_2$. Substituting this into (2.9), we obtain

$$\begin{aligned}\mathbf{D}(z) &= -z^2 \mathbf{R} (\mathbf{A}_1 - \mathbf{I}) - z^2 \mathbf{R}^2 \mathbf{A}_2 + z (\mathbf{A}_1 - \mathbf{I}) + \mathbf{A}_2 \\ &= (\mathbf{I} - z\mathbf{R}) (z (\mathbf{A}_1 - \mathbf{I}) + (\mathbf{I} + z\mathbf{R}) \mathbf{A}_2) \\ &= (\mathbf{I} - z\mathbf{R}) (z (\mathbf{U} - \mathbf{I}) + \mathbf{A}_2) \\ &= (\mathbf{I} - z\mathbf{R}) (\mathbf{U} - \mathbf{I}) \left(z\mathbf{I} - (\mathbf{I} - \mathbf{U})^{-1} \mathbf{A}_2 \right).\end{aligned}$$

The third equality follows from (2.22), with $\mathbf{U} = \mathbf{A}_1 + \mathbf{R}\mathbf{A}_2$. Since the determinant of a matrix product is the product of the individual determinants, we gather that the roots of $\det \mathbf{D}(z)$ are simply given by the roots of its factors. The stable modes of $\mathbf{\Pi}(z)$ correspond to the roots of $\det (\mathbf{I} - z\mathbf{R})$. Since $(\mathbf{I} - \mathbf{U})^{-1} \mathbf{A}_2$ is a stochastic ma-

trix, the Perron-Frobenius theorem asserts that $\det\left(z\mathbf{I}-(\mathbf{I}-\mathbf{U})^{-1}\mathbf{A}_2\right)$ has a root at unity, with any other root having a magnitude smaller than one. These remaining roots correspond to unstable modes of $\mathbf{\Pi}(z)$. Under partial fraction expansion, stability forces the coefficients associated with these latter roots to vanish. The remaining unknown is resolved through the normalization axiom of probability laws. This reconciles the two approaches, which necessarily lead to the same solution. The generating function method can give closed-form expressions if the channel has only two states, whereas the matrix geometric method gives rise to a numerical procedure that applies to any finite-state channel.

E. Performance Evaluation

The detailed characterization presented in the previous sections makes it possible to compute a number of performance criteria for the system under consideration, including the probability of decoding failure, average throughput and mean delay. In this chapter, we focus on two additional performance measures relevant to delay-sensitive communications. We consider the probability that the queue occupancy exceeds a certain threshold, $\Pr(Q > \tau)$. Furthermore, we examine the decay rate of the complementary cumulative distribution function of the queue, as discussed in Corollary D.4.

Throughout this numerical study, unless stated otherwise, we employ the following system parameters. The Gilbert-Elliott erasure channel is defined by $\alpha = 0.02$, $\beta = 0.005$, $\varepsilon_1 = 0.49$, and $\varepsilon_2 = 0.0025$. This yields an average bit-erasure probability of $\overline{\varepsilon} = 0.1$ and the channel memory decays at an exponential rate of $(1-\alpha-\beta) = 0.975$. During every codeword transmission attempt, a new packet arrives at the source with probability $\gamma = 0.25$, and the expected packet length is set to $\rho^{-1} = 195$

bits. The block length is fixed at $N = 114$ symbols. If codewords are transmitted every 4.615 ms, then this corresponds to a mean arrival rate of roughly 10.6 kbps and an ergodic channel capacity of roughly 22.2 kbps. These quantities are selected to loosely reflect the operation of a wireless GSM-based relay link. Collectively, these parameters dictate the evolution of the Markov process governing the queue.

The Shannon limit for the Gilbert-Elliott erasure channel channel is $1-\overline{\epsilon}$ bits per channel use and the capacity-achieving input distribution is i.i.d. uniform. In fact, this is true for an arbitrary (ergodic) erasure channel where the expected number of erasures is independent of the input sequence. Suppose X and Y are the input and output vectors of the erasure channel, respectively. Let O be a vector which indicates the observed (i.e., not erased) positions at the destination. Then,

$$I(X;Y) = H(Y) - H(Y|X) = H(O) + H(X_O) - H(O) = H(X_O),$$

where X_O is the subvector of X that includes only the observed positions. This quantity is simultaneously maximized, for all O, by choosing X to be i.i.d. uniform. Therefore, the maximum mutual-information rate is equal to the average number of unerased positions.

We continue our early analysis with simple performance criteria that are based solely on the evolution of the channel. They do not take into consideration the behavior of the queue at the transmitter. One such criterion is the probability of decoding failure at the receiver, which is equal to

$$\begin{aligned} P_{\mathrm{f}}(N-K) &= \sum_{e=0}^{N} P_{\mathrm{f}}(N-K,e)\Pr(E=e) \\ &= \sum_{e=0}^{N} P_{\mathrm{f}}(N-K,e)[\![x^e]\!]\left(\begin{bmatrix} \frac{\beta}{\alpha+\beta} & \frac{\alpha}{\alpha+\beta} \end{bmatrix}\mathbf{P}_x^N\mathbf{1}\right). \end{aligned}$$

A closely related measure of performance is the average throughput associated with

a saturated source; in the present setting, this is given by

$$T_{\mathrm{s}}(K,N) = \frac{K}{N}\left(1 - P_{\mathrm{f}}(N-K)\right).$$

The probability of decoding failure as a function of K appears in Fig. 3. The average

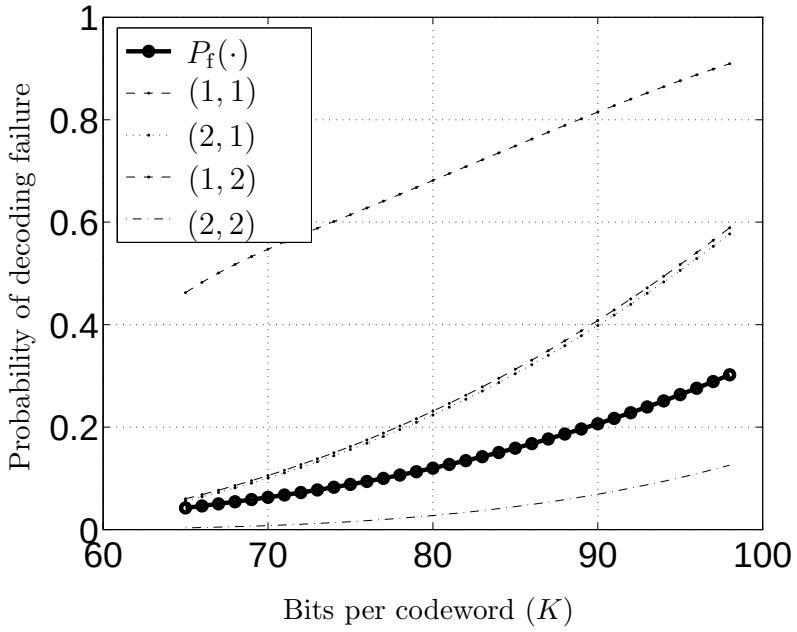

Fig. 3. Probability of decoding failure, $P_{\mathrm{f}}(N-K)$, as a function of the number of information bits per codeword. The conditional probabilities of decoding failure for various values of (c_1, c_{N+1}) are also included.

throughput associated with a saturated source is plotted as a function of the number of information bits per codeword in Fig. 4. These two figures illustrate well the natural tradeoff between data content and error protection. In particular, these competing considerations lead to the unimodal throughput function of Fig. 4, where optimal performance is achieved at $K = 87$. A naive conjecture would place $K = rN$ close to the rate implied by the Shannon limit $(1-\overline{\varepsilon})N = 0.9 \times 114 = 102.6$, but this is much larger than the optimal value of $K = 87$ for maximizing average throughput.

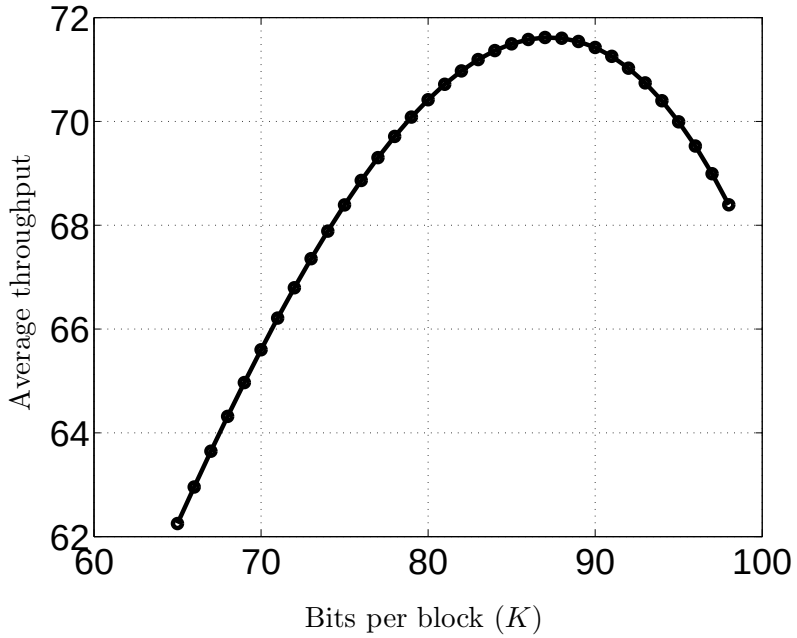

Fig. 4. Average throughput for a saturated source as a function of K, the number of information bits per codeword. The maximum throughput is obtained at $K = 87$.

For a short block length, two factors affect the optimal code-rate K for a prescribed queueing performance. The block length may be too small to ensure convergence of the empirical average number of erasures within a block. In addition, dependencies from block to block are non-negligible. Although the probability of decoding failure and the average throughput capture the effect of channel correlation within a block, they do not capture dependencies from block to block. This is a subtle yet important observation, especially for delay-sensitive traffic. The impact of these factors becomes more severe with increasing channel memory. This consideration underlies much of the queueing analysis presented in this chapter. Time-dependencies in the service process of a queue can alter system performance dramatically. We thus turn to queue-based performance criteria.

Figure 5 depicts $\Pr(Q > \tau)$, the complementary cumulative distribution function of the queue, as a function of K. Each curve represents the probability that, in steady-state, buffer occupancy exceeds a certain threshold τ, where $\tau \in \{5, 10, 15, 20, 25\}$. The low threshold values reflect the intended use of this methodology in the context of delay-sensitive applications. As expected, the probability of the queue exceeding

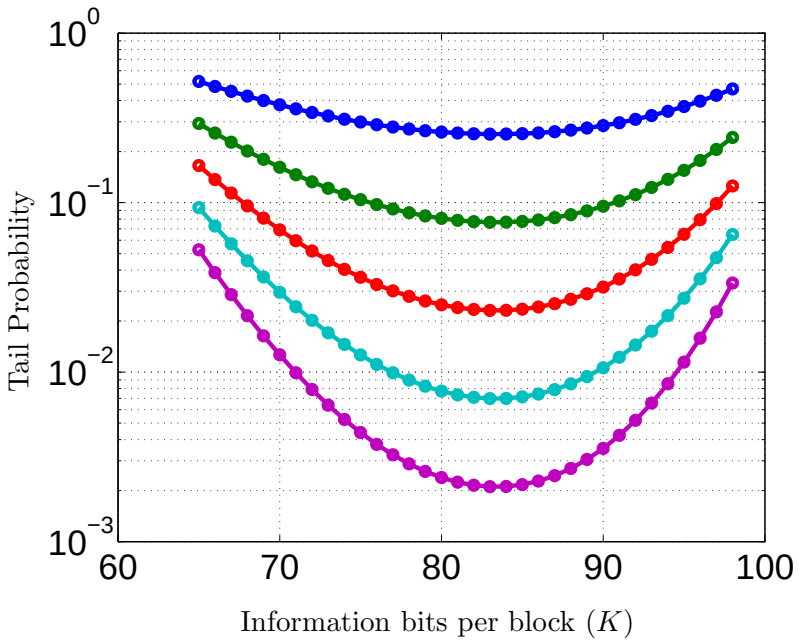

Fig. 5. Tail probabilities in the equilibrium packet distribution of the queue $\Pr(Q > \tau)$, for threshold values $\tau \in \{5, 10, 15, 20, 25\}$, as a function of the number of information bits K per codeword. The minimums occur uniformly at $K = 83$ for all threshold values.

a prescribed threshold decreases as τ increases. More interestingly, we note that $K = 83$ appears uniformly optimal for all values of τ. That is, the optimal code-rate seems robust to the choice of threshold value τ. This robustness property remains present for the other system parameters we tested. Further supporting evidence for this observation is offered by looking at the asymptotic decay rate in tail occupancy, displayed in Fig. 6. When the arrival rate $\gamma\rho^{-1}$ is between 47.5 and 60, one finds

that $K = 83$ is also optimal in terms of tail decay. The true optimum $K = 83$ is closer to the throughput maximizing code-rate $K = 87$ than to the naive conjecture.

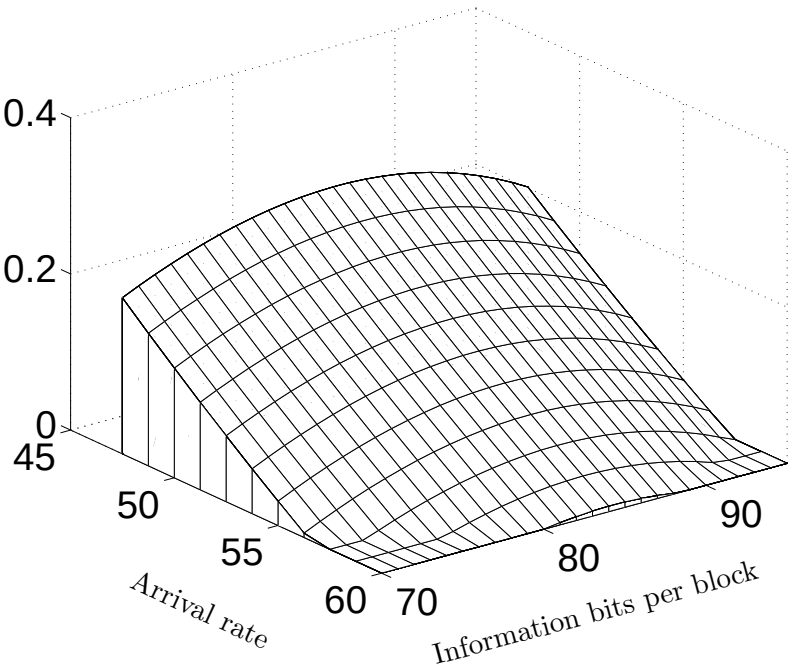

Fig. 6. Tail decay rate, $-\lim_{\tau\uparrow\infty} \frac{1}{\tau} \log \Pr(Q \geq \tau)$, as a function of the number of information bits K per codeword and the average arrival rate $\gamma\rho^{-1}$ in bits per codeword transmission interval.

We explore the impact of channel correlation on optimal code-rate in our next set of results. We fix $\beta : \alpha$ at a ratio of one to four, and vary the memory factor $(1-\alpha-\beta)$. When the channel is memoryless, the optimal K is 93. For comparison, the capacity is $1-\overline{\varepsilon} = 0.9$, which yields a K of roughly 103. As correlation increases, the optimal value of K initially decreases, thereby offering more protection against erasures. Yet, when the coherence time of the channel starts to approach the length of a codeword, $N = 114$, the error correcting code becomes ineffective as it fails to handle the increasingly likely long sequences of successive erasures. The optimal strategy then

progressively shifts to including more information bits in every packet, and hoping that the channel remains in its good state. In the limiting regime where $(1-\alpha-\beta)$ approaches one, the optimal strategy is to transmit uncoded data, $K = N$. Indeed, strong correlation is characterized by long strings of erasures followed by longer strings of reliable bits, and the best strategy is to send as many information bits as possible when the channel is good. At this point, the bit erasure channel essentially becomes a correlated packet erasure channel. Numerical results are summarized in Table I.

Table I. Optimal number of information bits per codeword and threshold violation probability as functions of channel memory.

$1-\alpha-\beta$	Optimal K	$\min \Pr(Q > 5)$
0	93	0.0359
0.5	91	0.0438
0.9	85	0.0982
0.98	85	0.2843
0.99	95	0.3169

F. Discussion and Concluding Remarks

This chapter presents a new framework to analyze the relation between code-rate and queueing behavior for communications over channels with memory. The simplicity of the erasure channel and its closed-form characterization of error events are instrumental in conducting our analysis. For short block lengths and channels with memory, the optimal code-rate appears to be linked to the relative size of a codeword compared to the coherence time of the channel. In certain circumstances, it

is beneficial to provide significant protection against erasures. However, as channel memory increases, performance may be improved by incorporating more data bits in every codeword. In this latter case, the transmitter resorts to a strategy where information is successfully sent when the channel starts in a good state, and it is lost otherwise. This is in stark contrast to information theoretic results obtained through asymptotically long codewords. Once the block length is selected, the optimal code-rate seems rather impervious to the queue occupancy threshold. This observation considerably simplifies system design because an optimal code-rate can be selected irrespective of the target queue length. The set of admissible arrival rates, on the other hand, will depend heavily on the queueing objective.

A distinguishing feature of this work is that it provides a rigorous approach linking queueing performance to the operation of a communication system at the physical layer. The methodology and results are developed for the Gilbert-Elliott erasure channel, but can be generalized to more intricate finite-state channels with memory. For example, the simple performance characterization of random codes over erasure channels may extend to hard-decision decoding of BCH codes over Gilbert-Elliott error channels. Possible avenues of future research include the study of alternative arrival processes, the ability to vary the rate and the length of codewords dynamically, and a more pragmatic feedback scheme.

CHAPTER III

QUEUEING ANALYSIS OF A BUTTERFLY NETWORK FOR COMPARING NETWORK CODING TO CLASSICAL ROUTING

Network coding has gained significant attention in recent years as a means to improve throughput, especially in multicast scenarios. These capacity gains are achieved by combining packets algebraically at various points in the network, thereby alleviating local congestion at the nodes. The benefits of network coding are greatest when the network is heavily utilized or, equivalently, when the sources are saturated so that there is data to send at every scheduling opportunity. Yet, when a network supports delay-sensitive applications, traffic is often bursty and congestion becomes undesirable. The lighter loads typical of real-time traffic with variable sources tend to reduce the returns of network coding. This work seeks to identify the potential benefits of network coding in the context of delay-sensitive applications. As a secondary objective, this chapter also studies the cost of establishing network coding in wireless environments. For a network topology to be suitable for coding, links need to possess a proper structure. The cost of establishing this structure may require excessive radio resources in terms of bandwidth and transmit power. Bursty traffic together with structural cost, tend to decrease the potential benefits of network coding. This chapter describes how, for real-time applications over wireless networks, there exist network topologies for which it may be best not to establish a network structure tailored to network coding.

A. Introduction

Network coding is a novel paradigm that has received much attention in the literature recently [40–44]. It has the potential to improve the throughput and robustness

of future communication networks. These performance gains are achieved by relaxing the restriction that data belonging to different information flows should remain separated. Indeed, network coding is a transmission strategy where packets are combined algebraically at intermediate nodes in the network. It can be viewed as an extension of traditional routing. In certain circumstances, network coding helps improve overall throughput; and it is known to achieve the min-cut flow in multicast scenarios [2].

The research enthusiasm generated by network coding can be explained, partly, by the ever expanding demand for Internet access and fast connectivity. Not only is network coding mathematically elegant, it seeks to improve network performance at a time when the number of data applications is rising furiously. The growing demand for network connectivity is felt both at the core of the Internet and at its periphery, where wireless systems are increasingly employed to provide flexibility to mobile users. One class of data connections that is rapidly gaining prominence on the Internet is the traffic generated by real-time applications. Delay-sensitive services including voice over Internet protocol (VoIP), video conferencing, gaming and electronic commerce are now commonly used by vanguardists on both wired and wireless devices. Future communication infrastructures are expected to carry much larger volumes of data with varying quality of service (QoS) requirements. As such, this chapter seeks to provide preliminary answers to two important questions related to delay-sensitive traffic and the efficient utilization of network coding.

First, are the potential benefits of network coding as substantial in the context of delay-sensitive applications? It seems intuitively clear that the gains of network coding are maximal when the links in the network are fully utilized. However, the bursty nature of many data sources and the service quality required of most real-time applications may force a network to operate much below its maximum throughput.

This phenomenon is captured by the concept of effective bandwidth, which identifies the data-rate needed by a source to fulfill its service requirement [3, 4]. In general, the effective bandwidth of a source can be much larger than the average throughput it produces. The bursty traffic generated by delay-sensitive applications combined with the gains associated with statistical multiplexing act to decrease the benefits of network coding. Therefore, it is not clear how much we gain by applying network coding in a communication system subject to QoS constraints. In this chapter, we provide quantitative results on the benefits of network coding for a simple butterfly network in the context of delay-sensitive applications.

Another pertinent observation about network coding is that it often requires a structured network topology. Coding benefits are optimum when the data-rates of the various links are integer multiples of one another. In a wireless environment, physical-layer resources can be allocated progressively to the different nodes. To maximize the coding gain, these resources must be assigned to create a suitable topology. While this enables efficient coding, there may be a non-negligible cost associated with creating such a structure. In other words, in a wireless environment, the performance of a system with network coding should be compared to the operation of the equivalent classic-routing system, with physical resources allocated optimally in both cases. This leads us naturally to the second question we seek to address. When is it relevant to create a topology suitable for network coding in a wireless environment?

These two important questions are not only related through the rising popularity of real-time applications and network coding, but also by their answers necessitating the development of analogous mathematical tools. This similarity motivates our joint treatment of these related topics. More specifically, we investigate the impact of network coding on the queueing behavior of wireless communication systems.

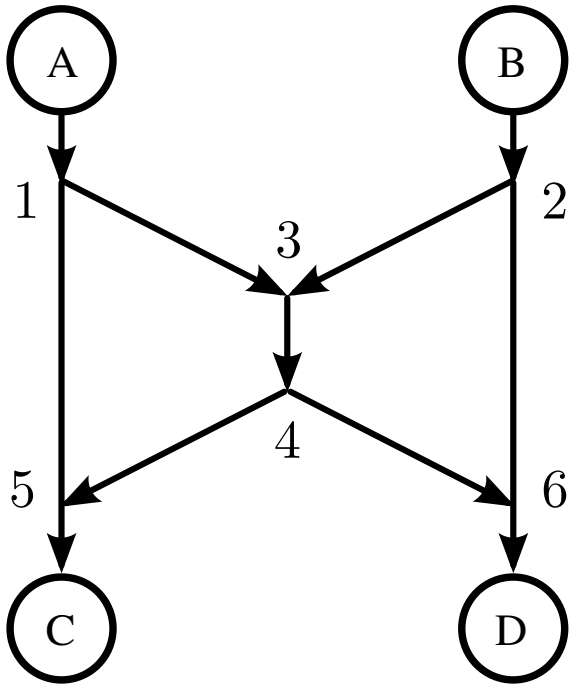

Fig. 7. Directed butterfly network.

We consider a simple scenario where two varying-rate sources communicate to multiple destinations through the notorious butterfly network, shown in Fig. 7. Every node is equipped with a data buffer where packets are stored prior to transmission. We analyze the performance of this system, and compute its achievable rate-region when the network operates under stringent service constraints. Due to the time-varying nature of typical arrival and service processes, it is difficult to provide deterministic delay guarantees for such systems. Accordingly, we adopt a popular statistical QoS criterion that captures the asymptotic decay-rate in buffer occupancy,

$$\theta = -\lim_{x\to\infty} \frac{\ln \Pr\{L > x\}}{x} \tag{3.1}$$

where L has the equilibrium distribution of the buffer at the transmitter. Parame-

ter θ reflects the perceived quality of the corresponding communication link: a larger θ implies a lower probability of violating a queue-length restriction and a tighter QoS constraint. This performance criterion is closely tied to large-deviations theory, and it forms a basis for the concept of effective bandwidth which has been studied extensively in the past [45–47, 47–50]. Given a specific arrival process, the effective bandwidth characterizes the minimum constant service-rate required for a communication system to meet its QoS requirements. Parameter θ is also related to the dual concept of effective capacity popularized by Guerin et al. [51], de Veciana et al. [52] and Wu and Negi [16]. Unlike wired networks, wireless links frequently feature time-varying service rates [53]. The effective capacity characterizes the maximum constant arrival rate that a wireless system can support, given a minimum buffer occupancy decay-rate θ_0. When the decay-rate θ_0 approaches zero, the effective capacity converges to the maximum throughput supported by the wireless channel.

To study the performance of a communication system subject to a buffer occupancy constraint akin to (3.1), we need to characterize the queueing performance of the network. In the mathematical framework under consideration, independent sources sharing a same link can be studied separately. This is one of the appealing properties of an analysis based on large deviations. The main challenge, as we will see, is to characterize the performance of the tandem network shown in Fig. 8. This network consists of two successive nodes where the output of the first node acts as an input to the second queue.

1. Contributions

Fig. 9 shows the butterfly network we want to study. For a multicast scenario, where stochastic sources A and B wish to communicate to destinations C and D, node 3 has the opportunity to employ packet combining. We consider two distinct versions

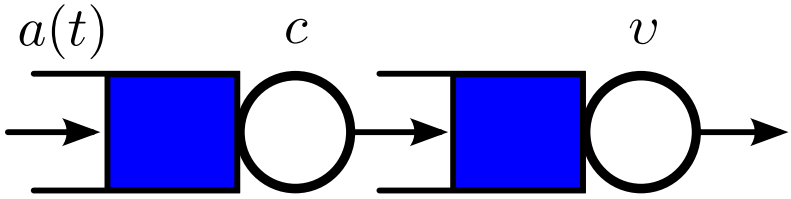

Fig. 8. Network with tandem queues.

of this simple butterfly network. First, we analyze a noise-limited network with constant and identical link capacities. This configuration is suitable for network coding and is the basis for our initial queueing comparison. Then, we examine a wireless network under a broadcast paradigm. In the latter scenario, we assume that physical resources can be allocated freely among the various nodes to create non-identical links, thereby enabling optimal operation within each configuration.

To compare the queueing performance of network coding versus classic routing, we characterize the achievable rate regions for both these cases under a QoS guarantee on the tail-asymptotics of the buffer-content distributions. Not too surprisingly, network coding outperforms classic routing for a network with identical link capacities. Although statistical multiplexing had the potential to offset some of the coding gain, classic routing remains a distant second to network coding for all QoS requirements. More interesting results come from the wireless butterfly analysis. Combining packets at an intermediate wireless relay does not necessarily yield performance gains, and may even be detrimental in some cases. This behavior depends on the topology of the butterfly network and the physical locations of the nodes. This peculiarity follows from the fact that network coding needs symmetric links between the sources and their destinations for maximal coding gains to be realized. If the link capacities are not identical, then packet combining entails delay and in-

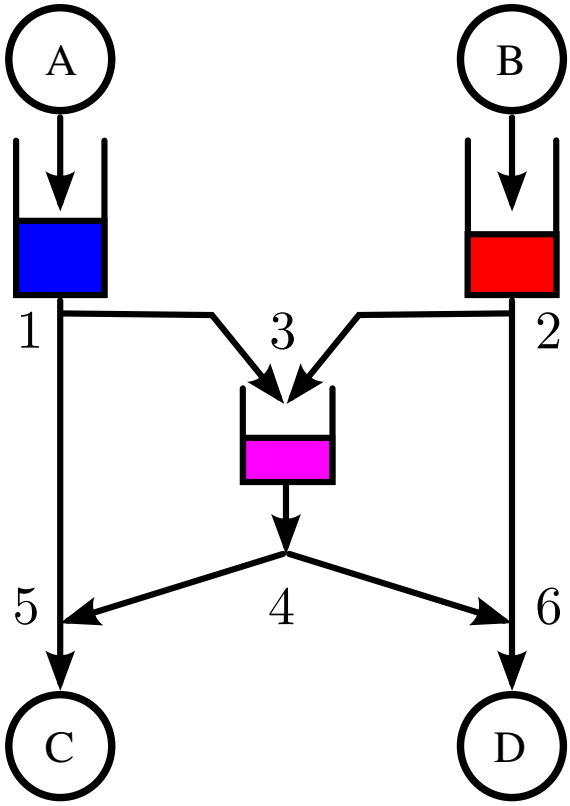

Fig. 9. Butterfly network of interest with corresponding buffers.

efficiencies. These results are detailed below, and they may be employed to provide guidelines on when to form a network suitable for combining packets algebraically at intermediate nodes. Analysis is limited to the simple butterfly network. However, the featured approach should generalized to topologies with constant service rates, using the same queueing methods employed in this chapter. In fact, our methodology for networks with two queues in tandem can be applied to several queues in cascade. Still, an exact analysis would get increasingly complex for large networks and these techniques may not lead to tractable expressions. In more complex topologies, it may be necessary to use approximation methods, which is a different topic altogether.

We would like to emphasize that our contribution is two-fold. We provide a quantitative analysis of the tradeoff between the cost of establishing a structure suitable for network coding and the ensuing returns associated with algebraic packet combining. Second, we offer an alternate proof for deriving the tail-exponent of the second buffer in a tandem queue.

2. Relevant Work

There have been many recent interesting contributions on queueing behavior when network coding is employed at the transmitter. We compare and contrast our work with the available literature below and summarize this discussion in Table II.

The random linear combining of data packets is considered in [54–58]. The authors focus on coding delay in [55], whereas decoding delay is studied in the remaining contributions. In [54], QoS is defined in terms of packet drop probability, and multiple flows are considered over an arbitrary network. In [55], the authors compare the performance of network coding versus scheduling for broadcast and multiple unicast scenarios; their work is based on average delay performance. In [56, 57], the authors explore the throughput-delay tradeoff with and without network coding. In [56], the coding scheme adapts to the underlying traffic conditions. Stability and delay performance of a multicast erasure channel with stochastic arrivals are studied in [57]. In [58], the authors propose a coding and queue management algorithm. Note that the random linear combining of packet transmissions is in effect a coding scheme which trades off delay and throughput over a single flow. In our work, we study the achievable rate region over a butterfly network when one employs either network coding or classic routing at the intermediate node, a distinct framework altogether.

Previous contributions differ from the framework presented below in many more respects. First, we are comparing network coding to routing at a given node in a

simple butterfly network. We consider a multicast scenario with two transmitters and two receivers. Second, we disregard coding/decoding delay and focus on the equilibrium queue-length buildup due to stochastic arrivals, limited service rate and feedback from the receiver. Third, we use pre-selected codes with fixed rates for reliable communication over the links. Furthermore, our QoS constraint on tail-asymptotics deviates significantly from average delay constraints, thereby offering better insights for delay-sensitive traffic. This methodology can be utilized to provide different QoS guarantees to flows with various requirements. In addition, we take into consideration the stochastic nature of arrivals, which may reduce the gains of network coding due to statistical multiplexing.

A two-way relay channel is considered in [59–62]. Note that two-way relays can be modeled as a butterfly network where sources are also destinations. However, this is not an equivalence relation. In the two-way relay model, direct links are absent and side-information about the received data is available at the destination, an advantage that is not present in our network. In [60], the authors characterize the end-to-end rate regions for MAC-XOR and PHY-XOR operations for two-way relay channel. They also present an opportunistic scheduling algorithm to show that the system can be stabilized for any bit-arrival rate pair within the Shannon rate region. In [61], the authors study the energy-delay tradeoff when network coding is used at the relay node. The energy is measured in terms of code-rates and channel conditions. In [62], authors show that, for asymmetric traffic, one needs to perform time-sharing between traditional slotted multi-hopping and network coding. In [63], the authors characterize the stability region for bursty traffic at multiple sources with and without network coding in the wireless network. In [59], the authors propose a framework to develop adaptive joint network coding and scheduling schemes. The authors show that, for asymmetric traffic, scheduling and network coding need to

be combined to maximize supportable throughputs over this network. In [64], the authors propose a dynamic routing-scheduling-coding strategy based on queue-state information for serving multiple unicast sessions. The authors study queues over a butterfly network for a discrete-time packet model, and they use fluid approximations to show the stability of the system for their proposed algorithm.

Again, we emphasize that the wireless butterfly network considered in our chapter is not equivalent to two-way relaying. We assume adaptive network coding in the sense that every time there is an opportunity to do network coding, the relay node combines packets algebraically. The servicing policy adopted throughout is taken to be first-come first-served. A fluid assumption allows us to bypass the scheduling problem. It is clear that, given a suitable network topology, the flexibility to switch between network coding and classical routing, when the need arises, is better than always using classical routing. What sets our work apart is accounting for the cost associated with making a network suitable for network coding. We study whether coding gains can offset the cost of facilitating a proper network structure. The rate regions achieved in [60, 62–64], correspond to the cases where the queue is stable for stochastic arrivals. There is no explicit guarantee on the buffer distribution. In our work, we find the achievable rate region under a tail-asymptote requirement on the decay-rate of the queue distribution. Therefore, the conclusions by the authors in the aforementioned papers would not necessarily apply in a framework akin to ours.

3. Organization

The remainder of this chapter is organized as follows. We introduce the system model in Section B. We list pertinent results on the performance of tandem queues in Section C. These tools are used to analyze symmetric butterfly networks in Section D. Wireless butterfly networks are studied in Section E. Key queueing results about

the second buffer of a tandem network are established in Section F. A complete characterization of the departure process at the output of a single queue with constant service and Markov-modulated arrivals is presented in Section 1. An expression for the equilibrium distribution of the buffer-occupancy is derived in Section 2. This enables us to obtain the corresponding asymptotic decay-rate of buffer-occupancy in Section G. In Section H, we determine the maximum achievable rate for the second buffer in the QoS constrained tandem network under consideration. Finally, we conclude with some relevant remarks and future directions in Section I.

B. Problem Statement

We study a communication system where two independent users wish to send their messages to two common destinations over a butterfly network, as shown in Fig. 9. A multicast scenario is considered where independent sources A & B store their respective information in buffers at nodes 1 & 2, and must transmit their data to both destinations C & D. To facilitate this process, node 1 sends its packets to nodes 3 & 5. Similarly, node 2 forwards its packets to nodes 3 & 6. Node 3 can take two courses of action; either it stores the received packets from the sources in a queue and then forwards them individually to node 4, or it combines the packets algebraically before transmitting the data.

The first setting will be called the *classic routing* case. In this scenario, node 4 duplicates the received packets from node 3 and forwards copies to nodes 5 & 6. These destination nodes disregard redundant information (they could potentially take advantage of redundancy to improve the reliability of previously received messages, but this is beyond the scope of this chapter) and retain new data. For the second scenario, we consider the network coding scheme where node 3 adds the two

streams of packets over *GF(2)* and relays the coded messages to node 4 [2]. The latter duplicates the received packets and transmits them to nodes 5 & 6. Node 5 can resolve the information received from node 2 by adding the packets obtained from node 1 to the corresponding packets received from node 4. In a similar fashion, node 6 can decode the information originating from node 1 by adding packets from node 2 to the corresponding packets from node 4. Service quality is captured by a global QoS constraint θ_0 on the system. That is, the asymptotic decay-rate of buffer occupancy must be greater than or equal to θ_0 for all the queues in the system.

For the sake of analysis, we assume that packets are infinitely divisible and hence the arrival and service processes are fluid in nature. Thus, it becomes possible to define instantaneous arrival and service rates. Under this assumption, every node in the network is equipped with a single fluid queue served by an individual transmitter. We also take the buffers in the system to be arbitrary large. A similar approach applies to the finite buffer case, albeit with additional boundary conditions on the buffer occupancy.

1. Source Model

Many real-time traffic sources can be accurately represented by *on-off* models [65]. This motivates our assumption of arrivals being two-state Markov-modulated fluid processes. In addition, there is a vast amount of literature available on the queueing behaviors of Markov-modulated fluid processes for wire-line networks [4, 45, 66]. We postulate that sources A & B are independent, and that they both satisfy the following assumption.

Assumption B.1. *During an* on *period, the source emits packets at a constant peak rate into its buffer; it remains idle otherwise. Moreover, the* on *and* off *times are*

independent and exponentially distributed.

Mathematically, the sources are defined through their underlying Markov chains. Let $\{I_1(t) \in \{0,1\} : t \geq 0\}$ and $\{I_2(t) \in \{0,1\} : t \geq 0\}$ be independent two-state continuous-time Markov chains (CTMC) modulating *on-off* sources A & B, respectively. State zero represents the *off* state and state one denotes the *on* state. Suppose that the peak-rate for the source at node $i \in \{1,2\}$ is taken as a_i. With a slight abuse of notation, we can write the arrival process at buffer i as

$$a(I_i(t)) = a_i \mathbf{1}_{\{I_i(t)=1\}}, \quad i \in \{1,2\}$$

where $\mathbf{1}_{\{\cdot\}}$ represents the standard set indicator function. We denote the mean *off* and *on* times by λ_i^{-1} and μ_i^{-1}, respectively. The generator matrix for the modulating two-state Markov process can then be written as

$$Q_i = \begin{bmatrix} -\lambda_i & \lambda_i \\ \mu_i & -\mu_i \end{bmatrix}, \quad i \in \{1,2\}.$$

2. Queueing Model

We denote the capacity of link i–j by c_{ij}. This capacity effectively limits the offered service rate at node i for transmission to node j. In particular, if there exists a link between nodes i & j and the buffer associated with node i is non-empty, then node i can transmit to node j at a maximum rate c_{ij}. For simplicity, we assume that $c_{34} = c_{45} = c_{46} = c_3$. The offered service-rates on links 4–5 and 4–6 are then equal to the maximum arrival-rate at node 4. As such, node 4 doesn't need to store data. It only facilitates the duplication and the forwarding of its received packets to nodes 5 & 6. In other words, the buffer associated with node 4 is always empty.

Node 1 sends the same information to both nodes 3 & 5, and therefore retains

data in its buffer until both receiving nodes have acquired the corresponding packet. Accordingly, the service rate at node 1 is $c_1 = \min\{c_{13}, c_{15}\}$. Similarly, the service offered to the buffer at node 2 is $c_2 = \min\{c_{23}, c_{26}\}$. Altogether, nodes 1, 2 and 3 transmit packets at rates c_1, c_2 and c_3 respectively, whenever their own buffers are non-empty. Observe that, by construction, congestion can only occur at these three nodes. We can therefore safely assume that there are no queues at the other nodes. We have depicted the fluid model of interest in Fig. 9 for the butterfly network under consideration. We represent the fluid level in the buffer at node i and time t by $L_i(t)$. The stochastic evolution of $L_i(t)$ depends on whether one opts for network coding or classic routing.

3. Network Coding

For network coding, packets originating from links 1–3 and 2–3 are combined algebraically over *GF(2)* and then stored in the buffer at node 3. From a fluid perspective, this is equivalent to both flows entering buffer 3 oblivious of each other. Buffer 3 can be serviced at a maximum rate c_3. However, to prevent decoding delays at the destinations, the service rates offered at nodes 1, 2 and 3 are made equal to $\hat{c}_1 = \min\{c_1, c_3\}, \hat{c}_2 = \min\{c_2, c_3\}$ and $\check{c}_3 = \max\{\hat{c}_1, \hat{c}_2\}$. In this scenario, there is no congestion at node 3 and hence $L_3(t) = 0$ for all times t. Furthermore, for the non-trivial case where $a_i > \hat{c}_i$, we can write the stochastic evolution of $L_i(t)$ as

$$\begin{aligned}\frac{d}{dt}L_i(t) = (a(I_i(t)) - \hat{c}_i)\mathbf{1}_{\{I_i(t)=1\}} \\ - \hat{c}_i\mathbf{1}_{\{I_i(t)=0, L_i(t)>0\}},\end{aligned} \tag{3.2}$$

where $i \in \{1, 2\}$. When the buffer at node 1 is non-empty, the net rate of fluid input is $a(I_1(t)) - \hat{c}_1$; this is called the drift rate. For $i \in \{1, 2\}$, we can define a drift matrix $D_i = \text{diag}(-\hat{c}_i, a_i - \hat{c}_i)$ for the buffer at node i, whose diagonal entries are

the drift rates corresponding to the state of the arrival process. In matrix form, we have

$$D_i = \begin{bmatrix} -\hat{c}_i & 0 \\ 0 & a_i - \hat{c}_i \end{bmatrix}, \quad i \in \{1, 2\}.$$

4. Classic Routing

For the case of classic routing, the service rate offered at node $i \in \{1, 2\}$ is c_i. The evolution of $L_i(t)$ for the non-trivial case of $a_i > c_i$, is then governed by

$$\begin{aligned} \frac{d}{dt} L_i(t) = (a(I_i(t)) - c_i)\mathbf{1}_{\{I_i(t)=1\}} \\ - c_i \mathbf{1}_{\{I_i(t)=0, L_i(t)>0\}}, \end{aligned} \tag{3.3}$$

where $i \in \{1, 2\}$. It will be shown in the later sections that the departure process at the output of buffer i is a two-state *on-off* process modulated by a countable-state Markov process $K_i(t)$. The departure process at node i can be represented as

$$c(K_i(t)) = c_i \mathbf{1}_{\{K_i(t)\neq 0\}}, \quad i \in \{1, 2\}.$$

The buffer at node 3 is fed by the aggregation of these two independent arrival processes, and it is serviced at a constant rate c_3. To initiate the analysis of this more complicated scenario, we study the simple case where the resources at buffer 3 are split between the flows of sources A & B. Consider two parallel buffers at node 3 with positive constant service-rates ν and $c_3 - \nu$, respectively. We assume that the flow from node 1 goes to the first parallel buffer; and the flow from node 2, to the second one. The aggregate fluid in both the buffers will be greater than or equal to the fluid level of a single-buffer system with incoming data from nodes 1 & 2 and service-rate c_3. Thus, the decay-rate of buffer occupancy for a system with a single buffer at node 3 must be no less than the exponential decay-rate of the

corresponding system with parallel buffers and partitioned service. Yet, it can be shown that these asymptotic values are equal for independent flows and optimal splitting rate [47, 49, 50, 67]. The two independent flows can therefore be decoupled and studied separately. If the shared buffer is constrained by a requirement θ_0 on the decay-rate of buffer occupancy, the queues in the decoupled system are constrained by the same parameter θ_0 as well. We denote by $L_{i3}(t)$ the fluid level in the buffer of the decoupled system holding data from node i. We can write the stochastic evolution of $L_{i3}(t)$ for the non-trivial case of $c_i > \max\{\nu, c_3 - \nu\}$, as

$$\begin{aligned}\frac{d}{dt}L_{13}(t) &= (c(K_1(t)) - \nu)\mathbf{1}_{\{K_1(t)\neq 0\}} \\ &\quad - \nu\mathbf{1}_{\{K_1(t)=0, L_{13}(t)>0\}} \\ \frac{d}{dt}L_{23}(t) &= (c(K_2(t)) - c_3 + \nu)\mathbf{1}_{\{K_2(t)\neq 0\}} \\ &\quad - (c_3 - \nu)\mathbf{1}_{\{K_2(t)=0, L_{23}(t)>0\}}.\end{aligned}$$

C. Key Results

In this section, we list the mathematical results needed to compute the achievable rate-regions for data multicast through the butterfly network, and under specific QoS requirements. Let $\ell_1(t)$ be the amount of fluid at time t in a queue being fed by an *on-off* source satisfying Assumption B.1, and serviced at a constant rate c. Let a denote the arrival-rate into this buffer when the source is *on*. The mean *off* and *on* times of the source are denoted by λ^{-1} and μ^{-1}, respectively. Furthermore, the output of this queue (also called departure process) is fed into another arbitrary large reservoir. This second queue is being serviced at a constant rate υ. The amount of fluid in the latter buffer at time t is denoted by $\ell_2(t)$.

We wish to find the maximum peak rate a, such that the QoS criterion of (3.1)

is no less than θ_0 for both queues. Let θ_1 and θ_2 be the asymptotic buffer decay-rate governing the first and second queues in the tandem network, i.e.,

$$\theta_j = -\lim_{x\to\infty} \frac{\ln \Pr\{\ell_j > x\}}{x} \tag{3.4}$$

where ℓ_j is the steady-state queue-length of buffer j. More specifically, we wish to identify set $\mathcal{A}$ (as a function of θ_0, c, and υ) defined by

$$\begin{aligned} \mathcal{A}(\theta_0, c, \upsilon) &= \{a \in \mathbb{R}^+ : \min\{\theta_1, \theta_2\} \geq \theta_0\} \\ &= \{a \in \mathbb{R}^+ : \theta_1 \geq \theta_0\} \cap \{a \in \mathbb{R}^+ : \theta_2 \geq \theta_0\}. \end{aligned} \tag{3.5}$$

If $a \leq c$, the first buffer always remains empty and the behavior of the tandem queue reduces to that of a system with a single queue. We therefore focus on the non-trivial case where $a > c$. Under this condition, Theorem F.1 in Section F asserts that if $a\lambda/(\lambda+\mu) < c < a$ then

$$\theta_1 = \frac{\mu}{a-c} - \frac{\lambda}{c}. \tag{3.6}$$

The lower bound on c ensures stability of the queue. This formula implicitly determines the maximum peak-rate a such that a QoS constraint of θ_0 is satisfied at buffer 1. We define

$$\begin{aligned} \mathcal{A}_1(\theta_0, c) &= \left\{a \in \mathbb{R}^+ : \theta_1 \geq \theta_0\right\} \\ &= \left\{a \in \mathbb{R}^+ : a \leq \bar{a}_1(\theta_0, c)\right\}, \end{aligned} \tag{3.7}$$

where we define $\bar{a}_1(\theta, c) = c + c\mu/(\lambda + c\theta)$. The second buffer always remains empty if $\upsilon \geq c$. Thus, we consider the situation where $\upsilon < c$. We identify the range of allowable rates such that buffer 2 satisfies QoS constraint θ_0 (see Section H) as

$$\begin{aligned} \mathcal{A}_2(\theta_0, c, \upsilon) &= \left\{a \in \mathbb{R}^+ : \theta_2 \geq \theta_0\right\} \\ &= \left\{a \in \mathbb{R}^+ : a \leq \bar{a}_2(\theta_0, c, \upsilon)\right\}, \end{aligned} \tag{3.8}$$

where the function $\bar{a}_2(\theta, c, \upsilon)$ is given by

$$\bar{a}_2(\theta, c, \upsilon) = \begin{cases} \bar{a}_1(\theta, \upsilon), & 0 < \upsilon \leq \upsilon^* \\ \bar{a}_3(\theta, c, \upsilon), & \upsilon^* < \upsilon < c. \end{cases}$$

Here, $\bar{a}_1(\theta, \upsilon)$ is as defined above and $\bar{a}_3(\theta, c, \upsilon)$ is given by the expression

$$\bar{a}_3(\theta, c, \upsilon) = \\ c + \frac{c\mu}{\lambda} \left(\frac{-1 + \sqrt{1 - \left((c-\upsilon)\frac{\theta}{\mu} - 1\right)\left((c-\upsilon)\frac{\theta}{\lambda} - 1\right)}}{(c-\upsilon)\frac{\theta}{\lambda} - 1} \right)^2$$

with parameter υ^* determined implicitly by

$$\frac{c}{\upsilon^*} - 1 = \frac{\theta \upsilon^* \mu}{\lambda\mu + (\lambda + \theta\upsilon^*)^2}.$$

Collecting these results, we obtain $\mathcal{A}(\theta_0, c, \upsilon) = \{a \in \mathbb{R}^+ : a \leq \min\{\bar{a}_1(\theta_0, c), \bar{a}_2(\theta_0, c, \upsilon)\}\}$. That is, a is admissible if and only if $a \in \mathcal{A}(\theta_0, c, \upsilon)$ with

$$a \leq \begin{cases} \bar{a}_1(\theta_0, \upsilon) & 0 < \upsilon \leq \upsilon^* \\ \min\{\bar{a}_1(\theta_0, c), \bar{a}_3(\theta_0, c, \upsilon)\} & \upsilon^* < \upsilon < c. \end{cases} \tag{3.9}$$

An intuitive explanation for the behavior of this tandem network is as follows. Build-ups in the system can occur at buffers 1 & 2. When the service rate of the second buffer is small ($\upsilon \leq \upsilon^*$), the behavior of the system is dominated solely by the action of the second queue. On the other hand, when $\upsilon > \upsilon^*$, deviations are caused by the combined behavior of the two queues with large build-ups occurring in the first queue when $\bar{a}_1(\theta_0, c) \geq \bar{a}_3(\theta_0, c, \upsilon)$, and in the second queue otherwise. The more complicated expression corresponding to this latter case follows from the fact that, being altered by the first buffer, the structure of the arrival process feeding the

second buffer is more intricate. We plot the boundary of $\mathcal{A}(\theta_0, c, \upsilon)$ with increasing $\upsilon \in (0, c)$ in Fig. 10 for various values of QoS constraint θ_0. It is clear from the figure that the achievable rate region $\mathcal{A}(\theta_0, c, \nu)$ shrinks with QoS constraint θ_0, as one would expect. The system parameters used in this example are $\lambda^{-1} = 0.65$ s, $\mu^{-1} = 0.352$ s and $c = 1$ Mbps.

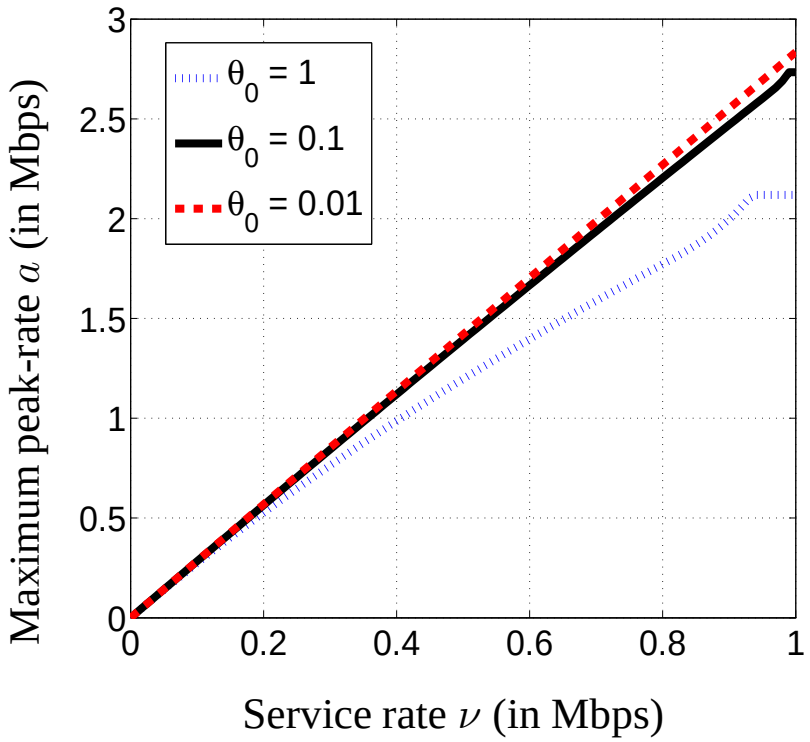

Fig. 10. Maximum supportable peak-rate of the on-off source $a(\theta_0, \nu) = \sup \mathcal{A}(\theta_0, c, \nu)$ versus ν under various values of QoS requirement θ_0 for a tandem queue. The first queue is serviced by a fixed constant service rate $c = 1$ Mbps and the second queue is serviced by a rate $\nu \in [0, c]$ Mbps. Parameter θ_0 denotes the target asymptotic exponential decay rate of the tail buffer occupancy.

Before we apply these results to compute achievable rate regions for the butterfly network, we point out that there are at least three different ways of obtaining the tail exponents of tandem queues. First, there is the transform method studied in [68–70].

Under this approach, the authors find the Laplace-Stieltjes transform of the desired distributions [69, 70]; still, these transforms are not straightforward to invert. A much stronger result, the joint distribution for tandem queues, is obtained in [68]. However, for our purposes it suffices to know the marginal content distributions of the two buffers. Furthermore, the marginal distribution for the second buffer in [68] is not expressed in its convenient reduced form.

A second approach would be to use sample-path large deviations as in the seminal paper by Chang-Zajic [47,49,50]. One can use the Lindley recursion and inverses of counting processes (Galois connections) to construct a discrete-time embedded process that would be closely related to the continuous-time process at hand. Thus, one can employ this methodology to study stationary random variables such as buffer content. However, to use these results [47,49,50], one needs to verify the general mixing conditions for the sample-path large deviations of the departure process. These conditions are highly technical and, to show they are satisfied, one requires expertise in probability theory and filtrations; this would lead to a more contrived exposition.

The third approach, which we adhere to in this chapter, can also be found in previous literature [71–73]. One can study two queues separately by first characterizing the stationary departure process from the first queue, and using it as an arrival process for the second queue. This gives us an explicit distribution for the marginal content of the second buffer in a tandem network. That is, we provide an alternate proof for the tail-exponent of the second buffer in a tandem queue (though it is a specialized result, it gives us the desired form) utilizing the Anick-Mitra-Sondhi approach [71–73] together with a characterization of the departure process by Aalto [74, 75].

D. Achievable Rate Regions

The results listed above can be employed to identify achievable rate regions for the butterfly network under consideration. With a slight abuse of notation, we let θ_i be the asymptotic decay-rate of buffer-occupancy for the queue at nodes $i \in \{1, 2, 3\}$ in Fig. 9, for a pair of peak-arrival rates (a_1, a_2) at sources A & B. We need to find the set of all two-tuples (a_1, a_2) such that the global QoS constraint θ_0 is satisfied, i.e.,

$$\mathcal{R} = \left\{(a_1, a_2) \in \mathbb{R}^+ \times \mathbb{R}^+ : \min\left\{\theta_1, \theta_2, \theta_3\right\} \geq \theta_0\right\}.$$

Network Coding As mentioned earlier, the effective service rates offered at nodes 1, 2 and 3 are $\hat{c}_1 = \min\{c_1, c_3\}$, $\hat{c}_2 = \min\{c_2, c_3\}$, and $\check{c}_3 = \max\{\hat{c}_1, \hat{c}_2\}$. This prevents undue decoding delays at the destinations. Using the notation of the previous section, we can write the achievable rate-region $\mathcal{R}$ for this system as

$$\mathcal{R}_{\text{nc}} = \mathcal{A}_1(\theta_0, \hat{c}_1) \times \mathcal{A}_1(\theta_0, \hat{c}_2),$$

where $\mathcal{A}_1$ is the set defined in (3.7).

Classic Routing For classic routing, consider two parallel buffers at node 3 with constant service-rates ν and $c_3 - \nu$, respectively. Assume that the flow from node 1 goes to the first buffer; and the flow from node 2, to the second one. Again, we emphasize that the aggregate fluid in both the buffers will be greater than or equal to the level of fluid in a single-buffer system with combined arrivals from nodes 1 & 2 and serviced at rate c_3. Using the aforementioned splitting property, the two independent flows can be decoupled and studied separately. If the shared buffer is constrained by a QoS requirement θ_0, the queues in the decoupled system are subject to the same criterion. For a fixed $0 < \nu < c_3$, there exists a unique peak-rate pair (a_1, a_2) such

that the QoS constraint θ_0 is satisfied by the system if the achievable rate region is $\mathcal{A}(\theta_0, c_1, \nu) \times \mathcal{A}(\theta_0, c_2, c_3 - \nu) = [0, a_1] \times [0, a_2]$. Accordingly, the achievable rate-region is equal to the union of the regions corresponding to all the possible values of ν. That is,

$$\mathcal{R}_{\text{cr}} = \bigcup_{0 \leq \nu \leq c_3} \mathcal{A}(\theta_0, c_1, \nu) \times \mathcal{A}(\theta_0, c_2, c_3 - \nu), \tag{3.10}$$

where $\mathcal{A}$ is the achievable rate region of (3.5).

E. Wireless Butterfly Network

In this section, we study a wireless butterfly network under a broadcast paradigm. We assume that the system operates in frequency division multiplexing (FDM) mode. We should point out that FDM operation is not necessarily optimal in terms of achievable rates or delays. More complex schemes provide future avenues of research. It is not clear though that a similar analytical framework can be used while considering alternative multiple-access schemes.

Consider the multicast scenario where two sources wish to communicate with two destinations. An additional node that acts as a relay is present to facilitate communication over the network. All the nodes have an identical power budget P, and the total spectral bandwidth available to the system is limited. The source nodes produce independent *on-off* traffic, as in our previous setting. The total available spectral bandwidth W is divided to make three non-interfering frequency bands.

Node 1 broadcasts its messages to nodes 3 & 5, and node 2 does the same to nodes 3 & 6 (see Fig. 11). Node 3 sends its messages to nodes 5 & 6 simultaneously. The packets transmitted by node 3 can be either multiplexed messages from nodes 1 & 2, or algebraic sums thereof. Again, we call these modes of operation classic routing and network coding, respectively. We note that this setup is closer

to the practical mesh networks being deployed today than it is to an information theoretic perspective seeking to identify fundamental limits of wireless systems.

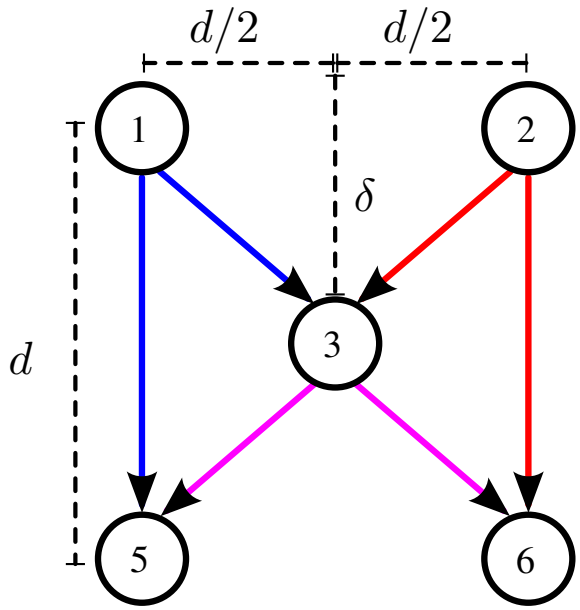

Fig. 11. Wireless butterfly network with two sources, two destinations and a relay node.

For simplicity, we assume that all the transmission links are time-invariant. We consider two cases: the additive white Gaussian noise (AWGN) case, and a scenario where the wireless channels are subject to path loss. Again, we suppose that every node is equipped with an arbitrary large buffer to store data packets that are awaiting transmission through the wireless medium. We also assume that a simple link layer acknowledgment scheme is present, so that data can be flushed out of the corresponding buffer once reception is confirmed.

1. Channel Model

For an AWGN channel, the maximum rate at which error-free data transfer is possible is given by

$$W \log_2 \left(1 + \frac{P}{N_0 W}\right) \tag{3.11}$$

where P represents the expected power of the signal, $N_0/2$ is the double-sided power spectral density of the noise process, and W is the spectral bandwidth. Recent developments in error-control coding allow operation near Shannon capacity with minimal error-rates and small delays. Therefore, the channel capacity expression of (3.11) can be viewed as an optimistic approximation of code performance. We assume that codes are designed to operate at a fixed rate, which is the constant service rate offered by the channel. In the case where all the links are AWGN limited, we allocate equal spectral bandwidth to the three nodes, and hence they become of equal capacity. We denote the constant service rate offered by each connection as $c = (W/3)\log_2(1 + 3P/(N_0 W))$, where $P/(N_0 W)$ is the observed signal-to-noise ratio (SNR).

In the second case, we assume that the received power decays exponentially in distance with exponential factor α. That is, the ratio of the transmit power to the received power is κd^{α}. Given a spectral bandwidth allocation of ηW and a distance of d meters, the capacity of the corresponding connection becomes

$$c(d, \eta) = \eta W \log_2 \left(1 + \frac{P}{\eta N_0 W \kappa d^{\alpha}}\right) \text{ bits/second.} \tag{3.12}$$

Once the spectral bandwidth allocation is completed, the capacity of each link stays fixed. We choose a code-rate to operate close to capacity, and this rate becomes the maximum allowable constant service rate for the corresponding queue.

2. AWGN Links

We begin the wireless analysis with the scenario where node 3 utilizes network coding. In this case, the data rate to be transmitted out of node 3 is the maximum of the rates from nodes 1 & 2, which is c. Therefore, there is no congestion at this node and the achievable rate region is limited by the QoS constraint at nodes 1 & 2. Under a QoS constraint θ_0, the maximum possible rate received by destinations C & D have identical functional form and are equal to a_1 from source A and a_2 from source B, where

$$a_i = c + \frac{c\mu_i}{\lambda_i + c\theta_0}, \quad i \in \{1, 2\}.$$

That is, the achievable rate region for the source rate-pair (a_1, a_2) is $\mathcal{R}_{\text{nc}} = \mathcal{A}_1^{(1)}(\theta_0, c) \times \mathcal{A}_1^{(2)}(\theta_0, c)$.

Next, we consider the situation where node 3 simply forwards packets from nodes 1 & 2 to the destinations. In this case, source rate-pairs (a_1, a_2) are also limited by the congestion at node 3; and, for all $\nu \in (0, c)$, we have $\mathcal{A}(\theta_0, c, \nu) \subset \mathcal{A}_1(\theta_0, c)$. The total achievable rate region for classic routing thus becomes

$$\mathcal{R}_{\text{cr}} = \bigcup_{0 \le \nu \le c} \mathcal{A}^{(1)}(\theta_0, c, \nu) \times \mathcal{A}^{(2)}(\theta_0, c, c - \nu).$$

The results for a specific value of θ_0 are shown in Fig. 12. The system parameters selected for this numerical study appear in Table III. Additionally, we chose an expected received power equal to 100 mW.

3. Links with Path Loss

To illustrate the effects of path attenuation, we consider an example where the sources and destinations are located on the vertices of a perfect square of side-length d. The relay node lies on the perpendicular bisector of the edges connecting the two sources

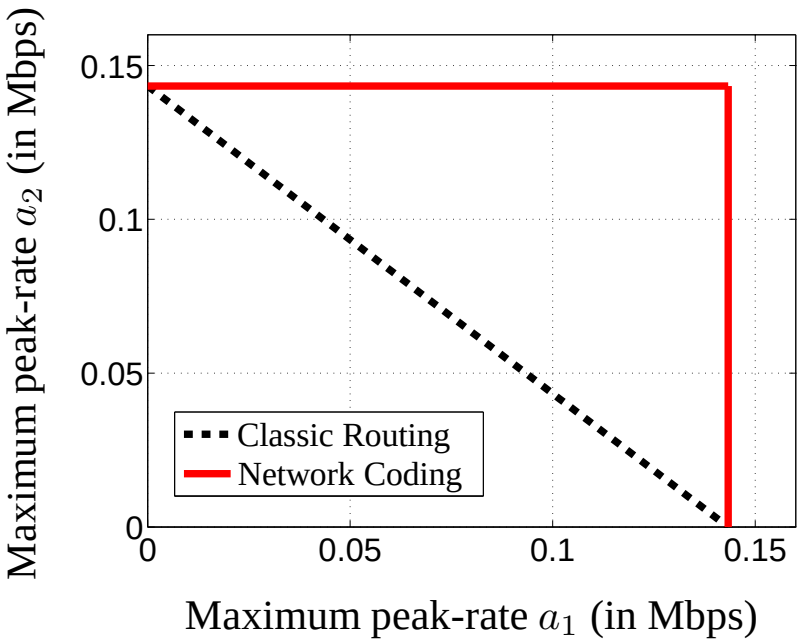

Fig. 12. Boundaries of achievable peak-rate regions for on-off sources when (a) classical routing or (b) network coding is employed at intermediate node 3 for the QoS constrained communication over butterfly network of Fig. 11, where each link is additive white Gaussian noise limited. The asymptotic exponential decay-rate of the buffer-occupancy is bounded below by $\theta_0 = 0.1$.

at a distance δ from the top of the square. The distance from the two sources to the relay node being identical, we assume that a fraction $\eta/2$ of the total bandwidth is allocated to each source; and the remaining $(1-\eta)W$, to the relay node.

To maximize the gains of network coding, we need to make the link capacities identical. The capacity of the link between a source and the relay is $c\,(d_{13}, \eta/2)$, where $d_{13} = d_{23} = \sqrt{(d/2)^2 + \delta^2}$. Similarly, the capacity of the link between the relay and a destination can be written as $c\,(d_{35}, (1-\eta))$, where the distance from the relay to a destination is $d_{35} = d_{36} = \sqrt{(d/2)^2 + (d-\delta)^2}$. Since the service-rate available to the source is limited by the minimum of the direct-link and relay-link

capacities, we allocate bandwidth parameter η in the following fashion:

$$\begin{aligned}\min\left\{c\left(d_{13},\frac{\eta}{2}\right),c\left(d,\frac{\eta}{2}\right)\right\} &= c\left(\max\{d_{13},d\},\frac{\eta}{2}\right)\\ &= c\left(d_{35},(1-\eta)\right).\end{aligned} \tag{3.13}$$

We denote this optimal bandwidth allocation parameter by η^*, which is evaluated numerically. We would like to point out that this η^* is unique, which follows from the monotonicity of $c(d,\eta)$ in η (see (3.12)).

Proposition E.1. *The optimal bandwidth allocation parameter η^* that satisfies* (3.13) *is unique.*

Proof. It is easy to see that $c(d,\eta)$ is continuous and even differentiable in the range $(0,1]$. We will quickly show that $c(d,\eta)$ is strictly increasing in η. Then, it will follow that $c(d,(1-\eta))$ is continuous and decreasing in η and hence η^* is unique since $c(d_1,0)=0<c(d_2,1)$ for any finite d_1,d_2. To that end, it suffices to show that $\frac{\partial c(d,\eta)}{\partial\eta}\geq 0$ for all $\eta\in[0,1]$. From (3.12), it is straightforward to verify that $\frac{\partial c(d,\eta)}{\partial\eta}$ is decreasing in η; hence we only need to check that $\frac{\partial c(d,\eta)}{\partial\eta}|_{\eta=1}=\frac{W}{\ln 2}\left(\ln(1+k_2)-\frac{k_2}{1+k_2}\right)$ is greater than zero. Here, we are denoting $\frac{P}{N_0 W\kappa d^\alpha}$ by k_2. Clearly, if $f(k_2)=\ln(1+k_2)-\frac{k_2}{1+k_2}\geq 0$ for all $k_2\geq 0$, we are done. This is easy to verify since $f(0)=0$ and $f'(k_2)=\frac{k_2}{(1+k_2)^2}\geq 0$. □

For the classic routing case, each source broadcasts its packets to the relay node where information is stored. The relay then forwards the received messages to the destinations using a first-come-first-serve scheduling policy. The links from the sources to the relay node have identical capacity $c(d_{13},\eta/2)=c(d_{23},\eta/2)$. Similarly, the link from the relay node to the destinations have capacity $c(d_{35},(1-\eta))$. Using the same queueing performance analysis as before and for a fixed δ, we can obtain the achievable rate-region under QoS constraint θ_0 for network coding and classic

routing (for different values of η in the latter case). The two regions are shown in Fig. 13 for a transmit power of $P = 40$ W and the system parameters of Table III. In this example, we have taken $d = 15$ m and $\alpha = 1.8$ (typical values of α range from 1.6 to 1.8 for line-of-sight communication in buildings [76, Page 139, Table 4.2]). The region enclosed by the thick solid line represents the achievable rate region achieved by network coding. The thin dashed lines characterize the regions corresponding to classic routing for different values of η. Clearly, classic routing outperforms network coding in this case. This is a scenario where the cost of establishing a network topology suitable for network coding exceeds the benefits of packet combining.

F. Queueing Results

Below, we list and derive queueing results related to the analysis of the tandem network introduced in Section C. Recall that the arrival process feeding the first queue is a Markov-modulated fluid process with *on*-rate a. This queue is serviced at a constant rate c, whereas the second queue is served at a rate υ. When $a \leq c$, the first buffer in the system remains empty at all times, and the analysis of the tandem network degenerates into a single-queue scenario. We therefore assume that $a > c$, which is the more interesting case. For this situation, the following theorem enables us to obtain the tail-asymptotics of the first buffer in a tandem network. The corresponding rate-region is governed by (3.4) and (3.6).

Theorem F.1 (Mitra [45, 71, 77])**.** *Let $\ell_1(t)$ be the amount of fluid in an arbitrary large reservoir being fed by an* on-off *source satisfying Assumption B.1, and serviced at a constant rate c. The* off *and* on *times of the source are exponentially distributed with means λ^{-1} and μ^{-1}, respectively. If the constant arrival-rate of the fluid in the*

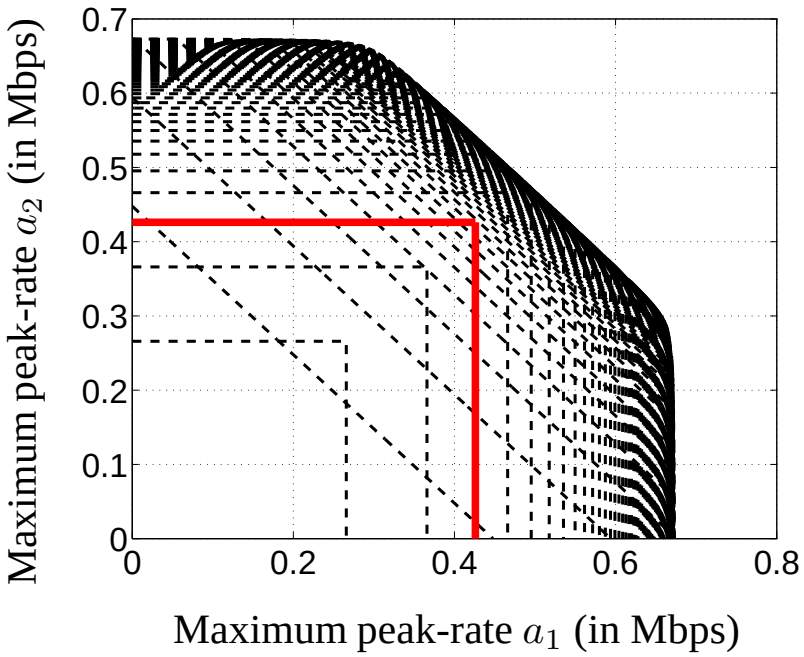

Fig. 13. Boundaries of achievable peak-rate regions for on-off sources when (a) classical routing (denoted by dashed lines for values of fraction $\eta \in [0.01, 0.99]$) and (b) network coding (denoted by the solid line) are employed at the intermediate node 3 for the QoS constrained communication over butterfly network of Fig. 11 where each link is limited by path loss. The asymptotic exponential decay-rate of the buffer occupancy is bounded below by $\theta_0 = 0.1$ and the relay node is at distance $\delta = 9$m.

on *state is such that*

$$\frac{a\lambda}{\lambda+\mu} < c < a,$$

then the limit $\lim_{t\to\infty} \Pr\{\ell_1(t) > x\}$ *exists for all* $x \geq 0$. *The corresponding asymptotic decay-rate of buffer occupancy can be identified by looking at the largest negative eigenvalue* ζ *that satisfies the matrix equation* $\zeta D\phi = Q\phi$. *Here,* Q *is the generator matrix for the modulating Markov chain and* D *is the drift matrix corresponding to*

the arrival process. More precisely, we have

$$\theta_1 = -\lim_{x\to\infty} \frac{\ln \Pr\{\ell_1 > x\}}{x} = \frac{\mu}{a-c} - \frac{\lambda}{c}$$

where ℓ_1 is a random variable whose distribution coincides with the equilibrium distribution of the queue.

This is a standard result and, as such, we state it without proof. For alternate treatments of this theorem and other pertinent queueing results on the decay-rate of buffer-occupancy under various conditions, see [3, 4, 45–47, 49, 50, 53, 66, 72].

1. Departure Process of a Fluid Buffer

We next consider the case where there are two queues in tandem and the departure process of the first queue serves as the input to the second buffer. The departure process at the output of the first buffer is characterized using a theorem first proved by Aalto [74, 75].

Theorem F.2 (Aalto)**.** *For the fluid queue described in Theorem F.1, the departure process can be viewed as an* on-off *source where packets are emitted at a constant rate during an* on *period, and the source is idle otherwise. The* off *times are exponentially distributed with mean λ^{-1}; while the* on *times have the same distribution as the duration of a busy period in an M/M/1 queue with arrival rate $(1-c/a)\lambda$ and service rate $c\mu/a$. Furthermore, the departure rate is c when the queue is non-empty. Mathematically, this departure process is modulated by a countable state birth-death*

process $\{K(t) : t \geq 0\}$ *with given transition rates*

$$\lambda_{0,1} = \lambda, \tag{3.14}$$

$$\lambda_{n-1,n} = \left(\frac{a-c}{a}\right)\lambda, \quad n = 2, 3, \ldots \tag{3.15}$$

$$\mu_{n,n-1} = \frac{c\mu}{a}, \quad n \in \mathbb{N}. \tag{3.16}$$

The modulating birth-death process is depicted in Fig. 14. The departure process is off*, when* $K(t) = 0$ *and it is* on *otherwise. That is, the departure rate is given by*

$$c(K(t)) = \begin{cases} 0, & K(t) = 0, \\ c, & K(t) \in \mathbb{N}. \end{cases} \tag{3.17}$$

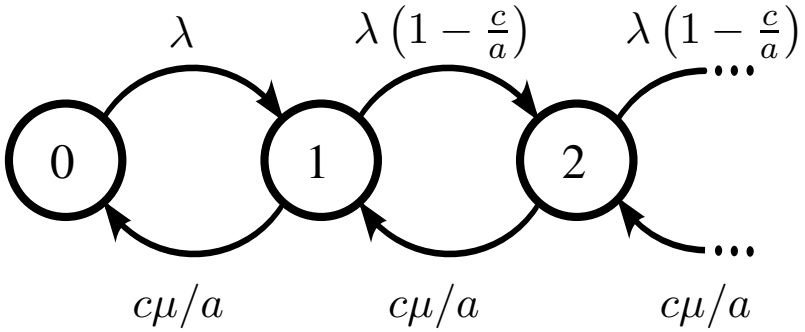

Fig. 14. Graphical representation of the modulating birth-death process.

Having characterized the departure process of the queue, we next present an expression for the equilibrium distribution of $K(t)$. Let

$$p_n = \lim_{t\to\infty} \Pr(K(t) = n), \quad n = 0, 1, 2, \ldots$$

and define $\rho = \left(\frac{a-c}{c}\right)\frac{\lambda}{\mu}$. Then, the stationary distribution of K is given by [78]

$$p_0 = \frac{(a-c)(1-\rho)}{a-c+c\rho},$$
$$p_n = \left(\frac{a}{a-c}\right)\rho^n p_0, \quad n = 1, 2, \ldots$$

2. Distribution of the Tandem Queue

For the tandem network described in Section C, we wish to find the equilibrium distribution of the second buffer. Note that we already have similar results for the first queue based on Theorem F.1. We also have knowledge of the departure process from the first queue, as afforded by Theorem F.2. The properties of the second queue can therefore be studied based on this departure process alone.

Consider a buffer that is being fed by a Markov modulated *on-off* source generating fluid at a rate $c(K(t))$, where $c(K(t))$ is the departure function of (3.17) and $K(t)$ is the modulating Markov process described in Theorem F.2. The arrival rate in the queue is c when $K(t) \geq 1$, and it is zero otherwise. Fluid is removed from the queue at constant rate υ, provided that it is non-empty. Paralleling the approach of Virtamo and Norros [73], we derive necessary and sufficient conditions for a non-trivial stationary probability distribution to exist. Then, we find the spectrum of the operator that governs the equilibrium distribution of the buffer. This allows us to present an explicit expression for the distribution of the queue.

Clearly, if $\upsilon \geq c$, there is no congestion in this buffer. We therefore examine the case where $\upsilon < c$. Writing the stochastic evolution equation for the buffer of interest, we get

$$\frac{d}{dt}\ell_2(t) = (c-\upsilon)\mathbf{1}_{\{K(t)\neq 0\}} - \upsilon\mathbf{1}_{\{K(t)=0,\ell_2(t)>0\}}.$$

It can be shown that the stochastic process governing the evolution of this buffer,

$\{\ell_2(t) : t \geq 0\}$, is positive recurrent if and only if [36, 66]

$$\left(\frac{a-\upsilon}{\upsilon}\right)\frac{\lambda}{\mu} < 1.$$

That is, there exists a stationary probability distribution for the buffer provided that the stability condition mentioned above is satisfied. Applying standard queueing arguments, we can derive the Chapman-Kolmogorov equations for the probability distribution of stationary buffer-occupancy ℓ_2 as follows. We use $\pi(x, n)$ to denote the probability that ℓ_2 exceeds x while the underlying birth-death process is in state n, i.e.,

$$\begin{aligned}\pi(x,n) &= \lim_{t\to\infty} \Pr(\ell_2(t) > x, K(t) = n) \\ &= \Pr(\ell_2 > x, K = n).\end{aligned}$$

The corresponding Chapman-Kolmogorov equations become

$$\begin{aligned}&\left(c\mathbf{1}_{\{n\geq 1\}} - \upsilon\right)\frac{d}{dx}\pi(x,n) \\ &= \lambda_{n-1,n}\mathbf{1}_{\{n\geq 1\}}\pi(x,n-1) + \mu_{n+1,n}\pi(x,n+1) \\ &- \left(\lambda_{n,n+1} + \mu_{n,n-1}\mathbf{1}_{\{n\geq 1\}}\right)\pi(x,n).\end{aligned} \tag{3.18}$$

for $n = 0, 1, 2, \ldots$ The constants $\{\lambda_{n-1,n}, \mu_{n,n-1} : n \in \mathbb{Z}^+\}$ are the transition rates of the modulating process $K(t)$ defined in (3.14)-(3.16). Additionally, we employ the convention that $\mu_{0,-1} = \lambda_{-1,0} = 0$. With some abuse of notation, we define the sequence $\pi(x)$ as

$$\pi(x) = \left\{\pi(x,n) : n \in \mathbb{Z}^+\right\}.$$

Let $\mathbf{H}$ denote the Hilbert space of square summable sequences indexed by $\mathbb{Z}^+$, and let $\mathcal{B}(\mathbf{H})$ be the space of bounded linear operators from $\mathbf{H}$ to itself. For sequence

$\mathbf{h} \in \mathbf{H}$, we define operators D and Q in $\mathcal{B}(\mathbf{H})$ by

$$(D\mathbf{h})_n = \left(c\mathbf{1}_{\{n\geq 1\}} - \upsilon\right) h_n, \tag{3.19}$$

$$\begin{aligned}(Q\mathbf{h})_n = {} & \lambda_{n-1,n}\mathbf{1}_{\{n\geq 1\}}h_{n-1} - \lambda_{n,n+1}h_n \\ & + \frac{c\mu}{a}\left(h_{n+1} - h_n\mathbf{1}_{\{n\geq 1\}}\right),\end{aligned} \tag{3.20}$$

where $\{\lambda_{n-1,n}\}$ are defined in (3.14)-(3.16). It is straightforward to check continuity of these two operators. In particular, for any $\mathbf{h} \in \mathbf{H}$, we have

$$\begin{aligned}\|D(\mathbf{h})\| &\leq \max\{\upsilon, c-\upsilon\}\|\mathbf{h}\|, \\ \|Q(\mathbf{h})\| &\leq 3\left(\lambda + \frac{c\mu}{a}\right)\|\mathbf{h}\|.\end{aligned}$$

For the aforementioned system, we can rewrite (3.18) in a compact form in terms of the sequence $\pi(x)$ and the bounded linear transformations Q and D,

$$D\frac{d}{dx}\pi(x) = Q\pi(x). \tag{3.21}$$

The transformation D can be expressed in terms of the identity operator I and the standard projection operator $\mathbf{e}_0^*$. For any $\mathbf{h} \in \mathbf{H}$, we have $I\mathbf{h} = \mathbf{h}$ and $\mathbf{e}_0^*\mathbf{h} = h_0$. We can then write

$$D = (c-\upsilon)\tilde{D} = (c-\upsilon)\left(I - \frac{c}{c-\upsilon}\mathbf{e}_0\mathbf{e}_0^*\right).$$

The linear transformation Q can be described in matrix form as

$$Q = \frac{c\mu}{a}E\tilde{Q}E^{-1},$$

where we have defined the operators $E = \text{diag}\left(\sqrt{\frac{a-c}{a}}, \rho^{1/2}, \rho, \rho^{3/2}, \ldots\right)$ and

$$\tilde{Q} = \begin{bmatrix} \frac{-a\rho}{a-c} & \sqrt{\frac{a\rho}{a-c}} & 0 & 0 & \cdots \\ \sqrt{\frac{a\rho}{a-c}} & -(1+\rho) & \sqrt{\rho} & 0 & \cdots \\ 0 & \sqrt{\rho} & -(1+\rho) & \sqrt{\rho} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

We emphasize that $\tilde{Q}$ is a symmetric matrix, which makes it very convenient for the spectral analysis carried out below. Note also that E, E^{-1} and D are diagonal and hence commutative operators.

It is easy to see that $\tilde{D}$ and $\tilde{Q}$ belong to $\mathcal{B}(\mathbf{H})$, and that $\tilde{D}$ is invertible. It follows that $\tilde{D}^{-1}\tilde{Q} \in \mathcal{B}(\mathbf{H})$. Since $\tilde{D}^{-1} = \left[I - \frac{c}{\upsilon}\mathbf{e}_0\mathbf{e}_0^*\right]$, we can represent $\tilde{D}^{-1}\tilde{Q}$ using a countable state matrix as

$$\begin{bmatrix} \frac{-a\rho}{a-c}\left(\frac{\upsilon-c}{\upsilon}\right) & \sqrt{\frac{a\rho}{a-c}}\left(\frac{\upsilon-c}{\upsilon}\right) & 0 & 0 & \cdots \\ \sqrt{\frac{a\rho}{a-c}} & -(1+\rho) & \sqrt{\rho} & 0 & \cdots \\ 0 & \sqrt{\rho} & -(1+\rho) & \sqrt{\rho} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{3.22}$$

It should also be apparent from the Chapman-Kolmogorov equation of (3.21) that

$$\pi(x) = E \exp\left(x\frac{c\mu}{a(c-\upsilon)}\tilde{D}^{-1}\tilde{Q}\right) E^{-1}\pi(0). \tag{3.23}$$

To evaluate $\pi(x)$, we need to identify boundary condition $\pi(0)$ and find a spectral representation for the operator

$$\exp\left(x\frac{c\mu}{a(c-\upsilon)}\tilde{D}^{-1}\tilde{Q}\right). \tag{3.24}$$

The equilibrium distribution $\pi(0)$ can be obtained from the natural boundary conditions on the buffer occupancy for a stable system; this is accomplished later. Rather,

we begin by deriving an expression for (3.24) through a two-step approach: first we compute the spectrum of $\tilde{D}^{-1}\tilde{Q}$, and then we obtain an equivalent representation for I in terms of the associated eigenfunctions.

a. Spectrum of Bounded Operator

In general, finding the spectrum of a non-compact bounded linear operator is a difficult task. However, in the present case, the relevant operator is a compact perturbation of a standard shift operator. It is well-known in the literature that the continuous spectrum of a shift operator remains invariant under compact perturbations [79, 80]. We present relevant results formally in the following theorem, which will be used to obtain a spectral representation for the operator of (3.24).

Let us denote the spectrum of an operator A by

$$\sigma(A) = \{\zeta \in \mathbb{C} : (A - \zeta I) \text{ is not invertible in } \mathcal{B}(\mathbf{H})\}.$$

This spectrum is composed of two parts. First, there is the discrete spectrum of A, also called the set of eigenvalues of A, which is defined as $\sigma_d(A) = \{\zeta \in \mathbb{C} : (A - \zeta I) \text{ is not injective}\}$. That is, ζ is an eigenvalue of A if and only if there exists a sequence $\mathbf{h} \in \mathbf{H}$ such that $(A - \zeta I)\mathbf{h} = \mathbf{0}$. On the other hand, the values of $\zeta \in \sigma(A)$ for which the operator $(A - \zeta I)$ is injective but not surjective belong to the continuous spectrum of A. If there exists $\mathbf{h} \in \mathbf{H}$ such that $(A - \zeta I)^{-1}\mathbf{h} \notin \mathbf{H}$, yet the range of operator $(A - \zeta I)^{-1}$ is dense in $\mathbf{H}$, then $\zeta \in \sigma_c(A)$.

Theorem F.3. *Let operators $D, Q \in \mathcal{B}(\mathbf{H})$ be as defined in* (3.19)-(3.20)*. Then, the*

continuous spectrum of $D^{-1}Q$ *is*

$$\sigma_c(D^{-1}Q) = \left[-\frac{c\mu}{a(c-\upsilon)}\left(1+\sqrt{\left(\frac{a-c}{c}\right)\frac{\lambda}{\mu}}\right)^2, \right. \\ \left. -\frac{c\mu}{a(c-\upsilon)}\left(1-\sqrt{\left(\frac{a-c}{c}\right)\frac{\lambda}{\mu}}\right)^2 \right] \tag{3.25}$$

and $D^{-1}Q$ *has a discrete eigenvalue* $\zeta_0 = 0$*. Furthermore, if*

$$\frac{a/c-1}{a/\upsilon-1} < \sqrt{\left(\frac{a-c}{c}\right)\frac{\lambda}{\mu}},$$

then there is an additional eigenvalue located at $\zeta_1 = -\frac{\mu}{a-\upsilon} + \frac{\lambda}{\upsilon}$*. In this latter case, we have* $\zeta_1 > \sup \sigma_c(D^{-1}Q)$*.*

Proof. Let $\mathbf{w} \in l^\infty(\mathbb{Z}^+)$ be such that, for some ζ, we have

$$\zeta D\mathbf{w} = Q\mathbf{w}. \tag{3.26}$$

Then, ζ belongs to the spectrum of the linear operator $D^{-1}Q$. Substituting $\xi = \frac{a(c-\upsilon)}{c\mu}\zeta$ and $\tilde{\mathbf{w}} = E^{-1}\mathbf{w}$, we obtain

$$\xi\tilde{D}\tilde{\mathbf{w}} = \tilde{Q}\tilde{\mathbf{w}}. \tag{3.27}$$

There is a one-to-one correspondence between the eigenvalues and eigenfunctions of (3.26) and (3.27). Therefore, it suffices to show that spectrum of $\tilde{D}^{-1}\tilde{Q}$ has a continuous part $[-(1+\sqrt{\rho})^2, -(1-\sqrt{\rho})^2]$ and an eigenvalue $\xi_0 = 0$. In addition, we need to show that operator $\tilde{D}^{-1}\tilde{Q}$ also has an additional eigenvalue $\xi_1 = -(1-c')(1-\frac{\rho}{c'})$ when $\sqrt{\rho} > c'$. Here we have implicitly defined $c' = \frac{a/c-1}{a/\upsilon-1}$. Now, we show that this is indeed true. We choose to solve (3.27) due to its symmetric structure

and greater simplicity. From (3.22), it follows that

$$\begin{aligned}
\xi\tilde{w}_0 &= -\frac{a\rho}{a-c}\left(\frac{\upsilon-c}{\upsilon}\right)\tilde{w}_0 + \sqrt{\frac{a\rho}{a-c}}\left(\frac{\upsilon-c}{\upsilon}\right)\tilde{w}_1,\\
\xi\tilde{w}_1 &= \sqrt{\frac{a\rho}{a-c}}\tilde{w}_0 - (1+\rho)\tilde{w}_1 + \sqrt{\rho}\tilde{w}_2,\\
\xi\tilde{w}_n &= \sqrt{\rho}\tilde{w}_{n-1} - (1+\rho)\tilde{w}_n + \sqrt{\rho}\tilde{w}_{n+1}, \quad n \geq 2.
\end{aligned}$$

Taking the z-transform, $\tilde{W}(z) = \sum_{n=0}^{\infty}\tilde{w}_n z^n$, we get

$$\begin{aligned}
\tilde{W}(z) = &\left(1 - \sqrt{\frac{a}{a-c}}\right)\tilde{w}_0\\
&+ \sqrt{\frac{a}{a-c}}\left[\frac{1 - \left(1 + \xi\frac{c/a - c/\upsilon}{1-c/\upsilon}\right)\frac{z}{\sqrt{\rho}}}{1 - \frac{1+\rho+\xi}{\sqrt{\rho}}z + z^2}\right]\tilde{w}_0.
\end{aligned} \tag{3.28}$$

We define $\gamma_0(\xi), \gamma_1(\xi)$ to be the roots of the characteristic polynomial $\gamma^2 - \frac{(1+\rho+\xi)}{\sqrt{\rho}}\gamma + 1$.

To solve for $\tilde{\mathbf{w}}$ from its z-transform, we break the problem into two separate cases. First, assume that $|\gamma_0(\xi)| \neq 1$. Then $\gamma_0(\xi), \gamma_1(\xi)$ are two different roots, and the z-transform can be written as

$$\begin{aligned}
\tilde{W}(z) = &\left(1 - \sqrt{\frac{a}{a-c}}\right)\tilde{w}_0\\
&+ \sqrt{\frac{a}{a-c}}\left(\frac{1-\alpha(\xi)}{1 - z\gamma_0(\xi)} + \frac{\alpha(\xi)}{1 - z\gamma_1(\xi)}\right)\tilde{w}_0,
\end{aligned} \tag{3.29}$$

where we have implicitly defined

$$\alpha(\xi) = \frac{1}{2}\left(1 + \frac{\rho - 1 - \frac{1+c'}{1-c'}\xi}{\sqrt{(1+\rho+\xi)^2 - 4\rho}}\right).$$

From this decomposition, we gather that

$$\tilde{w}_n(\xi) = \sqrt{\frac{a}{a-c}}\Big((1-\alpha(\xi))\gamma_0(\xi)^n + \alpha(\xi)\gamma_1(\xi)^n\Big)\tilde{w}_0,$$

where $n \in \mathbb{N}$. Note that $\tilde{\mathbf{w}} \in \mathbf{H}$ only when $\alpha(\xi)$ is 0 or 1 since $\gamma_0(\xi)\gamma_1(\xi) = 1$ (see [73]). The corresponding eigenvalues are $\xi_0 = 0$ and $\xi_1 = -(1-c')\left(1 - \frac{\rho}{c'}\right)$.

Since $\gamma_0(0) = \sqrt{\rho}$, zero is always an eigenvalue of (3.27). However, $\gamma_1(\xi_1) = \frac{c'}{\sqrt{\rho}}$ which implies that ξ_1 is an eigenvalue only when $c' < \sqrt{\rho}$.

On the other hand, we claim that ξ belongs to the continuous spectrum if and only if $|\gamma_1(\xi)| = 1$; this is equivalent to $\xi \in [-(1+\sqrt{\rho})^2, -(1-\sqrt{\rho})^2]$. To prove the claim, we note that $\tilde{D}^{-1}\tilde{Q}$ can be equivalently expressed in terms of the right shift operator S, the left shift operator T, a compact perturbation K and the identity operator I. In particular, we can write

$$\tilde{D}^{-1}\tilde{Q} = \sqrt{\rho}\left(S + T - \frac{1+\rho}{\sqrt{\rho}}I + K\right) \tag{3.30}$$

where $S = \sum_{n\in\mathbb{Z}^+} \mathbf{e}_{n+1}\mathbf{e}_n^*$, $T = \sum_{n\in\mathbb{Z}^+} \mathbf{e}_n\mathbf{e}_{n+1}^*$ and

$$\begin{aligned} K = &\left(\frac{1+\rho}{\sqrt{\rho}} - \sqrt{\rho}\left(\frac{av-ac}{av-cv}\right)\right)\mathbf{e}_0\mathbf{e}_0^* \\ &- \left(1 - \sqrt{\frac{a}{a-c}}\right)\mathbf{e}_1\mathbf{e}_0^* - \left(1 - \sqrt{\frac{a}{a-c}}\left(1-\frac{c}{v}\right)\right)\mathbf{e}_0\mathbf{e}_1^*. \end{aligned}$$

It is immediate that $S+T$ is self-adjoint with real continuous spectrum $\sigma_c(S+T)$, and $\|S\| = \|T\| = 1$. Also, it is a well-known result by Weyl that the continuous spectrum (plus limit points of point spectrum if any) of a self-adjoint operator remains unchanged under compact perturbations [79,80]. This fact, along with (3.30), gives us

$$\begin{aligned} \sigma_c\left(\tilde{D}^{-1}\tilde{Q}\right) &= \sqrt{\rho}\sigma_c\left(S + T - \frac{1+\rho}{\sqrt{\rho}}I\right) \\ &= \sqrt{\rho}\sigma_c(S+T) - (1+\rho). \end{aligned}$$

Next, we show that $\sigma_c(S+T) = [-2,2]$. Since $\|S+T\| \le 2$, it is clear that $\sigma_c(S+T) \subset [-2,2]$ [81, Proposition 7.19]. To establish set equality, we use Proposition 7.39 in [81] which states that

$$\sigma_c(S+T) = \left\{\lambda : \inf_{\|\mathbf{h}\|=1} \|(S+T-\lambda I)\mathbf{h}\| = 0\right\}.$$

We choose the following sequence $\{\mathbf{h}^{(N)}(\phi) : N \in \mathbb{Z}^+\}$ parametrized by $\phi \in [0,1]$,

$$\mathbf{h}^{(N)}(\phi) = \begin{cases} \frac{1}{\sqrt{N}} e^{j2\pi\phi n} & n \in \{0,1,\ldots,N-1\} \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, $\|\mathbf{h}^{(N)}\| = 1$. Furthermore, taking $\lambda = 2\cos(2\pi\phi)$, we get

$$\left\|(S+T-\lambda I)\mathbf{h}^{(N)}\right\| = \sqrt{\frac{3}{N}}.$$

This shows that $[-2,2] \subset \sigma_c(S+T)$. As a consequence, we have

$$\sigma_c\left(\tilde{D}^{-1}\tilde{Q}\right) = \left[-\left(1+\sqrt{\rho}\right)^2, -\left(1-\sqrt{\rho}\right)^2\right]$$

and hence (3.25) follows. We now show that $\zeta_1 > \sup \sigma_c\left(D^{-1}Q\right)$. It suffices to show that $\xi_1 > \sup \sigma_c\left(\tilde{D}^{-1}\tilde{Q}\right)$. Substituting the expressions for $\xi_1 = -(1-c')(1-\frac{\rho}{c'})$ and $\sup \sigma_c\left(\tilde{D}^{-1}\tilde{Q}\right) = -(1-\sqrt{\rho})^2$, and canceling common terms on both sides, we need to show that

$$\frac{c'}{\sqrt{\rho}} + \frac{\sqrt{\rho}}{c'} > 2.$$

This holds because $\gamma_1(\xi_1) = \frac{c'}{\sqrt{\rho}} < 1$. This completes the proof. □

b. Spectral Representation of Identity

For further analysis, we introduce variable y and constants p and q in the following way:

$$y = y(\xi) = \frac{1+\rho+\xi}{2\sqrt{\rho}}, \qquad p = \frac{2}{1-c'}, \qquad q = -\frac{c'+\rho}{(1-c')\sqrt{\rho}}.$$

We emphasize that there is a one-to-one correspondence between y and ξ; we can therefore express y as a function of ξ, and ξ as a function of y unambiguously. Hence, we use $y(\xi)$ and $\xi(y)$ interchangeably depending on the context. We define

$y_0 = y(\xi_0)$ and $y_1 = y(\xi_1)$ for eigenvalues ξ_0 and ξ_1. In addition, we assume without loss of generality that $\tilde{w}_0 = 1$ for eigen-sequence $\tilde{\mathbf{w}}(y(\xi))$ corresponding to $\xi \in \sigma(\tilde{D}^{-1}\tilde{Q})$. For convenience, we make the dependence of $\tilde{\mathbf{w}}$ on y (and hence ξ) explicit hereafter. For eigenvalues ξ_0 and ξ_1, we gather from proof of Theorem F.3 that the corresponding eigen-sequences are

$$\tilde{\mathbf{w}}(y_0) = \sqrt{\frac{a}{a-c}}\left(\sqrt{\frac{a-c}{a}}, \sqrt{\rho}, \rho, \dots\right)$$
$$\tilde{\mathbf{w}}(y_1) = \sqrt{\frac{a}{a-c}}\left(\sqrt{\frac{a-c}{a}}, \frac{c'}{\sqrt{\rho}}, \left(\frac{c'}{\sqrt{\rho}}\right)^2, \dots\right).$$

This leads to their z-transforms being

$$\tilde{W}(y_0, z) = \left(1 - \sqrt{\frac{a}{a-c}}\right) + \sqrt{\frac{a}{a-c}}\left(\frac{1}{1 - z\sqrt{\rho}}\right)$$
$$\tilde{W}(y_1, z) = \left(1 - \sqrt{\frac{a}{a-c}}\right) + \sqrt{\frac{a}{a-c}}\left(\frac{1}{1 - zc'/\sqrt{\rho}}\right).$$

We can also find $\tilde{W}(y_i, z)$ by substituting $\alpha(\xi_i) = i \in \{0, 1\}$ in (3.29). Furthermore, we can write the z-transform of the eigen-sequence corresponding to $\xi \in \sigma_c(\tilde{D}^{-1}\tilde{Q})$ by substituting y, p and q in (3.28) to obtain

$$\begin{aligned}\tilde{W}(y, z) &= \sum_{k \in \mathbb{Z}^+} \tilde{w}_k(y) z^k \\ &= \left(1 - \sqrt{\frac{a}{a-c}}\right) + \sqrt{\frac{a}{a-c}}\left(\frac{1 - (py+q)z}{1 - 2yz + z^2}\right).\end{aligned}$$

We notice that $\tilde{W}(y, \sqrt{\rho}) = 1 + \sqrt{\frac{a}{a-c}}\frac{c'}{(1-c')}$ is independent of y for $y \neq y_0$, and $\tilde{W}(y_0, \sqrt{\rho}) = 1 + \sqrt{\frac{a}{a-c}}\frac{\rho}{(1-\rho)}$.

For any eigen-sequence $\tilde{\mathbf{w}}$ and corresponding $\xi \in \sigma(\tilde{D}^{-1}\tilde{Q})$, we have $\tilde{Q}\tilde{\mathbf{w}} = \xi\tilde{D}\tilde{\mathbf{w}}$, which in turn implies

$$\tilde{\mathbf{w}}^*\tilde{Q} = \xi\tilde{\mathbf{w}}^*\tilde{D}. \tag{3.31}$$

Moreover, for $i \in \{0, 1\}$ and $y \in [-1, 1]$, we have $\xi(y_i) \neq \xi(y)$. Note that $\tilde{\mathbf{w}}(y) \notin \mathbf{H}$;

however, $\tilde{\mathbf{w}}(y) \in l^{\infty}(\mathbb{Z}^{+})$ and $\tilde{\mathbf{w}}(y_i) \in l^{1}(\mathbb{Z}^{+})$ (see proof of Theorem F.3). Therefore, the usual inner product on $\mathbf{H}$ (sum of point-wise product of sequences) is well defined for $\tilde{\mathbf{w}}(y)$ and $\tilde{\mathbf{w}}(y_i)$. Hence, using (3.31) and the eigen relationship, we get $\tilde{\mathbf{w}}^{*}(y_i)\xi(y)\tilde{D}\tilde{\mathbf{w}}(y) = \xi(y_i)\tilde{\mathbf{w}}^{*}(y_i)\tilde{D}\tilde{\mathbf{w}}(y)$, which in turn yields

$$(\xi(y) - \xi(y_i))\langle \tilde{D}\tilde{\mathbf{w}}(y), \tilde{\mathbf{w}}(y_i)\rangle = 0.$$

Thus, we obtain

$$\langle \tilde{D}\tilde{\mathbf{w}}(y), \tilde{\mathbf{w}}(y_i)\rangle = 0. \tag{3.32}$$

It is equally straightforward to see that

$$\langle \tilde{D}\tilde{\mathbf{w}}(y_1), \tilde{\mathbf{w}}(y_0)\rangle = 0. \tag{3.33}$$

Below, we use these orthogonal properties to derive the desired spectral representation of the identity operator.

Theorem F.4. *Let s_0, s_1 denote the weights corresponding to eigenvalues ξ_0 and ξ_1, and let $s(y)$ be the weighting function associated with the continuous spectrum of $\tilde{D}^{-1}\tilde{Q}$. Define these quantities as follows,*

$$\begin{aligned}
s_0 &= (\tilde{\mathbf{w}}(y_0)^{*}\tilde{D}\tilde{\mathbf{w}}(y_0))^{-1} = -\left(\frac{a-c}{a}\right)(1-\rho)\left(\frac{1-c'}{c'-\rho}\right) \\
s_1 &= \mathbf{1}_{\{c'<\sqrt{\rho}\}}\left(\tilde{\mathbf{w}}(y_1)^{*}\tilde{D}\tilde{w}(y_1)\right)^{-1} \\
&= \mathbf{1}_{\{c'<\sqrt{\rho}\}}\left(\frac{a-c}{a}\right)\left(\frac{\rho}{c'} - c'\right)\left(\frac{1-c'}{c'-\rho}\right) \\
s(y) &= \frac{2}{\pi}\left(\frac{a-c}{a}\right)\frac{\sqrt{1-y^2}}{1-y^2+((p-1)y+q)^2}.
\end{aligned}$$

Then, the identity operator I can be expressed in terms of $\tilde{D}$ as

$$I = \Bigg(s_0 \tilde{\mathbf{w}}(y_0)\tilde{\mathbf{w}}(y_0)^* + s_1 \tilde{\mathbf{w}}(y_1)\tilde{\mathbf{w}}(y_1)^* + \int_{-1}^{1} s(y)\tilde{\mathbf{w}}(y)\tilde{\mathbf{w}}(y)^* dy \Bigg) \tilde{D}. \tag{3.34}$$

Proof. It is easy to compute s_0 and s_1 using orthogonality by post-multiplying both sides of (3.34) by $\tilde{\mathbf{w}}(y_i)$. The proper weights are then obtained through (3.32)-(3.33). Getting $s(y)$ is slightly more involved. First, we right multiply (3.34) by $\tilde{D}^{-1}$, and then take the double z-transform of both sides. After rearranging terms, we deduce that it is equivalent to show that

$$\begin{aligned} &\frac{1}{1 - z_1 z_2} - \frac{c}{\upsilon} - s_0 \tilde{W}(y_0, z_1)\tilde{W}(y_0, z_2) \\ &- s_1 \tilde{W}(y_1, z_1)\tilde{W}(y_1, z_2) \\ &= \int_{-1}^{1} s(y)\tilde{W}(y, z_1)\tilde{W}(y, z_2) dy, \end{aligned} \tag{3.35}$$

where we know from the definition of $\tilde{D}$ that

$$\sum_{k=0}^{\infty}\sum_{l=0}^{\infty} z_1^k z_2^l \left[\tilde{D}^{-1}\right]_{kl} = \frac{1}{1 - z_1 z_2} - \frac{c}{\upsilon}.$$

It can be shown that the right-hand-side of (3.35) is the contour integral of a complex integrand over the unit circle [82]. To prove this, we denote the RHS of (3.35) by

$H(z_1, z_2)$ and substitute $y = \cos\theta$ to get

$$\begin{aligned}
&H(z_1, z_2) \\
&= \frac{2}{\pi}\int_0^{\pi}\left(1-\frac{c}{a}\right)\frac{\sin^2\theta}{\sin^2\theta + \left((p-1)\cos\theta+q\right)\right)^2} \\
&\times \tilde{W}\left(\cos\theta, z_1\right)\tilde{W}\left(\cos\theta, z_2\right)d\theta \\
&= \frac{1}{\pi}\int_0^{2\pi}\left(1-\frac{c}{a}\right)\mathrm{Re}\left[\frac{\sin\theta}{\sin\theta - i\left((p-1)\cos\theta+q\right)\right)}\right] \\
&\times \tilde{W}\left(\cos\theta, z_1\right)\tilde{W}\left(\cos\theta, z_2\right)d\theta \\
&= \frac{1}{\pi}\int_0^{2\pi}\left(1-\frac{c}{a}\right)\frac{\sin\theta}{\sin\theta - i\left((p-1)\cos\theta+q\right)\right)} \\
&\times \tilde{W}\left(\cos\theta, z_1\right)\tilde{W}\left(\cos\theta, z_2\right)d\theta.
\end{aligned}$$

The last equality follows from $\tilde{W}(\cos\theta, z)$ and $\mathrm{Im}\left(\frac{\sin\theta}{\sin\theta - i((p-1)\cos\theta+q))}\right)$ being respectively, even and odd in θ. Substituting $t = e^{j\theta}$ in the equation above, we get

$$\begin{aligned}
&H(z_1, z_2) \\
&= \frac{1}{i\pi}\oint\frac{dt}{t}\tilde{S}(t)\tilde{W}\left(\frac{t^2+1}{2t}, z_1\right)\tilde{W}\left(\frac{t^2+1}{2t}, z_2\right),
\end{aligned} \tag{3.36}$$

where we have used

$$\begin{aligned}
\tilde{S}(t) &= \left(1-\frac{c}{a}\right)\frac{t^2-1}{p\left(t-\sqrt{\rho}\right)\left(t-\frac{c'}{\sqrt{\rho}}\right)}, \\
\tilde{W}\left(\frac{t^2+1}{2t}, z\right) &= \left(1-\sqrt{\frac{a}{a-c}}\right) \\
&\quad + \sqrt{\frac{a}{a-c}}\left(\frac{\left(p(t^2+1)+2qt\right)z - 2t}{2(t-z)(zt-1)}\right).
\end{aligned}$$

We need to show that $H(z_1, z_2)$ is identical to the LHS of equation (3.35). To do so,

we employ complex integration, residue theory, and the fact that

$$\tilde{W}\left(\frac{t^2+1}{2t}, z\right)\Big|_{t=\sqrt{\rho}} = \tilde{W}(y_0, z),$$
$$\tilde{W}\left(\frac{t^2+1}{2t}, z\right)\Big|_{t=\frac{c'}{\sqrt{\rho}}} = \tilde{W}(y_1, z).$$

We find the residues pertinent to the integrand as

$$\text{Res}\left[\frac{1}{t}\tilde{S}(t)\tilde{W}\left(\frac{t^2+1}{2t}, z\right), t = z\right] = \frac{1}{2}\sqrt{\frac{a-c}{a}},$$
$$\text{Res}\left[\frac{1}{t}\tilde{S}(t), t = \sqrt{\rho}\right] = -\frac{s_0}{2},$$
$$\text{Res}\left[\frac{1}{t}\tilde{S}(t), t = \frac{c'}{\sqrt{\rho}}\right] = -\frac{s_1}{2}.$$

From the residue theorem, we obtain

$$\begin{aligned}
&H(z_1, z_2)\\
&= \lim_{t\to 0}\left(2\tilde{S}(0)\tilde{W}\left(\frac{t^2+1}{2t}, z_1\right)\tilde{W}\left(\frac{t^2+1}{2t}, z_2\right)\right)\\
&+ \sqrt{\frac{a-c}{a}}\left(\tilde{W}\left(\frac{z_1^2+1}{2z_1}, z_2\right) + \tilde{W}\left(\frac{z_2^2+1}{2z_2}, z_1\right)\right)\\
&- s_0\tilde{W}(y_0, z_1)\tilde{W}(y_0, z_2) - s_1\tilde{W}(y_1, z_1)\tilde{W}(y_1, z_2).
\end{aligned}$$

We can verify through algebraic manipulation that

$$\begin{aligned}
&\lim_{t\to 0}\left(2\tilde{S}(0)\tilde{W}\left(\frac{t^2+1}{2t}, z_1\right)\tilde{W}\left(\frac{t^2+1}{2t}, z_2\right)\right)\\
&+ \sqrt{\frac{a-c}{a}}\left(\tilde{W}\left(\frac{z_1^2+1}{2z_1}, z_2\right) + \tilde{W}\left(\frac{z_2^2+1}{2z_2}, z_1\right)\right)\\
&= \frac{1}{1-z_1z_2} - \frac{c}{v}
\end{aligned}$$

and, as such, the desired result follows. That is, (3.35) holds and hence the identity expression of (3.34) is valid. □

We are now ready to characterize the equilibrium distribution of the buffer

overflow probability for the second buffer in a tandem network.

Theorem F.5. *For the system described in Theorem F.1, let the departure process of the first queue serve as an input to the second buffer. The latter queue is assumed to be served at constant rate υ and its occupancy is denoted by $\ell_2(t)$. If*

$$\left(\frac{a}{\upsilon} - 1\right)\frac{\lambda}{\mu} < 1,$$

then we can express the equilibrium probability distribution of $\ell_2(t)$ exceeding a threshold x as

$$\lim_{t\to\infty} \Pr(\ell_2(t) > x) = K'\left(s_1 e^{\zeta_1 x} + \int_{-1}^{1} s(y) e^{\zeta(y)x} dy\right)$$

where K' is a constant.

Proof. Using the Chapman-Kolmogorov equation given in (3.21) and the identity expression of (3.34), we can rewrite (3.23) as

$$\pi(x) = E\exp\left(x\frac{c\mu}{a(c-\upsilon)}\tilde{D}^{-1}\tilde{Q}\right)\left(s_0\tilde{\mathbf{w}}(y_0)\tilde{\mathbf{w}}(y_0)^* + s_1\tilde{\mathbf{w}}(y_1)\tilde{\mathbf{w}}(y_1)^* + \int_{-1}^{1} s(y)\tilde{\mathbf{w}}(y)\tilde{\mathbf{w}}(y)^* dy\right)\tilde{D}E^{-1}\pi(0).$$

For this system to be stable, we need $\lim_{x\to\infty}\pi(x) = 0$ and $\pi(0,n) = p_n$ for $n \in \mathbb{N}$, since at steady state the probability of the buffer being empty is zero for the states where the input rate c exceeds the service rate υ. These boundary conditions imply that $\tilde{\mathbf{w}}(y_0)^*\tilde{D}E^{-1}\pi(0) = 0$, which can be employed to obtain $\pi(0,0)$,

$$\pi(0,0)\sqrt{\frac{a}{a-c}}\left(\frac{\upsilon}{c-\upsilon}\right) = \frac{ap_0}{a-c}\left(\tilde{W}\left(y_0, \sqrt{\rho}\right) - 1\right).$$

Since $\tilde{D}E^{-1}\pi(0)$ is almost a geometric sequence, the expression $\tilde{\mathbf{w}}(y)^*\tilde{D}E^{-1}\pi(0)$ is closely related to the z-transform of $\tilde{\mathbf{w}}(y)$. In view of the discussion at the beginning of Section b, it is not too surprising to find that $\tilde{\mathbf{w}}(y)^*\tilde{D}E^{-1}\pi(0)$ is constant for all

$y \in \{y_1\} \cup [-1,1]$; and it is equal to

$$\begin{aligned}
&\tilde{\mathbf{w}}(y)^* \tilde{D} E^{-1} \pi(0) \\
&= -\pi(0,0)\sqrt{\frac{a}{a-c}}\left(\frac{\upsilon}{c-\upsilon}\right) + \frac{a p_0}{a-c}\left(\tilde{W}(y,\sqrt{\rho}) - 1\right) \\
&= p_0 \left(\frac{a}{a-c}\right)^{\frac{3}{2}} \left(\frac{c'}{1-c'} - \frac{\rho}{1-\rho}\right) = -\frac{p_0}{s_0}\sqrt{\frac{a}{a-c}} > 0.
\end{aligned}$$

This, along with the fact that $\zeta = \frac{c\mu}{a(c-\upsilon)}\xi$, implies

$$\begin{aligned}
\pi(x) = &-\frac{p_0}{s_0}\sqrt{\frac{a}{a-c}} E\Bigg(e^{\zeta_1 x} s_1 \tilde{\mathbf{w}}(y_1) \\
&+ \int_{-1}^{1} e^{\zeta(y)x} s(y) \tilde{\mathbf{w}}(y) dy \Bigg).
\end{aligned}$$

We can get an expression for the probability of the buffer exceeding a fixed threshold x, using the relationship $\Pr(\ell_2 > x) = \sum_{n=0}^{\infty} \pi(x,n) = \langle \pi(x), \mathbf{1} \rangle$. Noting that

$$\langle E\tilde{\mathbf{w}}(y), \mathbf{1} \rangle = \sqrt{\frac{a-c}{a}} + \sqrt{\frac{a}{a-c}}\left(\frac{c'}{1-c'}\right) > 0$$

for all $y \in \{y_1\} \cup [-1,1]$, we obtain

$$\begin{aligned}
\Pr(\ell_2 > x) = &-\frac{p_0}{s_0}\left(1 + \frac{ac'}{(a-c)(1-c')}\right) \\
&\times \left(s_1 e^{\zeta_1 x} + \int_{-1}^{1} s(y) e^{\zeta(y)x} dy\right).
\end{aligned}$$

This is the desired expression. □

G. Tail Asymptotics for Buffer Occupancy

In this section, we characterize the exponential decay-rate of the complementary cumulative distribution function (also referred to as tail-asymptote) of the equilibrium buffer-occupancy random variable of the second buffer in a tandem-queue. We emphasize that finding this tail-asymptote is essentially the same as obtaining the

effective bandwidth of the departure process discussed in Section 1 (see also [74,75]). From the buffer distribution derived above, we can obtain the dominant exponential decay rate. An alternate way of finding the effective bandwidth of the departure process would be to first compute the moment generating function of the busy period of the fluid queue [83–85], and then use the method proposed in [86]. There is also literature available on finding tail-asymptote of the departure process in a discrete-time queue [47, 49, 50].

Recall that $\zeta(y) < \zeta_1 < 0$ for all $y \in [-1, 1]$, whenever $\sqrt{\rho} > c'$. That is, if the discrete eigenvalue ζ_1 exists, then it is larger than the supremum of the continuous spectrum. We characterize the tail-asymptote in the following theorem.

Theorem G.1. *For $\ell_2(t)$ described in Theorem F.5, the exponential decay rate associated with the steady-state probability of the buffer exceeding a threshold is given by θ_2, where*

$$\begin{aligned}\theta_2 &= -\lim_{x\to\infty} \frac{\ln \Pr(\ell_2 > x)}{x} \\ &= \begin{cases} \frac{\mu}{a-\upsilon} - \frac{\lambda}{\upsilon} & \sqrt{\rho} > c' \\ \frac{c\mu}{a(c-\upsilon)}(1-\sqrt{\rho})^2 & \sqrt{\rho} \le c.' \end{cases}\end{aligned} \tag{3.37}$$

Proof. We use the fact that $s(y)$ is non-negative, bounded and integrable. For $\sqrt{\rho} > c'$, the desired result follows from

$$K' s_1 e^{\zeta_1 x} \le \Pr(\ell_2 > x) \le K' e^{\zeta_1 x} \left(s_1 + \int_{-1}^{1} s(y) dy \right).$$

On the other hand, for the case where $\sqrt{\rho} \le c'$ and for some $\epsilon \in (0, 1)$, we have

$$\begin{aligned} K' e^{\zeta(1-\epsilon)x} \int_{1-\epsilon}^{1} s(y) dy &\le \Pr(\ell_2 > x) \\ &\le K' e^{\zeta(1)x} \int_{-1}^{1} s(y) dy. \end{aligned}$$

Taking logarithms on both sides, dividing by x, and taking limits; we get

$$\begin{aligned}\zeta(1-\epsilon) &\leq \liminf_{x\to\infty} \frac{\ln \Pr(\ell_2 > x)}{x} \\ &\leq \limsup_{x\to\infty} \frac{\ln \Pr(\ell_2 > x)}{x} \leq \zeta(1).\end{aligned}$$

Finally, letting $\epsilon \to 0$ and using the continuity of $\zeta(\cdot)$, we obtain the desired result.

□

H. Maximum Achievable Rate for Departure Process

In Theorem G.1, we found the tail-asymptote of the second buffer in a tandem queue as described in Theorem F.1. In this section, we find the achievable rate region $\mathcal{A}_2(\theta_0, c, \nu)$ for the tandem queue considered in Section C. While establishing this region, several cases must be considered. These cases are not individually difficult, but they are collectively tedious. Therefore, the proof appears in Appendix J. We have tried to make the presentation as clear and concise as possible.

Theorem H.1. *For exponential-decay rate θ_2 described in Theorem G.1, we define*

$$\begin{aligned}\mathcal{A}_2(\theta_0, c, \upsilon) &= \left\{a \in \mathbb{R}^+ : \theta_2 \geq \theta_0\right\} \\ &= \left\{a \in \mathbb{R}^+ : a \leq \bar{a}_2(\theta_0, c, \upsilon)\right\}\end{aligned}$$

where the function $\bar{a}_2(\theta, c, \upsilon)$ is given by

$$\bar{a}_2(\theta, c, \upsilon) = \begin{cases} \bar{a}_1(\theta, \upsilon), & 0 < \upsilon \leq \upsilon^* \\ \bar{a}_3(\theta, c, \upsilon), & \upsilon^* < \upsilon < c. \end{cases}$$

The first component, $\bar{a}_1(\theta, \upsilon)$, is equal to

$$\bar{a}_1(\theta, \upsilon) = \nu\left(1 + \frac{\mu}{\lambda + \nu\theta}\right)$$

and the second component $\bar{a}_3(\theta, c, \upsilon)$ *is given by the expression*

$$\bar{a}_3(\theta, c, \upsilon) =$$
$$c + \frac{c\mu}{\lambda}\left(\frac{-1 + \sqrt{1 - \left((c-\upsilon)\frac{\theta}{\mu} - 1\right)\left((c-\upsilon)\frac{\theta}{\lambda} - 1\right)}}{(c-\upsilon)\frac{\theta}{\lambda} - 1}\right)^2$$

with parameter υ^* *determined implicitly by*

$$\frac{c}{\upsilon^*} - 1 = \frac{\theta\upsilon^*\mu}{\lambda\mu + (\lambda + \theta\upsilon^*)^2}.$$

I. Conclusion

We compared network coding to classic routing for a QoS constrained communication system, and computed the achievable rate regions for both scenarios. For an AWGN model with identical link capacities, network coding significantly outperforms classic routing. This essentially implies that the benefits associated with network coding are far more important than the multiplexing gains achieved by routing for symmetric networks. However, we obtained more interesting results for the wireless butterfly network. In this case, allocating resources to form a network topology suitable for packet combining at intermediate nodes does not always offer gains and may even be detrimental at times. These results depend on the topology of the butterfly network. For network coding to be useful, we need symmetric direct links. It turns out that it is often better to route packets rather than trying to establish a direct link to the destination using excessive amounts of physical resources. A possible avenue of future research is to study networks with varying service rates.

J. Proof of Theorem H.1

If the second buffer in the tandem queue described in Theorem F.5 has a QoS constraint θ_0 on the asymptotic decay-rate of buffer-occupancy, then we must have

$$\theta_0 \leq \theta_2.$$

This condition enables us to determine the set $\mathcal{A}_2(\theta_0, c, \nu)$ of *on*-time arrival rates that can be supported by this queue under QoS constraint θ_0, and for a given service-rate ν.

First, we gather that, for $\nu_1 \leq \nu_2, \mathcal{A}_2(\theta_0, c, \nu_1) \subseteq \mathcal{A}_2(\theta_0, c, \nu_2)$. This fact follows directly from stochastic majorization of the buffer-content processes. Second, if the constant service-rate ν of the second queue is greater than or equal to the constant service-rate c of the first queue, then the second queue always remains empty (equivalently, $\theta_2 = \infty$). Under such circumstances, the second queue does not limit peak-rate a. We therefore focus on the case where $\nu \in [0, c)$. If $a \in [0, \nu)$, then both queues stay empty with $\theta_1 = \theta_2 = \infty$; thus, we have $[0, \nu) \subseteq \mathcal{A}_2(\theta_0, c, \nu)$ for all $\theta_0 \geq 0$ and all $0 \leq \nu < c$.

Let us define $\bar{a}_1(\theta_0, \nu) = \nu \left(1 + \frac{\mu}{\lambda + \theta_0 \nu}\right)$. Note that $\bar{a}_1(\theta_0, \nu)$ is increasing in ν. Let ν_c be the service-rate such that $\bar{a}_1(\theta_0, \nu_c) = c$. For $\nu \in [0, \nu_c]$, we can show that $[0, \bar{a}_1(\theta_0, \nu)] \subseteq \mathcal{A}_2(\theta_0, c, \nu)$. When $a \in [\nu, \bar{a}_1(\theta_0, \nu)] \subseteq [\nu, c]$, the first queue remains empty and the arrival processes at the first and second buffers are path-wise identical. Therefore, by Theorem F.1, we get $\theta_2 = \frac{\mu}{a - \nu} - \frac{\lambda}{\nu}$. Since $a \leq \bar{a}_1(\theta_0, \nu)$, it follows that $\theta_2 \geq \theta_0$ and hence $[0, \bar{a}_1(\theta_0, \nu)] \subseteq \mathcal{A}_2(\theta_0, c, \nu)$. It is clear that for any finite θ_0, $\nu_c < c$.

We can explicitly write ν_c as

$$\nu_c = \begin{cases} \frac{1}{2}\left(c - \frac{\lambda+\mu}{\theta_0}\right)\left(1 + \sqrt{1 + \frac{4c\lambda\theta_0}{(c\theta_0-\lambda-\mu)^2}}\right), & c \geq \frac{\lambda+\mu}{\theta_0} \\ \frac{1}{2}\left(c - \frac{\lambda+\mu}{\theta_0}\right)\left(1 - \sqrt{1 + \frac{4c\lambda\theta_0}{(c\theta_0-\lambda-\mu)^2}}\right), & c < \frac{\lambda+\mu}{\theta_0}. \end{cases}$$

Next, we consider the case where $a \geq c$. In this scenario, the arrival rate at the second queue is $c(K(t))$, as described in Section 1. In the previous section, we obtained the exponential decay-rate θ_2 governing the queue-distribution for this system,

$$\theta_2 = \begin{cases} \frac{\mu}{a-\nu} - \frac{\lambda}{\nu}, & \frac{a/c-1}{a/\nu-1} < \sqrt{\rho} \\ \frac{c\mu}{a(c-\nu)}(1-\sqrt{\rho})^2, & \frac{a/c-1}{a/\nu-1} \geq \sqrt{\rho}. \end{cases}$$

We note that $\bar{a}_1(\theta_0, \nu)$ denotes the maximum supportable rate when the discrete eigenvalue ζ_1 of $D^{-1}Q$ governs the decay-rate of the buffer overflow probability, i.e., when $c' = \frac{a/c-1}{a/\nu-1} < \sqrt{\rho}$. Furthermore, the condition $\sqrt{\rho} > c'$ is equivalent to the quadratic expression $\rho - \frac{\nu}{c}\sqrt{\rho} + \frac{\lambda}{\mu}\left(1 - \frac{\nu}{c}\right) > 0$. This equation specifies upper and lower bounds on the existence of ζ_1. That is, ζ_1 exists only when $\sqrt{\rho} \notin \left[\frac{\nu}{2c}\left(1 - \sqrt{1 - \frac{4\lambda c}{\mu\nu}\left(\frac{c}{\nu} - 1\right)}\right), \frac{\nu}{2c}\left(1 + \sqrt{1 - \frac{4\lambda c}{\mu\nu}\left(\frac{c}{\nu} - 1\right)}\right)\right]$. We call the endpoints of the interval $\sqrt{\rho_l}$ and $\sqrt{\rho_u}$, respectively. We have $\rho = \frac{\lambda}{\mu}\left(\frac{a}{c} - 1\right)$, which gives upper and lower bounds $a_{u1}(\nu) = c\left(1 + \frac{\mu}{\lambda}\rho_u\right)$ and $a_{l1}(\nu) = c\left(1 + \frac{\mu}{\lambda}\rho_l\right)$ in terms of ρ_u and ρ_l, respectively. Thus, when $a \in \mathcal{A}_2(\theta_0, c, \nu)$ lies between these two bounds, the supremum of the continuous spectrum dominates the tail-asymptotics and the discrete eigenvalue ζ_1 disappears. We can write $a_{u1}(\nu)$ and $a_{l1}(\nu)$ explicitly,

$$a_{u1}(\nu) = \nu\left(1 + \frac{\nu\mu}{2c\lambda}\left(1 + \sqrt{1 - 4\left(\frac{c}{\nu} - 1\right)\frac{c\lambda}{\nu\mu}}\right)\right), \tag{3.38}$$

$$a_{l1}(\nu) = \nu\left(1 + \frac{\nu\mu}{2c\lambda}\left(1 - \sqrt{1 - 4\left(\frac{c}{\nu} - 1\right)\frac{c\lambda}{\nu\mu}}\right)\right). \tag{3.39}$$

For the real interval $[a_{l1}, a_{u1}]$ to exist, the discriminant in (3.38)-(3.39) must be non-

negative. Defining

$$\nu' = \frac{2c}{1+\sqrt{1+\frac{\mu}{\lambda}}},$$

we deduce that a_{u1} and a_{l1} are well-defined real numbers provided that $\nu > \nu'$. That is, discrete eigenvalue ζ_1 exists for all $\nu \leq \nu'$. In this case, the achievable rate region is limited by $\bar{a}_1(\theta_0, \nu)$. Note that $\nu' < c$. Clearly,

$$1 - 4\left(\frac{c}{\nu} - 1\right)\frac{c\lambda}{\nu\mu} = (1 + \frac{\lambda}{\mu}) - \frac{\lambda}{\mu}\left(\frac{2c}{\nu} - 1\right)^2$$

is an increasing function that maps $\nu \in [\nu', c]$ to $[0, 1]$. Therefore, $\sqrt{\rho_u}$ is an increasing function of ν. Furthermore, because $\sqrt{\rho_u}\sqrt{\rho_l} = \frac{\lambda}{\mu}\left(1 - \frac{\nu}{c}\right)$ is a decreasing function of ν, it follows that $\sqrt{\rho_l}$ is a decreasing function of ν in $[\nu', c]$. Since ρ_u and ρ_l are both non-negative in this interval, it follows that a_{u1} monotonically increases and ranges over $\left[\nu'\left(1 + \frac{\nu'\mu}{2c\lambda}\right), c\left(1 + \frac{\mu}{\lambda}\right)\right]$, whereas a_{l1} decreases and ranges over $\left[\nu'\left(1 + \frac{\nu'\mu}{2c\lambda}\right), c\right]$.

When $a \in [a_{l1}(\nu), a_{u1}(\nu)]$, the continuous spectrum dominates the tail asymptotics. This implies,

$$\left((c-\nu)\frac{\theta_0}{\lambda} - 1\right)\rho + 2\sqrt{\rho} + \left((c-\nu)\frac{\theta_0}{\mu} - 1\right) \leq 0.$$

Therefore, since $\sqrt{\rho} \geq 0$, we know that $\sqrt{\rho}$ belongs to the interval

$$\left[0, \frac{-1 + \sqrt{1 - \left((c-\nu)\frac{\theta_0}{\lambda} - 1\right)\left((c-\nu)\frac{\theta_0}{\mu} - 1\right)}}{(c-\nu)\frac{\theta_0}{\lambda} - 1}\right]$$

when $(c-\nu)\frac{\theta_0}{\lambda} > 1$; and it belongs to the set

$$\left[0, \frac{1-\sqrt{1-\left((c-\nu)\frac{\theta_0}{\lambda}-1\right)\left((c-\nu)\frac{\theta_0}{\mu}-1\right)}}{1-(c-\nu)\frac{\theta_0}{\lambda}}\right] \bigcup \left[\frac{1+\sqrt{1-\left((c-\nu)\frac{\theta_0}{\lambda}-1\right)\left((c-\nu)\frac{\theta_0}{\mu}-1\right)}}{1-(c-\nu)\frac{\theta_0}{\lambda}}, \infty\right)$$

when $(c-\nu)\frac{\theta_0}{\lambda} < 1$. We emphasize that we require $(c-\nu)\frac{\theta_0}{\mu} \leq 1$ in the first case for a non-negative $\sqrt{\rho}$ to exist. Furthermore, we need $\nu \in [c - \frac{\lambda+\mu}{\theta_0}, c]$ for a real $\sqrt{\rho}$ to exist. For the second queue to be stable, we must have $a < \nu\left(1+\frac{\nu}{\lambda}\right)$, which implies $a \leq \bar{a}_3(\theta_0, c, \nu)$ where

$$\begin{aligned}
&\bar{a}_3(\theta_0, c, \nu) \\
&= \sup\left\{a \leq c\left(1+\frac{\mu}{\lambda}\right)\right. \\
&\left.: \theta_0 \leq \frac{c\mu}{a(c-\nu)}\left(1-\sqrt{\left(\frac{a-c}{c}\right)\frac{\lambda}{\mu}}\right)^2\right\}, \\
=&c + \frac{c\mu}{\lambda}\left(\frac{\sqrt{1-\left((c-\nu)\frac{\theta_0}{\mu}-1\right)\left((c-\nu)\frac{\theta_0}{\lambda}-1\right)}-1}{(c-\nu)\frac{\theta_0}{\lambda}-1}\right)^2.
\end{aligned}$$

It is clear from the explicit form of $\bar{a}_3(\theta_0, c, \nu)$ that it is increasing in ν. We can write the achievable rate region $\mathcal{A}_2(\theta_0, c, \nu)$ in terms of $\bar{a}_1, a_{u1}, a_{l1}$ and $\bar{a}_3$ as

$$\begin{aligned}
\mathcal{A}_2(\theta_0, c, \nu) = &\left\{a \in \mathbb{R}^+ / [a_{l1}(\nu), a_{u1}(\nu)] : a \leq \bar{a}_1(\theta_0, \nu)\right\} \\
&\cup \left\{a \in [a_{l1}(\nu), a_{u1}(\nu)] : a \leq \bar{a}_3(\theta_0, c, \nu)\right\}.
\end{aligned}$$

Let ν^* be the value of ν where $\bar{a}_1(\theta_0, \nu)$ intersects a_{l1} or a_{u1}. In other words, if we substitute $\bar{a}_1$ for a in the expression for ρ, the value of ν that equates $\sqrt{\rho}$ and c'

is ν^*. It can be seen that $\bar{a}_3(\theta_0, c, \nu^*) = \bar{a}_1(\theta_0, \nu^*)$ by substituting $\lambda = \frac{\mu(a/c-1)}{(a/\nu^*-1)^2}$ and $\sqrt{\rho} = \frac{a/c-1}{a/\nu^*-1}$ to obtain

$$\theta_0 = \frac{\bar{a}_1\mu(c-\nu^*)}{c(\bar{a}_1-\nu^*)^2} = \frac{\bar{a}_3\mu(c-\nu^*)}{c(\bar{a}_3-\nu^*)^2},$$

where $a = \bar{a}_1(\theta_0, \nu^*) = \bar{a}_3(\theta_0, c, \nu^*)$. The value ν^* can be written as the positive root of the following equation in the interval $[0, c]$,

$$\theta_0(c-\nu) = \frac{(\theta_0\nu)^2\,\mu}{\lambda\mu + (\lambda+\theta_0\nu)^2}.$$

The LHS is continuous and monotonically decreasing for $\nu \in [0, c]$ and ranges over $[c\theta_0, 0]$, whereas the RHS is continuous and monotonically increasing in the same interval and ranges over $[0, \frac{(c\theta_0)^2\mu}{\lambda\mu+(\lambda+c\theta_0)^2}]$. It is clear from the continuity and the monotonicity of these functions that there exists a unique real $\nu^* \in [0, c]$ for all values of $\lambda, \mu, c, \theta_0 > 0$. Furthermore, since $\nu > \nu'$ for a_{u1} and a_{l1} to exist, we have $\nu^* \geq \nu'$. In addition, since $\left((c-\nu)\frac{\theta_0}{\lambda} - 1\right)\left((c-\nu)\frac{\theta_0}{\mu} - 1\right)$ is decreasing in ν, we can prove that it is less than or equal to 1 for all $\nu \in [\nu^*, c]$ by showing it to be less than or equal to 1 at $\nu = \nu^*$. This is equivalent to showing $(c-\nu^*)\theta_0 \in [0, \lambda+\mu]$. Therefore, we only need to show that equation

$$\frac{(\theta_0\nu^*)^2\mu}{\lambda\mu + (\lambda+\theta_0\nu^*)^2} \leq \lambda + \mu$$

is valid, which is equivalent to $\lambda\left(\theta_0\nu^* + \lambda + \mu\right)^2 \geq 0$ and obviously true. Hence, the result holds. We know from Theorem F.3 that discrete eigenvalue ζ_1 is always greater than the supremum of the continuous spectrum of $D^{-1}Q$, whenever it exists. Therefore,

$$\frac{\mu}{a\left(1-\frac{\nu}{c}\right)}\left(1 - \sqrt{\left(\frac{a}{c}-1\right)\frac{\lambda}{\mu}}\right)^2 \geq \frac{\mu}{a-\nu} - \frac{\lambda}{\nu},$$

which in turn implies $c\left(1+\frac{\mu}{\lambda}\right) \geq \bar{a}_3(\theta_0, c, \nu) \geq \bar{a}_1(\theta_0, \nu)$.

If $\bar{a}_1$ intersects a_{u1} and $\bar{a}_3$ at ν^*, then $\bar{a}_1 > a_{u1} > a_{l1}$ for all $\nu < \nu^*$. Since, a_{l1} is decreasing, we have $\bar{a}_1 > a_{l1}$ for $\nu \in [\nu^*, c]$ as well. Also, since $\bar{a}_1$ and a_{u1} intersect uniquely at ν^* in $[\nu', c]$ and are both increasing, $\bar{a}_1 < a_{u1}$ within this interval. Thus, we have

$$\begin{aligned}\mathcal{A}_2(\theta_0, c, \nu) = &\{a \in \mathbb{R}^+ : a \leq \bar{a}_1(\theta_0, \nu), \nu < \nu^*\} \\ &\cup \{a \in \mathbb{R}^+ : a \leq \bar{a}_3(\theta_0, c, \nu), \nu \geq \nu^*\}.\end{aligned}$$

Otherwise, when $\bar{a}_1$ intersects a_{l1} and $\bar{a}_3$ at ν^*, $\bar{a}_3 < a_{l1} < a_{u1}$ for all $\nu < \nu^*$ because $\bar{a}_3$ is increasing in ν. Since a_{l1} is decreasing, $\bar{a}_1 > a_{l1}$ for $\nu \in (\nu^*, c]$. In this interval, $\bar{a}_1$ and a_{u1} do not intersect and are both increasing with $\bar{a}_1 < a_{u1}$. We conclude that

$$\begin{aligned}\mathcal{A}_2(\theta_0, c, \nu) = &\{a \in \mathbb{R}^+ : a \leq \bar{a}_1(\theta_0, \nu), \nu < \nu^*\} \\ &\cup \{a \in \mathbb{R}^+ : a \leq \bar{a}_3(\theta_0, c, \nu), \nu \geq \nu^*\}.\end{aligned}$$

From the discussion above, it is clear that we have $a_{l1}(\nu) < \bar{a}_1(\theta_0, \nu) < a_{u1}(\nu)$ for $\nu > \nu^*$. Also, it was shown that $\bar{a}_3(\theta_0, c, \nu^*) = \bar{a}_1(\theta_0, \nu^*)$, and therefore we conclude that $\mathcal{A}_2(\theta_0, c, \nu) = [0, \bar{a}_2(\theta_0, c, \nu)]$ where the maximum achievable rate $\bar{a}_2(\theta_0, c, \nu)$ is continuous in ν. This rate can therefore be characterized completely for all $\nu \in [0, c]$ as

$$\bar{a}_2(\theta_0, c, \nu) = \begin{cases} \bar{a}_1(\theta_0, \nu) & \nu \in [0, \nu^*] \\ \bar{a}_3(\theta_0, c, \nu) & \nu \in [\nu^*, c]. \end{cases}$$

This is the desired result.

Table II. Comparing our work with the literature

	[54]	[55]	[56]	[57]	[58]	[59]	[60]	[61]	[62]	[63]	[64]
Scheme	×U	B	U	M	B	×U	×U	×U	×U	×U	M, ×U
Arrivals	S	F	B	B	B	S	P	P	P	S	S
Service	C	S	B+V	IID	IID	V	C	C	C	C	V
Policy	WC	–	FC	FC	FC	FC	O	FC	O	FC	–
Aspect	D	C	D	C,D	D,Q	Q	Q	C	Q	Q	Q
QoS	PD	–	MD	S,D	MQ	S	S	MD	S	S	S
Network	L	D	P	D	D	A	2W	2W	2W	T	A

Schemes: broadcast (B), multicast (M), unicast (U), multiple unicast (×U). Arrivals: Bernoulli (B), fixed (F), Poisson (P), stochastic (S). Service: bulk (B), constant (C), independent and identically distributed (IID), stochastic (S), variable (V). Policy: first-come first-served (FC), opportunistic (O), work conserving (WC). Aspect: coding delay (C), decoding delay (D), queueing delay (Q). QoS: delay (D), mean delay (MD), mean queue-length (MQ), packet-drop probability (PD), stability (S). Network: arbitrary (A), downlink (D), large (L), point-to-point (P), two-way relay (2W), tandem (T).

Table III. System parameters

$N_0 = 10^{-6}$ W/Hz	Noise power spectral density
$W = 22$ MHz	Spectral bandwidth
$\lambda_1^{-1} = \lambda_2^{-1} = 650$ ms	Mean *off*-time
$\mu_1^{-1} = \mu_2^{-1} = 352$ ms	Mean *on*-time

CHAPTER IV

VALUE-AWARE RESOURCE ALLOCATION FOR SERVICE GUARANTEES IN NETWORKS

The traditional formulation of the total value of information transfer is a multi-commodity flow problem. Each data source is seen as generating a commodity along a fixed route, and the objective is to maximize the total system throughput under some concept of fairness, subject to capacity constraints of the links used. This problem is well studied under the framework of network utility maximization and has led to several different distributed congestion control schemes. However, this view of value does not capture the fact that flows may associate value, not just with throughput, but with link-quality metrics such as packet delay and jitter. In this work, the congestion control problem is redefined to include individual source preferences. It is assumed that degradation in link quality seen by a flow adds up on the links it traverses, and the total utility is maximized in such a way that the end-to-end quality degradation seen by each source is bounded by a value that it declares. Decoupling source-dissatisfaction and link-degradation through an effective capacity variable, a distributed and provably optimal resource allocation algorithm is designed to maximize system utility subject to these quality constraints. The applicability of the controller in different situations is supported by numerical simulations, and a protocol developed using the controller is simulated on ns-2 to illustrate its performance.

A. Introduction

Recent years have seen an enormous growth in demand for Internet access, with applications ranging from personal use to commercial and military operations. Sev-

eral of these applications are sensitive to a quality of packet delivery. For instance, archiving data transfers can tolerate long delays, while voice over Internet protocol (VoIP) is very sensitive to latency. Between these extreme examples lies a spectrum of applications with varying service requirements, e.g. video conferencing, electronic commerce and online gaming. All these applications require the allocation of enough network resources for satisfactory performance.

The design of efficient network control systems demands that end-user value be taken into consideration when allocating resources. The Internet architecture is built around the concept of a *flow*, which is a transfer of data between a fixed source-destination pair. How do we quantify the value of such a flow? The classical formulation of the total value of information transfer is a multi-commodity flow problem, in which each data source is seen as generating a commodity along a fixed route; the objective is to maximize the sum throughput under some concept of fairness, subject to capacity constraints on the links used [87–90]. If the flow from source r has a rate $x_r \geq 0$ and the system utility associated with such a flow is represented by a concave, increasing function $U_r(x_r)$, the objective can be stated as

$$\max \sum_{r \in \mathcal{S}} U_r(x_r) \quad \text{subject to} \quad y_l \leq c_l, \ \forall \ l \in \mathcal{L} \tag{4.1}$$

where $\mathcal{S}$ is the set of sources, $\mathcal{L}$ denotes the set of links, and c_l is the capacity of link $l \in \mathcal{L}$. Also, the load on link l is equal to $y_l = \sum_{r \in \mathcal{S}} R_{lr} x_r$ where R denotes the routing matrix of the network, with $R_{lr} = 1$ if the flow associated with source r is routed through link l. Note that we refer to flows and sources interchangeably; if there are multiple flows between a source and a destination, we simply give them different names. This is a convex optimization problem that is well studied under the framework of network utility maximization [87–90].

This approach to network resource allocation can often be used to decompose

the problem into several subproblems, each of which is amenable to a distributed solution. This so-called *optimization decomposition framework* has yielded a rich set of control schemes and protocols whose architectural implications are discussed in [91]. For example, there is a strong connection between the *primal solution* to the utility maximization problem and TCP-Reno [92, 93] characterized in [87, 88]. Similarly, one can identify connections between TCP-Vegas and the *dual solution* of the problem [94]. The same approach has been taken in the design of several new protocols such as Scalable TCP [95, 96] (which allows scaling of rate increases/decreases based on network characteristics), FAST-TCP [97] (meant for high bandwidth environments), TCP-Illinois [98] (which uses loss and delay signals to attain high throughput), and TRUMP [99] (a multipath protocol with fast convergence properties). A good tutorial on network utility maximization algorithms is [100].

Still, there is a growing realization that throughput cannot be considered as the sole value metric. As mentioned above, in applications such as voice calls, data is rendered useless after a certain delay threshold. Thus, simply ensuring that link capacities are not exceeded is not sufficient to capture value in this scenario. How do we ensure that the user is not dissatisfied with the quality of service? In many cases the quality of data transfer over a link decreases with load. For example, metrics such as the delay and the jitter experienced by packets as they pass through a router depend on the total load on the corresponding links. Such quality degradation may also add up over multiple hops. Indeed, the delay experienced by packets in a flow is the sum of the delays over each hop taken.

Once we have a clear conception of quality degradation as a function of link load, a pertinent question becomes: *Can we design a simple distributed algorithm for fair resource allocation under which each users' quality is no worse than a prescribed value?* To address this question, we need to redefine the traditional congestion

control problem to explicitly account for the tradeoff between throughput and quality degradation. We denote the degradation in the quality of link l with load y_l by a convex increasing function $V_l(y_l)$, and assume that degradation in link quality seen by a flow is additive over the links it traverses. Furthermore, we assume that the quality degradation is inherent to a link, and is identical for all flows sharing the link. Thus, there are no priorities assigned to particular flows. We maximize utility in such a way that the total degradation seen by source r is no greater than a pre-specified positive value σ_r. The modified objective becomes

$$\max \sum_{r \in \mathcal{S}} U_r(x_r) \quad \text{subject to} \quad \sum_{l \in \mathcal{L}} R_{lr} V_l\left(y_l\right) \leq \sigma_r, \tag{4.2}$$

where, again, $x_r \geq 0$, $y = Rx$ and under the assumption that $\lim_{y \to c_l} V_l(y) = \infty$. We emphasize that this convex optimization problem requires the quality degradation on each route to remain bounded. In this chapter, our objective is to design a distributed control scheme that can approach the optimal operating point which trades off throughput and quality, without maintaining per-flow information or prioritizing certain packets at intermediate hops. We overview our main contributions below, with details contained in subsequent sections.

Classical optimization-decomposition techniques typically yield a "source-rate responds to link-price" type of controller [87–91], wherein each link's price increases with the link-load in order to prevent the link-capacity from being exceeded. As the link-price increases, sources cut down their transmission rates, where the aggressiveness of the source controller is determined by its utility function. However, the solution to our delay-aware problem has remained elusive due to the strong coupling between the quality seen at source, which requires a hard guarantee, and *link quality degradation*, which depends on the link-loads along its route.

We present illustrative examples of what quality degradation functions may look

like in Section B, and discuss an example that we use later in the chapter. We then proceed to provide a centralized solution to our problem of interest in Section C. We develop two algorithms to this end. A primal algorithm is proposed in Section D. Our main contribution, a dual algorithm, is presented in Section E, and stems from the realization that it is possible to decouple the QoS guarantees at sources and link quality degradation using an *effective capacity* that is based on the link-price and user-dissatisfaction. Once we set the effective capacity for a link, the quality degradation depends solely on this decision and not on the actual link-load.

Each source declares its *dissatisfaction* to the links it uses based on the difference between the quality it sees and what it requires. The links select a *price* based on the difference between load and *effective capacity*, which in turn depends on link-load and aggregate dissatisfaction of users sharing the link. This decoupling of link-load and effective capacity appears to have the correct properties to allow a distributed solution. Finally, sources employ route-prices (the sum of all link-prices on a route) to determine their source-rates.

We prove that the algorithm is indeed capable of solving our resource allocation problem using Lyapunov techniques [101]. We report numerical results about the simulated operation of the controller in Section F to illustrate its performance. A new contribution over our earlier version [102] is the development of a realistic protocol based on the proposed controller. The protocol is presented in Section G, where we report simulation experiments on ns-2 to show that it performs as desired. We conclude with pointers to future work in Section H.

B. Examples of Quality Degradation Functions

We begin this section by discussing candidate measures of link quality degradation with load. We make several assumptions on the properties of link quality degradation functions. They can be expressed in the following manner; if the total sum-rate on link l is $y_l = \sum_{r \in \mathcal{S}} R_{lr} x_r$ then

- the quality degradation function $V_l(y_l)$ is non-negative, convex increasing in link-load y_l,
- the total quality degradation seen by flow r is $\beta_r(x) = \sum_{l \in \mathcal{L}} R_{lr} V_l(y_l)$ (i.e., quality degradation adds up over multiple hops), and
- the service process at one link does not impact the arrival process at the succeeding link.

The above assumptions ensure the analytical tractability of our optimization problem. We also believe that they provide acceptable models of quality degradations in communication systems with queues. Below, we support these assumptions with common examples of quality degradation functions.

For an M/M/1 queue with arrival rate y and service rate c, the expected waiting time in the queue is $y/(c\,(c-y))$ for a stable queue, that is when $y < c$. In this case, one can select the quality degradation function to be the expected waiting time for any packet in the queue,

$$V(y) = y/(c\,(c-y)).$$

We note that the quality degradation function is non-negative, convex, and increases from 0 to ∞ when x ranges in $[0, c)$.

As a second example, consider a single server fluid queue with constant-rate arrival y and a two-state on-off service process where on and off times are exponentially

distributed with rates μ and λ, respectively. When the service is on and the buffer is non-empty, it is serviced at a constant rate r such that $y < r\lambda/(\lambda+\mu)$. It can be shown [53] that the probability of buffer exceeding a threshold z is exponentially decreasing as z increases. A possible quality degradation function in this case is the inverse of this decay-rate. One can write the decay-rate explicitly in terms of the above parameters as

$$1/V(y) = -\lim_{z\to\infty} z^{-1}\log \Pr\left(L > z\right) = \lambda/y - \mu/(r-y).$$

If we denote $r\lambda/(\lambda+\mu)$ by c then, one can write

$$V(y) = y(c-y)^{-1}\left(c\lambda^{-1} - y(\mu+\lambda)^{-1}\right).$$

Again, note that the quality degradation function is non-negative, convex, and increases from 0 to ∞ when x ranges in $[0, c)$. Recent results [103] suggest that, under appropriate conditions, the delay seen in a queue is independent of other queues even though packets traverse the network along connected paths.

C. Centralized Resource Allocation

We begin by developing ideas on how to solve the resource allocation problem of (4.2) in a centralized fashion, and we create model networks that we will use as examples to illustrate the performance of various control loops throughout. Consider the scenario where the utility functions assume unbounded negative values when $x_r = 0$ and the quality degradation functions grow unbounded when sum-rates y_l approach c_l. In this case, an optimal solution is characterized by $x_r > 0$ and $y_l < c_l$. Let $x^* = \{x_r^* : r \in \mathcal{S}\}$ be a feasible point such that $y^* = Rx^*$, and suppose there exist constants $w_r \geq 0$

such that

$$\begin{aligned} U_r'(x_r^*) - \sum_{s \in \mathcal{S}} w_s \sum_{l \in \mathcal{L}} R_{ls} R_{lr} V_l'\left(y_l^*\right) = 0, \\ w_r \left(\beta_r(x^*) - \sigma_r\right) = 0, \end{aligned} \tag{4.3}$$

for all $r \in \mathcal{S}$, then x^* is a global maximum. Moreover, if $U_r(\cdot)$ is strictly concave, then x^* is the unique global maximum.

We illustrate by the following example how our model takes into consideration all of the desired properties of the quality degradation function and how they impact resource allocation with service guarantees.

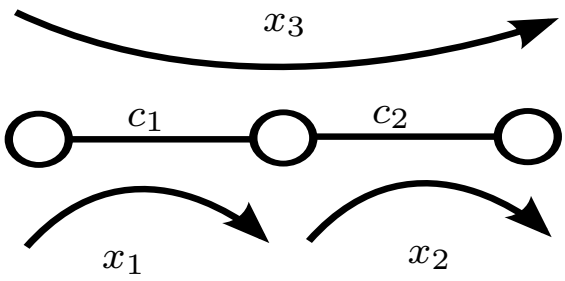

Fig. 15. Three flows sharing a two-link network.

Example [Three-Flows Two-Hop Network] Consider a network composed of three sources transmitting over two links, as shown on Fig. 15. Let link i have capacity c_i and suppose that we are using logarithmic utility and quality degradation functions. Then, the resource allocation problem becomes

$$\begin{aligned} &\max \ \sum_{i=1}^{n} a_i \log x_i \text{ subject to} \\ &-\log\left(1 - (x_i + x_3)/c_i\right) \leq \sigma_i, \quad i = 1, 2 \\ &-\sum_{i=1}^{2} \log\left(1 - \frac{x_i + x_3}{c_i}\right) \leq \sigma_3. \end{aligned} \tag{4.4}$$

Let w_i be the Lagrange multipliers corresponding to the quality degradation con-

straint of flow i. The *Lagrangian* is

$$L(x,w) = \sum_{i=1}^{3} a_i \log x_i + \sum_{i=1}^{2} w_i \left(\log \left(1 - \frac{x_i + x_3}{c_i} \right) + \sigma_i \right) + w_3 \left(\sum_{i=1}^{2} \log \left(1 - \frac{x_i + x_3}{c_i} \right) + \sigma_3 \right).$$

Therefore, we can derive the KKT conditions for this system,

$$\begin{aligned} a_i/x_i^* - (w_i + w_3)/(c_i - x_i^* - x_3^*) = 0, \quad i = 1,2 \\ a_3/x_3^* - \sum_{i=1}^{2} (w_i + w_3)/(c_i - x_i^* - x_3^*) = 0, \\ w_i \left(\log \left(1 - (x_i + x_3)/c_i \right) + \sigma_i \right) = 0, \quad i = 1,2 \\ w_3 \left(\sum_{i=1}^{2} \log \left(1 - (x_i + x_3)/c_i \right) + \sigma_3 \right) = 0. \end{aligned} \tag{4.5}$$

Let us consider the situation where $\sigma_3 < \min\{\sigma_1, \sigma_2\}$. In this case, $w_1 = w_2 = 0$ and

$$\sum_{i=1}^{2} \log \left(1 - \frac{x_i + x_3}{c_i} \right) + \sigma_3 = 0.$$

For the simple case where $a_i = 1$ and $c_i = c$ for $i = 1,2,3$, we get optimal rates

$$x_1^* = x_2^* = 2x_3^* = \frac{2c}{3} \left(1 - e^{-\sigma_3/2} \right). \tag{4.6}$$

This illustrative example provides supporting evidence that our modeling intuition is accurate for resource allocation with service guarantees. We list pertinent observations derived from this model:

- for any finite service requirement, the sum-rate is always less than the capacity of each link;
- throughputs decrease with the number of hops due to service requirements;
- when quality degradation is inherent to a link, the flow with the most stringent

service requirements limits the throughput of every neighboring flow.

D. Primal Algorithm

In this section, we develop an algorithm that can be employed to obtain an approximate solution to our optimization problem. The approach that we adopt is called the *Primal* method, as it follows from the primal formulation of the problem. The main idea is to relax the constraints by incorporating them as a cost into the objective. Essentially, the idea is that there is a price for violating the quality constraints and, as such, we can maximize the difference between utility and cost. We consider the objective

$$J(x) = \sum_{r \in \mathcal{S}} \left(U_r(x_r) - B_r\left(\beta_r(x)\right)\right), \tag{4.7}$$

where $B_r(\cdot)$ is a convex barrier function that increases from zero to infinity as its argument ranges from zero to σ_r. To maximize this function, we can use a gradient descent approach,

$$\begin{aligned} \dot{x}_r &= k_r(x_r)\left(U_r'(x_r) - q_r\right), \\ q_r &= \sum_{s \in \mathcal{S}} B_s'\left(\beta_s(x)\right) \sum_{l \in \mathcal{L}} R_{ls} R_{lr} V_l'\left(y_l\right). \end{aligned} \tag{4.8}$$

Since the problem is convex, it is straightforward to show using Lyapunov techniques [87, 88, 101] that this algorithm converges and leads to a maximizer of (4.7). To this end, note that $J(x)$, as defined in (4.7), is a strictly concave function. We denote its unique maximizer by $\hat{x}$. Then, $J(\hat{x}) - J(x)$ is non-negative and equals zero only at $x = \hat{x}$. This makes $W(x) \triangleq J(\hat{x}) - J(x)$ a natural candidate Lyapunov function; we use it in the following proposition, which has a similar proof to that found in [88].

Proposition D.1. *Consider a network in which all sources follow the primal control*

algorithm (4.8). *Suppose* $J(x)$ *is as defined in* (4.7) *and let functions* $U_r(\cdot)$, $k_r(\cdot)$, $V_l(\cdot)$ *and* $B_s(\cdot)$ *be such that* $W(x)$ *grows unbounded as* $\|x\| \to \infty$, *and* $\hat{x}_i > 0$ *for all* i. *Then, the controller in* (4.8) *is globally asymptotically stable and the equilibrium value maximizes* (4.7).

Proof. Differentiating $W(x)$ with time t, we get

$$\dot{W} = -\sum_{r\in\mathcal{S}} \frac{\partial J}{\partial x_r}\dot{x}_r = -\sum_{r\in\mathcal{S}} k_r(x_r)\left(U_r'(x_r) - q_r\right)^2 < 0,$$

for all $x \neq \hat{x}$, and $\dot{W} = 0$ for $x = \hat{x}$. Note that the second equality follows from (4.7) and (4.8). Thus, all the conditions of the Lyapunov theorem are satisfied, which ensures that the system state necessarily converges to $\hat{x}$. □

Convergence is a highly desirable attribute. Yet, the primal controller suffers from the following drawbacks. The approach is not optimal; the relaxation will produce an acceptable solution only if the barrier values at the optimal solution of our original objective (4.2) are small. Further, the above formulation does not allow for optimal points on the boundary of the constraint set. To overcome these limitations, we can approach the problem from a dual perspective and hope for better performance.

E. Dual Algorithm

We start by writing a *dual* formulation for the resource allocation problem defined in (4.2). This provides new insight on how to obtain a distributed means of achieving optimal resource allocation. The dual formulation corresponding to our optimization problem is given by

$$D(w) = \max_{x_s\geq 0} \sum_{s\in\mathcal{S}} U_s(x_s) - w_s\left(\beta_s(x) - \sigma_s\right).$$

Let x^* be the optimal maximizer, then for all $r \in \mathcal{S}$

$$U_r'(x_r^*) = \sum_{s \in \mathcal{S}} w_s \sum_{l \in \mathcal{L}} R_{ls} R_{lr} V_l' \left(y_l^*\right).$$

This gives us a system of equations that can be solved to find the optimal x^* for any vector w. However, it requires knowledge of the load on every link a flow traverses. Therefore, this approach is not completely distributed.

Nevertheless, this formulation gives us the hint that instead of link load and link-degradation being dependent on each other directly with load $y = Rx$ and degradation $V(y)$, we can break up their coupling. We do this by introducing a new variable $\tilde{y}_l$, we refer to as *effective capacity* of link l. This variable upper bounds the total link load y_l, and $V_l(\tilde{y}_l)$ upper bounds the link degradation. We then define "effective quality degradation" $\tilde{\sigma}_r \triangleq \sum_{l \in \mathcal{L}} R_{lr} V_l(\tilde{y}_l)$ seen by flow r. Then, the relaxed version of the resource allocation problem becomes

$$\begin{aligned} &\max \sum_{r \in \mathcal{S}} U_r(x_r) \\ &\text{subject to } y_l \leq \tilde{y}_l, \ \forall\ l \in \mathcal{L} \\ &\tilde{\sigma}_r \leq \sigma_r, \text{ and } x_r \geq 0 \ \forall\ r \in \mathcal{S}. \end{aligned} \tag{4.9}$$

Assuming that our concave utility and convex quality degradation functions ensure that values of x_r and $(c_l - \tilde{y}_l)$ are always positive, we can express the dual problem in terms of positive Lagrange multipliers p_l and w_r,

$$\min_{p,w \geq 0} D(p, w). \tag{4.10}$$

Here, $D(p, w)$ is the maximum of the *Lagrangian* function $L(x, \tilde{y}, p, w)$ with respect

to $x, \tilde{y}$, where

$$L = \sum_{s \in \mathcal{S}} U_s(x_s) - \sum_{l \in \mathcal{L}} p_l \left(y_l - \tilde{y}_l\right) - \sum_{s \in \mathcal{S}} w_s \left(\tilde{\sigma}_s - \sigma_s\right).$$

Let $x^*, \tilde{y}^*$ be the maximizers for L for any p, w, then

$$U_r'(x_r^*) = \sum_{l \in \mathcal{L}} R_{lr} p_l, \qquad\qquad p_l = V_l'(\tilde{y}_l^*) \sum_{r \in \mathcal{S}} R_{lr} w_r.$$

We find the partial derivatives of $D(p, w)$ with respect to variables p and w,

$$\frac{\partial D}{\partial p_l} = \tilde{y}_l^* - y_l^*, \qquad\qquad \frac{\partial D}{\partial w_r} = \sigma_r - \tilde{\sigma}_r^*, \tag{4.11}$$

where y_l^* and $\tilde{\sigma}_r^*$ are link-load and effective degradation for maximizing rate x^* and effective-capacity $\tilde{y}^*$. We employ gradient descent method for minimizing the dual of the relaxed problem. Therefore, the update equations for Lagrange multipliers p, w can be written as

$$\dot{p}_l = h_l(p_l) \left(-\frac{\partial D}{\partial p_l}\right)_{p_l}^{+}, \qquad\qquad \dot{w}_r = k_r(w_r) \left(-\frac{\partial D}{\partial p_l}\right)_{w_r}^{+}, \tag{4.12}$$

where $h_l(\cdot), k_r(\cdot)$ are positive functions and the notation $(z)_\rho^+$ is used to represent the function

$$(z)_\rho^+ = \begin{cases} z & \rho > 0 \\ \max\{z, 0\} & \rho = 0. \end{cases}$$

Function $(z)_\rho^+$ can be thought of as net input-rate into a fluid queue ρ. Clearly, when queue is non-empty fluid can enter at rate z, or leave at rate $-z$. However, when queue is empty fluid can only enter, but not leave.

We can now easily see that the above algorithm is distributed in nature. At any time during the evolution of our algorithm, we can treat Lagrange multipliers p_l and w_r as link-price and route-dissatisfaction, respectively. A flow r needs to "pay" link-price p_l for congesting link l if it uses the link (with the route-price being the

sum of all such p_l), and w_r is its end-to-end dissatisfaction under the current system state. The effective capacity of link l is $\tilde{y}_l$ and is decoupled from the actual load on this link, $y_l = \sum_{r \in \mathcal{S}} R_{lr} x_r$. We denote the sum of link-prices by $q_r = \sum_{l \in \mathcal{L}} R_{lr} p_l$ for any flow r, and the sum of route-dissatisfaction on a link l by $\nu_l = \sum_{r \in \mathcal{S}} R_{lr} w_r$ to yield the total dissatisfaction on that link. Note that such a total implies that there is no need to maintain per-flow information at the link. Notice again that due to decoupling through $\tilde{y}$, the perceived quality degradation is a function of the effective capacity, and not the actual link-load.

The algorithm is illustrated in Fig. 16. Although the diagram is reminiscent of traditional "source-rate responds to link-price" [87–90] corresponding to the congestion control problem defined in (4.1), the system is actually very different. The system may be described as follows:

- Each flow r, as it traverses its route, accrues the price q_r that it needs to "pay" for using each of the links l. Using this route-price, each source computes a feasible rate

$$x_r = U_r'^{-1}(q_r).$$

  Furthermore, each source declares its *dissatisfaction* w_r to the links it uses based on the difference between the quality degradation $\tilde{\sigma}_r$ that it sees and the degradation σ_r that it is willing to tolerate. The dissatisfaction is updated using

$$\dot{w}_r = k_r(w_r) \left(\tilde{\sigma}_r - \sigma_r\right)_{w_r}^{+}.$$

- Each link detects the total dissatisfaction ν_l of flows it accommodates, and it computes effective capacity

$$\tilde{y}_l = V_l'^{-1}(p_l/\nu_l)$$

and updates the link price by

$$\dot{p}_l = h_l(p_l)\,(y_l - \tilde{y}_l)^+_{p_l}\,.$$

The link ensures that the quality degradation inflicted on its users is $V_l(\tilde{y}_l)$ by increasing or reducing its aggregate flow as needed.

In summary, along with the two traditional elements of source-rate x_r and link-price p_l, we have two additional control variables: source-dissatisfaction w_r and effective capacity $\tilde{y}_l$ (with link-degradation $V_l(\tilde{y}_l)$) that provide two further dimensions of control that are required for a distributed solution. We next show that for admissible

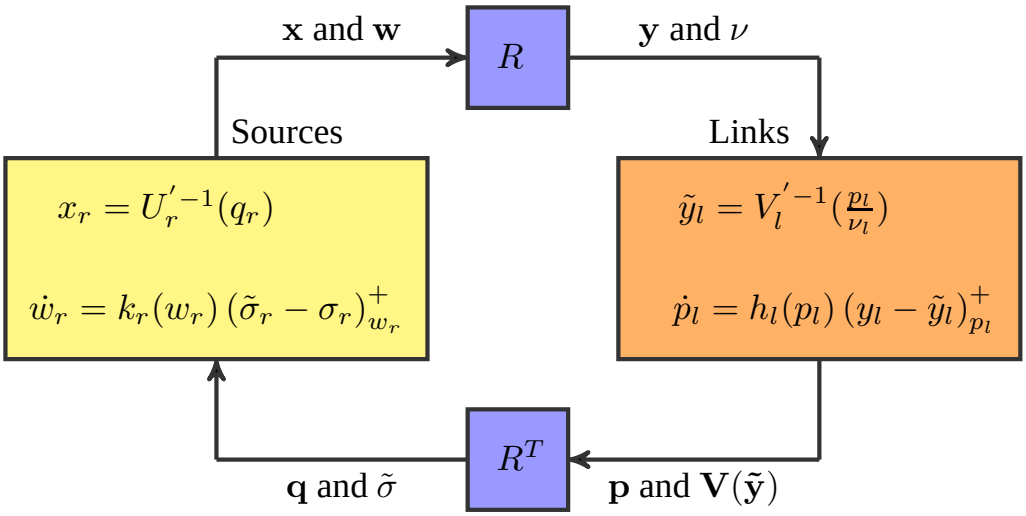

Fig. 16. Block diagram of value-aware resource allocation with decoupling of user-dissatisfaction on the source side, and quality on the link side.

functions $V_l(\cdot)$, the *effective capacity* is equal to the sum rate for all links for which the capacity constraint is binding.

Proposition E.1. *Assume that $V_l(\cdot)$ is strictly convex and increasing, and σ_r is finite for all $r \in \mathcal{S}$. Then, at equilibrium, for all $l \in \mathcal{L}_1 \triangleq \{l \in \mathcal{L} : \hat{y}_l < \hat{\tilde{y}}_l\}$, we have*

$\hat{p}_l = \hat{\nu}_l = 0$.

Proof. Our proof is by contradiction. Let us assume that at the equilibrium, there is a $l \in \mathcal{L}$ such that $\hat{y}_l < \hat{\hat{y}}_l$ and $\hat{\nu}_l \neq 0$. For this l, we have $\hat{p}_l = 0$ by KKT conditions. Note that $V_l'(\cdot)$ is a non-negative increasing function. That is, either $V_l'(0) = 0$ or $V_l'(z) > 0$ for all $z \in [0, c_l]$. For the former case, $0 \leq \hat{y}_l \leq \hat{\hat{y}}_l = 0 = V_l^{'-1}(0)$, i.e., $\hat{y}_l = \hat{\hat{y}}_l$. For the latter case, $\hat{p}_l$ cannot be zero since $V_l^{'-1}(0)$ is not in the feasible range of $\hat{y}_l$ and hence this will force $\hat{y}_l = \hat{\hat{y}}_l$ at the equilibrium. Therefore, the result holds. □

We have in effect, shown that the equilibrium conditions of our control loop satisfy the KKT conditions of the original optimization problem defined in (4.2). The conditions are easy to verify, and we may state this result as a corollary to Proposition E.1.

Corollary E.2. *The stationary point of* (4.12) *is a maximizer of the convex optimization problem described in* (4.2).

It is quite easy to show that the above algorithm is globally asymptotically stable. To show this, we choose our Lyapunov function to be

$$Q(p, w) = D(p, w) - D(\hat{p}, \hat{w}), \tag{4.13}$$

where $\hat{p}, \hat{w}$ are the unique minimizers of $D(p, w)$. It is clear that $Q(p, w) \geq 0$ for all values of p, w. Also, it is easily seen that $D(p, w)$ grows radially unbounded in p, w for our choice of $V_l(\cdot)$. Therefore, to prove that the above algorithm is stable it suffices to show $\dot{D}(p, w) \leq 0$, with equality if and only if $p = \hat{p}$ and $w = \hat{w}$. Note that at $\hat{p}, \hat{w}$, one would have $\dot{p}_l = \dot{w}_r = 0$.

Proposition E.3. *Let $Q(p,w)$ be as defined in* (4.13) *and functions $U_r(\cdot), V_l(\cdot), k_r(\cdot)$ and $h_l(\cdot)$ be such that $Q(p,w)$ grows unbounded with $\|p\|$ and $\|w\|$. The controller in* (4.12) *is globally asymptotically stable and the equilibrium value maximizes* (4.2)*.*

Proof. Differentiating D with respect to time, we get

$$\begin{aligned}\dot{D}(p,w) &= \sum_{l\in\mathcal{L}} \frac{\partial D}{\partial p_l}\dot{p}_l + \sum_{r\in\mathcal{S}} \frac{\partial D}{\partial w_r}\dot{w}_r \\ &= -\sum_{l\in\mathcal{L}} h_l(p_l)\left(y_l - \tilde{y}_l\right)\left(y_l - \tilde{y}_l\right)^+_{p_l} \\ &\quad -\sum_{r\in\mathcal{S}} k_r(w_r)\left(\tilde{\sigma}_r - \sigma_r\right)\left(\tilde{\sigma}_r - \sigma_r\right)^+_{w_r} < 0,\end{aligned}$$

for all $(p,w) \neq (\hat{p},\hat{w})$ and $\dot{D}(\hat{p},\hat{w}) = 0$. The second equality follows from equations (4.11) and (4.12). Thus, all the conditions of the Lyapunov theorem [101] are satisfied and we have proved that the Lagrange multipliers converge to $\hat{p},\hat{w}$. Hence, the system converges to the minimizer of (4.10). From the convexity of our original problem (4.2) and Corollary E.2, it follows that the stable point is the maximizer of (4.2). □

F. Numerical Studies

We utilize two realistic topologies presented in [99], illustrated in Fig. 17 and Fig. 18, to conduct numerical experiments. Our objective is to study the performance of our value-aware controller in different networking scenarios. We simulated our distributed resource allocation algorithm in Matlab using discrete-time evolution of *link-prices* and *end-to-end dissatisfaction*. Sources send packets at the rate generated by the controller, and links average this rate out using an exponential averaging factor α. Links base their decision on this average rate.

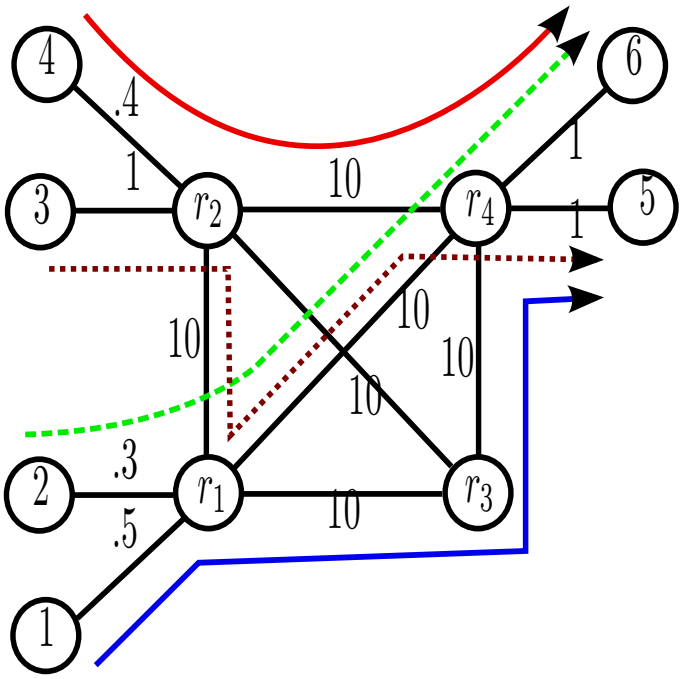

Fig. 17. Access-core topology for numerical experimentation.

1. Access-Core Topology

Our first network is an *Access-core* topology shown in Fig. 17. It represents a paradigm similar to commercially available Internet access, wherein users have a relatively small access bandwidth (from homes and businesses), connected together by a resource rich core-network. User bandwidth is constrained, either directly at the final hop into the home, or at a neighborhood head-end. Applications such as P2P file transfers (low quality constraint), as well as voice and video calls (higher quality constraints) result in end-to-end traffic on such a topology.

We consider the situation when nodes 1 and 3 wish to communicate to node 5; and similarly nodes 2 and 4, to node 6 over an Access-core network as shown in Fig. 17. The labels on the links denote their respective capacity. We refer to a flow

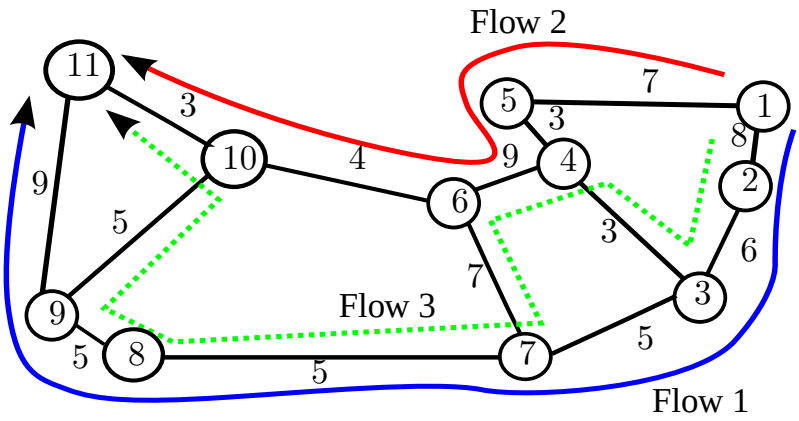

Fig. 18. Abilene topology for numerical experimentation.

by its origin node. The QoS constraints on quality of degradation for flows 1-4 are $2, 1, 3, 2$, respectively. We plot the convergence of the source rates in Fig. 19. We also plot load y_l and *effective capacity* $\tilde{y}_l$ for the diagonal core link used by flows 2 & 3 in Fig. 20. We also plotted the quality seen by the flow 2 and its constraint $\sigma_2 = 1$ in Fig. 21. Note, we have different rates of convergence of different parameters.

We assume that core links have a capacity of much higher order than that of access links. Thereby, every link on the core is taken to be of capacity 10, whereas access node 1 connects to the core with capacity 0.5. Similarly, capacities for nodes 2-4 are $0.3, 1, 0.4$ respectively. We chose nodes $5, 6$ to have identical access link capacities of 1.

2. Abilene Topology

Our second network represents the major nodes of the Abilene network topology [104], shown in Fig. 18. The network consists of high bandwidth links, and connects several universities and research labs. Traffic consists of large scale data transfers (low quality constraints) and distributed computation (where flows have strict delay con-

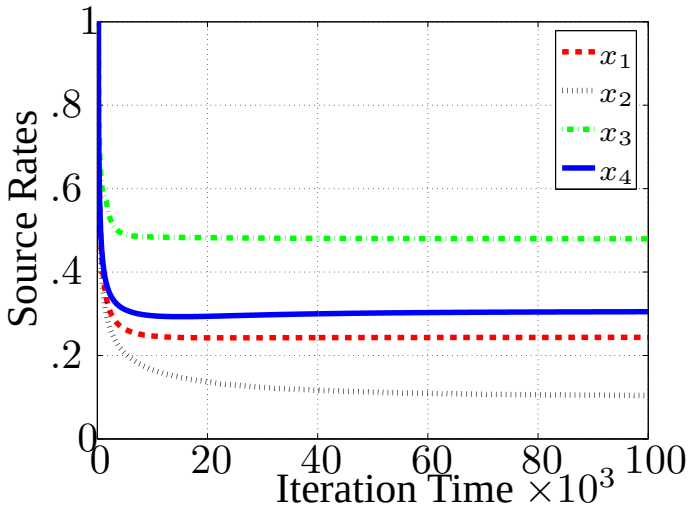

Fig. 19. Convergence of source-rate on Access-core topology.

straints).

We consider 3 flows over the Abilene network as shown in Fig. 18 with labels denoting the capacity of the corresponding link. Note, they are of same order. We call the flow on bottom to be flow 1 and one on the top, flow 2. These two flows have QoS constraint on dissatisfaction 8 and 15 respectively. Flow 3 has the zigzag path and has the most stringent QoS constraint of 1. We plot the convergence of flow rates in Fig. 22. We also plot load y_l and *effective capacity* $\tilde{y}_l$ for the link of capacity 3 shared by flows 2 and 3, in Fig. 23. We have also plotted the quality seen by the zigzag flow and it's acceptable constraint $\sigma_3 = 1$ in Fig. 24.

The conclusions that we draw from our simulations are (i) our value-aware resource allocation algorithm converges to a stable solution, (ii) user quality constraints are satisfied at equilibrium, i.e., the algorithm performs as designed, and (iii) the effective capacity is identical to the actual link load at equilibrium showing that our relaxation produces a tight solution.

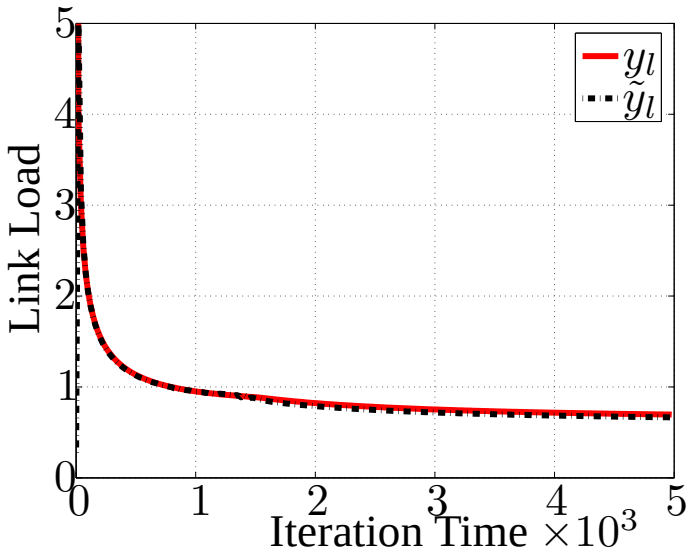

Fig. 20. Convergence of load & effective capacity on Access-core topology.

G. ns-2 Experiments

We adopt utility functions of the form $U_r(x_r) = a_r \log(x_r)$. We take $a_r = 100$ and normalize source-rates x_r to Megabits per second. The header of each packet transmitted by a source contains two additional fields, (i) source dissatisfaction w_r and (ii) route price q_r. Dissatisfaction is updated for each packet at the destination node as indicated by our dual algorithm. We selected the scaling function $k_r(w_r)$ to be a constant 10^3. Route price q_r is initialized to zero and is updated by the links. Depending on the current value, sources update their current rate as indicated by the algorithm.

As indicated by our algorithm, the price-update can be implemented by a virtual queue being served at rate equal to the *effective capacity* with an arrival rate of y_l. This virtual queue is implemented by traffic shaping (TS) queue, described below. We also need to compute the quality degradation at each link $V_l(\tilde{y}_l)$. We choose

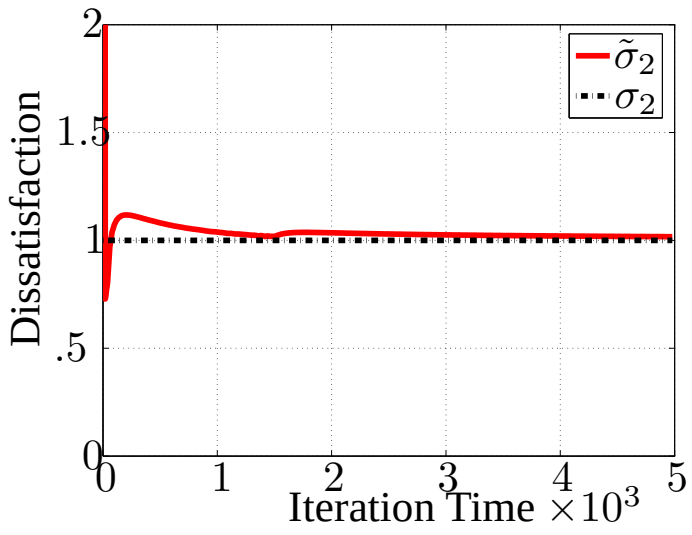

Fig. 21. Convergence of dissatisfaction for flow 2 on Access-core topology.

quality degradation per link to be the delay experienced by packets arriving in a queue with a rate $\tilde{y}_l$ and being served at link-capacity c_l. This queue is implemented by the router queue that buffers packets pending transmission over the link. The queues are shown in Fig. 25.

a. The Traffic Shaping Queue

The purpose of this queue is to shape traffic entering the router queue, and is illustrated by the queue on the left in Fig. 26. Since our system requires decoupling of the real load y_l on link l from the effective capacity $\tilde{y}_l$, we need to either add or subtract packets arriving at the link. For example, if two sources (S_1, S_2) are using link l, then a packet arrival from S_1 or S_2 is enqueued into the TS queue. The TS queue is drained at a rate $\tilde{y}_l$. The queue dynamics are implemented using a token generator at each link. Tokens are created at rate $\tilde{y}_l$. Each token observes the TS queue and if it is non-empty, it places the packet at the head of the TS queue into

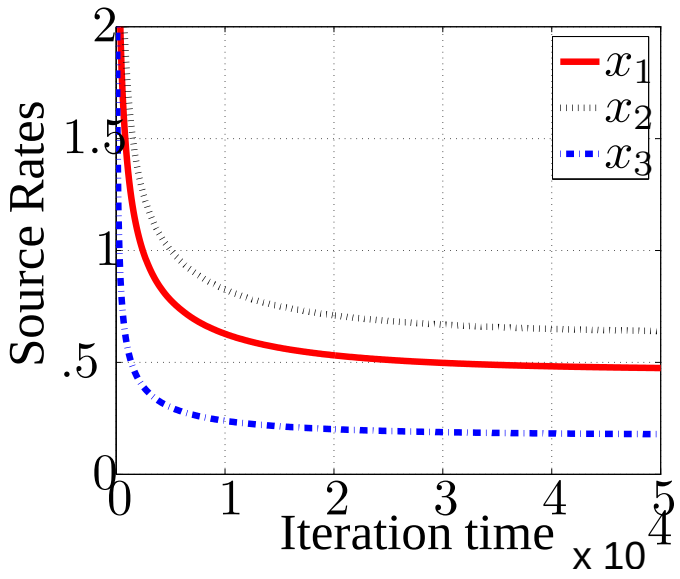

Fig. 22. Convergence of source-rate on Abilene topology.

the router queue. Otherwise, the token itself is placed in the router queue. Delays are created due to the time that packets spend in the TS queue, but we will see that these delays are small. We choose a large buffer size (10,000 packets) to ensure that very few packets are lost.

b. The Router Queue

We approximate the delay in router queue with a decreasing function of difference in link-capacity c_l and the effective capacity $\tilde{y}_l$. In particular, we take $V_l'(\tilde{y}_l) \approx K/(c_l - \tilde{y}_l)^2$. Therefore, we can update effective capacity $\tilde{y}_l$ at periodic time intervals according to the following equation

$$\tilde{y}_l = c_l - \sqrt{K\nu_l/p_l}. \tag{4.14}$$

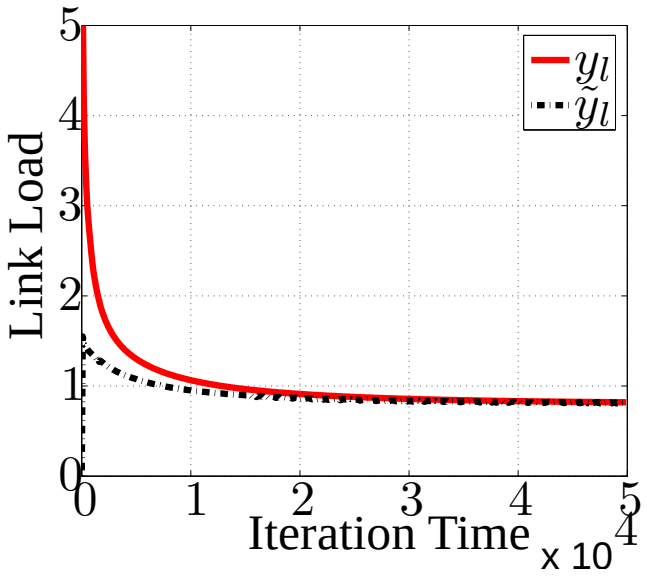

Fig. 23. Convergence of load & effective capacity on Abilene topology.

We choose $K = 10^9$ for fast convergence. Unlike the analytical result presented earlier, we note that in reality it takes a while for the impact of changing effective capacity $\tilde{y}_l$ to be felt on the end-to-end delay. Therefore, we do not change $\tilde{y}_l$ at the same time scale as source rates. Instead, we do so periodically at each link. The times at which each link changes its value of $\tilde{y}_l$ are not synchronized. The objective is to get the source rates to converge to effective capacity $\tilde{y}_l$. The value of $\tilde{y}_l$ is again changed after a time interval, assuming that the sum of source rates converged to that $\tilde{y}_l$, and hence the observed delay is $V_l(\tilde{y}_l)$.

We now study the performance of our protocol using a network simulator (ns-2) on the two realistic topologies presented that we saw earlier in Section F. Let T_i denote the propagation delay for flow i with no queueing delays. Our objective is to study the performance of our value-aware controller in different networking scenarios, and our example situations are shown in Fig. 27 and Fig. 28.

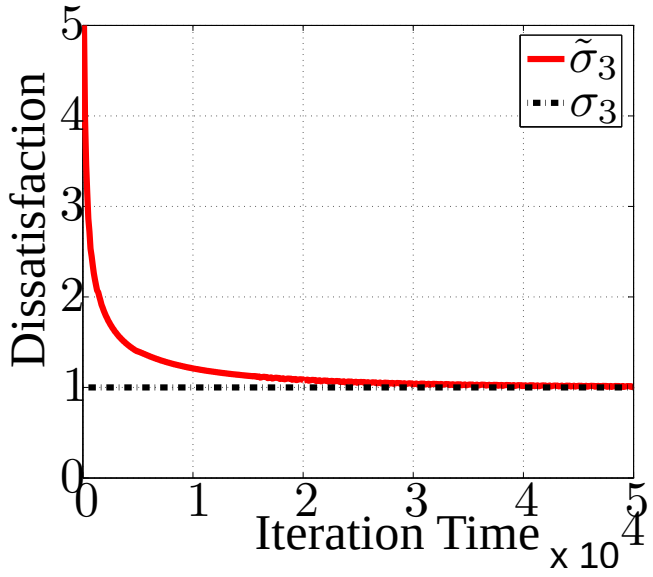

Fig. 24. Convergence of dissatisfaction for flow 3 on Abilene topology.

c. Abilene Topology

As shown in Fig. 27, the rates of the three flows are x_1, x_2 and x_3. In the system considered, $T_1 = 50$ ms, $T_2 = 45$ ms and $T_3 = 25$ms. We consider two cases. In the first case, flow 1 has stringent quality degradation constraint. In particular, respective degradation tolerances in terms of delays are $\sigma_1 = 55$ ms and $\sigma_2 = \sigma_3 = 1000$ s. Thus, we have set a very high delay tolerance for flows 2 and 3, whereas for

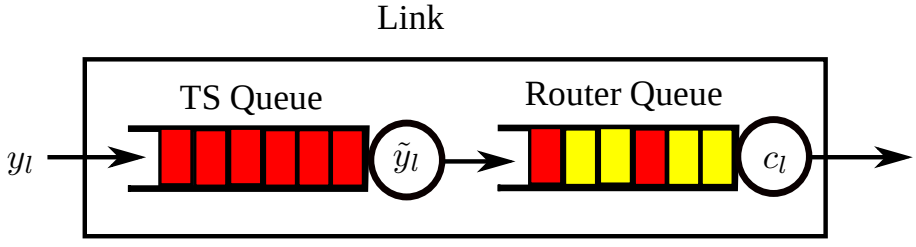

Fig. 25. Each node in our system contains a TS-queue and a router queue for each out-going link.

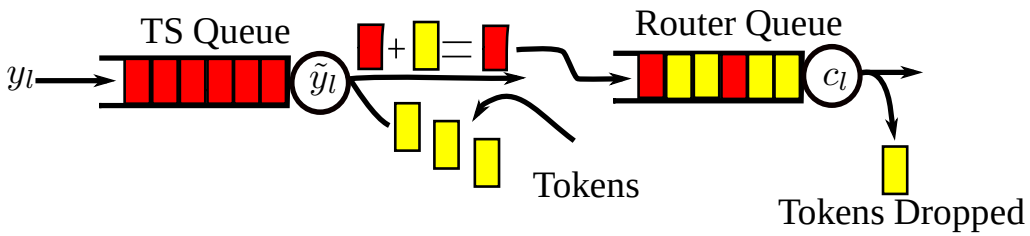

Fig. 26. The dynamics of the TS-queue and router queue in our system. Tokens are used in order to modulate the arrival rate into the router queue. Tokens are dropped when they reach the head of the router queue.

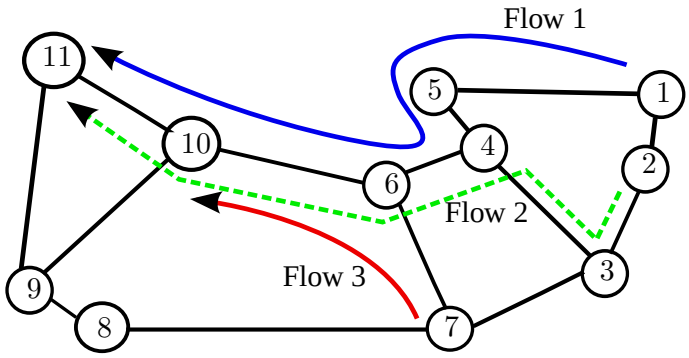

Fig. 27. Flows considered for ns-2 experimentation for Abilene topology.

flow 1 this tolerance is very low and nearly equals the propagation delay of 50 ms, and hence allows a queuing delay of only 5 ms for the five intermediate queues at routers 1, 5, 4, 6 and 10. The tight delay constraint on flow 1 has an effect on the other two flows which share the link between routers 6 and 10 with it. In Fig. 29 we plot the rates associated with individual flows. Figure 30 shows the acceptable delay for packets for flow 1, the delay through the router queues (the control delay), and the actual total delay (TS queue plus router queue) for packets to reach node 11 from node 1. Figure 31 plots the effective capacity $\tilde{y}$ for the link between nodes 6 and 10 (shared by all the three flows), which is less than the total link capacity of 25

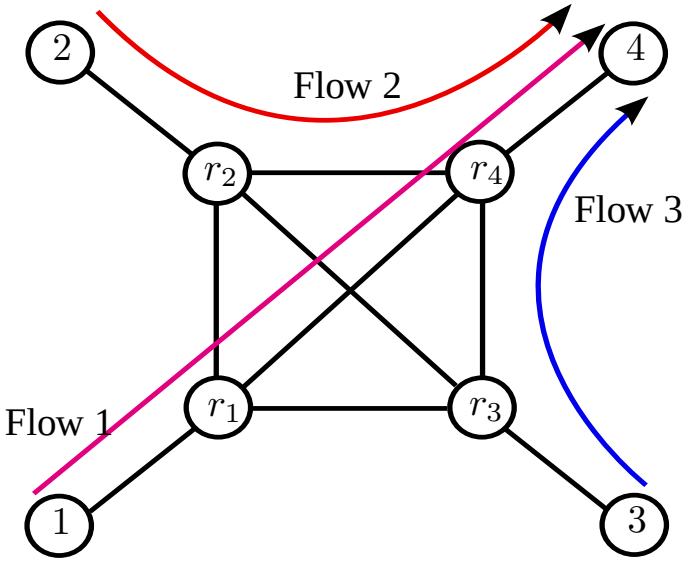

Fig. 28. Flows considered for ns-2 experimentation for Access-core topology.

Mbps. It is clear that the protocol is successful in ensuring a good delay performance for flow 1, at the expense of overall system throughput.

d. Access-Core Topology

Our other experiment involves the Access-core topology, with flows of interest shown in Fig. 28. In the system considered $T_1 = T_3 = 20$ ms and $T_2 = 25$ ms. Flow 1 has very stringent delay constraint. In particular, delay tolerances are $\sigma_1 = 25$ ms, $\sigma_2 = \sigma_3 = 1000$ s. In this case, flows 2 and 3 are highly delay tolerant, while flow 1 has low delay tolerance and it nearly equals the propagation delay of 20ms, and allows a queuing delay of 5ms for the three intermediate queues at routers 1, r_1 and r_4. The tight delay constraint on flow 1 has an effect on the other two flows which share the link between r_4 and 4 with it. In Fig. 32 we plot the three rates for flows 1, 2 and 3. Fig. 33 shows the acceptable delay for packets for flow 1, the delay through

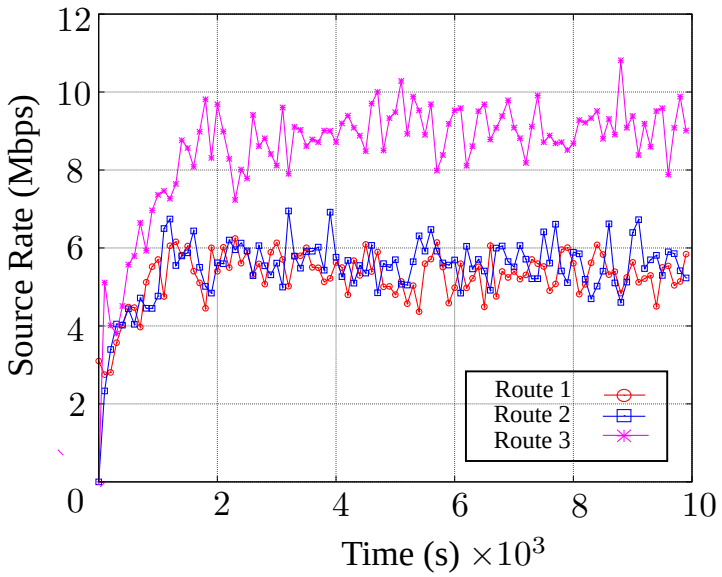

Fig. 29. Convergence performance of source-rates on Abilene network topology.

the router queue (control delay) and the actual total delay (TS queue plus router queue) for packets to reach node 4 from node 1. The total delay is close to the target. Fig. 34 plots the $\tilde{y}$ at router r_4 which is shared by all the three flows. Again, the throughput is less than the link capacity of 15 Mbps in the interest of reducing delay.

H. Conclusions

In this chapter we considered the design of a distributed resource allocation algorithm that would allow each individual flow to specify its measure of value. We assumed that every flow passing through a link suffers a certain quality-degradation due to the load on the link, and that such degradation adds up over the multiple links that the flow traverses. The objective is to ensure that the system throughput is maximized in a fair manner, subject to each flow's quality of service satisfying a hard constraint.

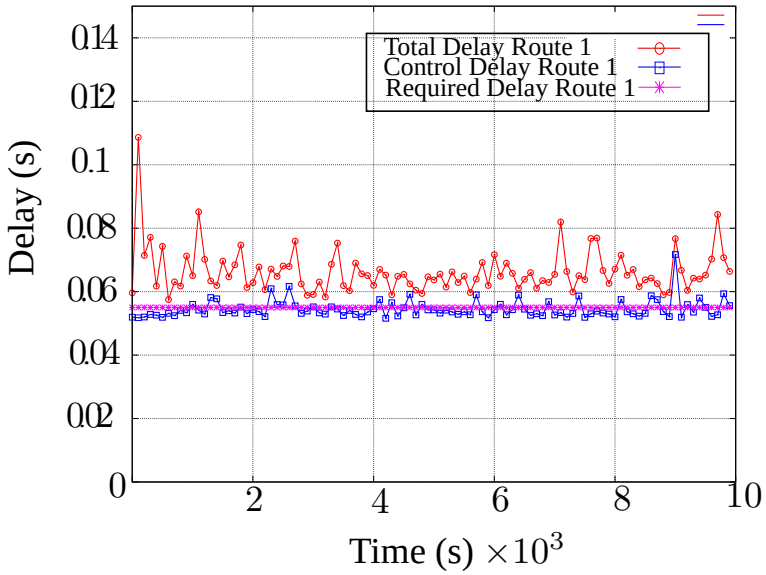

Fig. 30. Convergence performance of delay experienced by flow 1 on Abilene network topology.

Our aim was to ensure that the algorithm should be simple, use local information, and the relays need not maintain per-flow information.

We first showed that attempting to solve this problem by the usual optimization decomposition techniques using the primal formulation yields an approximate solution, and using the dual formulation yields a centralized solution. However, the observation that decoupling the link-load from the quality degradation using a secondary variable that we call *effective capacity*, allows us to design such a controller. Under our scheme, the source chooses its rate based on a *route price*, and it declares a *dissatisfaction* based on the quality of service that it sees. Links choose an effective capacity based on dissatisfaction and *link-price*, and modify the price as if the effective capacity were the actual capacity of the link. The control scheme only requires that links be aware of *aggregate* quantities of the flows using them, and the sources perform computations solely based on the parameters obtained from the links they

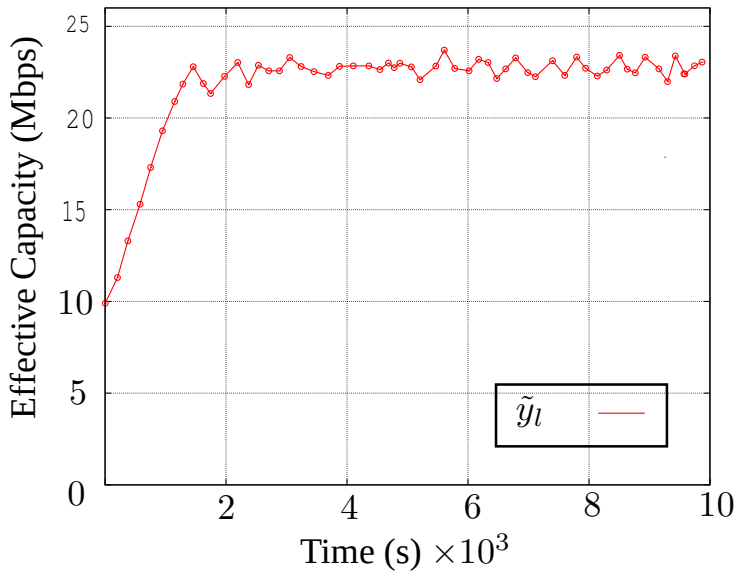

Fig. 31. Convergence performance of effective capacity on Abilene network topology.

traverse, hence satisfying our requirements.

We studied illustrative examples of quality-degradation functions that helped us gain insight into the working of the system. We performed simulations on realistic topologies to illustrate the performance of our algorithm, and used it as a basis for developing a delay-sensitive protocol. We showed that the protocol does indeed trade delay and throughput, so as to maximize the total utility of the system. In the future we would like to test out our ideas on a real network.

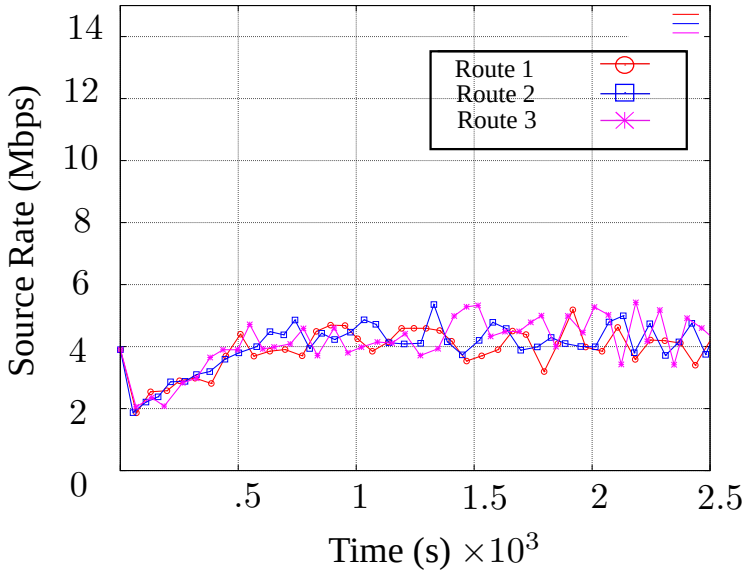

Fig. 32. Convergence performance of source-rates on Access-core network topology.

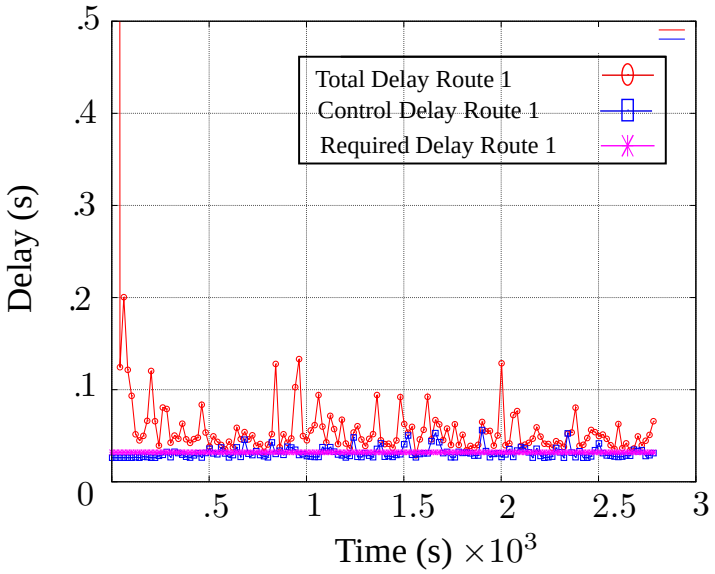

Fig. 33. Convergence performance of delay experienced by flow 1 on Access-core network topology.

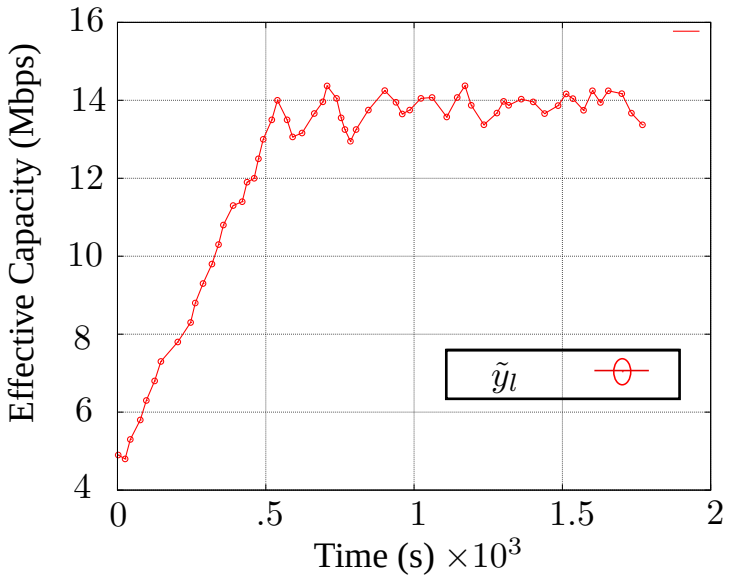

Fig. 34. Convergence performance of effective capacity on Access-core network topology.

CHAPTER V

CONCLUSION

We studied delay-sensitive communication over networks in this thesis. We looked at it at increasingly higher levels. First, we looked at point-to-point communication with service guarantees. Second, we studied the advantages of network coding for multicast over small networks, for communication with service guarantees. Third, we looked at a distributed algorithm to provide end-to-end service guarantees in a large network.

In the first chapter, we quantified the impact of channel correlation on optimal code-rate for finite-length block codes operating over a point-to-point communication link. To this end, we model the transmit-buffer occupancy by states of a Markov chain, where the transition matrix depends on the physical layer parameters. We notice that the occupancy of the transmit-buffer is dominated by the smallest negative eigenvalue of the transition matrix of this corresponding Markov chain. Therefore, the optimal code-rate depends on this large deviation parameter alone. Furthermore, the optimal coding-rate is much smaller than the Shannon capacity of lightly correlated channels. However, when the channel correlation becomes significant, the code-rate exceeds the channel capacity. In the extreme case, as the bit-erasure channel starts mimicing a packet-erasure channel, the code-rate approaches unity.

The results in this dissertation are for erasure channels. However, these results can be generalized to error channels to find the bounds of maximum mutual information rate with servic guarantees, utilizing thresholds on empirical mutual information for retransmission decision. We have also assumed feedback in one transmission opportunity at the transmitter. This assumption can be relaxed by considering parallel systems working in time-multiplexed fashion.

In the second chapter, we characterized the gains of network coding in case of multi-cast when each flow satisfies a service guarantee. We characterized the departure process at each intermediate node and thereby computed the exact buffer-occupancy distribution. We utilized the following facts for queueing analysis at intermediate nodes. The effective bandwidths of independent sources are additive in the routing case. In the coding case, flows are oblivious to one another in the bottleneck link. We also found the optimum physical resources to be allocated in a wireless environment to create a network suitable primarily for coding or routing. We showed that there exists a physical distribution of nodes, where creating a network suitable for coding performs worse than a network designed for routing alone.

In the third chapter, we proposed a distributed algorithm to provide per-flow based service guarantee in a large network. To this end, we employed an approximate "effective capacity" that is indicative of the upper limit on the link-rate, such that the service-guarantees for each flow sharing that link can be supported. In this algorithm, each link maintains a price queue with the arrival being the actual load on the link, and the service being the effective capacity of the link. Further, links can update their effective capacity based on the aggregate dissatisfaction declared by the flows sharing the link, and its current price. On the other hand, sources maintain a dissatisfaction queue, with the arrival being the flow-degradation, and the service being the accepted tolerance. Sources update their rates based on the route-price. We showed that this algorithm achieves the global objective of maximizing the sum-utility of flows while providing their individual service guarantee of maintaining their degradations below the declared acceptable thresholds.

REFERENCES

[1] C. Shannon, N. Petigara, and S. Seshasai, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.

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

[3] F. P. Kelly, "Effective bandwidths at multi-class queues," *Queueing Systems*, vol. 9, pp. 5–15, 1991.

[4] G. Kesidis, J. C. Walrand, and C.-S. Chang, "Effective bandwidths for multiclass markov fluids and other ATM sources," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 424–428, August 1993.

[5] A. Lapidoth and P. Narayan, *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2148–2177, October 1998.

[6] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY: Wiley, 1968.

[7] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., ser. Wiley Series in Telecommunications and Signal Processing. New York, NY: Wiley-Interscience, 2006.

[8] D. J. Costello Jr., J. Hagenauer, H. Imai, and S. B. Wicker, "Applications of error-control coding," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2531–2560, October 1998.

[9] A. Ephremides and B. Hajek, "Information theory and communication networks: an unconsummated union," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2416–2434, October 1998.

[10] R. A. Berry and R. G. Gallager, "Communication over fading channels with delay constraints," *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1135–1149, May 2002.

[11] S. Shakkottai, "Effective capacity and QoS for wireless scheduling," *IEEE Transactions on Automatic Control*, vol. 53, no. 3, pp. 749–761, April 2008.

[12] L. Ying, S. Yang, and R. Srikant, "Optimal delay-throughput tradeoffs in mobile ad hoc networks," *IEEE Transactions on Information Theory*, vol. 54, no. 9, pp. 4119–4143, September 2008.

[13] M. V. Burnashev, "Data transmission over a discrete channel with feedback: random transmission time," *Problems of Information Transmission*, vol. 12, no. 4, pp. 250–265, 1976.

[14] H. Yamamoto and K. Itoh, "Asymptotic performance of a modified Schalkwijk-Barron scheme for channels with noiseless feedback," *IEEE Transactions on Information Theory*, vol. 25, no. 6, pp. 729–733, November 1979.

[15] P. Berlin, B. Nakiboglu, B. Rimoldi, and E. Telatar, "A simple converse of Burnashev's reliability function," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3074–3080, July 2009.

[16] D. Wu and R. Negi, "Effective capacity: a wireless link model for support of quality of service," *IEEE Transactions on Wireless Communications*, vol. 2, no. 4, pp. 630–643, July 2003.

[17] S. Goel and R. Negi, "Analysis of delay statistics for the queued-code," in *IEEE International Conference on Communications*, Dresden, Germany, June 2009, pp. 1–6.

[18] Y.-C. Ko, M.-S. Alouini, and M. K. Simon, "Outage probability of diversity systems over generalized fading channels," *IEEE Transactions on Communications*, vol. 48, no. 11, pp. 1783–1787, November 2000.

[19] L. Li, N. Jindal, and A. Goldsmith, "Outage capacities and optimal power allocation for fading multiple-access channels," *IEEE Transactions on Information Theory*, vol. 51, no. 4, pp. 1326–1347, April 2005.

[20] J. G. Kim and M. M. Krunz, "Bandwidth allocation in wireless networks with guaranteed packet-loss performance," *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, pp. 337–349, June 2000.

[21] L. Liu, P. Parag, J. Tang, W.-Y. Chen, and J.-F. Chamberland, "Resource allocation and quality of service evaluation for wireless communication systems using fluid models," *IEEE Transactions on Information Theory*, vol. 53, no. 5, pp. 1767–1777, May 2007.

[22] W. Wu, A. Arapostathis, and S. Shakkottai, "Optimal power allocation for a time-varying wireless channel under heavy-traffic approximation," *IEEE Transactions on Automatic Control*, vol. 51, no. 4, pp. 580–594, April 2006.

[23] P. Wu and N. Jindal, "Coding versus ARQ in fading channels: how reliable should the PHY be?" in *IEEE Global Telecommunications Conference*, Honolulu, HI, December 2009, pp. 1–6.

[24] H. S. Wang and N. Moayeri, "Finite-state Markov channel – a useful model for

radio communication channels," *IEEE Transactions on Vehicular Technology*, vol. 44, no. 1, pp. 163–171, February 1995.

[25] P. Sadeghi, R. A. Kennedy, P. B. Rapajic, and R. Shams, "Finite-state Markov modeling of fading channels: a survey of principles and applications," *IEEE Signal Processing Magazine*, vol. 25, no. 5, pp. 57–80, September 2008.

[26] L. Wilhelmsson and L. B. Milstein, "On the effect of imperfect interleaving for the Gilbert-Elliott channel," *IEEE Transactions on Communications*, vol. 47, no. 5, pp. 681–688, May 1999.

[27] P. Parag, J.-F. Chamberland, H. D. Pfister, and K. R. Narayanan, "Code rate, queueing behavior and the correlated erasure channel," in *IEEE Information Theory Workshop on Information Theory*, Cairo, Egypt, January 2010, pp. 1–5.

[28] T. Richardson and R. Urbanke, *Modern Coding Theory*. New York, NY: Cambridge University Press.

[29] L. Kleinrock, *Queueing Systems. Volume 1: Theory*. New York, NY: Wiley-Interscience, 1975.

[30] F. Baccelli and P. Bremaud, *Elements of Queueing Theory*, 2nd ed., ser. Stochastic Modelling and Applied Probability. New York, NY: Springer, 2003.

[31] D. Gross, J. F. Shortle, J. M. Thompson, and C. M. Harris, *Fundamentals of Queueing Theory*, 4th ed., ser. Probability and Statistics. Hoboken, NJ: Wiley-Interscience, 2008.

[32] M. F. Neuts, *Structured Stochastic Matrices of M/G/1 Type and Their Applications*, ser. Probability: Pure and Applied. New York, NY: CRC Press,

1989.

[33] G. Latouche and V. Ramaswami, *Introduction to Matrix Analytic Methods in Stochastic Modeling*, ser. ASA-SIAM Series on Statistics and Applied Probability. Philadelphia, PA: Society for Industrial Mathematics, 1987.

[34] B. Hajek, "Birth-and-death processes on the integers with phases and general boundaries," *Journal of Applied Probability*, vol. 19, no. 3, pp. 488–499, September 1982.

[35] J. G. Kemeny, J. L. Snell, and A. W. Knapp, *Denumerable Markov Chains*, 2nd ed., ser. Graduate Texts in Mathematics. New York, NY: Springer-Verlag, 1976.

[36] J. R. Norris, *Markov Chains*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. New York, NY: Cambridge University Press.

[37] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. New York, NY: Cambridge University Press, 1994.

[38] M. F. Neuts, *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*, revised edition ed. Mineola, NY: Dover Publications, 1995.

[39] R. V. Evans, "Geometric distribution in some two-dimensional queuing systems," *Operations Research*, vol. 15, no. 5, pp. 830–846, September-October 1967.

[40] R. Ahlswede and I. Csiszár, "To get a bit of information may be as hard as to get full information," *IEEE Transactions on Information Theory*, vol. 27, no. 4, pp. 398–408, July 1981.

[41] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, October 2003.

[42] S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egner, K. Jain, and L. M. G. M. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1973–1982, June 2005.

[43] T. C. Ho, Y.-H. Chang, and K. J. Han, "On constructive network coding for multiple unicasts," in *44th Allerton Conference on Communication, Control and Computing*, September 2006.

[44] A. Khreishah, C.-C. Wang, and N. B. Shroff, "Optimization based rate control for communication networks with inter-session network coding," in *INFOCOM, 27th IEEE Conference on Computer Communications*, April 2008, pp. 81–85.

[45] A. I. Elwalid and D. Mitra, "Effective bandwidth of general markovian traffic sources and admission control of high speed networks," *IEEE/ACM Transactions on Networking*, vol. 1, no. 3, pp. 329–343, June 1993.

[46] C.-S. Chang, "Stability, queue length, and delay of deterministic and stochastic queueing networks," *IEEE Transactions on Automatic Control*, vol. 39, no. 5, pp. 913–931, May 1994.

[47] ——, "Sample path large deviations and intree networks," *Queueing Systems*, vol. 20, March 1995.

[48] W. Whitt, "Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues," *Telecommunication Systems*, vol. 2, no. 1, pp.

71–107, December 1993.

[49] C.-S. Chang and J. A. Thomas, "Effective bandwidth in high speed digital networks," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 1091–1100, August 1995.

[50] C.-S. Chang and T. Zajic, "Effective bandwidths of departure processes from queues with time varying capacities," in *INFOCOM (3)*, 1995, pp. 1001–1009.

[51] R. Guérin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high speed networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 7, pp. 968–981, September 1991.

[52] G. de Veciana, G. Kesidis, and J. C. Walrand, "Resource management in wide-area ATM networks using effective bandwiths," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 1081–1090, 1995.

[53] L. Liu, P. Parag, J. Tang, W.-Y. Chen, and J.-F. Chamberland, "Resource allocation and quality of service evaluation for wireless communication systems using fluid models," *IEEE Transactions on Information Theory*, vol. 53, no. 5, pp. 1767–1777, May 2007.

[54] S. Bhadra and S. Shakkottai, "Looking at large networks: Coding vs. Queueing," in *INFOCOM, 25th IEEE Conference on Computer Communications*, April 2006, pp. 1–12.

[55] A. Eryilmaz, A. Ozdaglar, and M. Médard, "On delay performance gains from network coding," in *CISS, 40th Annual Conference on Information Sciences and Systems*, March 2006, pp. 864–870.

[56] B. Shrader and A. Ephremides, "A queueing model for random linear coding," in *MILCOM, Military Communications Conference*, October 2007, pp. 1–7.

[57] ——, "On the queueing delay of a multicast erasure channel," in *ITW, IEEE Information Theory Workshop*, October 2006, pp. 423–427.

[58] J. K. Sundararajan, D. Shah, and M. Médard, "ARQ for network coding," in *ISIT, IEEE International Symposium on Information Theory*, July 2008, pp. 1651 – 1655.

[59] P. Chaporkar and A. Proutiere, "Adaptive network coding and scheduling for maximizing throughput in wireless networks," in *MobiCom, 13th Annual ACM International Conference on Mobile Computing and Networking*, September 2007, pp. 135 – 146.

[60] F. Xue, C.-H. Liu, and S. Sandhu, "MAC-layer and PHY-layer network coding for two-way relaying: Achievable rate regions and opportunistic scheudling," in *44th Annual Allerton Conference*, September 2007, pp. 396–402.

[61] X. He and A. Yener, "On the energy-delay trade-off of a two-way relay network," in *CISS, 42nd Annual Conference on Information Sciences and Systems*, March 2008, pp. 865–870.

[62] C.-H. Liu and F. Xue, "Network coding for two-way relaying: Rate region, sum rate and opportunistic scheduling," in *ICC, IEEE International Conference on Communications*, May 2008, pp. 1044–1049.

[63] Y. E. Sagduyu and A. Ephremides, "Cross-layer optimization of MAC and network coding in wireless queueing tandem networks," *IEEE Transactions on Information Theory*, vol. 54, no. 2, pp. 554–571, February 2008.

[64] A. Eryilmaz and D. S. Lun, "Control for inter-session network coding," in *NetCod, Workshop on Network Coding, Theory and Applications*, January 2007.

[65] C. Rago, P. Willett, and Y. Bar-Shalom, "Censoring sensors: A low-communication-rate scheme for distributed detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 2, pp. 554–568, April 1996.

[66] R. J. Gibbens and P. J. Hunt, "Effective bandwidths for the multi-type UAS channel," *Queueing Systems*, vol. 9, pp. 17–28, 1991.

[67] L. Liu, P. Parag, and J.-F. Chamberland, "Quality of service analysis for wireless user-cooperation networks," *IEEE Transactions on Information Theory*, vol. 53, no. 10, pp. 3833–3842, October 2007.

[68] D. P. Kroese and W. R. W. Scheinhardt, "Joint distributions for interacting fluid queues," *Queueing Systems*, vol. 37, no. 1-3, pp. 99–139, March 2001.

[69] S. Aalto and W. R. W. Scheinhardt, "Tandem fluid queues fed by homogeneous onoff sources," *Operations Research Letters*, vol. 27, no. 2, pp. 73–82, September 2000.

[70] N. Barbot and B. Sericola, "Exact stationary solution to tandem fluid queues," *International Journal of Simulation: Systems, Science & Technology*, vol. 4, no. 5-6, 2003.

[71] D. Mitra, "Stochastic theory of a fluid model of producers and consumers coupled by a buffer," *Advances in Applied Probability*, vol. 20, pp. 646–676, 1993.

[72] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell System Technical Journal*, vol. 61, no. 8,

pp. 1871–1894, October 1982.

[73] J. Virtamo and I. Norros, "Fluid queue driven by an M/M/1 queue," *Queueing Systems*, vol. 16, no. 3-4, pp. 373–386, September 1994.

[74] S. Aalto, "Characterization of the output rate process for a markovian storage model," *Journal of Applied Probability*, vol. 35, no. 1, pp. 184–199, March 1998.

[75] ——, "Output from an A-M-S type fluid queue," *Fundamental Role of Teletraffic in the Evolution of Telecommunication Networks*, pp. 421–430, 1997.

[76] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ: Prentice Hall PTR, 2001.

[77] T. E. Stern and A. I. Elwalid, "Analysis of separable markov-modulated rate models for information-handling systems," *Advances in Applied Probability*, vol. 23, pp. 105–139, 1991.

[78] S. Asmussen, *Applied Probability and Queues*, 2nd ed., ser. Stochastic Modelling and Applied Probability, Vol. 51. New York, NY: Springer, 2003.

[79] M. Reed and B. Simon, *Functional Analysis: Volume IV*, revised edition ed., ser. Methods of Modern Mathematical Physics. San Diego, CA: Academic Press, 1980.

[80] R. Carmona and J. Lacroix, *Spectral Theory of Random Schrödinger Operators*, ser. Probability and Its Applications. Cambridge, MA: Birkhäuser, 1990.

[81] M. Fabian, P. Habala, P. Hájek, V. M. Santalucía, J. Pelant, and V. Zizler, *Functional Analysis and Infinite-Dimensional Geometry*, ser. CMS Books in Mathematics. New York, NY: Springer-Verlag, 1990.

[82] R. V. Churchill and J. W. Brown, *Complex Variables and Applications*, 5th ed., ser. Mcgraw-Hill Series in Higher Mathematics. New York, NY: McGraw-Hill, 1989.

[83] M. Rubinovitch, "The ouput of a buffered data communication system," *Stochastic Processes and their Applications*, vol. 1, pp. 375–382, 1973.

[84] O. Kella and W. Whitt, "A storage model with a two-state random environment," *Operations Research*, vol. 40, no. S2, pp. 257–262, May-June 1992.

[85] O. J. Boxma and V. Dumas, "The busy period in the fluid queue," in *1998 ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*. ACM, 1998, pp. 100–110.

[86] V. G. Kulkarni, "Effective bandwidths for Markov regenerative sources," *Queueing Systems*, vol. 24, no. 1-4, pp. 137–153, March 1996.

[87] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, January-February 1997.

[88] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.

[89] R. Srikant, *The Mathematics of Internet Congestion Control*, ser. Systems & Control: Foundations & Applications. New York, NY: Birkhăuser, 2004.

[90] S. Shakkottai and R. Srikant, *Network Optimization and Control*, ser. Foundations and Trends in Networking. Hanover, MA: Now Publishers, 2008, vol. 2, no. 3.

[91] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, January 2007.

[92] W. R. Stevens, *TCP/IP Illustrated, Vol. 1: The Protocols*, 1st ed., ser. Addison-Wesley Professional Computing Series. Boston, MA: Addison-Wesley Professional, 1993.

[93] L. L. Peterson and B. S. Davie, *Computer Networks: A Systems Approach*, 2nd ed. Burlington, MA: Morgan-Kaufmann, 2000.

[94] S. H. Low and D. E. Lapsley, "Optimization flow control–I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, pp. 861–874, December 1999.

[95] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control," in *Proceedings of the IFAC World Congress*, Barcelona, Spain, 2002.

[96] T. Kelly, "Scalable TCP: improving performance in highspeed wide area networks," *ACM SIGCOMM Computer Communication Review*, vol. 33, pp. 83–91, April 2003.

[97] D. X. Wei, C. Jin, S. H. Low, and S. Hegde, "Fast tcp: Motivation, architecture, algorithms, performance," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1246–1259, December 2006.

[98] S. Liu, T. Başar, and R. Srikant, "TCP-Illinois: A loss- and delay-based congestion control algorithm for high-speed networks," *Performance Evaluation*, vol. 65, pp. 417–440, June 2008.

[99] J. He, M. Suchara, J. Rexford, and M. Chiang, "Rethinking internet traffic management: From multiple decompositions to a practical protocol," in *Proceedings of CoNEXT*, New York, December 2007.

[100] D. P. Palomar and M. Chiang, "A tutorial on decomposition methods for network utility maximization," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, August 2006.

[101] H. K. Khalil, *Nonlinear Systems*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 1996.

[102] P. Parag, S. Shakkottai, and J.-F. Chamberland, "Value-aware resource allocation for service guarantees in networks," in *Proceedings of INFOCOM*, March 2010, pp. 2730–2738.

[103] W. N. Kang, F. P. Kelly, N. H. Lee, and R. J. Williams, "State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy," *The Annals of Applied Probability*, vol. 19, no. 5, pp. 1719–1780, May 2009.

[104] "Internet2," 2009. [Online]. Available: http://www.internet2.edu/network/

VITA

Parimal Parag received a PhD in Electrical Engineering from Texas A&M University in December 2011. He finished his Dual Degree (B. Tech. and M. Tech.) in Electrical Engineering, with a specialization in communication systems, from Indian Institute of Technology Madras in 2004. He was at Stanford University and Los Alamos National Laboratory, in autumn of 2010 and summer of 2007, respectively. He is currently working at ASSIA Inc. His research interests include application of stochastic and optimization methods to communication networks. His other research interests lie in the areas of statistical signal processing, queueing theory, information theory, combinatorics, and probability theory.

He was a silver medalist at Indian Institute of Technology Madras. He is a recipient of prestigious Indian National Talent Search Scholarship and a final round participant of Indian National Mathematical Olympiad.

Address:
333 Twin Dolphin Dr, # 300
Redwood City, CA 94065-1449

The typist for this thesis was Parimal Parag
email: parimal@tamu.edu.