


On the tradeoff of average delay, average service cost, and average utility for single server queues with monotone policies

A THESIS
SUBMITTED FOR THE DEGREE OF
Doctor of Philosophy
IN THE FACULTY OF ENGINEERING

by

Vineeth Bala Sukumaran



Department of Electrical Communication Engineering
Indian Institute of Science
BANGALORE – 560 012

June 2013

To my family.

Acknowledgements	vi
Abstract	viii
Glossary of Abbreviations	x
Glossary of Conventions	xi
Glossary of Notation	xii
1 Introduction	1
1.1 Introduction to the queueing models and the tradeoff problem	3
1.1.1 A continuous time state dependent M/M/1 model	3
1.1.2 A discrete time queueing model	5
1.2 Literature survey	8
1.3 Overview of the thesis and contributions	13
2 On the tradeoff of average queue length, average service cost, and average utility for the state dependent M/M/1 queue: Part I	21
2.1 Introduction	21
2.1.1 Methodology	22
2.1.2 System model	22
2.1.3 A motivational example	25
2.1.4 Overview	26
2.2 Problem formulation	27
2.2.1 Problem	28
2.3 Analysis of FINITE- μ CHOICE	31
2.3.1 Asymptotic lower bounds	33
2.3.2 Asymptotic characterization of $Q_M^*(c_c)$	40
2.3.3 Asymptotic characterization of order optimal admissible policies	47
2.3.4 Numerical examples	56
2.3.5 An Application	65
2.4 Conclusions	67
Appendix	70

3	On the tradeoff of average queue length, average service cost, and average utility for the state dependent M/M/1 queue: Part II	73
3.1	Introduction	73
3.1.1	Overview	74
3.2	Analysis of INTERVAL- μ CHOICE	75
3.2.1	Asymptotic lower bounds	77
3.2.2	Asymptotic behaviour of the tradeoff curve	87
3.2.3	Asymptotic characterization of order-optimal admissible policies	91
3.2.4	Tradeoff problems which are similar to INTERVAL- μ CHOICE	94
3.3	Analysis of INTERVAL- $\lambda\mu$ CHOICE	102
3.3.1	Asymptotic lower bound	103
3.3.2	Asymptotic behaviour of the tradeoff curve	106
3.3.3	Tradeoff problems which are similar to INTERVAL- $\lambda\mu$ CHOICE	110
3.4	Conclusions	112
	Appendices	114
4	On the tradeoff of average queue length and average service cost for discrete time single server queues	117
4.1	Introduction	117
4.1.1	Methodology	119
4.1.2	System model - Integer valued queue evolution	119
4.1.3	System model - Real valued queue evolution	121
4.1.4	Overview	121
4.2	Problem formulation for I-model and R-model	122
4.2.1	A constrained Markov decision process formulation	122
4.3	Asymptotic bounds for I-model	123
4.3.1	An unconstrained MDP formulation	123
4.3.2	The tradeoff problem	125
4.3.3	Asymptotic analysis of TRADEOFF - Preliminaries	127
4.3.4	Asymptotic characterization of $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$	130
4.3.5	Asymptotic characterization of admissible policies for TRADEOFF	140
4.3.6	An asymptotic lower bound for the tradeoff problem (4.4)	142
4.3.7	Asymptotic lower bounds for ergodic $(A[m], m \geq 1)$	145
4.4	Asymptotic bounds for R-model	150
4.4.1	Asymptotic lower bound	151
4.4.2	R-model with a piecewise linear cost function	154
4.5	Conclusions	155
	Appendices	157
5	On the tradeoff of average power and average delay for fading wireless links	177
5.1	Introduction	177
5.1.1	Methodology	179
5.1.2	System model - Integer valued queue length evolution	179
5.1.3	System model - Real valued queue length evolution	181
5.1.4	An example	181
5.1.5	Overview	182
5.2	Problem formulation for I-model and R-model	183
5.3	Asymptotic bounds for I-model	186
5.4	Asymptotic bounds for R-model	191
5.5	A numerical example	192

5.6	Asymptotic lower bounds for ergodic $(A[m])$ and $(H[m])$	197
5.7	Queueing models with admission control	200
5.7.1	System model	200
5.7.2	Problem formulation for I-model-U and R-model-U	201
5.7.3	Asymptotic bounds	203
5.7.4	Discussion	208
5.8	Single hop networks	210
5.8.1	System model	210
5.8.2	Problem formulation	211
5.8.3	Asymptotic lower bound	212
5.8.4	Discussion	214
5.9	Conclusions	215
	Appendices	217
6	On the tradeoff of average error rate and average delay for point to point links	223
6.1	Introduction	223
6.1.1	System model	224
6.1.2	Related work	228
6.1.3	Overview	229
6.2	Problem formulation	230
6.2.1	The error cost $c_s(s, \tau)$	232
6.3	Asymptotic analysis for R-model-A	234
6.3.1	Problem Statement	234
6.3.2	Asymptotic analysis	234
6.3.3	Integer valued queue evolution	241
6.4	Asymptotic analysis for R-model-B	243
6.4.1	Problem Statement	243
6.4.2	Asymptotic analysis	245
6.5	Conclusions	251
	Appendices	253
7	Conclusions and Scope for Future Work	261
7.1	Scope for future work	262
7.1.1	Order optimality of policies obtained from a fluid model	263
7.1.2	Extensions to systems with service time control	264
	Bibliography	267

Acknowledgements

I am very much grateful to my wonderful family, without which this thesis would not have been possible.

I am indebted to Prof. Utpal Mukherji for his generous advice, guidance, and careful review of my work. I have benefitted a lot, both professionally and personally, by being his student.

I am very grateful to Prof. Anurag Kumar and Prof. Rajesh Sundaresan for their help and insightful comments regarding the work done during my PhD as well as for some of the nicest courses I have attended during my stay at IISc. I also thank Prof. Vinod Sharma and Prof. Vijay Kumar for their insightful comments and suggestions regarding my PhD work.

Akhil, Arjun, Birenjith, Divya, Deepak (L), Deepak (S), Jithin, Naveen, Nidhin, Sojan, Sreeram, Venu, and Vinodh are wonderful friends and have made my stay here at IISc memorable. I am grateful to Birenjith, Naveen, Venu, and Vinodh for their generous advice, great technical and even greater non-technical discussions, and being the sink nodes for my non-stop griping. My lab and department colleagues Arpan, Ashok, Avijeeth, Avijit, Avishek, Bharath, Chandramani, Deekshith, Jobin, Karthik, Krishna Chaitanya, Prakash, Prasad, Lalitha, Manoj, Prem, Rahul, Sayee, Srinidhi, Srinivasan, and Venkatesh have been very generous with their support along the way. I also thank Ashwin, Jaideep Sir, Manoj, Manu, Parameshwaran, Ravi, Rahul, Simil, Sunil, Sunilkumar Sir, Sandeep, and Vaisakh for the unforgettable times in A mess and C mess.

I am grateful to Govind and Nikhil for their support for all these years. I sincerely thank Hari, Sreeram, Sunil, Shri Indran Gurukkal, Shri Kumar, Shri Manjunath, Rahul, Guruprasad, Senthil, Srinidhi, Suresh, Midhun, Mahesh, Sreevalsa, Ravi, Mohanan Sir, and Naveen for encouraging me do things I would have never thought of doing before. I am grateful for the great food and hospitality of Divya, Birenjith, Lekha chechi, and Ajayan Sir. I also thank the members of SIMA for their companionship.

I owe a good deal to the staff of A and C messes, Prakurthi, Nesara, J.B., F.C., and IISc gymkhana for keeping me well fed and healthy throughout these years. I am very thankful for the help given by the staff of ERNET office, Network office, and DRDO-IISc office, especially Boregowda, Mrs. Chandrika, Mahesh, Priyanka, and Savitha. I am also grateful to the office staff of the Department of Electrical Communication Engineering, especially Mr. Srinivasa Murthy and Mr. Nagaraj, for their help and support. I also acknowledge the generous financial support from Indian Institute of Science, Defence Research and Development Organization (GoI), and the Ministry of Human Resource Development (GoI).

Abstract

In this thesis, we study the optimal tradeoff of average delay, average service cost, and average utility for single server queueing models, with and without admission control. The continuous time and discrete time queueing models that we consider are motivated by cross-layer models for noisy point-to-point links, with random packet arrivals. We study the above tradeoff problem for a class of admissible policies, which are monotone and stationary.

The solutions that we obtain for the above tradeoff problem are asymptotic in nature. For example, suppose we are interested in minimizing the average delay of packets, subject to a constraint on the average service cost of serving the packets. It is intuitive that to keep the queue *stable* the time average service rate of packets has to equal the time average arrival rate of packets. This in turn implies that queue stability requires a positive minimum average service cost expenditure. We obtain asymptotic bounds on the minimum average delay in the asymptotic regime \mathfrak{R} where the average service cost constraint is a small positive V more than the above minimum average service cost required for queue stability. We note that such asymptotic bounds can be used to obtain a first order characterization of the tradeoff curve, and are useful in identifying *good families* of scheduling policies, such as *buffer partitioning* policies.

In this thesis, we obtain asymptotic lower bounds on the minimum average delay in the regime \mathfrak{R} , for the cases for which lower bounds were previously not known, for admissible policies. The asymptotic characterization of the minimum average delay for admissible policies, for both continuous time and discrete time models, is obtained via new geometric bounds on the stationary probability distribution of the queue length, in the regime \mathfrak{R} . The restriction to admissible policies, also enables us to obtain an intuitive explanation for the behaviour of the asymptotic lower bounds, using the above geometric bounds on the stationary probability distribution of the queue length. We observe that the *shape* of the stationary probability distribution, in the regime \mathfrak{R} , determines the form of the asymptotic behaviour.

It is common practice to approximate a queueing model, where the queue length evolution is on the non-negative integers, with a queueing model where the queue length evolution is on the non-negative real numbers and the service cost function being strictly convex, for analytical tractability. We compare the asymptotic bounds which are obtained for the approximate real valued queue evolution model with that of the original integer valued queue evolution model. We observe that for some cases the average delay does not grow to infinity, in the regime \mathfrak{R} , although the real valued approximate queueing model, with a strictly convex cost function, suggests that the average delay should grow without bound in the regime \mathfrak{R} . In other cases where the average delay does grow to infinity in the regime \mathfrak{R} , our results illustrate that the approximate model strictly underestimates the behaviour of the tradeoff for the original model unless the service cost function is modelled as the piecewise linear lower convex envelope of the service cost function for the original integer valued queueing model.

The geometric bounds on the stationary probability distribution of the queue length also lead to asymptotic bounds on any optimal admissible policy, in the regime \mathfrak{R} . The asymptotic order bounds are independent of the exact service cost function, and are not available in previous work. For buffer partitioning policies, the bounds also show how buffer partitions have to scale with V .

We then apply the above asymptotic lower bounds to the motivating applications, discussed above. We develop geometric bounds on the stationary probability distribution for admissible policies to analyse the tradeoff problem in other scenarios, such as: (i) queueing models for N-user single hop communication networks, (ii) queueing models with non-convex service cost functions, and (iii) queueing models with general holding costs.

Glossary of Abbreviations

IID independent and identically distributed

DMC discrete memoryless channel

TPM transition probability matrix

MC or DTMC discrete time Markov chain

CTMC continuous time Markov chain

EMC embedded Markov chain

SMP semi Markov process

MDP Markov decision process

SMDP semi Markov decision process

CMDP constrained Markov decision process

ACOE average cost optimality equation

EXH exhaustive service policy

Glossary of Conventions

- An increasing (decreasing) function is a non-decreasing (non-increasing) function. The property of strictly increasing or decreasing is explicitly stated with the qualifier strict.
- Random variables are denoted by capital letters.
- Realizations of random variables are denoted by the corresponding small letters or by specifying a sample point/path ω . For example, X and x or $X(\omega)$ respectively.
- A random or deterministic vector is typeset in **bold** face.
- Derivatives will be explicitly shown, e.g. $\frac{d}{dx}$, rather than by using primes.
- All logarithms are natural logarithms, unless specified.

Glossary of Notation

\mathbb{R}_+	the set of all non-negative real numbers
\mathbb{Z}_+	the set of all non-negative integers
\mathbb{I}_A	indicator function for the event A
\mathbb{E}	expectation with respect to a distribution which is clear from the context
\mathbb{E}_p	expectation with respect to the distribution p
$Pr\{E\}$	probability of event E
t	continuous time variable
T	upper limit for $t, \in \mathbb{R}_+$
n	discrete time index
N	upper limit for $n, \in \mathbb{Z}_+$
m	index variable for special embedded epochs (such as decision instants)
M	upper limit for $m, \in \mathbb{Z}_+$
i	index variable for customers
I	upper limit for $i, \in \mathbb{Z}_+$
$P_A(a)$	distribution of the random variable A
$P_{A B}(a b)$	the distribution of A being a conditioned on B being b
\mathcal{O}	$f(x)$ is $\mathcal{O}(g(x))$ if there exists a $c > 0$ such that $\lim_{x \rightarrow 0} \frac{f(x)}{g(x)} \leq c; f(x), g(x) \geq 0$
o	$f(x)$ is $o(g(x))$ if for every $c > 0, \lim_{x \rightarrow 0} \frac{f(x)}{g(x)} \leq c; f(x), g(x) \geq 0$
Ω	$f(x)$ is $\Omega(g(x))$ if there exists a $c > 0$ such that $\lim_{x \rightarrow 0} \frac{f(x)}{g(x)} \geq c; f(x), g(x) \geq 0$
ω	$f(x)$ is $\omega(g(x))$ if for every $c > 0, \lim_{x \rightarrow 0} \frac{f(x)}{g(x)} \geq c$ or $g(x) = o(f(x)); f(x), g(x) \geq 0$
Θ	$f(x)$ is $\Theta(g(x))$ if $f(x) = \mathcal{O}(g(x))$ and $f(x) = \Omega(g(x)); f(x), g(x) \geq 0$

CHAPTER 1

Introduction

In this thesis, we study the optimal tradeoff of average delay, average service cost, and average utility for single server queueing models, with and without admission control. The continuous time and discrete time queueing models that we consider are motivated by cross-layer models for noisy point-to-point links. The features that we model are: (i) random packet arrivals, (ii) control of service and/or arrival rates, (iii) packet service costs and/or utility, and (iv) fading at slow and fast time scales. Our objective is to characterize the minimum average delay of the packets, under an upper bound constraint on the average service cost and/or a lower bound constraint on the average throughput for systems with admission control. We are also motivated by the problem of characterizing the minimum average delay of randomly arriving message symbols which are transmitted over a noisy point-to-point link with no admission control, under an upper bound constraint on the average error rate of the message symbols. Such tradeoff problems arise in the study of cross layer scheduling algorithms for wireless communication networks [45] or in the study of processor speed scaling [17].

In this thesis, we consider the performance of scheduling algorithms which optimally trade off average delay with other performance measures, as in [7], [45], and [29]. Related problems include the design of cross layer scheduling algorithms for: (a) *stabilizing* a communication network, as in [73], [74], [41], [26], [15], [38], [63], [66], [64], and [65], or (b) minimizing a delay measure, as in [79], [40], and [21].

We study the above tradeoff problem for a class of monotone policies, which we call admissible policies. Monotone policies are stationary policies, i.e., the service rate (and the number of packets

admitted, for queues with admission control) at a time is a function¹ only of the current state of the queue² rather than a function of the whole history of evolution as well as the current state of the queue. For monotone policies, the expected service rate for a queue length is non-decreasing as a function of the queue length. Intuitively, if the average delay is to be minimized subject to a constraint on the average service cost, for deterministic stationary policies, as the queue length increases, the service rate should also increase. We note that intuitively although the service rate should increase, the amount by which the service rate increases depends on the corresponding increase in the service cost. So in practice, monotone policies are usually used. Furthermore, in certain cases it can be shown that the *optimal* policy for the above tradeoff problem is in fact monotone³. Motivated by this reason, as well as the above intuition, we consider the tradeoff problem for the class of admissible policies only. The class of admissible policies is a subset of the class of monotone policies, possessing some additional properties, which makes their analysis more amenable.

We consider several variations of the above tradeoff problem. The solutions to the tradeoff problems that we address in this thesis are of the following form. Consider an example of an infinite buffer queueing model with service rate control, but no admission control. We are interested in minimizing the average delay subject to a constraint on the average service cost. It is intuitive that to keep the queue *stable* the time average service rate of packets has to equal the time average arrival rate of packets. This in turn implies that queue stability requires a positive minimum average service cost expenditure. We obtain asymptotic bounds on the minimum average delay in the asymptotic regime \mathfrak{R} where the average service cost constraint is a small positive V more than the above minimum average service cost required for queue stability. We note that such asymptotic bounds can be used to obtain a first order characterization of the tradeoff curve. Furthermore, such bounds can also be used to identify good families of scheduling policies, as in [7].

We note that asymptotic upper bounds (see [45]) as well as asymptotic lower bounds (see [7] and [43]) on the minimum average delay in the regime \mathfrak{R} are available for a variety of queueing models. However, asymptotic lower bounds are not known in many cases. We obtain asymptotic lower bounds for the minimum average delay in the regime \mathfrak{R} , for these cases, for admissible policies. Additionally, we also obtain *asymptotic bounds* on the *structure* of admissible policies which achieve the above asymptotic lower bounds. The method by which we derive these bounds, which is different from previous approaches, also leads to *geometric* bounds on the stationary probability of the queue length, in the regime \mathfrak{R} . These bounds have the added advantage of directly providing intuition for the behaviour of asymptotic lower bounds in the regime \mathfrak{R} .

Using the above asymptotic lower bounds on the minimum average delay and the already available

¹this function could be randomized.

²e.g., the current queue length or the current queue length and an auxiliary state variable such as the fade state.

³The monotonicity property is obtained using a Markov decision theoretic formulation of the tradeoff problem as in [70], [25], or [7]

asymptotic upper bounds we obtain a complete asymptotic characterization of the tradeoff between average delay and average service cost for several single server queueing models, in the regime \mathfrak{R} , for admissible policies. We start with an informal introduction to the queueing models and the tradeoff problems considered in this thesis.

1.1 Introduction to the queueing models and the tradeoff problem

We note that the primary features of the point-to-point links are: (i) random packet arrivals, (ii) control of service and/or arrival rates, (iii) packet service costs and/or utility, and (iv) fading at slow and fast time scales. These features are captured by two single server controlled queueing models in this thesis. Initial insights into the asymptotic behaviour of the tradeoff are obtained by studying a continuous time queueing model with exponential interarrival times and exponential service requirements, in Chapters 2 and 3. These insights are used in characterizing the asymptotic behaviour of the tradeoff for discrete time models in the following chapters.

For the purposes of this introductory discussion, a queue is defined to be stable (or more precisely, mean rate stable as in [45]) under a *policy* if the time average service rate equals the time average arrival rate, although in later chapters we use stronger notions of stability.

In the asymptotic regime \mathfrak{R} , asymptotic upper and lower bounds on the average delay can be obtained from asymptotic upper and lower bounds on the average queue length by applying Little's law with lower and upper bounds on the average throughput respectively. Therefore, we focus on the average queue length instead of average delay throughout this thesis.

1.1.1 A continuous time state dependent M/M/1 model

The state dependent M/M/1 queueing model is a birth death process with state being the queue length, and with state dependent birth (or arrival) rates and death (or service) rates, as shown in the transition diagram in Figure 1.1. We note that the control policy γ is the choice of the arrival rates $(\lambda(q), q \in \mathbb{Z}_+)$ and the service rates $(\mu(q), q \in \mathbb{Z}_+ \setminus \{0\})$, as a function of the queue length q , from sets \mathcal{X}_λ and \mathcal{X}_μ respectively.

The average queue length $\bar{Q}(\gamma)$ for a particular policy γ is the time average of the expectation of the queue length $Q(t)$, where $Q(t)$ is the state of the birth death process at time t under γ . We assume that utility is accrued at the rate of $u(\lambda(Q(t)))$ and service cost is incurred at the rate of $c(\mu(Q(t)))$, where $u(\lambda)$ is a non-decreasing concave function of λ and $c(\mu)$ is a non-decreasing convex function of μ . For the policy γ , the average utility rate $\bar{U}(\gamma)$ and average service cost rate $\bar{C}(\gamma)$ are defined as the time averages of the expectation of the utility rate $u(\lambda(Q(t)))$ and the expectation of the service cost rate $c(\mu(Q(t)))$ respectively. The general tradeoff problem that we consider is the minimization of $\bar{Q}(\gamma)$, subject to a lower bound constraint u_c on $\bar{U}(\gamma)$, and an upper




Figure 1.1: Transition diagram for the continuous birth-death queueing model

bound constraint c_c on $\bar{C}(\gamma)$, over all policies γ , i.e.,

$$\begin{aligned} & \underset{\gamma}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && \bar{U}(\gamma) \geq u_c \text{ and } \bar{C}(\gamma) \leq c_c. \end{aligned} \quad (1.1)$$

In the following, problems in the above constrained optimization form are called constrained tradeoff problems. We note that one way to analyse the above constrained tradeoff problem is to consider its unconstrained Lagrange dual, the dual function of which is as follows:

$$\underset{\gamma}{\text{minimize}} \quad \bar{Q}(\gamma) + \beta_1(\bar{C}(\gamma) - c_c) - \beta_2(\bar{U}(\gamma) - u_c), \quad (1.2)$$

where β_1 and β_2 are non-negative Lagrange multipliers. In the following, such unconstrained dual problems are called unconstrained tradeoff problems.

We primarily consider the case where $\lambda(q)$ is fixed to be a $\lambda \in \mathcal{X}_\lambda$ such that $u(\lambda) \geq u_c$, so that the tradeoff problem reduces to

$$\begin{aligned} & \underset{\gamma}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && \bar{C}(\gamma) \leq c_c. \end{aligned} \quad (1.3)$$

We note that as in the case of (1.1), we have the following unconstrained dual function:

$$\underset{\gamma}{\text{minimize}} \quad \bar{Q}(\gamma) + \beta_1(\bar{C}(\gamma) - c_c), \quad (1.4)$$

where $\beta_1 \geq 0$.

The optimal values of (1.1) and (1.3), as a function of their respective constraints, are referred to as the tradeoff curve in the following discussion. We note that in the context of communication networks, the function $u(\cdot)$ is usually assumed to be linear, so that $\bar{U}(\gamma)$ is the average throughput.

But for other applications, such as those in [5], $u(\cdot)$ could be a strictly concave function.

The optimization problem (1.2) and its variants (such as (1.4)) have been formulated as Markov decision problems (MDP) and analyzed by many authors, e.g. [5], [70], and [25]. They show that there exists a monotone optimal policy $\gamma^*(\beta_1, \beta_2)$ for (1.2) ($\gamma^*(\beta_1, \beta_2)$ is such that $\lambda(q)$ is a non-increasing function and $\mu(q)$ is a non-decreasing function of q). Also, from [35], we have that if $c_c = \bar{C}(\gamma^*(\beta_1, \beta_2))$ and $u_c = \bar{U}(\gamma^*(\beta_1, \beta_2))$, then $\gamma^*(\beta_1, \beta_2)$ is also optimal for (1.1). Therefore, for at least such values of c_c and u_c there exist monotone optimal policies for (1.1). This motivates us to consider (1.1) for a class of admissible policies, which are monotone.

The asymptotic regime \mathfrak{R} : We consider the tradeoff problems (1.1) and (1.3) in the asymptotic regime \mathfrak{R} where the average service cost constraint c_c is arbitrarily close to the minimum average service cost required for stability. It turns out that the minimum average service cost required for stability is $c(u^{-1}(u_c))$ (where the inverse $u^{-1}(\cdot)$ of $u(\cdot)$ is assumed to exist) for (1.1) and $c(\lambda)$ for (1.3). Therefore, the asymptotic regimes \mathfrak{R} for problems (1.1) and (1.3) are defined as the regime in which $c_c \downarrow c(u^{-1}(u_c))$ and $c_c \downarrow c(\lambda)$ respectively.

1.1.2 A discrete time queueing model

In Chapters 4, 5, and 6, we consider discrete time single server queueing models with random batch arrivals and batch service. We now introduce a general form of this model, shown in Figure 1.2. In each slot $m \geq 1$, we assume that a random number $R[m]$ of packets arrives into the system,




Figure 1.2: The discrete time single server queueing model with a single queue

with an arrival rate of λ per slot. In the m^{th} slot, $A[m] \leq R[m]$ arrivals are admitted into the queue, which is assumed to have infinite buffer space. We assume that there is a random process ($H[m]$) which models the environment in which the queue is evolving, e.g., this could be the fade state for a point-to-point wireless link, which stays constant for the duration of a slot. We assume that $H[m]$ is known at the start of every slot m . The queue length, at the slot boundaries, evolves

according to the evolution equation:

$$Q[m] = Q[m - 1] - S[m] + A[m], m \geq 1,$$

where $Q[0] = q_0$, $S[m] \leq \min(S_{max}, Q[m - 1])$, and $A[m] \leq R[m]$. We note that $S[m]$ and $A[m]$ are the control variables. The service batch size $S[m]$ in slot m is assumed to be chosen as a randomized function of the history

$$\sigma[m] = (q_0, H[1], S[1], R[1], Q[1], H[2], S[2], R[2], Q[2], \dots, Q[m - 2], H[m - 1]),$$

and the current queue length $Q[m - 1]$ and fade state $H[m]$. while the arrival batch size $A[m]$ in slot m is assumed to be chosen as a randomized function of the history $\sigma[m]$, $Q[m - 1]$, $H[m]$, and the current number of arrivals $R[m]$. The choice of the sequence $((S[m], A[m]), m \geq 1)$ constitutes the policy γ for this discrete time model.

The average queue length $\bar{Q}(\gamma)$ for a policy γ is defined to be the time average of the expectation of the queue length $Q[m]$. We assume that service cost is incurred at the rate of $P(H[m], S[m])$ in slot m , where $P(h, s)$ is a service cost function which is assumed to be non-decreasing and convex in s for every h . We note that the cost function $P(h, s)$ models the cost incurred in transmission of packets, e.g., $P(h, s)$ could be the expected number of packets that are received in error when a batch of s packets are jointly encoded and transmitted when the environment state is h , or $P(h, s)$ could be the power expended in transmission of s packets when the fade state is h . The average service cost $\bar{P}(\gamma)$ for γ is defined as the time average of the expectation of the service cost $P(H[m], S[m])$.

We define the average throughput $\bar{A}(\gamma)$ for γ as the time average of the expectation of the admitted arrival batch size $A[m]$. The performance measure that we are interested in is the utility of the average throughput, $u(\bar{A}(\gamma))$, achieved by γ , where $u(\cdot)$ is a non-decreasing and concave utility function. The general tradeoff problem that we consider is the minimization of $\bar{Q}(\gamma)$ subject to a lower bound constraint u_c on $u(\bar{A}(\gamma))$ and an upper bound constraint P_c on $\bar{P}(\gamma)$ over all policies γ , i.e.,

$$\begin{aligned} & \underset{\gamma}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && u(\bar{A}(\gamma)) \geq u_c \text{ and } \bar{P}(\gamma) \leq P_c. \end{aligned} \tag{1.5}$$

As for the state dependent M/M/1 model, we have the following unconstrained dual function:

$$\underset{\gamma}{\text{minimize}} \quad \bar{Q}(\gamma) + \beta_1 (\bar{P}(\gamma) - P_c) - \beta_2 (u(\bar{A}(\gamma)) - u_c), \tag{1.6}$$

where β_1 and β_2 are non-negative.

As for the state dependent M/M/1 model, we primarily consider the case where $A[m] = R[m]$, i.e., with no admission control. Then the tradeoff problem is:

$$\begin{aligned} & \underset{\gamma}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && \bar{P}(\gamma) \leq P_c, \end{aligned} \tag{1.7}$$

where we have assumed that u_c is such that $\lambda \geq u^{-1}(u_c)$ (where the inverse $u^{-1}(\cdot)$ of $u(\cdot)$ is assumed to exist), so that for stable policies the utility constraint is satisfied. Similar to (1.6), we have the unconstrained dual function:

$$\underset{\gamma}{\text{minimize}} \quad \bar{Q}(\gamma) + \beta_1 (\bar{P}(\gamma) - P_c), \tag{1.8}$$

where $\beta_1 \geq 0$. The optimal values of (1.5) and (1.7) as a function of their respective constraints are referred to as the tradeoff curve in the following discussion.

As for the state dependent M/M/1 model, for the discrete time model, again using a MDP formulation, it is possible to show that (e.g. [7], [29], and [1]) there exists a stationary monotone optimal policy for (1.6) for each pair of β_1 and β_2 . Again, if u_c and P_c are respectively equal to the utility and average service cost for the above optimal policy (for some β_1 and β_2), then this monotone policy is also optimal (see [35]) for (1.5). This motivates us to consider (1.5) for a class of admissible policies, which are monotone.

The asymptotic regime \mathfrak{R} : We consider the tradeoff problems (1.5) and (1.7) in the asymptotic regime \mathfrak{R} where the average service cost constraint P_c is arbitrarily close to the minimum average service cost required for stability. It turns out that this minimum average service cost required for stability is a function of u_c and λ for (1.5) and (1.7) respectively. Because of the similarities in its properties with those of $c(\cdot)$ for the state dependent M/M/1 model, the minimum average service cost required for stability is denoted as $c(\cdot)$ for (1.5) and (1.7) also. We note that if there is only one environment state, say h_0 , then $c(s) = P(h_0, s)$. The asymptotic regime \mathfrak{R} for problems (1.5) and (1.7) is defined as the regime in which $P_c \downarrow c(u^{-1}(u_c))$ and $P_c \downarrow c(\lambda)$ respectively.

A multiqueue model : We also consider a multiqueue single server queueing model, with N queues being served by a single server as shown in Figure 1.3. The model is a straightforward generalization of the single queue model discussed above. We assume that there is an environment variable $H_n[m]$ associated with the n^{th} queue. The vector of N environment variables is denoted as $\mathbf{H}[m] = (H_1[m], \dots, H_N[m])$. The vector of N queue lengths at the start of slot m is denoted as $\mathbf{Q}[m-1] = (Q_1[m-1], \dots, Q_N[m-1])$. We note that in this case the batch service vector, $\mathbf{S}[m]$, is a vector function of the history $\sigma[m]$ for the N queues, the current queue length vector $\mathbf{Q}[m-1]$, and the current environment state $\mathbf{H}[m]$ as in the single queue case. The batch of arrivals which are admitted, $\mathbf{A}[m]$, is a vector function of the history $\sigma[m]$ for the N queues, $\mathbf{Q}[m-1]$, $\mathbf{H}[m]$, and the vector of current arrivals $\mathbf{R}[m]$.

The service cost is a scalar function $P(\mathbf{h}, \mathbf{s})$ of the service vector \mathbf{s} and the environment vector \mathbf{h} . For a particular policy γ , we are interested in the total average queue length $\bar{Q}(\gamma)$, which is the sum of the average queue lengths for the N queues. The average service cost is the time average of the expectation of $P(\mathbf{H}[m], \mathbf{S}[m])$. We also assume that there are individual lower bound constraints on the $u(\bar{A}_n(\gamma))$ for each queue. Other definitions are straightforward extensions of the definitions for the single queue model. The asymptotic regime \mathfrak{R} for the multiqueue model is similar to that

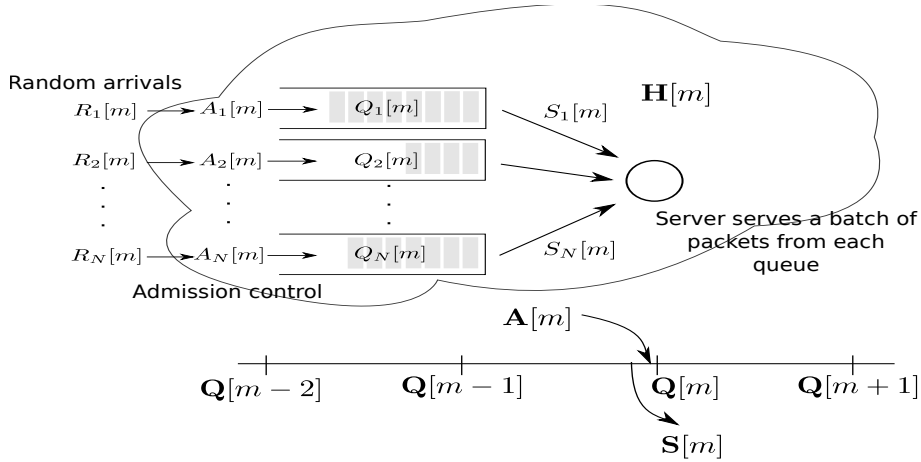


Figure 1.3: The discrete time single server queueing model with N queues

for the single queue case. We consider the problem of minimizing the average queue length in the asymptotic regime \mathfrak{R} , where the scalar cost constraint approaches the minimum average power required for mean rate stability.

1.2 Literature survey

Tradeoff problems for both continuous time and discrete time single server queueing models has been addressed by many researchers. To review what is known for such tradeoff problems, let us consider the problem (1.3). One of the first questions that can be asked is that of feasibility; for what values of c_c are there feasible solutions to (1.3)? Such questions are commonly answered using Lyapunov drift arguments as in, [45] and [41]. The next question that can be asked is that of the existence of an optimal policy for (1.3). Such questions are commonly answered by posing the problem (1.3) as a constrained Markov decision problem (CMDP) as in [68] or [2]. Under certain technical assumptions, the existence of an optimal policy which is also stationary can be shown using results as in [2], [28], or [27]. In certain cases, it may be possible to show that there is a stationary deterministic optimal policy for (1.3)⁴. In such cases, it is also possible to convert the CMDP (1.3) into an unconstrained Markov decision problem (MDP) using suitable Lagrange

⁴Stationary deterministic optimal policies are such that at a time, the service rate (and the arrival rate for models with admission control) is chosen as a deterministic function of the system state.

multipliers as in (1.4). For the state dependent M/M/1 model, the approach is then to find an optimal policy for the above MDP. We note that the optimal policy for (1.3) specifies the optimal service rate $\mu^*(q)$ and arrival rate $\lambda^*(q)$ as a function of queue length, to operate the system in order to minimize the time average of the single stage cost rate. Characterization of the optimal policy using the MDP approach yields structural properties, which are useful in reducing the search space for the optimal policy, for example see [70], [25], [3], and [5]. Surveys of the above approach can be found in [32] and [69].

We first review the results which are available for the continuous time model. In most cases, a monotonicity property of the optimal policy for (1.2) is obtained, i.e., $\mu^*(q)$ is a non-decreasing function of q and $\lambda^*(q)$ is a non-increasing function of q . Stidham and Weber [70] show that $\mu^*(q)$ is non-decreasing and $\lambda^*(q)$ is non-increasing, for a state dependent M/G/1 model, where the objective is to minimize the expected total cost from any initial queue length by serving customers until the queue length is zero, which is related to the average cost criterion. George and Harrison [25] show that $\mu^*(q)$ is non-decreasing in q , for a state dependent M/M/1 model with a Poisson arrival process of rate 1, for an MDP of the form in (1.4). Similar results have also been obtained by Ata in [3], and Ata and Shneorson in [5]. This monotonicity property of $\mu^*(q)$ and $\lambda^*(q)$ is used to motivate the definition of *admissible*⁵ policies in Chapters 2 and 3, which are policies with non-decreasing service rates $\mu(q)$ and non-increasing arrival rates $\lambda(q)$ as a function of q . We then obtain an asymptotic characterization of the tradeoff problems (1.1) and (1.3) for the class of admissible policies. We note that whenever the solution of (1.1) coincides with that of its dual (obtained by optimizing (1.2) over β_1 and $\beta_2 \geq 0$), then the optimal policy for (1.1) is admissible.

For the discrete time model, from a CMDP formulation for (1.7), it can be shown that there exists an optimal policy which chooses the service batch size $S[m]$ as a function $S^*(Q[m-1], H[m])$ of the current queue length and environment state, if the arrival process ($R[m]$) and the environment process ($H[m]$) are IID. The optimal rate or batch size $S^*(q, h)$ which has to be served as a function of the current queue length q and environment state h , so as to minimize the average queue length for a given constraint on the average service cost, can be characterized. Berry and Gallager [7], Collins and Cruz [18], Agarwal et al. [1], and Goyal et.al. [29] consider a fading point to point link with no admission control, where they study the tradeoff problem (1.7). They use a Lagrangian relaxation of the CMDP as in (1.8) to obtain that in many cases, $S^*(q, h)$ is a deterministic function $s^*(q, h)$ which is monotonically non-decreasing in q for every h . The monotonicity property of $s^*(q, h)$ is used in motivating the definition of admissible policies in Chapters 4, 5, and 6. We obtain an asymptotic characterization of the tradeoff problems (1.5) and (1.7) for the class of admissible policies. We note that whenever the solution of (1.5) coincides with that of its dual (obtained by

⁵Admissible policies are monotone policies. Since admissible policies are stationary, the queue evolution process under admissible policies is a Markov process. Admissible policies are monotone policies which are such that the Markov queue evolution process possesses *nice* properties such as aperiodicity, irreducibility, positive recurrence, and finite mean queue length.

optimizing (1.6) over β_1 and $\beta_2 \geq 0$), then the optimal policy for (1.5) is admissible. A similar observation holds for the solution of the tradeoff problem (1.7).

Going beyond the above monotonicity property of $s^*(q, h)$, Berry and Gallager [7] also characterized the asymptotic order behaviour of the tradeoff curve in the regime of large average delay. It was shown that when the average delay is allowed to be large, the average transmitter power can be made arbitrarily close to $c(\lambda)$. It was also shown that if the power $P(h, s)$ expended in transmission of s packets is a strictly convex function of s for fixed fade state h , then the average queue length is $\Omega\left(\frac{1}{\sqrt{V}}\right)$ if the average transmitter power constraint is V more than $c(\lambda)$ as $V \downarrow 0$ (this is known as the Berry-Gallager lower bound). This asymptotic characterization of the average transmitter power was motivated by the asymptotic characterization of the average distortion of an information source obtained by Tse [76].

We note that the order behaviour provides a first order characterization of the tradeoff curve. Furthermore, the order characterization provides a criterion to identify a *good* family⁶ of transmission policies. The authors in [7] suggest that the family of *buffer partitioning* policies achieves the $\frac{1}{D^2}$ order bound but were unable to prove this. Buffer partitioning policies partition the buffer into two regions, and use an average batch service rate less than λ in the lower region and an average batch service rate greater than λ in the higher region to drive the average queue length towards the partitioning value.

Neely [43] extended the Berry-Gallager lower bound to single hop networks and presented a back-pressure based (TOCA) algorithm which achieves the lower bound to within a logarithmic factor, i.e., the algorithm achieves an average delay of $\mathcal{O}\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)$ when the average transmitter power is at most V more than $c(\lambda)$. Neely [43] also observed that if the transmitter power $P(h, s)$ is a piecewise linear function of s for every h , then the above rate of increase of the average delay can be improved. We note that if $P(h, s)$ is piecewise linear in $s, \forall h$, then the function $c(\lambda)$ is a piecewise linear function of λ . If $(\lambda, c(\lambda))$ lies on a linear portion of the function $c(\cdot)$, it was observed [43] that there is a family of policies for which the average delay is $\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$ if the average transmitter power is V more than $c(\lambda)$. It was also observed [43] that for all values of λ and $P(h, s)$ convex in $s, \forall h$, there is a family of policies for which the average delay is $\mathcal{O}\left(\frac{1}{V}\right)$ if the average transmitter power is V more than $c(\lambda)$.

In [44], Neely considered the problem (1.5), with $u(x) = x$ and $P(h, s)$ a strictly convex function in s for every h . He observed that if the transmitter is allowed to drop a non-zero fraction of the customers arriving into the queue, such that $\bar{A}(\gamma) \geq \rho\lambda, 0 < \rho < 1$, then the average queue length grows only as $\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$ rather than $\mathcal{O}\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)$ if the average transmitter power is V more than $c(\lambda)$. Extensions to more general networks and other formulations can be found in [45].

⁶A family of transmission policies is a set of policies with common structure, e.g., a set of parametrized policies that do not serve below a threshold (parameter) while serving a particular batch size (another parameter) above the threshold.

We note that an order optimality result was obtained by Ramaiyan [50] for a particular birth death queueing model, for the average queue length at the relay node for a two-way relay link using network coding, in the regime \mathfrak{R} , for monotone policies. Ramaiyan et al. [51] also obtained the optimal tradeoff of average queueing delay and average transit delay for a two-hop vehicular relay network. Asymptotic upper bounds on the minimum average delay for general wireless networks using network coding, under an average power constraint, was obtained in [16].

An asymptotic upper bound for the tradeoff curve corresponding to (1.7) has been obtained for the case where $(A[m])$ and $(H[m])$ are ergodic Markov processes in [45, Section 4.9] and [30]. In [45, Theorem 4.12] and [30], it has been shown that if $(A[m])$ and $(H[m])$ are ergodic Markov processes, then for a sequence of Quadratic Lyapunov Algorithm (QLA) policies, parametrized by a sequence $V \downarrow 0$, the average queue length is $\mathcal{O}(\frac{1}{V})$ for an average power V more than $c(\lambda)$. Order optimality has also been explored for finite buffer systems. In [6, Chapter 6] it is shown that for a finite buffer discrete time queueing model, as the buffer size B goes to infinity, for any sequence of policies such that the buffer overflow probability is $o(\frac{1}{B^2})$, the average service cost is at least $\Omega(\frac{1}{B^2})$ more than $c(\lambda)$.

We note that, for the tradeoff problems (1.5) and (1.7), although asymptotic upper bounds on the minimum average queue length are known, asymptotic lower bounds are not available in many cases. We note that such asymptotic lower bounds are significant, since they may help in determining the best possible tradeoff. Let V be the difference between the average service cost constraint P_c and the minimum average service cost for stability, $c(\lambda)$, in the asymptotic regime \mathfrak{R} , where $V \downarrow 0$. The known asymptotic lower bounds on the minimum average queue length, along with the details of the models analysed, and the asymptotic upper bounds on the minimum average queue length are summarized in Table 1.1.

We note that several asymptotic lower bounds in Table 1.1 have been derived under the assumption that the queue length and service batch size take values in \mathbb{R}_+ and the service cost function is strictly convex. In certain cases, these real valued queueing models are used as approximate models for queueing models where the queue length and the service batch size take values in \mathbb{Z}_+ . We note that there are also scenarios, where modelling the queue length evolution to be on \mathbb{R}_+ is natural, such as when the queue is assumed to buffer a certain amount of *error exponent* as in [7].

An approximate solution to the tradeoff problem has been obtained by Ata et al. [4] by approximating $(Q[m])$ by a diffusion process, which enables them to find the optimal policy for the control of the approximating diffusion process in closed form. The complete characterization of the optimal admission control policy for a continuous time queueing model was obtained in [48].

Bettesh and Shamaï [10] obtain approximations for $s^*(q, h)$ for every h , in the regime of large q , by solving the average cost optimality equation associated with the MDP (1.8). For the MDP (1.8), Chen et al. [17] obtain approximations for $s^*(q)$ from a fluid approximation. However, we note that

	Model details	Asymptotic upper bound (Regime \mathfrak{R})	Asymptotic lower bound (Regime \mathfrak{R})
1	Berry-Gallager power delay tradeoff [7]; $P(h, s)$ strictly convex in $s, \forall h$	$\mathcal{O}\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)$	$\Omega\left(\frac{1}{\sqrt{V}}\right)$
2	Multuser Berry-Gallager power delay tradeoff [43]; $P(\mathbf{h}, \mathbf{s})$ strictly convex in $\mathbf{s}, \forall \mathbf{h}$	$\mathcal{O}\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)$	$\Omega\left(\frac{1}{\sqrt{V}}\right)$
3	Multuser Berry-Gallager power delay tradeoff [43]; piecewise linear $c(\cdot)$, λ is such that $c(\lambda)$ is on a piecewise linear portion of $c(\cdot)$	$\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$	$\Omega\left(\log\left(\frac{1}{V}\right)\right)$ shown for a specific example, not known in general
4	Multuser Berry-Gallager power delay tradeoff [43]; piecewise linear $c(\cdot)$, λ is any abscissa at which the slope of $c(\cdot)$ changes	$\mathcal{O}\left(\frac{1}{V}\right)$	Not known
5	Power delay tradeoff with lower bound constraint on average throughput [44]	$\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$	$\Omega\left(\log\left(\frac{1}{V}\right)\right)$ but with single fade state
6	Utility delay tradeoff [42]	$\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$	$\Omega\left(\log\left(\frac{1}{V}\right)\right)$
7	Power delay tradeoff with Markov arrival and fading process [30]	$\mathcal{O}\left(\frac{1}{V}\right)$	Not known

Table 1.1: Some of the available asymptotic bounds on the minimum average queue length; except for case 7 all other models assume that the arrival process and the fade process are IID, and except for cases 5 and 6 all models have $A[m] = R[m]$. Also, all lower bounds are derived under the assumption that the queue length can take real values.

the bounds on $s^*(q)$ depend on the form of the service cost functions.

Motivated by the above survey of known results, we ask and try to answer the following questions in this thesis:

1. We note that, in Table 1.1, there are several cases in which asymptotic lower bounds are not known. What are these asymptotic lower bounds?
2. From Table 1.1, we observe that asymptotic lower bounds have the form of $\log\left(\frac{1}{V}\right)$ or $\frac{1}{\sqrt{V}}$. What is the intuition behind such a behaviour?
3. As stated before, for certain cases, asymptotic lower bounds have been derived for an approximate queueing model, where the queue evolution is real valued. Are the asymptotic lower or upper bounds different for the original integer valued queue evolution model?
4. Bounds on $s^*(q, h_0)$ have been obtained, which are dependent on the service cost function. However, the asymptotic order bounds on the minimum average queue length are dependent

only on certain properties of the service cost function, rather than its exact form. Can we obtain asymptotic order bounds on the policy which are independent of the exact form of the service cost function?

In the next section, we briefly survey how these issues have been addressed in this thesis.

1.3 Overview of the thesis and contributions

This thesis consists of two parts. In the first part, comprising Chapters 2, 3, and 4, we consider the tradeoff problems (1.1), (1.3), (1.5), and (1.7) in their respective asymptotic regimes \mathfrak{R} , for the class of admissible policies. The second part, comprising Chapters 5 and 6, primarily illustrates the application of the results obtained in the first part to the motivating resource tradeoff problems arising in point-to-point communication links.

In this introductory discussion, the results are stated informally. We note that the results hold under further technical assumptions, which are stated in the respective chapters.

We study the continuous time state dependent M/M/1 queueing model in Chapters 2 and 3. The main results which are obtained in Chapters 2 and 3 are summarized in Table 1.2. For the first two cases, let V be the difference between c_c and $c(\lambda)$ in the asymptotic regime \mathfrak{R} , where $V \downarrow 0$. For the third case V is the difference between $u(\mu)$ and u_c , while for the fourth case V is the difference between c_c and $c(u^{-1}(u_c))$. We note that the first two cases are instances of (1.3) while the fourth case is an instance of (1.1). The third case is similar to (1.3), except that the roles of $\lambda(q)$ and $\mu(q)$ are interchanged, i.e., there is admission control with a fixed service rate. To the best of our knowledge, such asymptotic results for the tradeoff for the continuous time state dependent M/M/1 model, which may be of independent interest, are new. Thus, we obtain answers for the first question that we posed, for the continuous time queueing models.

The insights obtained from these two chapters are then used in deriving asymptotic lower bounds for a discrete time model in Chapter 4. The correspondence between discrete time models and the continuous time models can be achieved by the choice of \mathcal{X}_μ , \mathcal{X}_λ , and the form of the functions $c(\cdot)$ and $u(\cdot)$. The motivation behind the choice of \mathcal{X}_μ , \mathcal{X}_λ , and the form of the functions $c(\cdot)$ and $u(\cdot)$ is explained in Chapter 3.

From the analysis in Chapters 2 and 3, we obtain the following intuition for the behaviour of the asymptotic lower bounds, which partly answers the second question that we posed. We note that for the state dependent M/M/1 model, for admissible policies, it can be shown that the stationary probability distribution of the queue length exists. The intuition for the behaviour of the asymptotic lower bounds is based on the *shape* of this stationary probability distribution in the asymptotic regime \mathfrak{R} . We discuss the behaviour of the stationary probability distribution only for the cases where the minimum average queue length increases to infinity in the regime \mathfrak{R} .

Control	Service cost and utility functions	Results (in the regime \mathfrak{R} , for admissible policies)
\mathcal{X}_μ is discrete, $\mathcal{X}_\lambda = \{\lambda\}$	$c(\cdot)$ is piecewise linear	Depending on λ , minimum average queue length either increases to a finite value, is $\Theta(\log(\frac{1}{V}))$, or $\Theta(\frac{1}{V})$
\mathcal{X}_μ is a finite interval, $\mathcal{X}_\lambda = \{\lambda\}$	$c(\cdot)$ is strictly convex $c(\cdot)$ is piecewise linear	Minimum average queue length is $\Omega(\frac{1}{\sqrt{V}})$ Depending on λ , minimum average queue length either increases to a finite value, is $\Theta(\log(\frac{1}{V}))$, or $\Theta(\frac{1}{V})$
$\mathcal{X}_\mu = \{\mu\}$, \mathcal{X}_λ is a finite interval	$u(\cdot)$ is strictly concave $u(\cdot)$ is piecewise linear	Minimum average queue length is $\Omega(\frac{1}{\sqrt{V}})$ Depending on λ , minimum average queue length is $\Omega(\log(\frac{1}{V}))$, or $\Omega(\frac{1}{V})$
\mathcal{X}_μ is a finite interval, \mathcal{X}_λ is a finite interval	$c(\cdot)$ is strictly convex, $u(\cdot)$ is strictly concave	Minimum average queue length is $\Theta(\log(\frac{1}{V}))$

Table 1.2: List of asymptotic results derived for continuous time queueing models in Chapters 2 and 3.

The behaviour of the stationary probability distribution of the queue length is determined by the behaviour of the stationary probability distribution of the service rates, which is in turn decided by the nature of the function $c(\mu)$ at $\mu = \lambda$ or $u^{-1}(u_c)$ as the case may be. We consider non-idling admissible policies for the purpose of discussion. We find that in the regime \mathfrak{R} , the stationary probability of using a service rate of zero, goes to zero as $\mathcal{O}(V)$. Then, intuitively, since the probability of this queue being empty goes to zero, the stationary probability distribution *shifts to the right* as shown in Figure 1.4. Therefore, the average queue length has to increase. We note that this intuition has been used in the design of tradeoff optimal policies in [43].

We note that for monotone admissible policies, the stationary probability distribution has the following *shape* (as in Figure 1.4). The stationary probability distribution is monotonically increasing, then may or may not be constant for a set of queue lengths, and then is monotonically decreasing.

In the asymptotic regime \mathfrak{R} , the probability of using certain service rates decreases to zero as $\mathcal{O}(V)$, while the probability of using certain service rates increases. Consider the set of queue lengths, \mathcal{Q}_h , which are such that the stationary probability of using the service rates $\mu(q), q \in \mathcal{Q}_h$ does not decrease to zero. The different behaviours for the minimum average queue length, depends on (i) the shape of the stationary probability distribution for the set \mathcal{Q}_h , and (ii) the stationary probability of the smallest queue length in \mathcal{Q}_h .

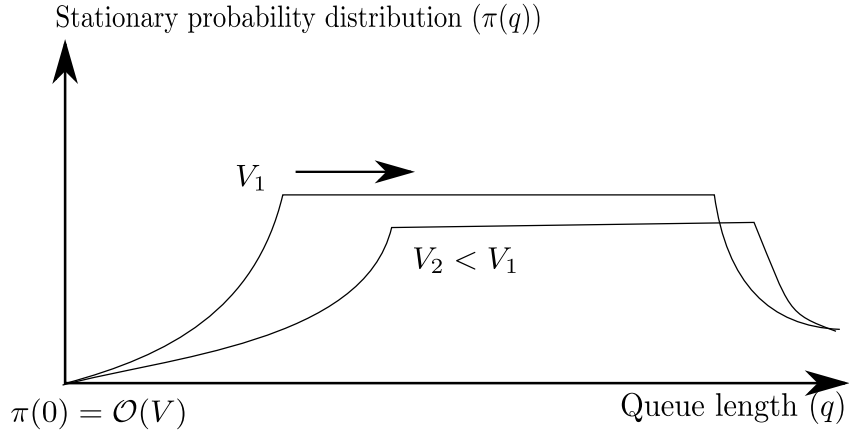


Figure 1.4: Intuition for the behaviour of the stationary probability distribution in the regime \mathfrak{R}




Figure 1.5: Possibilities for the behaviour of the stationary probability distribution in the regime \mathfrak{R}

In the asymptotic regime \mathfrak{R} , the shape of the stationary probability distribution for the set of queue lengths \mathcal{Q}_h , can be either (S1) monotonically increasing, constant, and then monotonically decreasing, or (S2) constant. We shall see that this is decided by the extent of freedom that we have in the choice of $\lambda(q)$ and $\mu(q)$. In the asymptotic regime \mathfrak{R} , the stationary probability of the smallest queue length in \mathcal{Q}_h is either (P1) $\mathcal{O}(V)$ or (P2) $\mathcal{O}(\sqrt{V})$. We shall see that this is decided by the form of the function $c(\cdot)$ at $\mu = \lambda$ or $u^{-1}(u_c)$ as the case may be. The various possibilities are illustrated in Figure 1.5.

Then, the $\log(\frac{1}{V})$ behaviour for the minimum average queue length arises because the stationary probability distribution of the queue length is (S1) and the stationary probability of the smallest queue length in \mathcal{Q}_h is (P1). The asymptotic $\frac{1}{V}$ behaviour arises with (S2) and (P1). The asymptotic $\frac{1}{\sqrt{V}}$ behaviour arises with (S2) and (P2). Although we discuss this in more detail in Chapters 2 and 3, here we provide an example to illustrate the derivation of the $\log(\frac{1}{V})$ behaviour. For a particular policy, we obtain a geometric upper bound of the form $\pi(0)\rho^q$ ($\rho > 1$) on the stationary probability $\pi(q)$ of the queue length. In the asymptotic regime \mathfrak{R} we also show that $\pi(0) = \mathcal{O}(V)$. Then

applying Markov inequality we obtain that the average queue length is a constant times $\log\left(\frac{1}{V}\right)$.

We note that these results are obtained using geometric upper bounds on the stationary probability of the queue length. Under the assumptions that we have made regarding the form of $c(\cdot)$ and $u(\cdot)$, we therefore observe that the possible forms for asymptotic lower bounds are $\log\left(\frac{1}{V}\right)$, $\frac{1}{V}$, or $\frac{1}{\sqrt{V}}$. By providing this intuition, we have successfully answered the second question, for admissible policies, for the continuous time model.

Using bounds on the stationary probability distribution of queue length, we also obtain asymptotic bounds on the structure of *order-optimal* admissible policies⁷ in the regime \mathfrak{R} . Since admissible policies are monotone, we note that corresponding to a set of service rates of S , there is a contiguous set or interval, \mathcal{Q}_S , of queue lengths such that $\mu(q) \in S, \forall q \in \mathcal{Q}_S$. In both Chapter 2 and Chapter 3, we obtain asymptotic bounds on the cardinality of \mathcal{Q}_S . These bounds are independent of the exact form of the functions $c(\cdot)$ and $u(\cdot)$. Thus, we obtain some answers for the fourth question that we have posed, for continuous time queueing models.

In Chapter 2, we also apply the analysis to a flow-level resource allocation model for a wireless downlink and obtain asymptotic bounds on the tradeoff of average power and average number of flows.

In Chapter 4, we consider the problem (1.7), with a single environment state (h_0), in the asymptotic regime \mathfrak{R} . We note that in this case the service cost $c(s) = P(h_0, s)$ is a function of the batch size s only. For admissible policies, it turns out that the stationary distribution of the queue length exists. Using the insights about the shape of the stationary distribution of the queue length in the regime \mathfrak{R} from Chapters 2 and 3, we obtain two upper bounds on the stationary probability distribution of the queue length, one of which is an extension of the bound on stationary probability distribution obtained in Bertsimas et al. [9] to the case where the service rate is dependent on the queue length. The bounds can be used to obtain the same intuition, as explained earlier for the continuous time model, for the discrete time model. The bounds are used to obtain the results that are summarized in Table 1.3. Thus, we answer the first and second questions that we have posed, for admissible policies. The bounds on the stationary probability distribution are also used to obtain asymptotic bounds on the cardinality of \mathcal{Q}_S , leading to some answers for the fourth question that we have posed.

We note that our asymptotic results apply to the tradeoff problems in Table 1.1 under the assumption that the optimal policy lies in the class of admissible policies, which is true for many cases⁸.

We also show that approximating the original integer valued queueing model by a real valued model with a strictly convex function, leads to the average queue length and average service cost being underestimated, in certain cases. We also analyse a case, which have not been hitherto identified,

⁷these are admissible policies which achieve the asymptotic lower bounds in Table 1.2.

⁸For example, whenever the solution to (1.7) coincides with that of its dual (obtained from (1.8)).

Model details	Service cost function	Results (in the regime \mathfrak{R} , for admissible policies)
$Q[m]$ evolves on \mathbb{Z}_+	$P(h_0, s)$ is piecewise linear in s	Depending on λ , minimum average queue length either increases to a finite value, is $\Theta(\log(\frac{1}{V}))$, or $\Theta(\frac{1}{V})$
$Q[m]$ evolves on \mathbb{Z}_+ ; $A[m]$ is ergodic	$P(h_0, s)$ is piecewise linear in s	Depending on λ , minimum average queue length either increases to a finite value or is $\Omega(\log(\frac{1}{V}))$
$Q[m]$ evolves on \mathbb{R}_+	$P(h_0, s)$ is strictly convex in s	Minimum average queue length is $\Omega(\frac{1}{\sqrt{V}})$ (previously known [7] but re-derived here using our method).

Table 1.3: List of asymptotic results derived for tradeoff problem (1.7) in Chapter 4 with single environment state h_0 .

where the average queue length increases to only a finite value in the asymptotic regime \mathfrak{R} , for the original integer valued queueing model. Thus, we obtain some answers to the third question that we have posed. We also show that a more appropriate real valued approximation is one in which the service cost function is approximated as the piecewise linear lower envelope of the service cost function for the original integer valued queueing model.

In Chapter 5, we consider problems (1.7) and (1.5), in the context of a point-to-point link with fast fading. In this chapter, the environment variable h models the fade state and $P(h, s)$ is the power expended in transmitting s packets in fade state h . Let V be the difference between the power constraint P_c and $c(\lambda)$ (or P_c and $c(u^{-1}(u_c))$) in the asymptotic regime \mathfrak{R} , where $V \downarrow 0$. Using the results in Chapter 4, we obtain an asymptotic characterization of the tradeoff. The main results obtained in this chapter are summarized in Table 1.4. We also comment on the extension of these asymptotic results to: (a) a N user single hop network model (as in Figure 1.3) with $Q_n[m]$ assumed to evolve on \mathbb{Z}_+ , (b) a model with admission control and ergodic arrival and fading processes, and (c) a model with no service cost, but for which we are interested in the tradeoff of utility and delay as in [42]. We note that except for the case where the minimum average queue length increases to only a finite value (which has been hitherto not identified in literature) we are able to obtain asymptotic lower bounds for all of the models in Table 1.1 for admissible policies.

We consider the tradeoff of average delay with average error rate for a point-to-point link in Chapter 6. The transmitter is assumed to use fixed or variable-length block coding. In this chapter, we interpret a packet as an information message symbol, which could be a bit. For fixed length block coding, we assume that s message symbols are encoded into a codeword of length N_c channel uses.

Model details	Service cost function	Results (in the regime \mathfrak{R} , for admissible policies)
$Q[m]$ evolves on \mathbb{Z}_+ , $A[m] = R[m]$, as in [43]	$P(h, s)$ is piecewise linear in $s, \forall h$	Depending on the arrival rate λ , minimum average queue length either increases to only a finite value, or is $\Theta(\log(\frac{1}{V}))$ or is $\Theta(\frac{1}{V})$
$Q[m]$ evolves on \mathbb{R}_+ , $A[m] = R[m]$, same as the Berry-Gallager tradeoff problem [7]	$P(h, s)$ is strictly convex in $s, \forall h$	Minimum average queue length is $\Omega(\frac{1}{\sqrt{V}})$ (previously known [7] but re-derived here using our method).
$Q[m]$ evolves on \mathbb{R}_+ , with admission control, same as the model in [44]	$P(h, s)$ is strictly convex in $s, \forall h$	Minimum average queue length is $\Theta(\log(\frac{1}{V}))$
$Q[m]$ evolves on \mathbb{R}_+ , $A[m] = R[m]$, same as the Berry-Gallager tradeoff problem [7], but with ergodic arrival and fading process	$P(h, s)$ is strictly convex in $s, \forall h$	Depending on the value of λ , minimum average queue length either increases to only a finite value or is $\Omega(\log(\frac{1}{V}))$
$Q_n[m]$ evolves on \mathbb{R}_+ for every user n , no admission control, same as the multiuser Berry-Gallager tradeoff problem [43]	$P(\mathbf{h}, \mathbf{s})$ is strictly convex in the vector $\mathbf{s}, \forall \mathbf{h}$	Minimum average total queue length is $\Omega(\frac{1}{\sqrt{V}})$, individual average queue length is also $\Omega(\frac{1}{\sqrt{V}})$ (previously known [43] but re-derived here using our method).
$Q_n[m]$ evolves on \mathbb{R}_+ for every user n , admission control	$P(\mathbf{h}, \mathbf{s})$ is strictly convex in the vector $\mathbf{s}, \forall \mathbf{h}$	Minimum average queue length is $\Theta(\log(\frac{1}{V}))$

Table 1.4: The main asymptotic results derived for discrete time queueing models with fading; unless stated otherwise the models assume that $(A[m])$ and $(H[m])$ are IID

We assume that N_c channel uses correspond to one slot. In the context of our discrete time model, the environment state is fixed (h_0) and $c(s) = P(h_0, s)$ is the expected number of message symbols which are decoded in error. The function $c(s)$ is approximated as s times the average block error probability when s message symbols are transmitted using a random block code of length N_c , where the average block error probability is further approximated by using Gallager's random coding upper bound [22, Chapter 5]. Asymptotic bounds to the optimal average delay for a given average error rate constraint are obtained as in Chapter 4, although in this case $c(s)$ is a non-convex function of s . The asymptotic lower bounds obtained in Chapter 4 can be applied to non-convex $c(s)$ through

the use of the lower convex envelope of $c(s)$.

In Chapter 6, we also consider a single server queueing model, where the codeword length N_c is a parameter for the policy, i.e., different policies can choose different codeword lengths, but every transmission uses codewords with the same length. For such models, it is intuitive that by using arbitrarily large block lengths the average message symbol error rate can be made arbitrarily close to zero. We show that the exponential decay rate of the average error rate with average queueing delay is at most $\frac{2}{3}E_R(\lambda)$ where $E_R(\lambda)$ is the Gallager random coding exponent and λ is the arrival rate of packets per slot. Furthermore for fixed length block codes, for λ sufficiently close to the capacity of the point-to-point link, a class of fixed rate service policies is shown to achieve the decay rate $\frac{2}{3}E_R(\lambda)$. We then consider a single server queueing model where the service time can also be varied, to model scenarios where variable length coding is used. For variable length block codes which constrain the average message symbol error rate by a constant bound on the block error probability, the class of exhaustive service policies, which transmit all the message symbols in the queue at a transmission instant, is shown to achieve the decay rate $\frac{2}{3}E_R(\lambda)$ for any λ .

We summarize the thesis in Chapter 7 and discuss some problems with scope for future work that are motivated by the analysis carried out in the thesis. The notation that is common to all the chapters in this thesis is summarized on page xii. The notation that is used in each chapter is summarized in each chapter.

On the tradeoff of average queue length, average service cost, and average utility for the state dependent M/M/1 queue: Part I

2.1 Introduction

We consider the tradeoff between average queue length, average service cost, and average utility for the continuous time single server queueing model in this chapter and the next. The mathematical model considered captures the problem of how a constrained/scarce resource should be dynamically allocated to randomly arriving demands, which may be subjected to admission control, in order that the system is operated *optimally*. Herein, this dynamic allocation problem is modelled using the simple state dependent M/M/1 model discussed in Chapter 1. Our primary motivation for modelling and studying this tradeoff problem as such, is the variety of tradeoff problems that arise in resource allocation problems in wireless networks.

The state dependent M/M/1 model that we consider in this chapter is a birth death process with the state corresponding to the queue length, as reviewed in Chapter 1. In this chapter, we consider the problem of making the optimal choice of the arrival rate and the service rate at each queue length, such that the time average queue length is minimized subject to constraints on both the time average service cost and the time average utility, associated with the service of customers. We recall that such problems have been analysed in [70], [25], [3], and [5], the results of which have been discussed in Chapter 1. But unlike the approach in these papers, in this chapter we obtain an asymptotic characterization of the tradeoff in the regime \mathfrak{R} . The asymptotic characterization of the tradeoff is discussed in this chapter and the next. The notation that we use in these two chapters are summarized in Table 2.1. We first summarize the methodology that is used for obtaining the

asymptotic bounds.

2.1.1 Methodology

We note that if $\pi(q)$ is the stationary distribution of the queue length for a policy γ , then $\bar{Q}(\gamma) = \mathbb{E}_\pi Q$. Suppose $Pr_u \{Q = q\}$ is any upper bound on $\pi(q)$. If \bar{q} is the largest q such that $Pr_u \{Q < q\} \leq \frac{1}{2}$, then $\mathbb{E}_\pi Q \geq \frac{\bar{q}}{2}$. We obtain $Pr_u \{Q = q\}$ for different cases to obtain \bar{q} . For example, if $c(\mu)$ is piecewise linear then it can be shown that $Pr_u \{Q < q\} = \pi(0)\rho^q$, where $\rho > 1$. We note that then \bar{q} is a function of $\pi(0)$. However, for the tradeoff problem that we consider, it can be shown that $\pi(0)$ is proportional to $c_c - c(\lambda)$, i.e. the difference between the average service cost constraint and the infimum of the average service cost. Then, $\frac{\bar{q}}{2} \geq \text{constant} \times \log\left(\frac{1}{c_c - c(\lambda)}\right)$ and therefore so is $\bar{Q}(\gamma)$. This leads to the $\Omega\left(\log\left(\frac{1}{c_c - c(\lambda)}\right)\right)$ asymptotic lower bound for a particular case.

For deriving asymptotic upper bounds, we consider a sequence of admissible policies. For a particular policy in the sequence, the derivation of asymptotic upper bounds on average service cost rate and average utility rate uses upper bounds on stationary probability distribution of the queue length, whereas in all except one case asymptotic upper bounds on the average queue length are derived using the Lyapunov comparison theorem [36, Theorem A.4.3]. For several cases, we identify sequences of order-optimal admissible policies γ_k , for which the asymptotic growth rates of $\bar{Q}(\gamma_k)$ matches with the corresponding asymptotic lower bounds. Now we will discuss the model that we consider in this chapter.

2.1.2 System model

The queue evolves in continuous time, which is denoted by $t \in \mathbb{R}_+$. The number of customers in the queue at time t (including the one in service, if any) is denoted by $Q(t) \in \mathbb{Z}_+$. The state dependent M/M/1 model for the process $Q(t)$ is a birth death process with birth rate $r_{q,q+1} = \lambda(q)$, death rate $r_{q,q-1} = \mu(q)$ for $q \in \{1, \dots\}$, and birth rate when there are zero customers in the queue, $r_{0,1} = \lambda(0)$. The state transition diagram of the birth-death process for a policy γ is shown in Figure 2.1. A policy γ is the sequence $(\mu(0) = 0, \lambda(0), \mu(1), \lambda(1) \dots)$ ¹. The set of all policies is denoted as Γ .

We associate an utility rate function $u(\cdot)$ with the arrival of customers and a cost rate function $c(\cdot)$ with their service. The utility rate function models the benefit in serving customers, while the cost rate function models the cost incurred in serving customers. We assume that utility is accrued

¹We note that we are restricting to policies which are stationary. Such a restriction is reasonable for the class of tradeoff problems that we are interested in.

Symbol	Description
t	time index
$Q(t)$	queue length at time t
$\mu(q)$	service rate used at queue length q
$\lambda(q)$	arrival rate used at queue length q
r_{q_1, q_2}	transition rate from state q_1 to q_2 for a CTMC
$u(\cdot)$	utility rate function
$c(\cdot)$	service cost rate function
\mathcal{X}_μ	set of all possible service rates
\mathcal{X}_λ	set of all possible arrival rates
$[r_{a, \min}, r_{a, \max}]$	range of values for $\lambda(q)$
$[0, r_{\max}]$	range of values for $\mu(q)$
γ	a policy
Γ	set of all policies
$\bar{Q}(\gamma)$	average queue length
$\bar{C}(\gamma)$	average service cost rate
$\bar{U}(\gamma)$	average utility rate
Γ_a	set of admissible policies
π_γ	stationary probability distribution for policy γ
$\Gamma_{a, M}$	set of all mixtures of admissible policies
γ_M	a mixture policy
u_c	constraint on average utility rate
c_c	constraint on average service cost rate
\mathfrak{R}	asymptotic regime in which $c_c \rightarrow u^{-1}(u_c)$
$Q_M^*(c_c, u_c)$	minimum average queue length over $\Gamma_{a, M}$ under constraints c_c and u_c
$Q^*(c_c, u_c)$	minimum average queue length over Γ_a under constraints c_c and u_c
β_c, β_u	non-negative Lagrange multipliers corresponding to service cost and utility constraints
\mathcal{O}^u	set of constraint value pairs (c_c, u_c) for which admissible policies are optimal
$\gamma^*(c_c, u_c)$	optimal policy for constraints c_c and u_c
$Q^*(c_c)$	minimum average queue length for the set Γ_a service cost constraint c_c
$\pi_\mu(k)$	stationary probability of service rate μ_k
μ_l	largest service rate $\leq \lambda$ at which the slope of $c(\cdot)$ changes
μ_u	smallest service rate $\geq \lambda$ at which the slope of $c(\cdot)$ changes
$Q_M^*(c_c)$	minimum average queue length for the set $\Gamma_{a, M}$ under service cost constraint c_c
$P_l \{.\}$	lower bound on probability of an event
$P_u \{.\}$	upper bound on probability of an event

Table 2.1: Notation used in Chapters 2 and 3.




Figure 2.1: The birth death process under a particular policy γ

at the rate of $u(\lambda(Q(t)))$ at time t and cost is incurred at the rate of $c(\mu(Q(t)))$ at time t . The functions $u(\cdot)$ and $c(\cdot)$ are assumed to satisfy the following properties :

U1 : The function $u(\lambda) : \mathcal{X}_\lambda \rightarrow \mathbb{R}_+$ is strictly increasing and concave in λ , with $u(0) = 0$ and \mathcal{X}_λ the set of arrival rates.

C1 : The function $c(\mu) : \mathcal{X}_\mu \rightarrow \mathbb{R}_+$ is strictly increasing and convex in μ , with $c(0) = 0$ and \mathcal{X}_μ the set of service rates.

The set \mathcal{X}_λ is assumed to be either a) a finite set of discrete points $(\lambda_0, \lambda_1, \dots, \lambda_K)$ or b) an interval $[r_{a,min}, r_{a,max}]$ of the real line. Similarly the set \mathcal{X}_μ is assumed to be either a) a finite set of discrete points $(\mu_0, \mu_1, \dots, \mu_K)$ or b) an interval $[r_{min}, r_{max}]$ of the real line. If the set \mathcal{X}_λ is a set of discrete points, then we extend the definition of $u(\lambda)$ to $[r_{a,min} = \lambda_0, r_{a,max} = \lambda_K]$ by linear interpolation. A similar extension is done for the function $c(\mu)$. Note that the linear interpolation preserves the concavity of $u(\cdot)$ and the convexity of $c(\cdot)$. Let $u^{-1} : \mathbb{R}_+ \rightarrow [r_{a,min}, r_{a,max}]$ and $c^{-1} : \mathbb{R}_+ \rightarrow [r_{min}, r_{max}]$ be the inverse functions of $u(\cdot)$ and $c(\cdot)$ respectively.

The average service cost for the policy γ , $\bar{C}(\gamma)$ is defined as

$$\bar{C}(\gamma) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T c(\mu(Q(t))) dt \middle| Q(0) = q_0 \right]. \quad (2.1)$$

The average utility for the policy γ , $\bar{U}(\gamma)$ is defined as

$$\bar{U}(\gamma) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T u(\lambda(Q(t))) dt \middle| Q(0) = q_0 \right]. \quad (2.2)$$

The average queue length for the policy γ , $\bar{Q}(\gamma)$ is defined as

$$\bar{Q}(\gamma) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_0^T Q(t) dt \middle| Q(0) = q_0 \right]. \quad (2.3)$$

In this chapter, we restrict attention to policies for which the above three performance measures are

independent of the initial state q_0 , hence in the above definitions the dependence of these quantities on q_0 is not made explicit. We note that the above definition of average utility (as in [5]) is much more general and encompasses scenarios where the utility of average throughput is of interest (e.g. as in [44]).

We note that the state dependent M/M/1 model can be directly applied to study resource allocation in modern high rate data networks. We consider such a motivational example in the next section. Thus, the tradeoff problem for the state dependent M/M/1 model can be studied in its own right. Furthermore, in Chapters 3 and 4, we shall see that the ideas developed for this simple state dependent M/M/1 model can be used in the study of discrete time queueing models, which in some cases are more representative of the resource allocation problems in wireless networks.

2.1.3 A motivational example

We discuss a motivating example in this section, which is based on the problem considered by Borst [13]. We consider the downlink of a base station, operating in slotted time, with each slot of duration $1.67ms$. Flows, each a file of size $480Kb$, arrive at the downlink scheduler queue for transmission to different users. We assume that at most one flow arrives in a slot, and the flow arrival process is an IID Bernoulli process. Each flow is destined to an user, which is one of two types, T1 or T2 with uniform probability. The base station uses a round robin scheduler, which transmits bits from each flow, in order of their arrival instants. We assume that, if the transmitter is transmitting at a power level of $10W$, then the transmission rate is $50Kb/s$ for T1 flows, while for T2 flows, it is $150Kb/s$. Let $r_1 = 50Kb/s$ and $r_2 = 150Kb/s$. The transmitter may dynamically vary its power to change a multiplier m of the transmission rate whenever a flow arrives or a flow leaves the system. The transmitter may choose $m \in \mathcal{M} = \{0, 0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2\}$. The multiplier m is used to model the constraint that each flow may choose a rate corresponding to the choice of a codebook from a finite set of codebooks. Then the transmission rate is mr_1 and mr_2 for T1 and T2 flows. As in Borst [13], we assume that the transmission rate as a function of the received SNR is $800 \log_{10}(1 + SNR) Kb/s$. We note that this formula for the transmission rate models a case where the fading gain is fixed or a case where transmission is done at a fixed rate only if the fading gain is above a certain threshold. Then the transmitter power is $P^1(m) = 64.6 \left(10^{\left(\frac{50m}{800}\right)} - 1 \right) W$ and $P^2(m) = 18.5 \left(10^{\left(\frac{150m}{800}\right)} - 1 \right) W$, when transmitting to receivers with T1 and T2 flows respectively. We note that as a function of the number of T1 and T2 flows, and therefore the total number of flows in the system, the actual transmitter power used varies within a round-robin scheduling cycle. We are interested in dynamically controlling m as a function of the current total number of flows to minimize the average number of flows in the system (or for a fixed arrival rate, the average flow transfer latency) subject to a constraint on the average transmitter power.

As in Borst [13] we note that, as the minimum flow transfer time ($\frac{480Kb}{300Kb/s}$) for any flow is of the

order of seconds while the slot duration is of the order of milliseconds, we can model the system as a M/G/1-PS (processor sharing) queue, but with control on the total rate of service, through the choice of m as a function $m(q)$ of the current number of flows q in the system. The Bernoulli arrival process is approximated as a Poisson process of rate λ . Let us consider the normalized service requirement for T1 and T2 flows. If a T1-flow is the only flow present in the system, then it requires a time of $\frac{480}{50}s$ if $m = 1$. Similarly a T2-flow requires a time of $\frac{480}{150}s$ if $m = 1$. So the normalized service requirement of the flows arriving are distributed as $\frac{480}{50}$ with probability 0.5 and $\frac{480}{150}$ with probability 0.5. The time sharing amongst users manifests itself as processor sharing in the continuous time model. At a time t , if there are q flows in the M/G/1-PS model, the remaining service requirement of each flow is reduced at the rate $\frac{m(q)}{q}$. From Bonald [11, Theorem 2], we note that for any policy γ , the stationary probability of the M/G/1-PS queue with control on the service rate, is independent of the service requirement distribution. Therefore, following Borst [13], we consider a M/M/1-PS queue with control on the service rate, where the service requirement distribution is exponential, but with the same mean as the service requirement distribution in the M/G/1-PS queue, i.e, a mean service requirement of $\frac{1}{2} \left(\frac{480}{50} + \frac{480}{150} \right) = 6.4s$. At a time t , if there are q flows in this model, the remaining service requirement of each flow is reduced at the rate $\frac{m(q)}{q}$. Therefore the rate at which a flow leaves the system is $q \frac{m(q)}{q} \frac{1}{6.4}$. We then note that the M/M/1-PS queue with state dependent total service rate, determined by the policy γ , is the same as the model considered in Section 2.1.2. The number of flows in the system is modelled by the state of the birth-death process. The birth rate is λ , while the death rate which is dependent on q can take values in $\mathcal{S} = \frac{1}{6.4} \times \{0, 0.25, 0.50, 0.75, 1, 1.25, 1.50, 1.75, 2\}$.

We note that the average transmitter power used when there are q flows in the system depends on the proportion of T1 flows and T2 flows. If there are q_1 T1 flows when there are q flows in the system, then the average transmitter power is $\frac{q_1 P^1(m) + (q - q_1) P^2(m)}{q}$. To obtain *good* policies which tradeoff average power with average number of flows for the discrete time system, we analyse the M/M/1-PS model, under the assumption that the transmitter power is $P(m) = P^1(m)$, irrespective of the flow which is being transmitted. The model obtained under this assumption is denoted as M/M/1-PS (I). We also repeat the analysis under the assumption that $P(m) = P^2(m)$, in which case the model is denoted as M/M/1-PS (II). The results obtained from the analysis of both M/M/1-PS (I) and M/M/1-PS (II) models are found to be useful in designing policies for the discrete time system.

2.1.4 Overview

We formulate the tradeoff problem in Section 2.2 for a restricted class of admissible policies. A non-idling property of any optimal admissible policy is also shown in the same section. We consider

three cases of the tradeoff problem: FINITE- μ CHOICE, INTERVAL- μ CHOICE, and INTERVAL- $\lambda\mu$ CHOICE, which correspond to different choices of the sets \mathcal{X}_λ and \mathcal{X}_μ . The analysis of FINITE- μ CHOICE in the asymptotic regime \mathfrak{R} is carried out in Section 2.3 while INTERVAL- μ CHOICE, and INTERVAL- $\lambda\mu$ CHOICE are analysed in Chapter 3. For FINITE- μ CHOICE, where $\mathcal{X}_\lambda = \{\lambda\}$, we identify three cases based on the value of λ and the set of available service rates $\{\mu_0, \dots, \mu_K\}$, for which the asymptotic behaviour of the tradeoff curve in the asymptotic regime \mathfrak{R} is different. Asymptotic lower bounds and upper bounds to the tradeoff problem for these three cases are then obtained in Section 2.3.1 and Section 2.3.2 respectively. An asymptotic characterization of optimal policies is presented in Section 2.3.3. We then numerically illustrate the asymptotic behaviour of the solution to FINITE- μ CHOICE for several examples in Section 2.3.4. Asymptotic bounds to the tradeoff curve for the example in Section 2.1.3 are then presented in Section 2.3.5 using the results derived in this chapter.

2.2 Problem formulation

In this chapter and the next, we consider the tradeoff problems (1.1) and (1.3) for a restricted class of admissible policies Γ_a . The set of admissible policies is defined as follows.

Stability : A policy γ is defined to be stable if the birth death process $Q(t)$ under policy γ is irreducible and positive recurrent with stationary distribution π_γ .

Admissibility : A policy γ is admissible, if

G1 : it is stable,

G2 : the sequence $(\mu(0), \mu(1), \mu(2), \dots)$ is non-decreasing, and,

G3 : the sequence $(\lambda(0), \lambda(1), \lambda(2), \dots)$ is non-increasing.

Then we define the set of admissible policies as

$$\Gamma_a \triangleq \{\gamma : \gamma \in \Gamma, \gamma \text{ is admissible}\}.$$

Remark 2.2.1. We note that restricting attention to Γ_a is reasonable, as the optimal policy which minimizes the average queue length subject to constraints on the average service cost and average utility possesses the properties G1, G2, and G3 in many cases (see Chapter 1) ².

²We note that there exists an admissible policy which achieves the minimum for the constrained optimization problems (1.1) or (1.3) for certain values of the constraints c_c and u_c . For example, for (1.1), these values of c_c and u_c are such that there exists Lagrange multipliers β_1 and β_2 for which average service cost and average utility of any admissible optimal policy for the dual problem (1.2) are equal to c_c and u_c .

We note that for any admissible policy, we have that

$$\begin{aligned}\bar{C}(\gamma) &= \mathbb{E}_{\pi_\gamma} c(\mu(Q)), \\ \bar{U}(\gamma) &= \mathbb{E}_{\pi_\gamma} u(\lambda(Q)), \text{ and} \\ \bar{Q}(\gamma) &= \mathbb{E}_{\pi_\gamma} Q,\end{aligned}$$

where $Q \sim \pi_\gamma$ and the performance measures are independent of the initial state q_0 . We note that for a policy γ if G2 and G3 holds and if at any finite q' , $\mu(q') - \lambda(q') = \epsilon > 0$, then γ is stable and therefore admissible.

We note that a larger set of policies can be obtained by mixing the *pure* policies in Γ_a . We note that a mixture policy corresponds to time sharing of pure policies, with the time period, in which a particular pure policy is used, tending to infinity. The set of policies which are obtained by a finite mixture of the policies in Γ_a is denoted as $\Gamma_{a,M}$. We note that associated with a $\gamma_M \in \Gamma_{a,M}$ we have a set $\Gamma(\gamma_M) \subseteq \Gamma_a$, which is the set of policies which are mixed according to a probability mass function p_γ , for $\gamma \in \Gamma(\gamma_M)$. For a $\gamma_M \in \Gamma_{a,M}$, $\bar{Q}(\gamma_M) = \sum_\gamma p_\gamma \bar{Q}(\gamma)$. The average service cost rate and average utility rate are defined similarly for $\gamma \in \Gamma_{a,M}$.

2.2.1 Problem

Our objective is to solve the following optimization problem, TRADEOFF-M:

$$\begin{aligned}\text{minimize } & \gamma \in \Gamma_{a,M} \quad \bar{Q}(\gamma) \\ \text{such that } & \bar{C}(\gamma) \leq c_c, \\ & \text{and } \bar{U}(\gamma) \geq u_c,\end{aligned} \tag{2.4}$$

where c_c and u_c are constraints on the average service cost and average utility respectively. The optimal value of the above problem is denoted by $Q_M^*(c_c, u_c)$. We note that above constrained minimization of the average queue length corresponds to the constrained minimization of average delay if the average arrival rate is fixed.

In the following lemma, we show that TRADEOFF-M can be solved, only for certain values of c_c and u_c .

Lemma 2.2.2. If TRADEOFF-M has any feasible solutions, then $u^{-1}(u_c) \leq c^{-1}(c_c)$.

Proof. Assume that there is an policy $\gamma_M \in \Gamma_{a,M}$ which is feasible for TRADEOFF. Then from Jensen's inequality we have that $c(\mathbb{E}_{p_\gamma} [\mathbb{E}_{\pi_\gamma} [\mu(Q)]]) \leq \mathbb{E}_{p_\gamma} [\mathbb{E}_{\pi_\gamma} [c(\mu(Q))]]$ and $\mathbb{E}_{p_\gamma} [\mathbb{E}_{\pi_\gamma} [u(\lambda(Q))]] \leq u(\mathbb{E}_{p_\gamma} [\mathbb{E}_{\pi_\gamma} [\lambda(Q)]])$. For brevity let us denote $\mathbb{E}_{p_\gamma} [\mathbb{E}_{\pi_\gamma} [\cdot]]$ by just $\mathbb{E}[\cdot]$ in this proof. Therefore $\mathbb{E}\mu(Q) \leq c^{-1}(\mathbb{E}c(\mu(Q)))$ and $u^{-1}(\mathbb{E}u(\lambda(Q))) \leq \mathbb{E}\lambda(Q)$. As $\mathbb{E}_{\pi_\gamma} Q < \infty$, $\mathbb{E}_{\pi_\gamma} \mu(Q) =$

$\mathbb{E}_{\pi_{\text{gamma}}} \lambda(Q)$, $\forall \gamma \in \Gamma(\gamma_M)$. Therefore for γ , $u^{-1}(\mathbb{E}u(\lambda(Q))) \leq c^{-1}(\mathbb{E}c(\mu(Q)))$. From the non-decreasing properties of $c(\cdot)$ and $u(\cdot)$ we have that $c^{-1}(\cdot)$ and $u^{-1}(\cdot)$ are also non-decreasing. Hence if there is any one feasible policy γ , $u^{-1}(u_c) \leq c^{-1}(c_c)$. \square

If $c^{-1}(c_c) > u^{-1}(u_c)$, then we show that there exists a feasible policy $\gamma \in \Gamma_a$ on a case by case basis in the following discussion. We note that if $c^{-1}(c_c) > u^{-1}(u_c)$, then it is not guaranteed that an optimal policy $\gamma^*(c_c, u_c) \in \Gamma_{a,M}$ exists for the above problem. However, in the following discussion we identify a set of (c_c, u_c) for which the existence of an optimal policy in $\Gamma_{a,M}$ is guaranteed.

Remark 2.2.3. Let β_c and $\beta_u \in \mathbb{R}_+$. Consider an unconstrained MDP denoted as $MDP(\beta_c, \beta_u)$ as in [5] which is obtained by uniformization at rate r_u with single stage cost $\frac{q + \beta_c c(\mu) - \beta_u u(\lambda)}{r_u}$. Then from [5] we know that an optimal policy $\gamma^*(\beta_c, \beta_u) \in \Gamma_a$ exists for $MDP(\beta_c, \beta_u)$. Let $\Gamma^*(\beta_c, \beta_u)$ be the set of all optimal Γ_a policies for $MDP(\beta_c, \beta_u)$. Also let $\Gamma_M^*(\beta_c, \beta_u)$ be the set of all mixed policies obtained by a finite mixture of $\gamma \in \Gamma^*(\beta_c, \beta_u)$. Let $\mathcal{O}^u = \{(\overline{C}(\gamma), \overline{U}(\gamma)), \gamma \in \Gamma_M^*(\beta_c, \beta_u), \forall \beta_c, \beta_u \geq 0\}$. Then from [35], if $(c_c, u_c) \in \mathcal{O}^u$, then there exists an optimal policy in $\Gamma_{a,M}$ for TRADEOFF-M.

Remark 2.2.4. Suppose $(c_c, u_c) \notin \mathcal{O}^u$, but $c^{-1}(c_c) > u^{-1}(u_c)$. We note that then for all γ_M which are feasible for TRADEOFF-M, $\overline{Q}(\gamma_M) > 0$. Hence, for every $\epsilon > 0$, there exists some feasible γ_M such that $\overline{Q}(\gamma_M) < Q_M^*(c_c, u_c) + \epsilon$. We call such policies ϵ -optimal for c_c .

We now show that any optimal policy for TRADEOFF-M is *non-idling*, if it exists.

Lemma 2.2.5. Any optimal mixed policy $\gamma^*(c_c, u_c)$ for TRADEOFF-M, has $\mu(q) > 0$ for every $q \geq 1$, for every $\gamma \in \Gamma(\gamma^*(c_c, u_c))$.

Proof. Let $\gamma \in \Gamma(\gamma^*(c_c, u_c))$ be an admissible policy with service rate and arrival rate given by $\mu(q)$ and $\lambda(q)$ respectively for $q \geq 0$. Let $q_0 = \max\{q : \mu(q) = 0\}$. Assume that for γ , $q_0 > 0$. As γ is admissible, $\mu(q) = 0$, for all $q \leq q_0$. Then the states $\{0, \dots, q_0 - 1\}$ are transient under policy γ . Let γ' be another policy such that at a queue length q the service rate and arrival rate are $\mu'(q)$ and $\lambda'(q)$ respectively. For $\forall q$, let $\mu'(q) = \mu(q + q_0)$ and $\lambda'(q) = \lambda(q + q_0)$. We note that the birth-death process under γ' is obtained by a relabelling of the states under the policy γ . And γ' is admissible as γ is admissible. It is clear that $\overline{U}(\gamma') = \overline{U}(\gamma)$ and $\overline{C}(\gamma') = \overline{C}(\gamma)$, but $\overline{Q}(\gamma') = \overline{Q}(\gamma) - q_0$. Thus any γ such that $q_0 > 0$ cannot be an element of $\Gamma(\gamma^*(c_c, u_c))$. \square

Thus in the following we need only consider non-idling admissible policies.

In the following discussion we consider the problem TRADEOFF,

$$\begin{aligned} & \text{minimize}_{\gamma \in \Gamma_a} && \overline{Q}(\gamma) \\ & \text{such that} && \overline{C}(\gamma) \leq c_c, \\ & && \text{and } \overline{U}(\gamma) \geq u_c, \end{aligned} \tag{2.5}$$

where we minimize over the set Γ_a only. The asymptotic bounds on the optimal value for TRADEOFF-M can be obtained easily from the analysis of TRADEOFF. We now consider three special cases of TRADEOFF, which are either representative of the problems that arise in the context of communication networks or are useful in understanding the tradeoff for discrete time queues.

FINITE- μ CHOICE:

1. We restrict to policies γ such that $\lambda(q) = \lambda, \forall q \in \mathbb{Z}_+$.
2. For any such policy γ , $\bar{U}(\gamma) = u(\lambda)$. We choose λ such that $u(\lambda) \geq u_c$.
3. We also restrict to policies γ such that $\mu(q) \in \mathcal{X}_\mu = \{\mu_0 = 0, \mu_1, \mu_2, \dots, \mu_K\}$, where $\mu_i < \mu_{i+1}$, $\mu_K = r_{max} < \infty$. Thus the available service rates take values from a finite discrete set. We assume that $\lambda < r_{max}$.
4. The optimal value of the tradeoff problem is denoted by $Q^*(c_c)$.

INTERVAL- μ CHOICE:

1. We restrict to policies γ such that $\lambda(q) = \lambda, \forall q \in \mathbb{Z}_+$.
2. For any such policy γ , $\bar{U}(\gamma) = u(\lambda)$. We choose λ such that $u(\lambda) \geq u_c$.
3. We restrict to policies γ such that $\mu(q) \in [0, r_{max}]$. Thus the available service rates take values in a finite interval. We assume that $\lambda < r_{max}$.
4. The optimal value of the tradeoff problem is denoted by $Q^*(c_c)$.

INTERVAL- $\lambda\mu$ CHOICE:

1. We restrict to policies γ such that $\lambda(q) \in [r_{a,min}, r_{a,max}]$, where $r_{a,min} \geq 0, r_{a,max} < \infty$.
2. We restrict to policies γ such that $\mu(q) \in [0, r_{max}]$. We assume that $r_{a,min} < r_{max}$.

We note that for FINITE- μ CHOICE and INTERVAL- μ CHOICE, the constraint on the average utility in TRADEOFF is satisfied by the choice of λ , and therefore this constraint is not explicitly mentioned (as in (1.3)). In the following, we obtain an asymptotic characterization of $Q^*(c_c)$ for FINITE- μ CHOICE, while INTERVAL- μ CHOICE and INTERVAL- $\lambda\mu$ CHOICE are analysed in Chapter 3.

2.3 Analysis of FINITE- μ CHOICE

In the following, we state the motivation for considering FINITE- μ CHOICE.

Remark 2.3.1. FINITE- μ CHOICE is motivated by the tradeoff problem (1.7) for the following discrete time queueing model. Customers arrive in a batch of random size, in every slot, into an infinite length queue. All the customers which arrive in a slot are admitted into the queue. The number of customers, which are served in each slot, or the service batch size, is chosen as a deterministic function, of the current queue length. This feature of the discrete time queue is modelled by the choice of the service rate, $\mu(q)$, as a function of q in FINITE- μ CHOICE. We assume that for the discrete time model, the queue evolves on the set of non-negative integers. Hence, the service batch size also takes values in the set of non-negative integers. The essential feature here is that the set of batch sizes is discrete and therefore we assume that $\mu(q)$ takes values in a finite discrete set. As there is no admission control in the discrete time model, we assume that the arrival rate is a fixed λ for every q for FINITE- μ CHOICE. In each slot, assume that there is a service cost incurred in serving the customers. This is modelled by the service cost rate function $c(\cdot)$ in FINITE- μ CHOICE. By analysing FINITE- μ CHOICE we illustrate the basic techniques which are used in the asymptotic analysis of TRADEOFF, which turn out to be useful in the analysis of problem (1.7).

We now present an asymptotic analysis of FINITE- μ CHOICE in the regime \mathfrak{R} . We note that part of this analysis was presented in [62]. For brevity, we use π rather than π_γ to denote the stationary distribution corresponding to an admissible policy γ in places where there is no source for confusion. The stationary probability of queue length being q is denoted by $\pi(q)$. The stationary probability of using a rate μ_k is denoted by $\pi_\mu(k)$. We note that $\pi_\mu(k) = \sum_{\{q: \mu(q)=\mu_k\}} \pi(q)$, and $\bar{C}(\gamma) = \sum_{k=0}^K \pi_\mu(k) c(\mu_k)$. Since $c(\mu)$ is convex in μ , by Jensen's inequality, we have that the average service cost $\bar{C}(\gamma) \geq c(\lambda)$, for any admissible policy γ . We note that $c(\lambda)$ is the minimum average service cost which has to be expended for the average service rate to be equal to the average arrival rate, as noted in Chapter 1.

We first obtain an asymptotic lower bound to $Q^*(c_c)$ in the regime \mathfrak{R} as $c_c \downarrow c(\lambda)$, by finding a lower bound on $\bar{Q}(\gamma_k)$, as a function of $\bar{C}(\gamma_k) - c(\lambda)$, for any sequence of feasible admissible policies γ_k with $\bar{C}(\gamma_k) \downarrow c(\lambda)$. Subsequently, we show that there exists a sequence of admissible policies γ_k for which $\bar{C}(\gamma_k)$ approaches $c(\lambda)$ arbitrarily closely, so that $c(\lambda) = \inf_{\gamma \in \Gamma_a} \bar{C}(\gamma)$.

The asymptotic behaviour of $Q^*(c_c)$ for FINITE- μ CHOICE depends on the behaviour of $c(\mu)$ in the neighbourhood of $\mu = \lambda$. We now define quantities μ_u, k_u, μ_l , and k_l , which are related to this

behaviour.

$$\mu_u = \begin{cases} \min \left\{ \mu_k : k \leq K - 1, \mu_k \geq \lambda, \frac{c(\mu_{k+1}) - c(\mu_k)}{\mu_{k+1} - \mu_k} > \frac{c(\mu_k) - c(\mu_{k-1})}{\mu_k - \mu_{k-1}} \right\} & \text{if this set is non-empty,} \\ \mu_K & \text{otherwise.} \end{cases}$$

$$\mu_l = \begin{cases} \max \left\{ \mu_k : k \geq 1, \mu_k \leq \lambda, \frac{c(\mu_{k+1}) - c(\mu_k)}{\mu_{k+1} - \mu_k} > \frac{c(\mu_k) - c(\mu_{k-1})}{\mu_k - \mu_{k-1}} \right\} & \text{if this set is non-empty,} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mu_{k_u} = \mu_u$ and $\mu_{k_l} = \mu_l$. In words, μ_u is the service rate μ_{k_u} greater than or equal to λ at which the slope of $c(\cdot)$, considered for service rates in \mathcal{X}_μ , changes. A similar interpretation can be given for μ_l . Depending upon the value of μ_l , μ_u , and their relationship with λ , there are three different cases that need to be considered (also see Figure 2.2):

FINITE- μ CHOICE-1 : $\mu_l = 0, \lambda < \mu_u$,

FINITE- μ CHOICE-2 : $\mu_l \geq \mu_1, \mu_l < \lambda < \mu_u$, and

FINITE- μ CHOICE-3 : $\mu_l \geq \mu_1, \mu_l = \lambda = \mu_u$.




Figure 2.2: Illustration of the relationship between λ , μ_l , and μ_u along with the minimum average service cost $c(\lambda)$ and the line $l(\mu)$ for the three cases that arises for FINITE- μ CHOICE problem

We now state the motivation for this classification. We note that as the constraint c_c approaches $c(\lambda)$, if there exists a feasible policy γ with $\bar{C}(\gamma) \leq c_c$, then for that policy the stationary probability of certain service rates should go to zero. The classification is based on the set of service rates, whose stationary probability goes to zero. For example, for FINITE- μ CHOICE-3, as $(\lambda, c(\lambda))$ is a *corner* point, as $c_c \downarrow c(\lambda)$, the stationary probability that any service rate other than λ is used approaches zero. For FINITE- μ CHOICE-1 and FINITE- μ CHOICE-2, as $c_c \downarrow c(\lambda)$, the stationary probability that any service rate which is less than μ_l or greater than μ_u is used, approaches zero. We note that $\mu_u > \lambda$ in the case of FINITE- μ CHOICE-1 and FINITE- μ CHOICE-2, and as $c_c \downarrow c(\lambda)$ the service rate $\mu_u > \lambda$ could be used, unlike in the case of FINITE- μ CHOICE-3 where only $\mu = \lambda$ can be used. Furthermore we note that for FINITE- μ CHOICE-1, $\mu_l = 0$ and for both FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3, $\mu_l > 0$. Then as $c_c \downarrow c(\lambda)$, for FINITE- μ CHOICE-1, for a non-idling $\gamma \in \Gamma_a$, the queue becomes empty infinitely often, while this cannot happen for both FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3.

2.3.1 Asymptotic lower bounds

For an admissible policy γ , for obtaining an asymptotic lower bound on $\overline{Q}(\gamma)$ we: a) obtain an upper bound on $\pi_\mu(k)$ for certain values of k in terms of $\overline{C}(\gamma)$ and $c(\lambda)$, b) relate the stationary probability $\pi_\mu(k)$ to the stationary probability of the queue $\pi(q)$, and c) obtain a lower bound on $\overline{Q}(\gamma)$ in terms of $\pi(q)$. For the cases FINITE- μ CHOICE-1 and FINITE- μ CHOICE-2, define the line $l(\mu)$ as the line through the points $(\mu_l, c(\mu_l))$ and $(\mu_u, c(\mu_u))$. For the case FINITE- μ CHOICE-3, let $l(\mu)$ be any line through $(\lambda, c(\lambda))$ with slope greater than $\frac{c(\lambda) - c(\mu_{k_l-1})}{\lambda - \mu_{k_l-1}}$ and less than $\frac{c(\mu_{k_l+1}) - c(\lambda)}{\mu_{k_l+1} - \lambda}$. The line $l(\mu)$ is illustrated for the three cases in Figure 2.2. We note that $c(\lambda) = l(\lambda)$ by construction. Furthermore $\mathbb{E}_{\pi_\mu} l(\mu(Q)) = l(\lambda)$ as l is linear.

We now present an upper bound on the stationary probability of certain service rates as the average service cost approaches $c(\lambda)$.

Lemma 2.3.2. Let $\mathcal{R}_k = \{k : \mu_k < \mu_l \text{ or } \mu_k > \mu_u\}$. For an admissible policy γ , for all $k \in \mathcal{R}_k$,

$$\pi_\mu(k) \leq \frac{\overline{C}(\gamma) - c(\lambda)}{c(\mu_k) - l(\mu_k)}.$$

Proof. We have that $\overline{C}(\gamma) - c(\lambda) = \sum_{k=0}^K \pi_\mu(k) [c(\mu_k) - l(\mu_k)]$. For $k \in \mathcal{R}_k^c$, $c(\mu_k) - l(\mu_k) = 0$. Since $c(\mu_k) > l(\mu_k)$, for all $k \in \mathcal{R}_k$, we have that $\pi_\mu(k) \leq \frac{\overline{C}(\gamma) - c(\lambda)}{c(\mu_k) - l(\mu_k)}$. \square

A non-idling admissible policy γ is specified by the sequence $(q_0 = 0, q_1, q_2, \dots, q_{K-1}, q_K = \infty)$, $q_k \leq q_{k+1}$, which is such that

$$\begin{aligned} \mu(q_0) &= \mu(0) = 0 \\ \mu(q) &= \mu_k, \text{ if } q \in \{q_{k-1} + 1, \dots, q_k\}, \text{ for } k \in \{1, \dots, K\}. \end{aligned}$$

For any k , if $q_{k-1} = q_k$, then the rate μ_k is not used by the policy γ . From the definition of $\pi_\mu(k)$ we have that

$$\begin{aligned} \pi_\mu(0) &= \pi(0), \\ \pi_\mu(k) &= \sum_{q=q_{k-1}+1}^{q_k} \pi(q), \forall k \geq 1. \end{aligned} \tag{2.6}$$

In the following we obtain lower bounds on the average queue length $\overline{Q}(\gamma)$ as a function of the upper bounds in Lemma 2.3.2 on stationary probabilities of service rates, when $c_c - c(\lambda) = V \downarrow 0$ for the three cases FINITE- μ CHOICE-1, FINITE- μ CHOICE-2, and FINITE- μ CHOICE-3 respectively.

We first consider the case FINITE- μ CHOICE-1, where $k_l = 0$ and $\mu_l = 0$.

Remark 2.3.3. If $\mu_u = \mu_K$, then we note that $c(\mu)$ is a linear function of μ . Then the admissible policy γ , with $\mu(0) = 0$, and $\mu(q) = \mu_K$ for all $q \geq 1$, has $\bar{C}(\gamma) = \mathbb{E}c(\mu(Q)) = c(\mathbb{E}\mu(Q)) = c(\lambda)$. Furthermore, γ has the minimum average queue length $\frac{\lambda}{\mu_K - \lambda}$. Hence in this case there is no tradeoff between the average queue length and average service cost. So in the following, we assume $c(\mu)$ is such that $\mu_u < \mu_K$.

We note that an admissible policy γ , which uses only the service rates $\{\mu_k, k \in \{0, \dots, k_u\}\}$ has $\bar{C}(\gamma) = \sum_{k=0}^{k_u} \pi_\mu(k) c(\mu_k) = \sum_{k=0}^{k_u} \pi_\mu(k) l(\mu_k) = c(\lambda)$. Furthermore, $\bar{C}(\gamma) = c(\lambda)$ is achieved only by admissible policies which uses only the service rates $\{\mu_k, k \in \{0, \dots, k_u\}\}$. Hence, as $\lambda < \mu_u$, the policy γ_u , that uses $\mu(0) = 0$ and $\mu(q) = \mu_u$ for $q \geq 1$ has a service cost of $c(\lambda)$. We note that γ_u has the minimum average queue length $\frac{\lambda}{\mu_u - \lambda}$, among all policies γ for which $\bar{C}(\gamma) = c(\lambda)$. Thus the minimum average queue length, among policies γ for which $\bar{C}(\gamma) \leq c_c$, where $c_c > c(\lambda)$ is at most $\frac{\lambda}{\mu_u - \lambda}$. We note that if $\bar{C}(\gamma) > c(\lambda)$, then service rates μ_k with $k > k_u$ could be used, which could yield an average queue length less than $\frac{\lambda}{\mu_u - \lambda}$. In the following lower bound on $\bar{Q}(\gamma)$, we observe that if $c_c > c(\lambda)$, then the average queue length can be less than $\frac{\lambda}{\mu_u - \lambda}$, but has the limit $\frac{\lambda}{\mu_u - \lambda}$, as $\bar{C}(\gamma) \rightarrow c(\lambda)$.

Lemma 2.3.4. For any sequence of non-idling admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that

$$\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_k) = \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right).$$

In the proof, for any sequence of non-idling admissible policies γ_k , we first show that as $V_k = \bar{C}(\gamma_k) - c(\lambda) \downarrow 0$, the largest queue length q_{k_u} at which service rate μ_u is used, increases to infinity as $\log\left(\frac{1}{V_k}\right)$. This asymptotic lower bound is obtained by showing that a lower bound $q_{k_u, l}$ to q_{k_u} increases as $\log\left(\frac{1}{V_k}\right)$. For a V_k , the lower bound $q_{k_u, l}$ is used to define the policy γ'_k , which has $\mu(0) = 0, \mu(q) = \mu_u$ for $1 \leq q \leq q_{k_u, l}$, and $\mu(q) = \mu_K$ for $q > q_{k_u, l}$. We note that for any policy γ_k with $\bar{C}(\gamma) - c(\lambda) = V_k$, $\bar{Q}(\gamma_k) \geq \bar{Q}(\gamma'_k)$. The sequence $\bar{Q}(\gamma'_k)$, obtained as $V_k \downarrow 0$, is shown to have the asymptotic behaviour in the above lemma.

Proof. Let us consider a particular policy γ in the sequence γ_k with $\bar{C}(\gamma) - c(\lambda) = V$. As $Q(t)$ is a birth-death process we have that

$$\pi(q-1) \leq \pi(q) \frac{\mu_u}{\lambda}, \forall q \in \{1, \dots, q_{k_u}\}.$$

By induction, we obtain that

$$\pi(q_{k_u} - m) \leq \pi(q_{k_u}) \left(\frac{\mu_u}{\lambda}\right)^m \leq \pi(q_{k_u} + 1) \frac{\mu_K}{\lambda} \left(\frac{\mu_u}{\lambda}\right)^m, \forall m \in \{0, \dots, q_{k_u}\}.$$

Now we note that

$$\sum_{q=0}^{q_{k_u}} \pi(q) = \sum_{k=0}^{k_u} \pi_\mu(k) = 1 - \sum_{k=k_u+1}^K \pi_\mu(k)$$

$$\text{but } \sum_{q=0}^{q_{k_u}} \pi(q) = \sum_{m=0}^{q_{k_u}} \pi(q_{k_u} - m) \leq \pi(q_{k_u} + 1) \frac{\mu_K}{\lambda} \sum_{m=0}^{q_{k_u}} \left(\frac{\mu_{k_u}}{\lambda}\right)^m,$$

$$\text{and from Lemma 2.3.2, } 1 - \sum_{k=k_u+1}^K \pi_\mu(k) \geq 1 - \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)}.$$

Hence

$$1 - \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)} \leq \pi(q_{k_u} + 1) \frac{\mu_K}{\lambda} \sum_{m=0}^{q_{k_u}} \left(\frac{\mu_{k_u}}{\lambda}\right)^m,$$

$$\text{or } 1 - \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)} \leq \pi(q_{k_u} + 1) \mu_K \frac{\left(\frac{\mu_{k_u}}{\lambda}\right)^{q_{k_u}+1} - 1}{\mu_{k_u} - \lambda},$$

$$\text{or } \frac{\mu_u - \lambda}{\mu_K} \frac{1}{\pi(q_{k_u} + 1)} \left[1 - \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)} \right] + 1 \leq \left(\frac{\mu_u}{\lambda}\right)^{q_{k_u}+1}.$$

Therefore

$$q_{k_u} \geq \log\left(\frac{\mu_u}{\lambda}\right) \left[\frac{\mu_u - \lambda}{\mu_K} \frac{1}{\pi(q_{k_u} + 1)} \left[1 - \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)} \right] + 1 \right] - 1.$$

But we note that $\pi(q_{k_u} + 1) \leq \sum_{k=k_u+1}^K \pi_\mu(k) \leq \sum_{k=k_u+1}^K \frac{V}{c(\mu_k) - l(\mu_k)} \triangleq \frac{V}{c_1}$. So that

$$q_{k_u} \geq \log\left(\frac{\mu_u}{\lambda}\right) \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1\right) + 1 \right] - 1. \quad (2.7)$$

Therefore, for any non-idling admissible policy γ we have that $q_{k_u} \geq q_{k_u, l}$, where

$$q_{k_u, l} \triangleq \left\lceil \log\left(\frac{\mu_u}{\lambda}\right) \left[\frac{\mu_u - \lambda}{\mu_K} \frac{c_1}{V} \left[1 - \frac{V}{c_1} \right] + 1 \right] - 1 \right\rceil.$$

Now we note that for the policy γ under consideration, $\mu(q) \leq \mu_u$ for $q \in \{0, \dots, q_{k_u}\}$, and $\mu(q) \leq \mu_K$ for $q \in \{q_{k_u} + 1, \dots\}$. Let policy γ' be defined as follows :

$$\mu'(q) = \begin{cases} 0 & \text{if } q = 0, \\ \mu_u & \text{if } q \in \{1, \dots, q_{k_u, l}\}, \\ \mu_K & \text{otherwise.} \end{cases}$$

We note that $\bar{Q}(\gamma') \leq \bar{Q}(\gamma)$. Let the stationary distribution of CTMC under γ' be denoted as

$\pi'(q)$. Let $a \triangleq \frac{\lambda}{\mu_u}$ and $b \triangleq \frac{\lambda}{\mu_K}$. Then we have that

$$\pi'(q) = \begin{cases} \pi'(0)a^q & \text{if } q \in \{0, \dots, q_{k_u,l}\} \\ \pi'(0)a^{q_{k_u,l}}b^{q-q_{k_u,l}} & \text{otherwise.} \end{cases}$$

We note that γ' is also admissible. As $\sum_{q=0}^{\infty} \pi'(q) = 1$, we have that

$$\begin{aligned} \pi'(0) \left[1 + \sum_{q=1}^{q_{k_u,l}} a^q + \sum_{q=q_{k_u,l}+1}^{\infty} a^{q_{k_u,l}} b^{q-q_{k_u,l}} \right] &= 1, \\ \pi'(0) \left[1 + \frac{a}{1-a} (1 - a^{q_{k_u,l}}) + a^{q_{k_u,l}} \frac{b}{1-b} \right] &= 1. \end{aligned} \quad (2.8)$$

We note that $\pi'(0)$ can be obtained from (2.8) in terms of $q_{k_u,l}$. We have

$$\bar{Q}(\gamma') = \sum_{q=0}^{\infty} \pi'(q)q = \pi'(0) \left[\sum_{q=1}^{q_{k_u,l}} qa^q + \sum_{q=q_{k_u,l}+1}^{\infty} qa^{q_{k_u,l}} b^{q-q_{k_u,l}} \right].$$

Simplification leads to

$$\begin{aligned} \bar{Q}(\gamma') &= \pi'(0) \left[\frac{a}{(1-a)^2} ((1-a)(1 - (q_{k_u,l}+1)a^{q_{k_u,l}}) + a(1 - a^{q_{k_u,l}})) + a^{q_{k_u,l}} \left(q_{k_u,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right) \right], \\ &= \frac{\left[\frac{a}{(1-a)^2} ((1-a)(1 - (q_{k_u,l}+1)a^{q_{k_u,l}}) + a(1 - a^{q_{k_u,l}})) + a^{q_{k_u,l}} \left(q_{k_u,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right) \right]}{\left[1 + \frac{a}{1-a} (1 - a^{q_{k_u,l}}) + a^{q_{k_u,l}} \frac{b}{1-b} \right]} \end{aligned} \quad (2.9)$$

where $\pi'(0)$ was obtained from (2.8). At this point, we note that as $V \rightarrow 0$, $q_{k_u,l} \rightarrow \infty$ and this lower bound $\bar{Q}(\gamma')$ to $\bar{Q}(\gamma)$ approaches $\frac{a}{1-a}$, which is what we expect. However, in order to ascertain the behaviour of average queue length as V approaches zero, we need to lower bound the right-hand side of (2.9) for $V > 0$.

The denominator of (2.9) can be bounded above as follows :

$$\begin{aligned} &1 + \frac{a}{1-a} - \frac{a}{1-a} a^{q_{k_u,l}} + \frac{b}{1-b} a^{q_{k_u,l}} \\ &= 1 + \frac{a}{1-a} + a^{q_{k_u,l}} \left[\frac{1}{\frac{1}{b}-1} - \frac{1}{\frac{1}{a}-1} \right]. \end{aligned}$$

As $b < a$, $\frac{1}{\frac{1}{b}-1} < \frac{1}{\frac{1}{a}-1}$, so that the denominator of (2.9) $\leq \frac{1}{1-a}$. After substituting this upper

bound for the denominator in (2.9), we have

$$\begin{aligned}
\overline{Q}(\gamma') &\geq (1-a) \left[\frac{a}{(1-a)^2} \{1 - (1-a)(q_{k_u,l} + 1)a^{q_{k_u,l}} - a \cdot a^{q_{k_u,l}}\} + a^{q_{k_u,l}} \left\{ q_{k_u,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right\} \right], \\
&= \frac{a}{1-a} + (1-a) \left[a^{q_{k_u,l}} \left\{ q_{k_u,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right\} - \frac{a}{(1-a)^2} \{(1-a)(q_{k_u,l} + 1)a^{q_{k_u,l}} + a \cdot a^{q_{k_u,l}}\} \right], \\
&\geq \frac{a}{1-a} - \frac{a}{1-a} a^{q_{k_u,l}} [1 + (1-a)q_{k_u,l}]. \tag{2.10}
\end{aligned}$$

From (2.7), with $a < 1$, we have that

$$\begin{aligned}
a^{q_{k_u,l}} &\leq \frac{1}{a} \log_{\frac{1}{a}} \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1 \right], \\
&= \frac{1}{a} \left[a^{\log_a \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1 \right]} \right]^{\frac{1}{\log_a(1/a)}}, \\
&= \frac{1}{a} \left[\frac{\frac{\mu_u - \lambda}{\mu_K} (c_1 - V) + V}{V} \right]^{-1}, \\
&= \frac{1}{a} \frac{1}{\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1}.
\end{aligned}$$

From the definition of $q_{k_u,l}$ we have that

$$q_{k_u,l} \leq \log_{\left(\frac{\mu_u}{\lambda}\right)} \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1 \right]$$

Therefore, $a^{q_{k_u,l}}(1 + (1-a)q_{k_u,l}) \leq$

$$\frac{1}{a} \frac{1}{\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1} \left(1 + (1-a) \left\{ \log_{\left(\frac{\mu_u}{\lambda}\right)} \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1 \right] \right\} \right)$$

Substituting this in (2.10) we have that $\frac{a}{1-a} - \overline{Q}(\gamma') \leq$

$$\frac{a}{1-a} \left\{ \frac{V}{a} \frac{1}{\frac{\mu_u - \lambda}{\mu_K} c_1 \left(1 - \frac{V}{c_1} \right) + V} \left(1 + (1-a) \left\{ \log_{\left(\frac{\mu_u}{\lambda}\right)} \left[\frac{\mu_u - \lambda}{\mu_K} \left(\frac{c_1}{V} - 1 \right) + 1 \right] \right\} \right) \right\}. \tag{2.11}$$

Thus, for a sequence γ_k such that $V_k \downarrow 0$, we have that $\frac{a}{1-a} - \overline{Q}(\gamma_k) = \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right)$. \square

Corollary 2.3.5. For any sequence of non-idling admissible policies $\gamma_{M,k} \in \Gamma_{a,M}$ with $\overline{C}(\gamma_{M,k}) - c(\lambda) = V_k \downarrow 0$, we have that

$$\frac{\lambda}{\mu_u - \lambda} - \overline{Q}(\gamma_{M,k}) = \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right).$$

Proof. For a k , if $\gamma_{M,k}$ is such that $\overline{C}(\gamma_{M,k}) - c(\lambda) = V_k$, then we have that for every $\gamma_k \in \Gamma(\gamma_{M,k})$,

$\bar{C}(\gamma_k) - c(\lambda) \leq \frac{V_k}{p_{\gamma_k}} = U_{\gamma_k}$. Then

$$\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_k) = \mathcal{O}\left(U_{\gamma_k} \log\left(\frac{1}{U_{\gamma_k}}\right)\right).$$

We note that $\bar{Q}(\gamma_{k,M}) = \mathbb{E}_{p_{\gamma_k}} \bar{Q}(\gamma_k)$. Then applying $\mathbb{E}_{p_{\gamma_k}}$ to LHS and RHS of the above equation we obtain that

$$\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_{k,M}) = \sum p_{\gamma_k} \left[\mathcal{O}\left(U_{\gamma_k} \log\left(\frac{1}{U_{\gamma_k}}\right)\right) \right].$$

We note that by definition, there exists some constant $c > 0$ such that the RHS is

$$\leq \sum p_{\gamma_k} \left[c \left[U_{\gamma_k} \log\left(\frac{1}{U_{\gamma_k}}\right) \right] \right].$$

Since the function $u \log\left(\frac{1}{u}\right)$ is concave, we have that

$$\sum p_{\gamma_k} \left[c \left[U_{\gamma_k} \log\left(\frac{1}{U_{\gamma_k}}\right) \right] \right] \leq c \sum p_{\gamma_k} [U_{\gamma_k}] \log\left(\frac{1}{\sum p_{\gamma_k} [U_{\gamma_k}]}\right).$$

Since $\sum p_{\gamma_k} [U_{\gamma_k}] = V_k$ we have that

$$\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_{M,k}) = \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right).$$

□

We now present asymptotic lower bounds for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3. For an admissible policy γ , to relate the stationary probability distribution $\pi(q)$ to the average queue length $\bar{Q}(\gamma)$, as noted in Section 2.1.1 we make use of the fact that if \bar{q} is such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$, then $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2}$. The choice of $\frac{1}{2}$ here is arbitrary. The best lower bound on $\bar{Q}(\gamma)$ is given by the largest \bar{q} such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$.

Lemma 2.3.6. For any sequence of non-idling admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that

$$\bar{Q}(\gamma_k) = \begin{cases} \Omega\left(\log\left(\frac{1}{V_k}\right)\right) & \text{for FINITE-}\mu\text{CHOICE-2,} \\ \Omega\left(\frac{1}{V_k}\right) & \text{for FINITE-}\mu\text{CHOICE-3.} \end{cases} \quad (2.12)$$

Proof. We note that $k_l \geq 1$ for the cases FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3. Consider a particular policy γ in the given sequence γ_k , with $\bar{C}(\gamma) - c(\lambda) = V$. Then we note that $\sum_{k=0}^{k_l-1} \pi_{\mu}(k) \leq \sum_{k=0}^{k_l-1} \frac{V}{c(\mu_k) - l(\mu_k)} \triangleq \frac{V}{c'_1}$, from the upper bound in Lemma 2.3.2. Therefore, $\sum_{q=0}^{q_{k_l-1}} \pi(q) = \sum_{k=0}^{k_l-1} \pi_{\mu}(k) \leq \frac{V}{c'_1}$. Also, for every $q \in \{0, \dots, q_{k_l-1}\}$, $\pi(q) \leq \frac{V}{c'_1}$. Now we intend to find the largest \bar{q} such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$. But as $\sum_{q=0}^{q_{k_l-1}} \pi(q) \leq \frac{V}{c'_1}$ and

$\pi(q_{k_l-1} + 1) \leq \pi(q_{k_l-1}) \frac{\lambda}{\mu_l} \leq \frac{\lambda V}{\mu_l c'_1}$, the largest such \bar{q} satisfies

$$\sum_{q=q_{k_l-1}+1}^{\bar{q}} \pi(q) \leq \frac{1}{2} - \frac{V}{c_1}$$

In the following, we use an upper bound on $\pi(q)$ which leads to a lower bound \bar{q}_1 on \bar{q} . Since $\pi(q)\lambda = \pi(q+1)\mu(q)$, and for $q > q_{k_l}$, $\frac{\lambda}{\mu(q)} < \frac{\lambda}{\mu_l}$ (from the admissibility of γ), we have that

$$\sum_{q=q_{k_l-1}+1}^{\bar{q}} \pi(q) < \pi(q_{k_l-1}) \sum_{m=1}^{\bar{q}-q_{k_l-1}} \left(\frac{\lambda}{\mu_l}\right)^m.$$

If \bar{q}_1 is the largest integer such that

$$\pi(q_{k_l-1}) \sum_{m=1}^{\bar{q}_1-q_{k_l-1}} \left(\frac{\lambda}{\mu_l}\right)^m \leq \frac{1}{2} - \frac{V}{c'_1}, \quad (2.13)$$

then $\sum_{q=0}^{\bar{q}_1} \pi(q) \leq \frac{1}{2}$ and $\bar{q} \geq \bar{q}_1$.

In the case FINITE- μ CHOICE-2, $\lambda > \mu_l$, so that summing the geometric series in (2.13), we have that

$$\begin{aligned} \left(\frac{\lambda}{\mu_l}\right)^{\bar{q}_1-q_{k_l-1}} - 1 &\leq \frac{\lambda - \mu_l}{\lambda\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1}\right) \text{ or} \\ \bar{q}_1 &\leq q_{k_l-1} + \left\lceil \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{\lambda - \mu_l}{\lambda\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1}\right)\right) \right\rceil, \end{aligned} \quad (2.14)$$

and, in fact,

$$\bar{q}_1 = \left\lceil q_{k_l-1} + \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{\lambda - \mu_l}{\lambda\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1}\right)\right) \right\rceil.$$

Since $q_{k_l-1} \geq 0$, we have that

$$\bar{q}_1 \geq \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{\lambda - \mu_l}{\lambda\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1}\right)\right) - 1$$

Now since $\pi(q_{k_l-1}) \leq \pi_{\mu}(k_l) \leq \frac{V}{c(\mu_{k_l-1}) - l(\mu_{k_l-1})}$ we have that

$$\bar{q}_1 \geq \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{(\lambda - \mu_l)(c(\mu_{k_l-1}) - l(\mu_{k_l-1}))}{\lambda V} \left(\frac{1}{2} - \frac{V}{c'_1}\right)\right) - 1$$

Then we have that

$$\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{1}{2} \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{(\lambda - \mu_l)(c(\mu_{k_l-1}) - l(\mu_{k_l-1}))}{\lambda V} \left(\frac{1}{2} - \frac{V}{c'_1}\right)\right) - \frac{1}{2} \quad (2.15)$$

Thus for FINITE- μ CHOICE-2, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

For FINITE- μ CHOICE-3, $\lambda = \mu_l$, so that from (2.13), instead of (2.14) we have that

$$\bar{q}_1 - q_{k_l-1} \leq \frac{1}{\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1} \right),$$

and, in fact,

$$\bar{q}_1 = \left\lfloor q_{k_l-1} + \frac{1}{\pi(q_{k_l-1})} \left(\frac{1}{2} - \frac{V}{c'_1} \right) \right\rfloor.$$

Proceeding as for FINITE- μ CHOICE-2, we have that

$$\bar{q}_1 \geq \frac{c(\mu_{k_l-1}) - l(\mu_{k_l-1})}{V} \left(\frac{1}{2} - \frac{V}{c'_1} \right).$$

Therefore we have that

$$\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{1}{2} \left[1 + \frac{c(\mu_{k_l-1}) - l(\mu_{k_l-1})}{V} \left(\frac{1}{2} - \frac{V}{c'_1} \right) \right]. \quad (2.16)$$

Hence for FINITE- μ CHOICE-3, we conclude that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$. \square

Corollary 2.3.7. For any sequence of non-idling admissible policies $\gamma_{k,M} \in \Gamma_{a,M}$ with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that

$$\bar{Q}(\gamma_{k,M}) = \begin{cases} \Omega\left(\log\left(\frac{1}{V_k}\right)\right) & \text{for FINITE-}\mu\text{CHOICE-2,} \\ \Omega\left(\frac{1}{V_k}\right) & \text{for FINITE-}\mu\text{CHOICE-3.} \end{cases} \quad (2.17)$$

The proof of this corollary is very similar to that of Corollary 2.3.5, except that the convexity property of the functions $\log\left(\frac{1}{u}\right)$ and $\frac{1}{u}$ are used instead of the concavity of the function $u \log\left(\frac{1}{u}\right)$.

2.3.2 Asymptotic characterization of $Q_M^*(c_c)$

In this section we obtain asymptotic upper bounds for TRADEOFF-M. The sequence of policies that is constructed for FINITE- μ CHOICE-1 is motivated by the policy γ' that was used in the proof of the lower bound. We shall see that a sequence of policies with q_{k_u} scaling as $\log\left(\frac{1}{V}\right)$ gives the correct asymptotic upper bound.

Lemma 2.3.8. For FINITE- μ CHOICE-1, there exists a sequence of non-idling admissible policies γ_k with a sequence $V_k \downarrow 0$ such that $\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_k) = \Theta\left(V_k \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. We first consider a policy γ in the sequence of policies γ_k . The policy γ is defined as follows:

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \mu_u, \text{ for } q \in \{1, \dots, q_{k_u}\}, \\ \mu(q) &= \mu_K, \text{ for } q \in \{q_{k_u} + 1, \dots\}.\end{aligned}$$

The sequence of policies $\gamma_k, k \geq 1$ is obtained by choosing $q_{k_u} = k$. Consider the policy γ . As $\sum_{q=0}^{\infty} \pi(q) = 1$ we have that

$$\pi(0) \left[1 + \frac{\lambda}{\mu_u - \lambda} \left(1 - \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \right) + \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \frac{\lambda}{\mu_K - \lambda} \right] = 1$$

Therefore, we obtain that

$$\begin{aligned}\pi(0) &= \frac{1}{1 + \frac{\lambda}{\mu_u - \lambda} + \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \left[\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda} \right]}, \text{ and} \\ \pi_\mu(K) &= \sum_{q=q_{k_u}+1}^{\infty} \pi(q) = \frac{\left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \frac{\lambda}{\mu_K - \lambda}}{1 + \frac{\lambda}{\mu_u - \lambda} + \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \left[\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda} \right]}.\end{aligned}\tag{2.18}$$

From the definition of $\bar{C}(\gamma)$ we have that

$$\begin{aligned}\bar{C}(\gamma) &= \pi_\mu(k_u)c(\mu_u) + \pi_\mu(K)c(\mu_K), \\ &= \pi_\mu(k_u)(c(\lambda) + m_1(\mu_u - \lambda)) + \pi_\mu(K)(c(\lambda) + m_2(\mu_K - \lambda)), \\ &= c(\lambda)(1 - \pi(0)) + m_1(\mu_u - \lambda)\pi_\mu(k_u) + m_2(\mu_K - \lambda)\pi_\mu(K), \\ &= c(\lambda) - \pi(0)c(\lambda) + m_1(\mu_u - \lambda)(1 - \pi_\mu(K) - \pi(0)) + m_2(\mu_K - \lambda)\pi_\mu(K). \\ \bar{C}(\gamma) - c(\lambda) &= m_1(\mu_u - \lambda) + \pi_\mu(K)[m_2(\mu_K - \lambda) - m_1(\mu_u - \lambda)] - \pi(0)[c(\lambda) + m_1(\mu_u - \lambda)],\end{aligned}\tag{2.19}$$

where $m_1 = \frac{c(\mu_u) - c(\lambda)}{\mu_u - \lambda}$ and $m_2 = \frac{c(\mu_K) - c(\lambda)}{\mu_K - \lambda}$. We denote $\bar{C}(\gamma) - c(\lambda)$ by V . We note that $c(\lambda) + m_1(\mu_u - \lambda) = c(\mu_u)$ and $m_2(\mu_K - \lambda) - m_1(\mu_u - \lambda) = c(\mu_K) - c(\mu_u)$. Then (2.19) can be written as

$$\begin{aligned}V &= m_1(\mu_u - \lambda) + \pi_\mu(K)(c(\mu_K) - c(\mu_u)) - \pi(0)c(\mu_u), \\ V - m_1(\mu_u - \lambda) &= \frac{\left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \frac{\lambda}{\mu_K - \lambda} (c(\mu_K) - c(\mu_u)) - c(\mu_u)}{1 + \frac{\lambda}{\mu_u - \lambda} + \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u}} \left[\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda} \right]}.\end{aligned}$$

Simplifying, we obtain, with $c(\lambda) = m_1\lambda$,

$$\frac{[V - m_1(\mu_u - \lambda)] \left[1 + \frac{\lambda}{\mu_u - \lambda}\right] + c(\mu_u)}{\frac{\lambda}{\mu_K - \lambda} (c(\mu_K) - c(\mu_u)) - \left[\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda}\right] (V - m_1(\mu_u - \lambda))} = \left(\frac{\lambda}{\mu_u}\right)^{q_{k_u}},$$

$$\frac{\bar{C}(\gamma) \frac{\mu_u}{\mu_u - \lambda} - \frac{\lambda}{\mu_u - \lambda} (c(\lambda) + m_1(\mu_u - \lambda))}{\lambda(m_2 - m_1) - \left(\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda}\right) (\bar{C}(\gamma) - c(\lambda))} = \left(\frac{\lambda}{\mu_u}\right)^{q_{k_u}},$$

so that,

$$q_{k_u} = \log\left(\frac{\mu_u}{\lambda}\right) \left[\frac{\lambda(m_2 - m_1) - \left(\frac{\lambda}{\mu_K - \lambda} - \frac{\lambda}{\mu_u - \lambda}\right) V}{V \frac{\mu_u}{\mu_u - \lambda}} \right].$$

We note that the average queue length for the policy γ is of the form given in (2.9), wherein $a = \frac{\lambda}{\mu_u}$ and $b = \frac{\lambda}{\mu_K}$. Simplifying this expression, we obtain that

$$\bar{Q}(\gamma) = \frac{\frac{a}{(1-a)^2} + q_{k_u} a^{q_{k_u}} \left[\frac{b}{1-b} - \frac{a}{1-a}\right] + a^{q_{k_u}} \left[\frac{b}{(1-b)^2} - \frac{a}{(1-a)^2}\right]}{\frac{1}{1-a} + a^{q_{k_u}} \left[\frac{b}{1-b} - \frac{a}{1-a}\right]}. \quad (2.20)$$

We are interested in only order approximations of $\bar{Q}(\gamma)$; so we proceed by considering large k , so that $V \downarrow 0$ and $a^{q_{k_u}} \downarrow 0$. We obtain that (only the dominant terms as $V \downarrow 0$)

$$\bar{Q}(\gamma) \approx \left(\frac{a}{1-a} + q_{k_u} a^{q_{k_u}} (1-a) \left[\frac{b}{1-b} - \frac{a}{1-a} \right] + a^{q_{k_u}} (1-a) \left[\frac{b}{(1-b)^2} - \frac{a}{(1-a)^2} \right] \right) \times \left(1 + a^{q_{k_u}} (1-a) \left[\frac{a}{1-a} - \frac{b}{1-b} \right] \right).$$

Expanding, we obtain

$$\bar{Q}(\gamma) \approx \frac{a}{1-a} + \frac{a}{1-a} a^{q_{k_u}} (1-a) \left[\frac{a}{1-a} - \frac{b}{1-b} \right] + q_{k_u} a^{q_{k_u}} (1-a) \left[\frac{b}{1-b} - \frac{a}{1-a} \right] \quad (2.21)$$

$$+ a^{q_{k_u}} (1-a) \left[\frac{b}{(1-b)^2} - \frac{a}{(1-a)^2} \right] \quad (2.22)$$

$$- q_{k_u} a^{2q_{k_u}} (1-a)^2 \left[\frac{a}{1-a} - \frac{b}{1-b} \right]^2 + a^{2q_{k_u}} (1-a)^2 \left[\frac{a}{1-a} - \frac{b}{1-b} \right] \left[\frac{b}{(1-b)^2} - \frac{a}{(1-a)^2} \right] \quad (2.23)$$

We note that the second term in (2.21) is positive and $\Theta(V)$, while the third term is negative and $\Theta(V \log(\frac{1}{V}))$. The term in (2.22) is negative and $\Theta(V)$. The first term in (2.23) is negative and $\Theta(V^2 \log(\frac{1}{V}))$, while the second term is negative and $\Theta(V^2)$. Thus, the dominating term in $\frac{a}{1-a} - \bar{Q}(\gamma)$ is positive and $\Theta(V \log(\frac{1}{V}))$ as $V \downarrow 0$. Hence we have that $\frac{a}{1-a} - \bar{Q}(\gamma) = \Theta(V \log(\frac{1}{V}))$.

We note that every γ_k constructed, by choosing $q_{k_u} = k$, is non-idling and admissible. So, there

exists a sequence of non-idling admissible policies γ_k such that $\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_k) = \Theta\left(V_k \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$. \square

Using the asymptotic lower bound on $\bar{Q}(\gamma_k)$ from Corollary 2.3.5, and the above asymptotic upper bound, we have the following result.

Proposition 2.3.9. For FINITE- μ CHOICE-1, we have that the optimal value of the tradeoff problem $Q_M^*(c_{c,k})$ is $\frac{\lambda}{\mu_u - \lambda} - \Theta\left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$, for a sequence $c_{c,k} = \bar{C}(\gamma_k)$ for the sequence of policies γ_k in Lemma 2.3.8.

Proof. Consider the sequence $c_{c,k} = V_k + c(\lambda)$. Let us choose a sequence $\epsilon_V \triangleq V_k \log\left(\frac{1}{V_k}\right)$ which decreases to zero as $V_k \downarrow 0$. Let $\gamma_k \in \Gamma_{a,M}$ be a sequence of ϵ_V -optimal policies for FINITE- μ CHOICE such that $\bar{Q}(\gamma_k) \leq Q_M^*(c_{c,k}) + \epsilon_V$. Then, applying Corollary 2.3.5 we have that $\frac{\lambda}{\mu_u - \lambda} - \bar{Q}(\gamma_k) = \mathcal{O}\left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$. Then, as $V_k \downarrow 0$, there exists some constant c_1 such that

$$\bar{Q}(\gamma_k) \geq \frac{\lambda}{\mu_u - \lambda} - c_1 \left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right) \right).$$

Since $\bar{Q}(\gamma_k) \leq Q_M^*(c_{c,k}) + \epsilon_V$, we have that

$$Q_M^*(c_{c,k}) + \epsilon_V \geq \frac{\lambda}{\mu_u - \lambda} - c_1 \left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right) \right).$$

Then we have that

$$Q_M^*(c_{c,k}) \geq \frac{\lambda}{\mu_u - \lambda} - (c_1 + 1) \left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right) \right).$$

We have that $Q_M^*(c_{c,k}) \leq \bar{Q}(\gamma_k)$, where γ_k is the sequence of policies constructed in Lemma 2.3.8. Therefore $Q_M^*(c_{c,k}) = \bar{Q}(\gamma_k) = \frac{\lambda}{\mu_u - \lambda} - \Theta\left((c_{c,k} - c(\lambda)) \log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$. \square

Remark 2.3.10. We note that the asymptotic characterization of $Q^*(c_c)$ has been obtained only for a particular sequence $V_k = \bar{C}(\gamma_k) - c(\lambda)$, where γ_k is as in Lemma 2.3.8. The set of average service cost values that can be achieved depends upon the set of service rates, $\{0, \mu_1, \dots, \mu_K\}$, available for control. For example, if the set of service rates available for control is $\{0, \mu_u, \mu_K\}$, then the average service cost always corresponds to the set of values V_k in Proposition 2.3.9. In fact, the asymptotic Θ characterization of $Q_M^*(c_c)$ can be obtained for any sequence of $c_{c,k}$ such that there exists a sequence of non-idling admissible γ_k such that $\bar{C}(\gamma_k) - c(\lambda) = \Theta(c_{c,k} - c(\lambda))$.

Remark 2.3.11. In this thesis, any sequence of admissible policies γ_k , which achieve the asymptotic lower bound is called an *order-optimal* sequence of policies. For example, the sequence of policies γ_k in Lemma 2.3.8 is order-optimal.

For obtaining bounds on the average queue length for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3, we use a result, presented in Appendix 2.A.1, that uses a quadratic Lyapunov function to obtain bounds on the average queue length.

Lemma 2.3.12. For FINITE- μ CHOICE-2, there exists a sequence of non-idling admissible policies γ_k , with a sequence $V_k \downarrow 0$ such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. Consider a policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \mu_l, \text{ for } q \in \{1, \dots, q_{k_l}\}, \\ \mu(q) &= \mu_u, \text{ for } q \in \{q_{k_l} + 1, \dots\}.\end{aligned}$$

where $q_{k_l} \triangleq \left\lceil \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{\lambda - \mu_l}{\lambda} \frac{1}{U}\right) \right\rceil$, with $U > 0$. The sequence of policies γ_k is obtained by choosing U from a sequence U_k that decreases to zero.

Now we note that for $q \in \{1, \dots, q_{k_l}\}$,

$$\pi(q) = \pi(0) \left(\frac{\lambda}{\mu_l}\right)^q,$$

and for $q \in \{q_{k_l} + 1, \dots\}$

$$\pi(q) = \pi(0) \left(\frac{\lambda}{\mu_l}\right)^{q_{k_l}} \left(\frac{\lambda}{\mu_u}\right)^{q - q_{k_l}}.$$

As $\sum_{q=0}^{\infty} \pi(q) = 1$, we have that

$$\begin{aligned}\pi(0) \left[1 + \sum_{q=1}^{q_{k_l}} \left(\frac{\lambda}{\mu_l}\right)^q + \left(\frac{\lambda}{\mu_l}\right)^{q_{k_l}} \sum_{q=1}^{\infty} \left(\frac{\lambda}{\mu_u}\right)^q \right] &= 1, \\ \pi(0) \left[1 + \frac{\lambda}{\lambda - \mu_l} \left[\left(\frac{\lambda}{\mu_l}\right)^{q_{k_l}} - 1 \right] + \left(\frac{\lambda}{\mu_l}\right)^{q_{k_l}} \frac{\lambda}{\mu_u - \lambda} \right] &= 1.\end{aligned}\tag{2.24}$$

We note that $q_{k_l} \geq \log_{\left(\frac{\lambda}{\mu_l}\right)} \left(1 + \frac{\lambda - \mu_l}{\lambda} \frac{1}{U}\right)$. Since $\frac{\lambda}{\mu_l} > 1$ we have that

$$\left(\frac{\lambda}{\mu_l}\right)^{q_{k_l}} \geq 1 + \frac{\lambda - \mu_l}{\lambda} \frac{1}{U}.$$

Substituting the above lower bound, in (2.24), we have that

$$\begin{aligned}\pi(0) &\leq \frac{U}{U \left(\frac{\mu_u}{\mu_u - \lambda} \right) + 1 \left(\frac{\mu_u - \mu_l}{\mu_u - \lambda} \right)}, \\ \text{or, } \pi(0) &\leq \frac{U}{\left(\frac{\mu_u - \mu_l}{\mu_u - \lambda} \right)}.\end{aligned}\tag{2.25}$$

We note that for γ , $\bar{C}(\gamma) = \pi(0).0 + \pi_\mu(k_l)c(\mu_l) + \pi_\mu(k_u)c(\mu_u)$. Also $c(\mu_l) = c(\lambda) + (\mu_l - \lambda)m$ and $c(\mu_u) = c(\lambda) + (\mu_u - \lambda)m$, where $m = \frac{c(\mu_u) - c(\mu_l)}{\mu_u - \mu_l}$. Then we have that

$$\begin{aligned}\bar{C}(\gamma) &= \pi_\mu(k_l)(c(\lambda) + (\mu_l - \lambda)m) + \pi_\mu(k_u)(c(\lambda) + (\mu_u - \lambda)m), \\ &\leq c(\lambda) + m(\pi_\mu(k_l)(\mu_l - \lambda) + \pi_\mu(k_u)(\mu_u - \lambda)).\end{aligned}$$

We note that γ is admissible, therefore, we have that $\pi_\mu(k_l)\mu_l + \pi_\mu(k_u)\mu_u = \lambda$. Hence,

$$\pi_\mu(k_l)(\mu_l - \lambda) + \pi_\mu(k_u)(\mu_u - \lambda) = \pi(0)\lambda$$

Then,

$$\bar{C}(\gamma) \leq c(\lambda) + m\pi(0)\lambda\tag{2.26}$$

From (2.25), we have that

$$\bar{C}(\gamma) - c(\lambda) = \mathcal{O}(U).$$

Let $V \triangleq \bar{C}(\gamma) - c(\lambda)$, then $V = \mathcal{O}(U)$.

From Proposition 2.A.1, with $q_\epsilon = q_{k_l} + 1$ and $\epsilon = \mu_u - \lambda$, we obtain that

$$\bar{Q}(\gamma) \leq \frac{(q_{k_l} + 1)\mu_u}{\mu_u - \lambda} + \frac{\lambda + r_{max}}{2(\mu_u - \lambda)}.\tag{2.27}$$

Now, for the sequence of policies γ_k with $U_k \downarrow 0$, $q_{k_l} = \mathcal{O}\left(\log\left(\frac{1}{U_k}\right)\right)$. Hence $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{U_k}\right)\right)$ and, since $V_k = \mathcal{O}(U_k)$, $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$. So there exists a sequence of policies γ_k such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$. \square

Using the asymptotic lower bound on $\bar{Q}(\gamma)$ from Corollary 2.3.7, and the asymptotic upper bound above, and proceeding as for Proposition 2.3.9 (except that ϵ_V is a constant $\epsilon > 0$) we obtain the following result.

Proposition 2.3.13. For FINITE- μ CHOICE-2, we have that the optimal value of the tradeoff problem $Q_M^*(c_{c,k})$ is $\Theta\left(\log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$, for a sequence $c_{c,k} = V_k + c(\lambda)$, where $V_k = \bar{C}(\gamma_k) - c(\lambda)$

for the sequence of policies γ_k in Lemma 2.3.12.

The following asymptotic upper bound for FINITE- μ CHOICE-3, is obtained using a procedure similar to that for FINITE- μ CHOICE-2 in Lemma 2.3.12.

Lemma 2.3.14. For FINITE- μ CHOICE-3, there exists a sequence of non-idling admissible policies γ_k with a sequence $V_k \downarrow 0$ such that $\overline{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$ with $\overline{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. Consider a policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \lambda, \text{ for } q \in \{1, \dots, q_\lambda\}, \\ \mu(q) &= \mu', \text{ for } q \in \{q_\lambda + 1, \dots\},\end{aligned}$$

where $\mu' = \min\{\mu_k : \mu_k > \lambda\}$ and q_λ is chosen as $\lceil \frac{1}{U} \rceil$, with $U > 0$. The sequence of policies γ_k is obtained by choosing U from a sequence $U_k \downarrow 0$.

We note that for $q \in \{1, \dots, q_\lambda\}$, as $\mu(q) = \lambda$, we have that $\pi(q) = \pi(0)$. And for $q \in \{q_\lambda + 1, \dots\}$, we have that $\pi(q) = \pi(0) \left(\frac{\lambda}{\mu'}\right)^{q-q_\lambda}$. As $\sum_{q=0}^{\infty} \pi(q) = 1$, we have that

$$\begin{aligned}\pi(0) \left(1 + q_\lambda + \frac{\lambda}{\mu' - \lambda}\right) &= 1, \\ \pi(0) &\leq \frac{U(\mu' - \lambda)}{U\mu' + \mu' - \lambda} \leq U.\end{aligned}\tag{2.28}$$

We note that $\overline{C}(\gamma) = \pi(0).0 + \pi_\mu(\lambda)c(\lambda) + \pi_\mu(\mu')c(\mu')$. We have that $c(\mu') = c(\lambda) + m(\mu' - \lambda)$ where $m = \frac{c(\mu') - c(\lambda)}{\mu' - \lambda}$. We note that γ is admissible. Therefore, $\pi(0).0 + \pi_\mu(\lambda)\lambda + \pi_\mu(\mu')\mu' = \lambda$, or $\pi_\mu(\mu')(\mu' - \lambda) = \pi(0)\lambda$. Thus

$$\begin{aligned}\overline{C}(\gamma) &= \pi_\mu(\lambda)c(\lambda) + \pi_\mu(\mu')(c(\lambda) + m(\mu' - \lambda)), \\ &= c(\lambda) + m\lambda\pi(0) - c(\lambda)\pi(0), \\ &\leq c(\lambda) + (m\lambda - c(\lambda))U.\end{aligned}\tag{2.29}$$

Using (2.28) we obtain that, for the policy γ , $\overline{C}(\gamma) - c(\lambda) = \mathcal{O}(U)$. Let $V \triangleq \overline{C}(\gamma) - c(\lambda)$. Then $V = \mathcal{O}(U)$.

To use Proposition 2.A.1, we set $q_\epsilon = q_\lambda + 1$ and $\epsilon = \mu' - \lambda$. We obtain that

$$\overline{Q}(\gamma) \leq \frac{(q_\lambda + 1)\mu'}{\mu' - \lambda} + \frac{\lambda + r_{max}}{2(\mu' - \lambda)}.\tag{2.30}$$

For the policy γ_k , as $q_\lambda = \mathcal{O}\left(\frac{1}{U_k}\right)$, we have that $\overline{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{U_k}\right)$. Since $V_k = \mathcal{O}(U_k)$, we have

that $\overline{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. Hence there exists a sequence of non-idling admissible policies γ_k , with $\overline{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$ and $\overline{C}(\gamma_k) - c(\lambda) = V_k$. \square

Using the asymptotic lower bound on $\overline{Q}(\gamma_k)$ from Corollary 2.3.7, and the asymptotic upper bound above, and proceeding as in Proposition 2.3.9 (except that ϵ_V is a constant $\epsilon > 0$), we obtain the following result.

Proposition 2.3.15. For FINITE- μ CHOICE-3, we have that the optimal value of the tradeoff problem $Q_M^*(c_{c,k})$ is $\Theta\left(\frac{1}{c_{c,k} - c(\lambda)}\right)$, for a sequence $c_{c,k} = V_k + c(\lambda)$, where $V_k = \overline{C}(\gamma_k) - c(\lambda)$ for the sequence of policies γ_k in Lemma 2.3.14.

2.3.3 Asymptotic characterization of order optimal admissible policies

The above approach that characterizes the asymptotic behaviour of the average queue length, via an upper bound on the stationary probability distribution of the queue length for admissible policies, is used throughout this thesis. In this section, we discuss an advantage that this approach has over methods proposed in [7] or [44]. We illustrate how the above approach can be used to obtain an *asymptotic characterization* of any sequence of order-optimal admissible policies γ_k for which $\overline{C}(\gamma_k) \downarrow c(\lambda)$. Such an asymptotic characterization may lead to a reduction in the search space for the ϵ -optimal admissible policy for (2.4) as $c_c \approx c(\lambda)$.

We note that an admissible policy is equivalently described by the sequence $(q_0 = 0, q_1, \dots, q_{K-1}, q_K = \infty)$. In this section, we discuss how asymptotic bounds on q_k can be obtained for an admissible policy with $\overline{C}(\gamma) \approx c(\lambda)$. We consider only FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3 in this section, since these are the cases for which the design of policies is more critical (since $Q^*(c_c)$ increases to infinity as $c_c \downarrow c(\lambda)$).

We note that for special classes of admissible policies, such as buffer-partitioning policies proposed in [7], intuition about q_k can be obtained from the asymptotic characterization of $\overline{Q}(\gamma)$ for a γ such that $\overline{C}(\gamma) \approx c(\lambda)$. We note that buffer partitioning policies use only a specific set of rates, e.g., for FINITE- μ CHOICE-2, only the rates $\{0, \mu_l, \mu_u\}$ are used. Since $\mu_l < \lambda < \mu_u$, intuitively, we expect that the partition $q_{k_l} \approx \overline{Q}(\gamma)$, which implies that $q_{k_l} = \Omega\left(\log\left(\frac{1}{c_c - c(\lambda)}\right)\right)$ for any feasible policy as $c_c \downarrow c(\lambda)$. We note that the above asymptotic behaviour for the partition can be surmised from the asymptotic behaviour of $\overline{Q}(\gamma)$, which may be obtained via methods as in [43] or [44]. However, in the following, we discuss how bounds on the stationary probability distribution of the queue length are useful in obtaining a much more refined asymptotic characterization of the policy.

Two inequalities:

In this section, we discuss two inequalities, which form the basis for the asymptotic characterization of any sequence of order-optimal policies. Let $q^1 \leq q^2$. Let $P_l \{q^1 \leq Q \leq q^2\}$ and $P_u \{q^1 \leq Q \leq q^2\}$ be any lower bound and upper bound to $Pr \{q^1 \leq Q \leq q^2\}$, i.e., $P_l \{q^1 \leq Q \leq q^2\} \leq Pr \{q^1 \leq Q \leq q^2\} \leq P_u \{q^1 \leq Q \leq q^2\}$. Also let $\pi_l(q)$ and $\pi_u(q)$ be any lower bound and upper bound to $\pi(q)$, i.e., $\pi_l(q) \leq \pi(q) \leq \pi_u(q)$. The asymptotic characterization of any policy is obtained using the following two observations. If q_l is the largest integer such that

$$\sum_{q=q^1}^{q^1+q_l} \pi_u(q) \leq P_l \{q^1 \leq Q \leq q^2\}, \quad (2.31)$$

then $q_l \leq q^2 - q^1$. If q_u is the smallest integer such that

$$\sum_{q=q^1}^{q^1+q_u} \pi_l(q) \geq P_u \{q^1 \leq Q \leq q^2\}, \quad (2.32)$$

then $q_l \geq q^2 - q^1$. We note that bounds on $q_k - q_{k-1} + 1$, which is the set of queue lengths for which the service rate μ_k is used, can be obtained by using $q^1 = q_{k-1} + 1$ and $q^2 = q_k$.

The bounds - $P_l \{.\}$, $P_u \{.\}$, $\pi_l(\cdot)$, and $\pi_u(\cdot)$:

In the following we discuss how $P_l \{.\}$, $P_u \{.\}$, $\pi_l(\cdot)$, and $\pi_u(\cdot)$ can be obtained. We note that we consider only cases where $q^1 = q_{k'} + 1$ and $q^2 = q_k$, where $k' < k$. Then $Pr \{q^1 \leq Q \leq q^2\} = Pr \{\mu_{k'+1} \leq \mu(Q) \leq \mu_k\} = \sum_{n=k'+1}^k \pi_\mu(n)$.

Consider any admissible policy γ . Then for any k , we can obtain a lower bound on $\pi_\mu(k)$ as the optimal value of the linear program:

$$\min \quad \pi_\mu(k) \quad (2.33)$$

$$\text{such that} \quad \sum_{k=0}^K \pi_\mu(k) = 1, \quad (2.34)$$

$$\sum_{k=0}^K \pi_\mu(k) \mu_k = \lambda, \quad (2.35)$$

$$\sum_{k=0}^K \pi_\mu(k) (c(\mu_k) - l(\mu_k)) = \bar{C}(\gamma) - c(\lambda). \quad (2.36)$$

An upper bound on $\pi_\mu(k)$ can be obtained by maximising $\pi_\mu(k)$ in the above linear program. However, we note that finding the above bounds analytically is difficult. Hence, in the following we find other bounds on $\pi_\mu(k)$ or $\sum_{n=k'+1}^k \pi_\mu(n)$, which can be expressed analytically. We note that

these analytical bounds are obtained from the constraints in the above linear program.

We note that for $k < k_l$ and $k > k_u$, from Lemma 2.3.2 we have that $\pi_\mu(k) \leq \frac{\bar{C}(\gamma) - c(\lambda)}{c(\mu_k) - l(\mu_k)}$. Proceeding as in the proof of Lemma 2.3.2 we can also show that if $k < k_l$ or $k' \geq k_u$, then

$$\sum_{n=k'+1}^k \pi_\mu(k) \leq \frac{\bar{C}(\gamma) - c(\lambda)}{\min_{n \in \{k'+1, \dots, k\}} (c(\mu_n) - l(\mu_n))}. \quad (2.37)$$

If $k_l = k_u$, as in FINITE- μ CHOICE-3, then from (2.34) and (2.36), we have that

$$\pi_\mu(k_l) = 1 - \sum_{k \neq k_l} \pi_\mu(k) \geq 1 - \frac{\bar{C}(\gamma) - c(\lambda)}{\min_{k \neq k_l} (c(\mu_k) - l(\mu_k))}. \quad (2.38)$$

If $k_l < k_u$, as in FINITE- μ CHOICE-2, then again from (2.34) and (2.36), we have that

$$\sum_{k=k_l}^{k_u} \pi_\mu(k) \geq 1 - \frac{\bar{C}(\gamma) - c(\lambda)}{\min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \quad (2.39)$$

Suppose we are interested in $\sum_{k=k_l}^{\tilde{k}} \pi_\mu(k)$. Then we have that

$$\sum_{k=k_l}^{\tilde{k}} \pi_\mu(k) + \sum_{k=\tilde{k}+1}^{k_u} \pi_\mu(k) \geq 1 - \frac{\bar{C}(\gamma) - c(\lambda)}{\min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \quad (2.40)$$

If $\tilde{k} = k_u$, then we note that the RHS is a lower bound. If $\tilde{k} < k_u$, then we proceed as follows

$$\sum_{k=k_l}^{\tilde{k}} \pi_\mu(k) \geq 1 - \sum_{k=\tilde{k}+1}^{k_u} \pi_\mu(k) - \frac{\bar{C}(\gamma) - c(\lambda)}{\min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \quad (2.41)$$

From (2.35), we have that

$$\begin{aligned} \sum_{k=k_l}^{\tilde{k}} \pi_\mu(k) \mu_k + \sum_{k=\tilde{k}+1}^{k_u} \pi_\mu(k) \mu_k &\leq \lambda, \\ \mu_l \sum_{k=k_l}^{\tilde{k}} \pi_\mu(k) + \mu_{\tilde{k}} \sum_{k=\tilde{k}+1}^{k_u} \pi_\mu(k) &\leq \lambda, \end{aligned}$$

Or,

$$\sum_{k=\tilde{k}+1}^{k_u} \pi_\mu(k) \leq \frac{\lambda - \mu_l \sum_{k=k_l}^{\tilde{k}} \pi_\mu(k)}{\mu_{\tilde{k}}}.$$

Therefore, we have that

$$\begin{aligned} \sum_{k=k_l}^{\tilde{k}} \pi_{\mu}(k) &\geq 1 - \frac{\lambda - \mu_l \sum_{k=k_l}^{\tilde{k}} \pi_{\mu}(k)}{\mu_{\tilde{k}}} - \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \\ \sum_{k=k_l}^{\tilde{k}} \pi_{\mu}(k) &\geq \frac{\mu_{\tilde{k}} - \lambda}{\mu_{\tilde{k}} - \mu_l} - \mu_{\tilde{k}} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \end{aligned} \quad (2.42)$$

We note that this lower bound is non-negative only if $\mu_{\tilde{k}} > \lambda$. We note that

$$\begin{aligned} \sum_{k < k_l, k > k_u} \pi_{\mu}(k) (c(\mu_k) - l(\mu_k)) &= \overline{C}(\gamma) - c(\lambda), \\ \sum_{k < k_l, k > k_u} \pi_{\mu}(k) &\geq \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \end{aligned} \quad (2.43)$$

We note that for FINITE- μ CHOICE-2, the best lower bound that can be obtained for $\sum_{k < k_l} \pi_{\mu}(k)$ and $\sum_{k > k_u} \pi_{\mu}(k)$ separately is zero. For example, if we try to obtain a lower bound on $\sum_{k < k_l} \pi_{\mu}(k)$, since the constraints (2.34), (2.35), and (2.36) can be met by assigning positive probability to $\mu_k \geq k_l$ only, we obtain the trivial lower bound that $\sum_{k < k_l} \pi_{\mu}(k) \geq 0$. However, for FINITE- μ CHOICE-3 we have that

$$\begin{aligned} \sum_{k < k_l} \pi_{\mu}(k) (\mu_k - \lambda) + \sum_{k > k_u} \pi_{\mu}(k) (\mu_k - \lambda) &= 0, \text{ from (2.35),} \\ \sum_{k > k_u} \pi_{\mu}(k) (\mu_k - \lambda) &= \sum_{k < k_l} \pi_{\mu}(k) (\lambda - \mu_k). \end{aligned} \quad (2.44)$$

To obtain a lower bound on $\sum_{k > k_u} \pi_{\mu}(k)$ we proceed as follows. From (2.44)

$$\begin{aligned} (\mu_K - \lambda) \sum_{k > k_u} \pi_{\mu}(k) &\geq (\lambda - \mu_{k_l-1}) \sum_{k < k_l} \pi_{\mu}(k), \\ \frac{\mu_K - \lambda}{\lambda - \mu_{k_l-1}} \sum_{k > k_u} \pi_{\mu}(k) + \sum_{k > k_u} \pi_{\mu}(k) &\geq \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \text{ from (2.43),} \end{aligned}$$

Or, we have that

$$\sum_{k > k_u} \pi_{\mu}(k) \geq \frac{\lambda - \mu_{k_l-1}}{\mu_K - \mu_{k_l-1}} \left(\frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))} \right). \quad (2.45)$$

To obtain a lower bound on $\sum_{k < k_l} \pi_{\mu}(k)$ we proceed as follows from (2.44)

$$\begin{aligned} (\mu_{k_u+1} - \lambda) \sum_{k > k_u} \pi_{\mu}(k) &\leq \lambda \sum_{k < k_l} \pi_{\mu}(k), \\ \frac{\lambda}{\mu_{k_u+1} - \lambda} \sum_{k < k_l} \pi_{\mu}(k) + \sum_{k < k_l} \pi_{\mu}(k) &\geq \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \text{ from (2.43),} \end{aligned}$$

Or, we have that

$$\sum_{k < k_l} \pi_\mu(k) \geq \frac{\mu_{k_u+1} - \lambda}{\mu_{k_u+1}} \left(\frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))} \right). \quad (2.46)$$

We now consider a method to obtain an upper bound on $\sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k)$. We note that from (2.34) we have

$$\begin{aligned} \sum_{k < k_l, k > k_u} \pi_\mu(k) + \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) + \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) &= 1, \\ \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) + \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) &\leq 1 - \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \text{ from (2.43)}. \end{aligned} \quad (2.47)$$

We note that RHS is the upper bound if $k_l = k_u$ or $\tilde{k} = k_u$. Suppose $\tilde{k} < k_u$, then we proceed as follows. From (2.35), we have that

$$\begin{aligned} \sum_{k < k_l} \pi_\mu(k) \mu_k + \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) \mu_k + \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) \mu_k + \sum_{k_u < k} \pi_\mu(k) \mu_k &= \lambda, \\ \mu_{k_l-1} \sum_{k < k_l} \pi_\mu(k) + \mu_{\tilde{k}} \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) + \mu_u \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) + \mu_K \sum_{k_u < k} \pi_\mu(k) &\geq \lambda, \end{aligned}$$

Or we have that

$$\begin{aligned} \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) &\geq \frac{\lambda}{\mu_u} - \frac{\mu_{\tilde{k}}}{\mu_u} \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) - \frac{\mu_{k_l-1}}{\mu_u} \sum_{k < k_l} \pi_\mu(k) - \frac{\mu_K}{\mu_u} \sum_{k_u < k} \pi_\mu(k), \\ \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) &\geq \frac{\lambda}{\mu_u} - \frac{\mu_{\tilde{k}}}{\mu_u} \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) - \frac{\mu_{k_l-1}}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n < k_l} (c(\mu_n) - l(\mu_n))} - \frac{\mu_K}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n > k_u} (c(\mu_n) - l(\mu_n))}. \end{aligned}$$

From (2.47) we have that

$$\begin{aligned} \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) &\leq 1 - \sum_{\tilde{k} < k \leq k_u} \pi_\mu(k) - \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \\ &\leq 1 - \frac{\lambda}{\mu_u} + \frac{\mu_{\tilde{k}}}{\mu_u} \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) + \frac{\mu_{k_l-1}}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n < k_l} (c(\mu_n) - l(\mu_n))} + \\ &\quad \frac{\mu_K}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n > k_u} (c(\mu_n) - l(\mu_n))} - \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \end{aligned}$$

Or, we have that

$$\begin{aligned}
\left(1 - \frac{\mu_{\tilde{k}}}{\mu_u}\right) \sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) &\leq 1 - \frac{\lambda}{\mu_u} + \frac{\mu_{k_l-1}}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n < k_l} (c(\mu_n) - l(\mu_n))} + \\
&\quad \frac{\mu_K}{\mu_u} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n > k_u} (c(\mu_n) - l(\mu_n))} - \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}, \\
\sum_{k_l \leq k \leq \tilde{k}} \pi_\mu(k) &\leq \frac{\mu_u - \lambda}{\mu_u - \mu_{\tilde{k}}} + \mu_{k_l-1} \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n < k_l} (c(\mu_n) - l(\mu_n))} + \\
&\quad \mu_K \frac{\overline{C}(\gamma) - c(\lambda)}{\min_{n > k_u} (c(\mu_n) - l(\mu_n))} - \mu_u \frac{\overline{C}(\gamma) - c(\lambda)}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))}. \tag{2.48}
\end{aligned}$$

We note that the above bound is less than one, in the limit as $\overline{C}(\gamma) \downarrow c(\lambda)$ only if $\mu_{\tilde{k}} < \lambda$.

We note that the above bounds can be used to obtain upper and lower bounds on $\sum_{k=0}^{\tilde{k}} \pi_\mu(k)$ in many cases. We then obtain the lower and upper bounds $P_l\{\cdot\}$ and $P_u\{\cdot\}$ using $\sum_{k=0}^{\tilde{k}} \pi_\mu(k)$. Now we discuss how $\pi_l(q)$ and $\pi_u(q)$ can be obtained.

We note that if $q^1 = q_{k'} + 1$ and $q^2 = q_k$ for $k' < k$, then for any q such that $q^1 \leq q \leq q^2$, we have that

$$\begin{aligned}
\pi(q) &\leq \pi(q_{k'}) \left(\frac{\lambda}{\mu_{k'+1}}\right)^{q-q_{k'}}, \text{ or,} \\
\pi(q) &\leq \pi(q_k) \left(\frac{\mu_k}{\lambda}\right)^{q_k-q}.
\end{aligned}$$

We can then bound $\pi(q_{k'})$ by $\pi_\mu(k')$ or $\pi(q_k)$ by $\pi_\mu(k)$, which leads to an upper bound $\pi_u(q)$ for $q^1 \leq q \leq q^2$. We note that a similar upper bound on $\pi(q)$ has been used in the asymptotic analysis of $Q^*(c_c)$ above.

We are only able to obtain asymptotic lower bounds on $\pi(0)$, for any sequence of non-idling order optimal admissible policies γ_k . The asymptotic lower bounds are obtained using the same method as in the proof of [44, Theorem 2]. We consider the DTMC $Q_d[m]$ which is obtained by uniformizing $Q(t)$ at rate r_u , as in Appendix 2.A. Then, we note that the stationary distribution of $Q_d[m]$ is the same as that of $Q(t)$ under the policy γ . We proceed as in [44] by assuming that the process is stationary at $m = 0$. Then, from Markov inequality we have that the probability that $Q_d[0] \leq \lceil 2\overline{Q}(\gamma) \rceil$ and there are no arrivals (or up-transitions for $Q_d[m]$) in k successive slots is at least $\frac{1}{2} \left(1 - \frac{\lambda}{r_u}\right)^k$. Suppose $k = \lceil 2\overline{Q}(\gamma) \rceil + 1$. Then we have that the probability that there is no service (or down-transition for $Q_d[m]$) in k slots is at least $\frac{1}{2} \left(1 - \frac{\lambda}{r_u}\right)^k$. Therefore, we have that $\pi(0) \geq \frac{\left(1 - \frac{\lambda}{r_u}\right)^k}{2k}$. With our choice of k , we have that

$$\pi(0) \geq \frac{\left(1 - \frac{\lambda}{r_u}\right)^{\lceil 2\overline{Q}(\gamma) \rceil + 1}}{2 \left(\lceil 2\overline{Q}(\gamma) \rceil + 1\right)}. \tag{2.49}$$

We note that for the policy under consideration, for FINITE- μ CHOICE-2 $\bar{Q}(\gamma) = \Theta\left(\log\left(\frac{1}{V}\right)\right)$ and for FINITE- μ CHOICE-3 $\bar{Q}(\gamma) = \Theta\left(\frac{1}{V}\right)$, as $V \downarrow 0$. Therefore for small enough V , we have that

$$\begin{aligned} \pi(0) &\geq \frac{\left(1 - \frac{\lambda}{r_u}\right)^{(2\bar{Q}(\gamma)+2)}}{4\bar{Q}(\gamma)}, \\ \pi(0) &\geq \pi_l(0) = \begin{cases} \Omega\left(\frac{V}{\log\left(\frac{1}{V}\right)}\right), & \text{for FINITE-}\mu\text{CHOICE-2,} \\ \Omega\left(V\left(1 - \frac{\lambda}{r_u}\right)^{\frac{1}{V}}\right), & \text{for FINITE-}\mu\text{CHOICE-3.} \end{cases} \end{aligned} \quad (2.50)$$

We note that an asymptotic lower bound $\pi_l(q)$ can then be obtained since $\pi(q) \geq \pi_l(q) = \pi_l(0) \left(\frac{\lambda}{\mu_K}\right)^q$, but this bound is very weak in most cases.

Asymptotic characterization:

In this section, we obtain asymptotic bounds on q_k . We note that since we are not able to obtain analytical forms for $P_l\{\cdot\}$, $P_u\{\cdot\}$, $\pi_l(\cdot)$, and $\pi_u(\cdot)$ in all cases, we are not able to obtain asymptotic bounds on $q_k, \forall k \in \{1, \dots, K\}$. The asymptotic bounds on q_k are obtained using the methodology described in Section 2.3.3.

Proposition 2.3.16. For FINITE- μ CHOICE-2, for any sequence of non-idling order-optimal admissible policies γ_k , with $\bar{C}(\gamma_k) - c(\lambda) = V_k$, we have that

$$q_{\tilde{k}} = \begin{cases} \mathcal{O}\left(\log\left(\log\left(\frac{1}{V_{\tilde{k}}}\right)\right)\right), & \text{if } 1 \leq \tilde{k} \leq k_l - 1, \\ \mathcal{O}\left(\log\left(\frac{1}{V_{\tilde{k}}}\right)\right), & \text{if } k_l \leq \tilde{k} \text{ and } \mu_{\tilde{k}} \leq \lambda, \\ \Omega\left(\log\left(\frac{1}{V_{\tilde{k}}}\right)\right), & \text{if } \tilde{k} \text{ is such that } \mu_{\tilde{k}} > \lambda. \end{cases}$$

Proof. Consider any policy γ in the sequence with $V_{\tilde{k}} = V$. Consider any \tilde{k} such that $1 \leq \tilde{k} \leq k_l - 1$. From (2.37), we have that

$$\sum_{k=0}^{\tilde{k}} \pi_{\mu}(k) \leq \frac{V}{\min_{n \in \{0, \dots, \tilde{k}\}} (c(\mu_n) - l(\mu_n))} \triangleq \frac{V}{c_1}.$$

From (2.32), we have that if $q_{\tilde{k}, u}$ is the smallest integer such that

$$\sum_{q=0}^{q_{\tilde{k}, u}} \pi_l(q) \geq \frac{V}{c_1},$$

then $q_{\tilde{k}} \leq q_{\tilde{k},u}$. Substituting $\pi_l(q) = \pi_l(0) \left(\frac{\lambda}{\mu_{\tilde{k}}}\right)^q$, we have that

$$\begin{aligned} \pi_l(0) \sum_{q=0}^{q_{\tilde{k},u}} \left(\frac{\lambda}{\mu_{\tilde{k}}}\right)^q &\geq \frac{V}{c_1}, \\ \pi_l(0) \left(\left(\frac{\lambda}{\mu_{\tilde{k}}}\right)^{q_{\tilde{k},u}+1} - 1 \right) &\geq \frac{\lambda - \mu_{\tilde{k}}}{\mu_{\tilde{k}}} \frac{V}{c_1}. \end{aligned}$$

Or we have that $q_{\tilde{k},u}$ is the smallest integer such that

$$q_{\tilde{k},u} \geq \log \left(\frac{\lambda}{\mu_{\tilde{k}}} \left[1 + \frac{\lambda - \mu_{\tilde{k}}}{\mu_{\tilde{k}}} \frac{V}{c_1 \pi_l(0)} \right] \right).$$

Since $\pi_l(0) = \Omega \left(\frac{V}{\log(\frac{1}{V})} \right)$ from (2.50), we have that $q_{\tilde{k}} \leq q_{\tilde{k},u} = \mathcal{O}(\log(\log(\frac{1}{V})))$. We do not have any asymptotic lower bounds for $q_{\tilde{k}}$, $1 \leq \tilde{k} \leq k_l - 1$.

Let us now consider \tilde{k} such that $k_l \leq \tilde{k}$ and $\mu_{\tilde{k}} < \lambda$. Then from (2.37) and (2.48), we have that

$$\begin{aligned} \Pr \{0 \leq Q \leq q_{\tilde{k}}\} &= \sum_{k=0}^{\tilde{k}} \pi_{\mu}(k) = \sum_{k=0}^{k_l-1} \pi_{\mu}(k) + \sum_{k=k_l}^{\tilde{k}} \pi_{\mu}(k), \\ &\leq \frac{V}{\min_{n \in \{0, \dots, k_l-1\}} (c(\mu_n) - l(\mu_n))} + \frac{\mu_u - \lambda}{\mu_u - \mu_{\tilde{k}}} + \mu_{k_l-1} \frac{V}{\min_{n < k_l} (c(\mu_n) - l(\mu_n))} + \\ &\mu_K \frac{V}{\min_{n > k_u} (c(\mu_n) - l(\mu_n))} - \mu_u \frac{V}{\max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))} = P_u \{0 \leq Q \leq q_{\tilde{k}}\}. \end{aligned}$$

We also have that $\forall q \leq q_{\tilde{k}}$, $\pi_l(q) = \pi_l(0) \left(\frac{\lambda}{\mu_{\tilde{k}}}\right)^q$. Then, from (2.32), if $q_{\tilde{k},u}$ is the smallest integer such that

$$\sum_{q=0}^{q_{\tilde{k},u}} \pi_l(q) \geq \frac{\mu_u - \lambda}{\mu_u - \mu_{\tilde{k}}} + \mathcal{O}(V),$$

then $q_{\tilde{k}} \leq q_{\tilde{k},u}$. Then, we obtain that $q_{\tilde{k},u} = \mathcal{O}(\log(\frac{1}{V}))$. Therefore, for \tilde{k} such that $k_l \leq \tilde{k} < \lambda$, $q_{\tilde{k}} = \mathcal{O}(\log(\frac{1}{V}))$.

Now consider \tilde{k} such that $\tilde{k} \leq k_u$ and $\lambda < \mu_{\tilde{k}}$. From (2.42), we have that

$$\sum_{k=k_l}^{\tilde{k}} \pi_{\mu}(k) \geq \frac{\mu_{\tilde{k}} - \lambda}{\mu_{\tilde{k}} - \mu_l} - \frac{\mu_{\tilde{k}}}{c_1} V,$$

where $c_1 = \min_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))$. Then

$$\sum_{k=0}^{\tilde{k}} \pi_{\mu}(k) \geq P_l \{0 \leq Q \leq q_{\tilde{k}}\} \triangleq \frac{\mu_{\tilde{k}} - \lambda}{\mu_{\tilde{k}} - \mu_l} - \frac{\mu_{\tilde{k}}}{c_1} V,$$

For $0 \leq q \leq q_{\tilde{k}}$, we have that $\pi(q) \leq \pi_u(q) = \pi_u(0) \left(\frac{\lambda}{\mu_1}\right)^q$. From (2.31) if $q_{\tilde{k},l}$ is the largest integer such that

$$\sum_{q=0}^{q_{\tilde{k},l}} \pi_u(q) \leq P_l \{0 \leq Q \leq q_{\tilde{k}}\},$$

then $q_{\tilde{k},l} \leq q_{\tilde{k}}$. Substituting for $\pi_u(q)$ and $P_l \{0 \leq Q \leq q_{\tilde{k}}\}$, we have that

$$\sum_{q=0}^{q_{\tilde{k},l}} \pi_u(0) \left(\frac{\lambda}{\mu_1}\right)^q \leq \frac{\mu_{\tilde{k}} - \lambda}{\mu_{\tilde{k}} - \mu_l} - \frac{\mu_{\tilde{k}}}{c_1} V.$$

Since $\pi_u(0) = \mathcal{O}(V)$, we have that $q_{\tilde{k},l} = \Omega\left(\log\left(\frac{1}{V}\right)\right)$. \square

Proposition 2.3.17. For FINITE- μ CHOICE-3, for any sequence of non-idling order-optimal admissible policies γ_k , with $\bar{C}(\gamma_k) - c(\lambda) = V_k$, we have that

$$q_{\tilde{k}} = \begin{cases} \mathcal{O}\left(\frac{1}{V_k}\right), & \text{if } \mu_{\tilde{k}} < \lambda, \\ \Omega(1), & \text{if } \mu_{\tilde{k}} = \mu_{k_l-1}, \\ \mathcal{O}\left(\frac{1}{V\left(1-\frac{\lambda}{r_u}\right)^{\frac{1}{V}}}\right), & \text{if } \mu_{\tilde{k}} = \lambda, \\ \Omega\left(\frac{1}{V_k}\right), & \text{if } \mu_{\tilde{k}} \geq \lambda. \end{cases}$$

Proof. The methods used in this proof are similar to that used for the proof of Proposition 2.3.16. We first consider \tilde{k} such that $\mu_{\tilde{k}} < \lambda$. Since

$$\sum_{k=0}^{\tilde{k}} \pi_{\mu}(k) \leq \frac{V}{\min_{n \in \{0, \dots, \tilde{k}\}} (c(\mu_n) - l(\mu_n))} \triangleq \frac{V}{c_1},$$

from (2.37) and $\pi_l(0) = \Omega\left(V\left(1-\frac{\lambda}{r_u}\right)^{\frac{1}{V}}\right)$ from (2.50), we have that $q_{\tilde{k},u} = \mathcal{O}\left(\frac{1}{V}\right)$. From (2.46), for $\tilde{k} = k_l - 1$, we have that

$$\sum_{0 \leq k \leq k_l-1} \pi_{\mu}(k) \geq \frac{\mu_{k_u+1} - \lambda}{\mu_{k_u+1}} \left(\frac{V}{c_2}\right),$$

where $c_2 = \max_{k < k_l, k > k_u} (c(\mu_k) - l(\mu_k))$. Then using $\pi_u(q) = \pi_u(0) \left(\frac{\lambda}{\mu_1}\right)^q$, we have that $q_{\tilde{k},l} = \Omega(1)$, so that $q_{\tilde{k}} = \Omega(1)$.

Consider \tilde{k} such that $\mu_{\tilde{k}} = \lambda$. We note that $\pi_{\mu}(\lambda) \leq 1$, then $\pi_l(q) = \pi_l(0), \forall q \leq q_{\tilde{k}}$. Then we obtain that $q_{\tilde{k},u} = \mathcal{O}\left(\frac{1}{V\left(1-\frac{\lambda}{r_u}\right)^{\frac{1}{V}}}\right)$.

Ex. Identifier	λ	$c(\lambda)$	Type	Tradeoff plot
E-111	0.25	0.125	FINITE- μ CHOICE-1	2.3
E-112	0.50	0.250	FINITE- μ CHOICE-3	2.5
E-113	0.75	0.625	FINITE- μ CHOICE-2	2.7

Table 2.2: Numerical examples with $c(\mu) = \mu^2$ for $\mu \in \{0, 0.5, 1\}$.

Ex. Identifier	λ	$c(\lambda)$	Type	Tradeoff plot
E-121	0.25	0.015625	FINITE- μ CHOICE-1	2.4
E-122	0.50	0.031250	FINITE- μ CHOICE-3	2.6
E-123	0.75	0.515600	FINITE- μ CHOICE-2	2.8

Table 2.3: Numerical examples with $c(\mu) = \mu^5$ for $\mu \in \{0, 0.5, 1\}$.

Now consider any \tilde{k} such that $\mu_{\tilde{k}} \geq \lambda$. We have that $\sum_{k=0}^{\tilde{k}} \pi_{\mu}(k) \geq \pi_{\mu}(\lambda) \geq 1 - \frac{V}{c_1}$, where $c_1 = \min_{k < k_l, k > k_l} \{c(\mu_k) - l(\mu_k)\}$. Since $\pi_u(q) = \pi_u(q_{k_l-1} + 1)$, for all $q > q_{k_l-1}$, we have that $q_{k_l, l} = \Omega\left(\frac{1}{V}\right)$. \square

2.3.4 Numerical examples

In this section, we consider some numerical examples for FINITE- μ CHOICE. In the examples, we compare the bounds on $Q^*(c_c)$ which were obtained above, with the optimal tradeoff curve for FINITE- μ CHOICE, which is obtained by the numerical solution of an MDP, obtained by uniformization as in [5]. We now state, the chosen parameters for each numerical example and the quantities plotted in the corresponding figures. Each numerical example is identified by “E-abc”, where a, b, and c are numbers. For each numerical example, we choose the set of service rates and the service cost function $c(\mu)$. Then we consider a set of arrival rates, for each of which the bounds and the optimal tradeoff are plotted. We note that for each value of λ , we obtain a corresponding minimum average service cost $c(\lambda)$. All the numerical examples that we consider in this section, along with their parameters, are given in Tables 2.2, 2.3, 2.4, and 2.5, along with references to their corresponding plots.

Ex. Identifier	λ	$c(\lambda)$	Type	Tradeoff plot
E-211	0.10	0.02	FINITE- μ CHOICE-1	2.9
E-212	0.20	0.04	FINITE- μ CHOICE-3	2.11
E-213	0.25	0.07	FINITE- μ CHOICE-2	2.13
E-214	0.70	0.50	FINITE- μ CHOICE-2	2.15

Table 2.4: Numerical examples with $c(\mu) = \mu^2$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.

Ex. Identifier	λ	$c(\lambda)$	Type	Tradeoff plot
E-221	0.10	0.00016	FINITE- μ CHOICE-1	2.10
E-222	0.20	0.00032	FINITE- μ CHOICE-3	2.12
E-223	0.25	0.00280	FINITE- μ CHOICE-2	2.14
E-224	0.70	0.20272	FINITE- μ CHOICE-2	2.16

Table 2.5: Numerical examples with $c(\mu) = \mu^5$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.

The *Optimal* set of points in each plot is obtained by considering a MDP with the single stage cost given by $q + \beta c(\mu)$, and solving for the optimal infinite horizon average cost policy. The state of the MDP corresponds to the queue length q , and the set of actions taken at each state q is the set of service rates. Here β is a positive Lagrange multiplier. The MDP is obtained by uniformization at rate $r_u = \lambda + r_{max}$. The transitions in the uniformized MDP are as follows : a) for $q = 0$, the Markov chain moves to $q = 1$ with probability $\frac{\lambda}{r_u}$, and stays in $q = 0$ with probability $1 - \frac{\lambda}{r_u}$, and b) for $q > 0$, the Markov chain moves to $q + 1$ with probability $\frac{\lambda}{r_u}$, to $q - 1$ with probability $\frac{\mu(q)}{r_u}$, and stays in q with probability $1 - \frac{\lambda + \mu(q)}{r_u}$. We note that the state space of the MDP is truncated at a maximum queue length, which is such that the optimal value does not change appreciably with further increase in this maximum queue length. The Optimal points are obtained by varying β . The x -coordinate of a point corresponding to a value of β is the difference between the average service cost, for the β -optimal policy for the MDP, and $c(\lambda)$ while the y -coordinate is the average queue length for the β -optimal policy. The *Lower bound (Analytical)* curve in each plot is: a) (2.11) for FINITE- μ CHOICE-1, b) (2.15) for FINITE- μ CHOICE-2, and c) (2.16) for FINITE- μ CHOICE-3. The *Upper bound (Analytical)* curve in each plot is obtained as follows: a) for FINITE- μ CHOICE-1, we choose the sequence q_{k_u} to be a sequence of increasing positive integers, and for each q_{k_u} obtain the bound on the average queue length from (2.20), and the average service cost from (2.19), b) for FINITE- μ CHOICE-2, we choose the sequence U_k to be a decreasing sequence, and for each U_k obtain the bound on the average queue length from (2.27) and the average service cost from (2.26), c) for FINITE- μ CHOICE-3, we choose the sequence U_k to be a decreasing sequence, and for each U_k obtain the bound on the average queue length and the average service cost from (2.30) and (2.29) respectively. We note that the bounds on the average queue length which are obtained using the Lyapunov drift method in Proposition 2.A.1 are usually weak (although they give the correct order behaviour). Therefore, for the sequence of policies which we have considered for *Upper bound (Analytical)*, we also evaluate the actual average service cost and average queue length, by obtaining the stationary probability of the queue length for the system with truncated state space. This curve is denoted as *Upper bound* in the plots. We note that discrete points are obtained while varying β , time sharing of the policies corresponding to these points leads to the continuous curves shown in the figures.

We note that in all the cases that we have considered, both the analytical upper and lower bounds




Figure 2.3: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.25$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.4: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.25$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.5: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.50$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.6: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.50$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.7: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.75$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.8: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.75$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.5, 1\}$.




Figure 2.9: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.10$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.10: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.10$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.11: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.20$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.12: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.20$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.13: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.25$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.14: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.25$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.15: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.70$ and $c(\mu) = \mu^2$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.16: $Q^*(c_c)$ as a function of V , where $c_c = c(\lambda) + V$, for $\lambda = 0.70$ and $c(\mu) = \mu^5$ for $\mu \in \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$.




Figure 2.17: The optimal tradeoff curve for the system with $\mathcal{X}_\mu = \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$, $c(\mu) = \mu^2, \forall \mu \in \mathcal{S}$, and $\lambda = 0.39, 0.40$, and 0.41 . The minimum average service cost rates are $c(0.39) = 0.154, c(0.40) = 0.160$, and $c(0.41) = 0.169$

are very loose.

The difference in the asymptotic behaviour of $Q^*(c_c)$ for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3 is illustrated by the following example. Consider the following example : we choose the set of service rates $\mathcal{X}_\mu = \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$, and $c(\mu) = \mu^2, \forall \mu \in \mathcal{S}$. In Figure 2.17 we plot the tradeoff curve, numerically obtained from a suitably truncated MDP, for $\lambda = 0.39, 0.40$ and 0.41 . The minimum average service cost rates corresponding to $\lambda = 0.39, 0.40$ and 0.41 are $0.154, 0.160$, and 0.169 . We note that the difference between the average service cost and $c(\lambda)$ increases when λ is changed from 0.39 to 0.40 and then decreases when λ is increased, since at $\lambda = 0.40$, the average queue length increases at the rate $\frac{1}{\lambda}$.

In Figures 2.18 and 2.19 we illustrate the stationary probability mass functions $\pi(q)$ for optimal policies for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3 respectively.

2.3.5 An Application

In this section we discuss the application of the above asymptotic results to the example considered in Section 2.1.3. We obtain a sequence of optimal policies for both the M/M/1-PS (I) and the M/M/1-PS (II) models which trade off the average transmitter power with the average number of flows. Both of these sequences are obtained by the numerical evaluation of the optimal policy for a suitably truncated MDP, with single stage cost $q + \beta P(m)$ and with state transitions as shown in Figure 2.1, for a sequence of $\beta > 0$. The optimal tradeoff for both M/M/1-PS (I) and (II) models are shown in Figure 2.20. We then obtain via simulation, the average power and average number of flows for the discrete time system, for the policies obtained from the M/M/1-PS (I) and M/M/1-PS (II) models, to obtain a possible (sub-optimal) tradeoff curve as shown in Figure

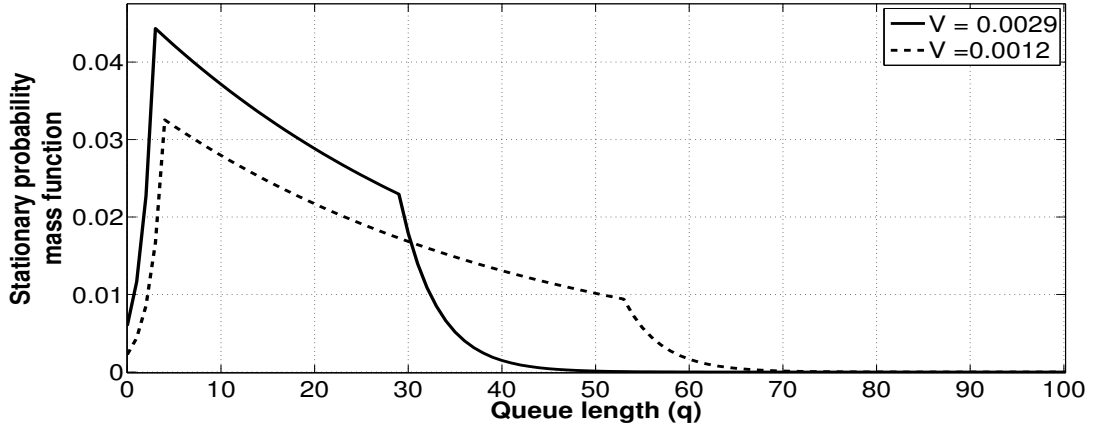


Figure 2.18: Stationary probability mass functions for optimal policies for the system with $\mathcal{X}_\mu = \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$, $c(\mu) = \mu^2, \forall \mu \in \mathcal{S}$, and $\lambda = 0.39$ (corresponding to FINITE- μ CHOICE-2). The $\Omega(\log(\frac{1}{V}))$ asymptotic lower bound arises due to the geometrically increasing and decreasing nature of $\pi(q)$ as $V \downarrow 0$.




Figure 2.19: Stationary probability mass functions for optimal policies for the system with $\mathcal{X}_\mu = \{0, 0.2, 0.4, 0.5, 0.6, 0.8, 1\}$, $c(\mu) = \mu^2, \forall \mu \in \mathcal{S}$, and $\lambda = 0.40$ (corresponding to FINITE- μ CHOICE-3). The $\Omega(\frac{1}{V})$ asymptotic lower bound arises due to the constant nature of $\pi(q)$ as $V \downarrow 0$.

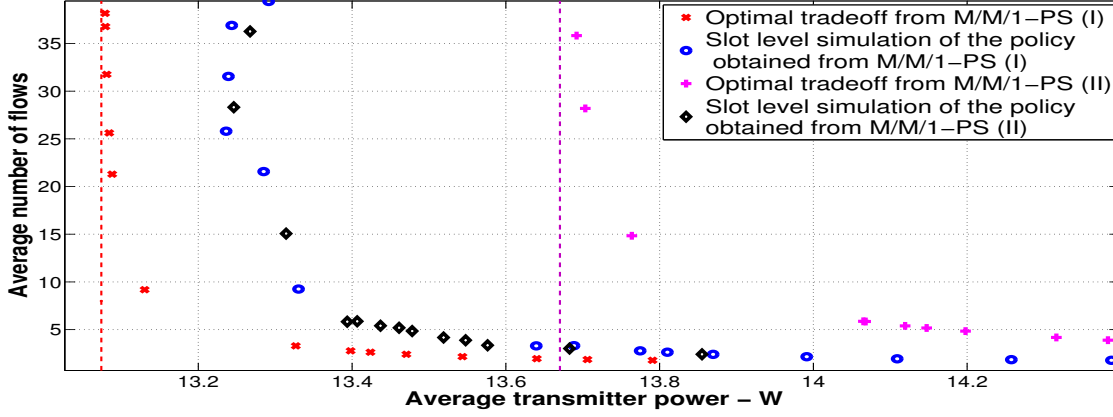


Figure 2.20: Tradeoff of average number of flows in the downlink scheduler queue with the average transmitter power for the motivating problem considered in Section 2.1.3, for a flow arrival rate 0.2 flows/sec

2.20. We note that the sequence of policies suggested by both (I) and (II) models have similar performance. For this example, the tradeoff curve obtained from the simulation of the policies suggested by the analysis, demonstrates that with a $1W$ increase in transmitter power, a 36 fold decrease in latency, i.e., from 3 mins to 5 secs, is possible. We note that $P^1(m) \leq P^2(m)$ for $m \in \{1.25, 1.50, 1.75, 2\}$, i.e., for $\mu \geq \mu_l$ given that $\lambda = 0.2$ (we recall that actual service rates are $\frac{m}{6.4}$). Then for each of the optimal policies for the M/M/1-PS (I) (or (II)) model, for which the average number of flows is large, $P^1(m(Q)) \leq P(m(Q)) \leq P^2(m(Q))$ with high probability. Thus, the tradeoff curves obtained from the M/M/1-PS (I) and (II) models are approximate upper and lower bounds to tradeoff curve for the discrete time system, especially when the average number of flows are large. Then the asymptotic analysis in the above sections, leads to asymptotic upper and lower bounds on the average transmitter power as a function of the average number of flows for the discrete time system. Let $c^1(m)$ and $c^2(m)$ be the lower convex envelopes of $P^1(m)$ and $P^2(m)$ respectively. Since $\mu_l = 1.25/6.4 < \lambda = 0.2 < \mu_u = 1.50/6.4$, by considering the sequence of optimal policies from the M/M/1-PS (II) model, the average transmitter power for the discrete time system can be bounded below by $c^1(6.4\lambda) + \Omega(e^{-\bar{q}})$ and bounded above by $c^2(6.4\lambda) + \mathcal{O}(e^{-\bar{q}})$, for large \bar{q} , where \bar{q} is the average number of flows.

2.4 Conclusions

The main purpose of this chapter is the illustration of the techniques involved in: (i) the derivation of asymptotic bounds on $Q^*(c_c)$ and $Q_M^*(c_c)$ for admissible policies and (ii) asymptotic bounds on order-optimal admissible policies in the regime \mathfrak{R} , using a simple queueing model.

For FINITE- μ CHOICE, we observe that the constraint on the average service cost leads to a restriction on the stationary probability of service rates which in turn restricts the behaviour of the

stationary probability of the queue length in the asymptotic regime \mathfrak{R} , where $V = c_c - c(\lambda) \downarrow 0$. For FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3, we note that the stationary probability of service rates less than μ_l or greater than μ_u goes to zero as $V \downarrow 0$. Since, state 0 uses a service rate of 0, the above fact implies that $\pi(0) \downarrow 0$ as $V \downarrow 0$. Hence, it is intuitive that the average queue length has to increase. For FINITE- μ CHOICE-3, as $\mu_l = \mu_u = \lambda$, as $V \downarrow 0$, the stationary probability that any service rate other than λ is used approaches zero. Then we observe that the stationary probability $\pi(q)$ becomes equal and is $\mathcal{O}(V)$ for all q which occur with high probability, as $V \downarrow 0$, since $\pi(q)\lambda = \pi(q+1)\mu(q)$, $q \geq 0$ and $\mu(q) = \lambda$ for all q occurring with high probability. The above constant nature of $\pi(q)$ leads to the $\Omega\left(\frac{1}{V}\right)$ asymptotic lower bound. For FINITE- μ CHOICE-2, as $V \downarrow 0$, service rates $\mu_l < \lambda$ and $\mu_u > \lambda$ could be used. Therefore, intuitively, one expects that the rate at which $Q^*(c_c)$ increases is less for FINITE- μ CHOICE-2 compared to FINITE- μ CHOICE-3. Furthermore, since $\pi(q)\lambda = \pi(q+1)\mu(q)$, $q \geq 0$, we observe that the stationary probability $\pi(q)$ has a geometric growth and decay with growth and decay rates at most $\frac{\lambda}{\mu_l}$ and at least $\frac{\lambda}{\mu_u}$ respectively, for the set of queue lengths occurring with high probability, as $V \downarrow 0$. The above geometric growth and decay of $\pi(q)$ leads to the $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ asymptotic lower bound in Lemma 2.3.6. The above behaviour of the stationary probability in the different cases motivates us to analyse the tradeoff problem for the discrete time queue in Chapter 4, by constructing bounds on the stationary probability of the queue length, which have the same behaviour in the respective cases.

We note that the sequence of policies which achieve the asymptotic order behaviour in Section 2.3.2 for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3 are similar to buffer partitioning policies. We note that the buffer partitions that were used for these sequences of policies scaled as $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ and $\Omega\left(\frac{1}{V}\right)$ for FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3 respectively. This scaling of the buffer partitions was suggested by the asymptotic lower bounds that were derived in Section 2.3.1, and we shall see that similar ideas can be used in the design of buffer partitioning policies for discrete time systems. Furthermore, we have also derived asymptotic bounds on any sequence of order-optimal policies in Section 2.3.3.

We note that FINITE- μ CHOICE-1, where the average queue length increases only to a finite value, even when the average service cost rate is the minimum possible $c(\lambda)$, has been hitherto unidentified in the literature. To the best of our knowledge, the asymptotic characterizations of the optimal tradeoff curve obtained in this chapter, for all the three cases, were previously not known for the state dependent M/M/1 model.

The development of the asymptotic results in this chapter, partly motivates the definition of admissible policies for the discrete time queueing model in Chapter 4. The non-idling nature of the optimal policy which has been obtained in Lemma 2.2.5, motivates us to consider whether the optimal policy for the discrete time queueing model has the same property also. In Chapter 4, we shall show that in fact it does. The non-idling property of the optimal policy and its relation with the *place-holder bit* scheduling policies [45] are discussed in more detail in Chapter 4. We have also

illustrated the utility of the simple state dependent M/M/1 queueing model in the analytical study of scheduling schemes for next generation wireless systems using the example in Section 2.1.3.

Appendix

2.A Uniformization and a bound on the average queue length

Let $r_u = r_{max} + r_{a,max}$. Consider a discrete time Markov chain $(Q_d[m])$ which is obtained by uniformization of $(Q(t))$ at rate r_u . The transition probabilities, $p_{q_1, q_2} = P(Q_d[m+1] = q_2 | Q_d[m] = q_1), m \in \mathbb{Z}_+$, of the DTMC are as follows :

$$\begin{aligned} p_{0,0} &= 1 - \frac{\lambda(0)}{r_u}, \\ p_{0,1} &= \frac{\lambda(0)}{r_u}, \\ p_{q,q+1} &= \frac{\lambda(q)}{r_u}, \forall q \geq 1, \\ p_{q,q-1} &= \frac{\mu(q)}{r_u}, \\ p_{q,q} &= 1 - \frac{\lambda(q) + \mu(q)}{r_u}. \end{aligned}$$

We note that for an admissible policy γ , the stationary distribution π is the same for both the CTMC $Q(t)$ and the DTMC $Q_d[m]$. Thus $\mathbb{E}_\pi Q_d = \bar{Q}(\gamma)$, $\mathbb{E}_\pi u(\lambda(Q_d)) = \bar{U}(\gamma)$, and $\mathbb{E}_\pi c(\mu(Q_d)) = \bar{C}(\gamma)$. The following proposition states an upper bound on $\bar{Q}(\gamma)$ for an admissible policy γ subject to an assumption about the structure of the policy.

Proposition 2.A.1. Assume that the admissible policy γ is such that there exists a q_ϵ such that $\mu(q_\epsilon) - \lambda(q_\epsilon) \geq \epsilon$, for some $\epsilon > 0$. Then

$$\bar{Q}(\gamma) \leq \frac{q_\epsilon(\epsilon + r_{a,max})}{\epsilon} + \frac{r_u}{2\epsilon} \quad (2.51)$$

Proof. Let $L(q) = q^2$. We use $L(q)$ as a Lyapunov function to derive the above upper bound. The expected Lyapunov drift $\Delta(q) =$

$$\mathbb{E}[L(Q_d[m+1]) - L(Q_d[m]) | Q_d[m] = q] \quad (2.52)$$

We have that

$$\Delta(q) = \frac{-2q}{r_u}(\mu(q) - \lambda(q)) + \frac{\mu(q) + \lambda(q)}{r_u}, \forall q.$$

Note that from the admissibility of γ , $\forall q \geq q_\epsilon$, $\mu(q) - \lambda(q) \geq \epsilon$. So for $q \geq q_\epsilon$ we have that

$$\Delta(q) \leq \frac{-2q\epsilon}{r_u} + \frac{\mu(q) + \lambda(q)}{r_u}.$$

For $q < q_\epsilon$,

$$\Delta(q) = \frac{-2q\epsilon}{r_u} + \frac{2q\epsilon}{r_u} - \frac{2q}{r_u}(\mu(q) - \lambda(q)) + \frac{\mu(q) + \lambda(q)}{r_u}. \quad (2.53)$$

For $q < q_\epsilon$, $\mu(q) - \lambda(q) < \epsilon$. Therefore

$$\begin{aligned} \Delta(q) &\leq \frac{-2q\epsilon}{r_u} + \frac{2q}{r_u}(\epsilon - \mu(q) + \lambda(q)) + \frac{\mu(q) + \lambda(q)}{r_u}, \\ &\leq \frac{-2q\epsilon}{r_u} + \frac{2q\epsilon}{r_u}(\epsilon + r_{a,max}) + 1. \end{aligned}$$

For all q , we therefore have that

$$\Delta(q) \leq \frac{-2q\epsilon}{r_u} + \frac{2q\epsilon}{r_u}(\epsilon + r_{a,max}) + 1.$$

Hence from [36, Theorem A.4.3] we have that

$$\bar{Q}(\gamma) = \mathbb{E}_\pi Q_d \leq \frac{q_\epsilon(\epsilon + r_{a,max})}{\epsilon} + \frac{r_u}{2\epsilon}.$$

□

On the tradeoff of average queue length, average service cost, and average utility for the state dependent M/M/1 queue: Part II

3.1 Introduction

We continue our analysis of the tradeoff problem for the state dependent M/M/1 model in this chapter. We note that FINITE- μ CHOICE was primarily motivated by the wireless network problem in Section 2.1.3, for which the set of possible service rates was a finite discrete set. In this chapter, we consider the INTERVAL- μ CHOICE and INTERVAL- $\lambda\mu$ CHOICE problems, which we study with the objective of understanding the tradeoff problem for the discrete time queueing model. We note that in this chapter, \mathcal{X}_μ and \mathcal{X}_λ are chosen to be finite intervals.

The method of analysis for INTERVAL- μ CHOICE and INTERVAL- $\lambda\mu$ CHOICE is similar to that in Chapter 2. We again obtain bounds on the stationary probability distribution of the queue length for admissible policies, leading to an asymptotic characterization of the average queue length as well as order-optimal admissible policies, in the asymptotic regime \mathfrak{R} . We note that the analysis for INTERVAL- μ CHOICE can be used to obtain the results for FINITE- μ CHOICE. However, in Chapter 4, we will see that some of the steps used in the analysis for INTERVAL- μ CHOICE and INTERVAL- $\lambda\mu$ CHOICE, which are different from that for FINITE- μ CHOICE, are essential in the analysis of the discrete time models. We recall that the consideration of the stationary probability distribution of the queue length in the asymptotic regime \mathfrak{R} , for admissible policies (which are monotone), as a method for understanding the asymptotic behaviour of the average queue length, underlies most of the results obtained in this thesis. We note that in this chapter, since we consider queueing models with arrival rate control as well as other forms of $c(\mu)$ (other than the piecewise

linear form in Chapter 2), we are able to obtain new insights in this direction.

The question arises as to how insights about the asymptotic behaviour of the minimum average queue length in the regime \mathfrak{R} for the discrete time model can be obtained from the state dependent M/M/1 model. It is clear that all the model features for the discrete time model, even with a single environment state, cannot be captured by the state dependent M/M/1 model. The model features for the state dependent M/M/1 model are the sets \mathcal{X}_μ and \mathcal{X}_λ , and the functions $c(\mu)$ and $u(\lambda)$, with the restriction that $\mu(q)$ and $\lambda(q)$ are deterministic functions of the queue length. For the discrete time model with a single environment state, we note that in addition to the sets of possible service batch sizes and possible admitted arrival batch sizes, and the cost and utility functions, we also have that the service batch size and the amount of arrivals admitted in a slot are randomized functions of the history of the process, as discussed in Chapter 1.

Suppose we consider only the set of stationary policies for the discrete time model, which choose the service batch size S (or the amount A of arrivals admitted) as a randomized function of the current queue length q only, say with probability distribution $P_{s|q}$ (or distributed as $P_{a|q,r}$ which is a function of the current queue length and number of actual arrivals r). Even then, we have to reduce $P_{s|q}$ (or $P_{a|q,r}$) to a real value, which can then be modelled by $\mu(q)$ (or $\lambda(q)$). In the following, we take $\mathbb{E}_{P_{s|q}} S$ as the quantity which represents $P_{s|q}$. Thus $\mu(q)$ is assumed to correspond to $\mathbb{E}_{P_{s|q}} S$ and therefore \mathcal{X}_μ is the set of all values that $\mathbb{E}_{P_{s|q}} S$ can take. A similar assumption is made for $\lambda(q)$.

We note that for the discrete time model, the service cost incurred in a slot is a random variable, since the batch size S itself is random. At a queue length q , the expected service cost is $\mathbb{E}_{P_{s|q}} c(S)$. Since we have already chosen $\mu(q)$ to correspond to $\mathbb{E}_{P_{s|q}} S$, a possibility is to choose $c(\mu)$ to correspond to $c(\mathbb{E}_{P_{s|q}} S)$. A similar assumption is made for $u(\lambda)$. In retrospect, this turns out to be a good choice of model features for the state dependent M/M/1 model (e.g., we obtain the asymptotic Berry Gallager lower bound for the M/M/1 model which, with the above choice of model features, corresponds to the discrete time model in [7] with admissible policies). We can also then surmise that one of the reasons for the asymptotic behaviour of the stationary queue length for admissible policies in the regime \mathfrak{R} for the discrete time model are the behaviours of asymptotic probability distributions of $\mathbb{E}_{P_{s|q}} S(Q)$ and $\mathbb{E}_{P_{a|q,r}} A(Q, R)$, in the regime \mathfrak{R} , since for the M/M/1 model the behaviours of these quantities are significant.

3.1.1 Overview

We present the analysis of INTERVAL- μ CHOICE in Section 3.2. The motivating discrete time problem for INTERVAL- μ CHOICE is discussed in the same section. We recall that in defining INTERVAL- μ CHOICE we considered admissible policies for which $\lambda(q)$ is a constant λ for all $q \in \mathbb{Z}_+$. We obtain asymptotic lower and upper bounds on the tradeoff problem in Sections 3.2.1

and 3.2.2 respectively. Asymptotic bounds on order-optimal policies for INTERVAL- μ CHOICE, are presented in Section 3.2.3. We consider the counterpart INTERVAL- λ CHOICE of INTERVAL- μ CHOICE, for which $\mu(q)$ is constant over all $q \in \mathbb{Z}_+$ and $\lambda(q)$ is the control variable, in Section 3.2.4. The analysis of INTERVAL- $\lambda\mu$ CHOICE is presented in Section 3.3, along with the motivating discrete time problem. We present the main conclusions in Section 3.4, where we discuss the main ideas obtained in the asymptotic analysis of INTERVAL- μ CHOICE, INTERVAL- λ CHOICE, and INTERVAL- $\lambda\mu$ CHOICE, and how these ideas can be applied to the discrete time model in Chapters 4, 5, and 6.

3.2 Analysis of INTERVAL- μ CHOICE

We recall that for INTERVAL- μ CHOICE we restrict to admissible policies γ such that $\lambda(q) = \lambda$ and $\mu(q) \in [0, r_{max}]$, $\forall q \in \mathbb{Z}_+$. The tradeoff problem for INTERVAL- μ CHOICE is

$$\begin{aligned} & \text{minimize}_{\gamma \in \Gamma_a} && \overline{Q}(\gamma) \\ & \text{such that} && \overline{C}(\gamma) \leq c_c, \end{aligned}$$

whose optimal value is denoted as $Q^*(c_c)$. We also note that a tradeoff problem can be defined where we minimize over the set of policies, which includes finite mixtures of pure policies in Γ_a . The optimal value of this problem is denoted as $Q_M^*(c_c)$. Then, as for the case of FINITE- μ CHOICE (e.g. as in Corollary 2.3.5), asymptotic lower bounds for $Q^*(c_c)$ can be used to obtain asymptotic lower bounds for $Q_M^*(c_c)$. In fact, the asymptotic lower bounds in the regime \mathfrak{R} are the same for $Q^*(c_c)$ and $Q_M^*(c_c)$. Furthermore, the asymptotic upper bounds that we derive for $Q^*(c_c)$ are by definition upper bounds for $Q_M^*(c_c)$. Hence, in the following, we present the results for $Q^*(c_c)$ only.

The study of INTERVAL- μ CHOICE is classified into:

INTERVAL- μ CHOICE-1 : $c(\mu)$ is strictly convex for $\mu \in [0, r_{max}]$, and

INTERVAL- μ CHOICE-2 : $c(\mu)$ is piecewise linear. That is, (a) there exists a minimal partition of $[0, r_{max}]$ into intervals $\{[a_i, b_i], i \in \{1, \dots, P\}\}$ with $a_1 = 0$, $b_P = r_{max}$, and $b_i = a_{i+1}$, and (b) there are linear functions f_i such that $\forall \mu \in [a_i, b_i], f_i(\mu) = c(\mu)$.

Remark 3.2.1. We first discuss the motivation for INTERVAL- μ CHOICE-1. INTERVAL- μ CHOICE-1 corresponds to the tradeoff problem for the following discrete time queueing model. Work arrives in a batch, of random size, in every slot, into an infinite buffer queue. The state of the queue is the amount of unfinished work. We note that for the discrete time model, the amount of unfinished work or the queue state evolves on the set of non-negative real numbers. This is approximated by an integer-valued queue evolution process in INTERVAL- μ CHOICE-1. The amount of work completed in each slot, or the service batch size, can be chosen as a function (possibly randomized), of the

current backlog of unfinished work. The choice of the amount of work completed in each slot as a function of the current backlog for the discrete time queue, is modelled by the control of the service rate, $\mu(q)$, in INTERVAL- μ CHOICE-1. Hence, the service batch size and the average service batch size (or rate) also takes values in the set of non-negative real numbers. As discussed in Section 3.1 we then assume that $\mu(q)$ takes values in an interval, which for technical reasons is assumed to be finite. We assume that there is no admission control in the discrete time model, therefore we assume that the arrival rate is λ for every q for INTERVAL- μ CHOICE-1. For the discrete time queue, we assume that there is a service cost associated with the amount of work done in each slot. This is modelled by the service cost rate function $c(\cdot)$ in INTERVAL- μ CHOICE-1. We note that since the amount of work done in a slot can be any real value, the service cost for the discrete time model could be a strictly convex function defined on an interval, which provides the motivation for assuming $c(\mu)$ to be strictly convex for INTERVAL- μ CHOICE-1. We note that this discrete time model is similar to the model considered by Berry and Gallager [7], but with a single fade state.

The motivating discrete time queueing model for INTERVAL- μ CHOICE-2 is very similar to the discrete time model discussed above, except that the queue evolution is assumed to be on integers. But for stationary randomized policies, the average service batch size could still take any real value in a finite interval, and in light of the discussion in Section 3.1, $\mu(q)$ is again assumed to take any value from a finite interval. In Chapter 4, we shall see that then $c(\mathbb{E}_{P_{s|q}} S)$ is piecewise linear, which is the motivation for the piecewise linear assumption on $c(\mu)$ for INTERVAL- μ CHOICE-2.

For any admissible policy γ , from Jensen's inequality, we have that $\overline{C}(\gamma) \geq c(\lambda)$. We study INTERVAL- μ CHOICE-1 and INTERVAL- μ CHOICE-2 in the asymptotic regime \mathfrak{R} where the service cost constraint c_c approaches $c(\lambda)$, since it can be shown that $c(\lambda) = \inf_{\gamma \in \Gamma_a} \overline{C}(\gamma)$.

Similar to FINITE- μ CHOICE, since the asymptotic behaviour of $Q^*(c_c)$ for INTERVAL- μ CHOICE-2 depends on the behaviour of $c(\mu)$ in a neighbourhood of $\mu = \lambda$, we consider the following cases for INTERVAL- μ CHOICE-2:

INTERVAL- μ CHOICE-2-1 : $\lambda \in (0, b_1 = b_\lambda)$,

INTERVAL- μ CHOICE-2-2 : $\lambda \in (a_i = a_\lambda, b_i = b_\lambda)$ for some $i \in \{2, \dots, P\}$, and,

INTERVAL- μ CHOICE-2-3 : $\lambda = a_i = a_\lambda$ for some $i \in \{2, \dots, P\}$.

The motivation for classifying INTERVAL- μ CHOICE-2 into the three cases is the same as that for FINITE- μ CHOICE. The different cases are illustrated in Figure 3.1. We note that for INTERVAL- μ CHOICE-1, the function $c(\mu)$ is strictly convex for every $\mu \in [0, r_{max}]$ and therefore it has no subcases. We note that the analysis of INTERVAL- μ CHOICE-2 is similar to that of FINITE- μ CHOICE except that now the set of service rates is not a given finite set. We now present the asymptotic lower bounds for $Q^*(c_c)$ in the regime \mathfrak{R} for the above cases.

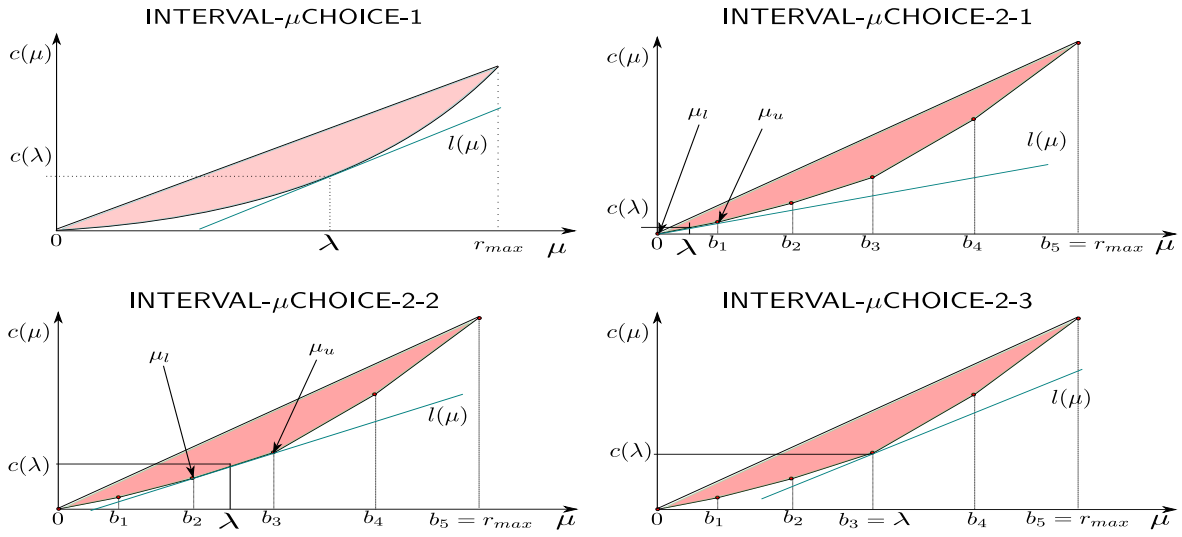


Figure 3.1: Illustration of the relationship between λ , μ_l , and μ_u along with the minimum average cost $c(\lambda)$ and the line $l(\mu)$ for the four cases of the INTERVAL- μ CHOICE problem.

3.2.1 Asymptotic lower bounds

Similar to the definition used for FINITE- μ CHOICE, line $l(\mu)$ in the case of INTERVAL- μ CHOICE is defined as: (a) The tangent to the $c(\mu)$ curve at λ for INTERVAL- μ CHOICE-1, (b) the line passing through $(0, 0)$ and $(b_1, c(b_1))$ for INTERVAL- μ CHOICE-2-1, (c) the line passing through $(a_\lambda, c(a_\lambda))$ and $(b_\lambda, c(b_\lambda))$ for INTERVAL- μ CHOICE-2-2, and (d) any line that passes through $(\lambda, c(\lambda))$ with a slope m , such that $\frac{dc(\mu)^-}{d\mu}|_{\mu=\lambda} < m < \frac{dc(\mu)^+}{d\mu}|_{\mu=\lambda}$ (the left and right derivatives respectively) for INTERVAL- μ CHOICE-2-3.

We note that like in the case of FINITE- μ CHOICE, here we find an asymptotic lower bound on $\bar{Q}(\gamma)$ by (a) obtaining an upper bound on the stationary probability for a certain set of service rates in terms of $\bar{C}(\gamma)$ and $c(\lambda)$, (b) relating the stationary probability of this set of service rates to the stationary probability $\pi(q)$, of a set of queue lengths, and (c) obtaining a lower bound on $\bar{Q}(\gamma)$ in terms of $\pi(q)$.

We first consider INTERVAL- μ CHOICE-1, for which $c(\mu)$ is a strictly convex function of $\mu \in [0, r_{max}]$. We make the following assumption regarding $c(\mu)$ at $\mu = \lambda$.

C2 : For INTERVAL- μ CHOICE-1, the second derivative of $c(\mu)$ is non-zero at $\mu = \lambda$.

The above assumption has been used in [7]. We note that since $c(\mu)$ is strictly convex, the second derivative of $c(\mu)$ is non-zero for all $\mu \in [0, r_{max}]$ except for μ in a countable set. We note that even though $\mu(q) \in [0, r_{max}]$, the set $\{\mu(q), q \in \mathbb{Z}_+\}$ is only countable. Let $\{\mu_0 = 0, \dots, \mu_k, \dots\}$ denote the set of service rates that is used by a policy γ .

Lemma 3.2.2. For INTERVAL- μ CHOICE-1, for any sequence of non-idling admissible policies γ_k such that $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$ we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$.

Proof. Consider a particular policy γ in the sequence γ_k with $V_k = V$. Let $\mu^* \triangleq \lambda - \epsilon_V$, where $\epsilon_V > 0$ is a function of V . The functional form of ϵ_V will be chosen later. We have that

$$\begin{aligned} V &= \sum_{k=0}^{\infty} (c(\mu_k) - l(\mu_k)) \pi_{\mu}(k) = \sum_{q=0}^{\infty} (c(\mu(q)) - l(\mu(q))) \pi(q), \\ &= \sum_{q=0}^{\infty} \left(c(\lambda) + \frac{dc(\mu)}{d\mu} \Big|_{\mu=\lambda} (\mu(q) - \lambda) - l(\mu(q)) + G(\mu(q) - \lambda) \right) \pi(q), \end{aligned}$$

where $G(x)$ is a strictly convex function in x as in [7, Proposition 4.2]. We note that $c(\lambda) + \frac{dc(\mu)}{d\mu} \Big|_{\mu=\lambda} (\mu(q) - \lambda) = l(\mu(q))$. Thus we have that $V = \sum_{q=0}^{\infty} G(\mu(q) - \lambda) \pi(q)$. As $G(x)$ is strictly convex in x and $\mu(q) - \lambda$ is bounded, we have from Proposition 3.A.1 that $G(\mu(q) - \lambda) \geq a_1(\mu(q) - \lambda)^2$ for some constant $a_1 > 0$. Thus we have that

$$V \geq a_1 \sum_{q=0}^{\infty} (\mu(q) - \lambda)^2 \pi(q). \quad (3.1)$$

Define $q_{\mu^*} = \inf \{q : \mu(q) \geq \mu^*\}$. From the non-decreasing property of $\mu(q)$ for γ , we have that $Pr \{\mu(Q) < \mu^*\} = Pr \{Q < q_{\mu^*}\}$. Then,

$$Pr \{\mu(Q) < \mu^*\} = \sum_{q=0}^{q_{\mu^*}-1} \pi(q) \leq \frac{V}{a_1 \epsilon_V^2}, \quad (3.2)$$

where we have used the upper bound (3.1). We choose ϵ_V as $a_2 \sqrt{V}$, so that $\alpha \triangleq \frac{2V}{a_1 \epsilon_V^2} = \frac{2}{a_1 a_2^2}$. We choose a_2 such that $\alpha < 1$. In fact, we note that α can be made arbitrarily close to zero by the choice of a_2 . Therefore, $Pr \{\mu(Q) \geq \mu^*\} \geq 1 - \frac{\alpha}{2}$, which can be made arbitrarily close to one.

We note that for any $q < q_{\mu^*}$, $\pi(q) \leq \frac{\alpha}{2}$. Therefore $\pi(q_{\mu^*}) \leq \pi(q_{\mu^*} - 1) \frac{\lambda}{\mu^*} \leq \frac{\lambda \alpha}{2 \mu^*}$. In order to obtain a lower bound on $\bar{Q}(\gamma)$, we intend to find the largest \bar{q} such that $Pr \{Q \leq \bar{q}\} \leq \frac{1}{2}$. But we note that $Pr \{Q < q_{\mu^*}\} \leq \frac{V}{a_1 \epsilon_V^2} = \frac{\alpha}{2}$. Therefore the largest \bar{q} satisfies

$$\sum_{q=0}^{q_{\mu^*}-1} \pi(q) + \sum_{q=q_{\mu^*}}^{\bar{q}} \pi(q) \leq \frac{1}{2}$$

If \bar{q}_1 satisfies

$$\sum_{q=q_{\mu^*}}^{\bar{q}_1} \pi(q) \leq \frac{1}{2} - \frac{\alpha}{2},$$

then $\bar{q}_1 \leq \bar{q}$, for α sufficiently small. As $Q(t)$ is a birth-death process, we have that $\pi(q)\lambda = \pi(q+1)\mu(q+1)$. Furthermore, if $q \geq q_{\mu^*}$ we have that $\pi(q-1)\lambda \geq \pi(q)\mu^*$. By induction, we obtain that for $q \in \{q_{\mu^*}, \dots\}$

$$\pi(q) \leq \pi(q_{\mu^*}) \left(\frac{\lambda}{\mu^*} \right)^{q-q_{\mu^*}} \leq \pi(q_{\mu^*}-1) \left(\frac{\lambda}{\mu^*} \right)^{q-q_{\mu^*}+1}, \quad (3.3)$$

$$\text{and for any } q' \geq q_{\mu^*}, \sum_{q=q_{\mu^*}}^{q'} \pi(q) \leq \pi(q_{\mu^*}-1) \sum_{m=1}^{q'-q_{\mu^*}+1} \left(\frac{\lambda}{\mu^*} \right)^m. \quad (3.4)$$

Using the above upper bound on $\sum_{q=q_{\mu^*}}^{q'} \pi(q)$, we obtain a lower bound \bar{q}_2 to \bar{q}_1 . If \bar{q}_2 is the largest integer such that

$$\pi(q_{\mu^*}-1) \sum_{m=1}^{\bar{q}_2-q_{\mu^*}+1} \left(\frac{\lambda}{\mu^*} \right)^m \leq \frac{1}{2} - \frac{\alpha}{2}, \quad (3.5)$$

then $\sum_{q=0}^{\bar{q}_2} \pi(q) \leq \frac{1}{2}$ and $\bar{q}_1 \geq \bar{q}_2$.

Now we obtain an upper bound on $\pi(q_{\mu^*}-1)$, which is tighter than the upper bound $\frac{\alpha}{2}$ derived before. From (3.1) we have that

$$\begin{aligned} \frac{V}{a_1} &\geq \sum_{q=0}^{\infty} (\mu(q) - \lambda)^2 \pi(q) \geq \sum_{q < q_{\mu^*}} (\mu(q) - \lambda)^2 \pi(q) \\ &= \sum_{q < q_{\mu^*}} (\mu(q) - \lambda)^2 \pi(q) + 0 \sum_{q \geq q_{\mu^*}} \pi(q), \\ &\geq \left(\sum_{q < q_{\mu^*}} (\mu(q) - \lambda) \pi(q) \right)^2 \quad (\text{using Jensen's inequality as in [7]}). \end{aligned} \quad (3.6)$$

But, as $\pi(q)\mu(q) = \pi(q-1)\lambda$, we obtain that

$$\begin{aligned} \sum_{q < q_{\mu^*}} (\mu(q) - \lambda) \pi(q) &= -\lambda \pi(0) + \sum_{1 \leq q \leq q_{\mu^*}-1} (\lambda \pi(q-1) - \lambda \pi(q)) = -\lambda \pi(q_{\mu^*}-1), \\ \text{or } \frac{V}{a_1} &\geq \lambda^2 \pi(q_{\mu^*}-1)^2. \end{aligned} \quad (3.7)$$

Now we find a lower bound \bar{q}_3 on \bar{q}_2 by using the above upper bound on $\pi(q_{\mu^*}-1)$ in (3.5). Let \bar{q}_3 be the largest integer such that

$$\frac{1}{\lambda} \sqrt{\frac{V}{a_1}} \sum_{m=1}^{\bar{q}_3-q_{\mu^*}+1} \left(\frac{\lambda}{\mu^*} \right)^m \leq \frac{1}{2} - \frac{\alpha}{2}.$$

Then $\bar{q}_3 \leq \bar{q}_2$. We have that \bar{q}_3 satisfies

$$\begin{aligned} \frac{1}{\lambda} \sqrt{\frac{V}{a_1}} \sum_{m=1}^{\bar{q}_3 - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*} \right)^m &\leq \frac{1}{2} - \frac{\alpha}{2}, \\ \frac{\left(\frac{\lambda}{\mu^*} \right)^{\bar{q}_3 - q_{\mu^*} + 1} - 1}{\lambda - \mu^*} &\leq \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right], \\ \left(\frac{\lambda}{\mu^*} \right)^{\bar{q}_3 - q_{\mu^*} + 1} &\leq 1 + (\lambda - \mu^*) \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right] \\ \bar{q}_3 - q_{\mu^*} + 1 &\leq \log_{\frac{\lambda}{\mu^*}} \left[1 + (\lambda - \mu^*) \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right] \right]. \end{aligned}$$

Since $q_{\mu^*} > 0$, we note that \bar{q}_3 is at least

$$\left\lceil \log_{\frac{\lambda}{\mu^*}} \left[1 + (\lambda - \mu^*) \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right] \right] - 1 \right\rceil.$$

Therefore,

$$\begin{aligned} \bar{q}_3 &\geq \log_{\frac{1}{1 - \frac{\epsilon_V}{\lambda}}} \left[1 + \epsilon_V \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right] \right] - 2, \\ &= \frac{\log \left[1 + \epsilon_V \sqrt{\frac{a_1}{V}} \left[\frac{1 - \alpha}{2} \right] \right]}{-\log \left(1 - \frac{\epsilon_V}{\lambda} \right)} - 2. \end{aligned}$$

Since $\epsilon_V = a_2 \sqrt{V}$, we have

$$\bar{q}_3 \geq \frac{\log \left[1 + a_2 \sqrt{a_1} \left[\frac{1 - \alpha}{2} \right] \right]}{-\log \left(1 - \frac{a_2 \sqrt{V}}{\lambda} \right)} - 2.$$

Since $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{\bar{q}_2}{2} \geq \frac{\bar{q}_3}{2}$ we have that

$$\bar{Q}(\gamma) \geq \frac{1}{2} \left[\frac{\log \left[1 + a_2 \sqrt{a_1} \left[\frac{1 - \alpha}{2} \right] \right]}{-\log \left(1 - \frac{a_2 \sqrt{V}}{\lambda} \right)} - 2 \right].$$

As $V \downarrow 0$, we note that $\log \left(1 - \frac{a_2 \sqrt{V}}{\lambda} \right) = \Theta \left(\sqrt{V} \right)$. Hence, for the sequence γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega \left(\frac{1}{\sqrt{V_k}} \right)$. \square

Remark 3.2.3. We note that as $V \downarrow 0$, $\pi(0) \rightarrow 0$. We also note that as $V \downarrow 0$, there exists a set of queue lengths, \mathcal{Q}_h , occurring with high probability $(1 - V)$, such that $\mu(q) \rightarrow \lambda, \forall q \in \mathcal{Q}_h$. If $\pi(0) \rightarrow 0$, then $|\mathcal{Q}_h| \rightarrow \infty$. Furthermore, for each $q \in \mathcal{Q}_h$, $\pi(q) = \mathcal{O}(\sqrt{V})$. We also note that the stationary probability for each $q \in \mathcal{Q}_h$ become equal as $V \downarrow 0$. Then the average queue length is $\Omega \left(\frac{1}{\sqrt{V}} \right)$.

Remark 3.2.4. We note that the lower bounding technique in [7] can be used to obtain the $\Omega \left(\frac{1}{\sqrt{V}} \right)$

lower bound by considering a uniformized version of $Q(t)$. We outline this method in Appendix 3.B. Using the stationary probability of the queue length has its advantages, since it gives us additional insights into the form of the optimal policy.

Lemma 3.2.5. For INTERVAL- μ CHOICE-2-1, for any sequence of non-idling admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$ we have that $\bar{Q}(\gamma_k) = \frac{\lambda}{b_\lambda - \lambda} - \mathcal{O}(V^{1-\delta} \log(\frac{1}{V}))$, for $0 < \delta < 1$.

We note that INTERVAL- μ CHOICE-2-1 is very similar to FINITE- μ CHOICE-1 for which we recall that the asymptotic order was $\mathcal{O}(V \log(\frac{1}{V}))$. However, for INTERVAL- μ CHOICE-2-1, we are only able to show that the order is $\mathcal{O}(V^{1-\delta} \log(\frac{1}{V}))$, where δ can be made arbitrarily close to zero.

Proof. We note that in this case there exists a policy γ , for which $\mu(q) = b_1 = b_\lambda, \forall q > 0$, with $\bar{C}(\gamma) = c(\lambda)$ and $\bar{Q}(\gamma) = \frac{\lambda}{b_\lambda - \lambda}$. For $V = 0$, we note that the above policy is optimal. The solution to the tradeoff problem is similar to FINITE- μ CHOICE-1 in which the average queue length increases but only to a finite limit as $V_k \downarrow 0$. Consider a particular policy γ in the sequence γ_k with $V_k = V$. We have that

$$V = \sum_{k=0}^{\infty} (c(\mu_k) - l(\mu_k)) \pi_\mu(k) = \sum_{\mu_k > b_\lambda} (c(\mu_k) - l(\mu_k)) \pi_\mu(k).$$

Let $\mu^* = b_\lambda + \epsilon_V$, where ϵ_V is a function of V to be chosen later. Then we have that

$$\begin{aligned} V &\geq \sum_{\mu_k \geq \mu^*} (c(\mu_k) - l(\mu_k)) \pi_\mu(k) \\ &\geq m_a \sum_{\mu_k \geq \mu^*} (\mu_k - b_\lambda) \pi_\mu(k) \\ &\geq m_a \epsilon_V \sum_{\mu_k \geq \mu^*} \pi_\mu(k) \\ \frac{V}{m_a \epsilon_V} &\geq \sum_{\mu_k \geq \mu^*} \pi_\mu(k), \end{aligned} \tag{3.8}$$

where m_a is the tangent of the angle made by the line passing through $(b_\lambda, c(b_\lambda))$ and $(b_2, c(b_2))$ with the line $l(\mu)$. We proceed as in the proof of the asymptotic lower bound for problem FINITE- μ CHOICE-1. But unlike in FINITE- μ CHOICE-1, we note here that any service rate arbitrarily close to b_λ might be used by a policy γ . Intuitively, since $\sum_{\mu_k \geq \mu^*} \pi_\mu(k)$ should approach 0 as $V \downarrow 0$, we require that the choice of ϵ_V should be such that $\frac{V}{\epsilon_V} \downarrow 0$ as $V \downarrow 0$.

Let $q_{\mu^*} = \inf \{q : \mu(q) \geq \mu^*\}$. For $q < q_{\mu^*}$, $\mu(q) < \mu^*$ and therefore $\pi(q)\lambda < \pi(q+1)\mu^*$. Hence, by induction we obtain that

$$\pi(q_{\mu^*} - m) < \pi(q_{\mu^*}) \left(\frac{\mu^*}{\lambda} \right)^m, \text{ for } m \in \{1, \dots, q_{\mu^*}\}. \tag{3.9}$$

From (3.8), we have

$$\sum_{q < q_{\mu^*}} \pi(q) = 1 - \sum_{q \geq q_{\mu^*}} \pi(q) \geq 1 - \frac{V}{m_a \epsilon_V}.$$

Now from (3.9) we have

$$\begin{aligned} \sum_{q < q_{\mu^*}} \pi(q) &\leq \sum_{q=0}^{q_{\mu^*}-1} \pi(q_{\mu^*}) \left(\frac{\mu^*}{\lambda}\right)^{q_{\mu^*}-q}, \text{ hence we have that,} \\ 1 - \frac{V}{m_a \epsilon_V} &\leq \pi(q_{\mu^*}) \sum_{m=1}^{q_{\mu^*}} \left(\frac{\mu^*}{\lambda}\right)^m, \\ \frac{1}{\pi(q_{\mu^*})} \left(1 - \frac{V}{m_a \epsilon_V}\right) &\leq \sum_{m=1}^{q_{\mu^*}} \left(\frac{\mu^*}{\lambda}\right)^m = \frac{\mu^*}{\mu^* - \lambda} \left[\left(\frac{\mu^*}{\lambda}\right)^{q_{\mu^*}} - 1 \right]. \end{aligned}$$

We note that as $\pi(q_{\mu^*}) \leq \sum_{q \geq q_{\mu^*}} \pi(q) \leq \frac{V}{m_a \epsilon_V}$, we have that $\frac{1}{\pi(q_{\mu^*})} \geq \frac{m_a \epsilon_V}{V}$ and therefore

$$\begin{aligned} \frac{m_a \epsilon_V}{V} \left(1 - \frac{V}{m_a \epsilon_V}\right) &\leq \frac{\mu^*}{\mu^* - \lambda} \left[\left(\frac{\mu^*}{\lambda}\right)^{q_{\mu^*}} - 1 \right], \\ \log_{\frac{\mu^*}{\lambda}} \left[\frac{\mu^* - \lambda}{\mu^*} \frac{m_a \epsilon_V}{V} \left(1 - \frac{V}{m_a \epsilon_V}\right) + 1 \right] &\leq q_{\mu^*}. \end{aligned}$$

By definition, for every $q < q_{\mu^*}$, $\mu(q) < \mu^*$, and for every $q \geq q_{\mu^*}$, $\mu(q) \leq r_{max}$. Let us define

$$q_{\mu^*,l} = \left\lceil \log_{\frac{\mu^*}{\lambda}} \left[\frac{\mu^* - \lambda}{\mu^*} \frac{m_a \epsilon_V}{V} \left(1 - \frac{V}{m_a \epsilon_V}\right) + 1 \right] \right\rceil,$$

which is the smallest possible value for q_{μ^*} for any policy γ . Consider another policy γ' defined as follows :

$$\begin{aligned} \mu(0) &= 0, \\ \mu(q) &= \mu^*, \text{ for } 1 \leq q \leq q_{\mu^*,l}, \\ \mu(q) &= r_{max}, \text{ for } q > q_{\mu^*,l}. \end{aligned}$$

Then $\bar{Q}(\gamma') \leq \bar{Q}(\gamma)$. We now obtain a lower bound on $\bar{Q}(\gamma')$ as in FINITE- μ CHOICE-1.

Recall that for γ' we have (using the sequence of steps leading to (2.10))

$$\begin{aligned} \bar{Q}(\gamma') &\geq (1-a) \left[\frac{a}{(1-a)^2} \{1 - (1-a)(q_{\mu^*,l} + 1)a^{q_{\mu^*,l}} - a.a^{q_{\mu^*,l}}\} + a^{q_{\mu^*,l}} \left\{ q_{\mu^*,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right\} \right], \\ &= \frac{a}{1-a} + (1-a) \left[a^{q_{\mu^*,l}} \left\{ q_{\mu^*,l} \frac{b}{1-b} + \frac{b}{(1-b)^2} \right\} - \frac{a}{(1-a)^2} \{(1-a)(q_{\mu^*,l} + 1)a^{q_{\mu^*,l}} + a.a^{q_{\mu^*,l}}\} \right], \\ &\geq \frac{a}{1-a} - \frac{a}{1-a} a^{q_{\mu^*,l}} [1 + (1-a)q_{\mu^*,l}], \end{aligned}$$

where $a = \frac{\lambda}{\mu^*}$ and $b = \frac{\lambda}{r_{max}}$. We note that for $V \downarrow 0$, the term $\frac{a}{1-a} = \frac{\lambda}{b_\lambda - \lambda} \left(1 - \frac{\epsilon_V}{b_\lambda - \lambda} + o(\epsilon_V)\right)$. If $V \downarrow 0$, since we require that $\frac{V}{\epsilon_V} \downarrow 0$, $q_{\mu^*,l} \uparrow \infty$ and therefore the second term in the lower bound for $\bar{Q}(\gamma')$ is $aq_{\mu^*,l}a^{q_{\mu^*,l}}$. We note that at $V = 0$, since we require that the lower bound is tight, we only consider ϵ_V such that $\epsilon_V \downarrow 0$ as $V \downarrow 0$. Then it can be shown that $\bar{Q}(\gamma) \geq \bar{Q}(\gamma') = \frac{\lambda}{b_\lambda - \lambda} - \mathcal{O}\left(\epsilon_V + \frac{V}{\epsilon_V} \log\left(\frac{\epsilon_V}{V}\right)\right)$, for any sequence $\epsilon_V \downarrow 0$ and $\frac{V}{\epsilon_V} \downarrow 0$ as $V \downarrow 0$. By choosing $\epsilon_V = V^{1-\delta}$, where $0 < \delta < 1$, we obtain that

$$\bar{Q}(\gamma) = \frac{\lambda}{b_\lambda - \lambda} - \mathcal{O}\left(V^{1-\delta} \log\left(\frac{1}{V}\right)\right).$$

For the sequence γ_k , we therefore obtain that $\bar{Q}(\gamma_k) = \frac{\lambda}{b_\lambda - \lambda} - \mathcal{O}\left(V^{1-\delta} \log\left(\frac{1}{V}\right)\right)$. \square

Lemma 3.2.6. For INTERVAL- μ CHOICE-2-2, for any sequence of non-idling admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

Proof. Consider a particular policy γ in the sequence γ_k with $V_k = V$. Let $\mu^* \triangleq a_\lambda - \epsilon_V$. We have that

$$\begin{aligned} V &= \sum_{\mu_k < a_\lambda} (c(\mu_k) - l(\mu_k))\pi_\mu(k) + \sum_{\mu_k > b_\lambda} (c(\mu_k) - l(\mu_k))\pi_\mu(k), \\ &\geq \sum_{\mu_k < a_\lambda - \epsilon_V} (c(\mu_k) - l(\mu_k))\pi_\mu(k), \\ &= m_a \sum_{\mu_k < a_\lambda - \epsilon_V} (a_\lambda - \mu_k)\pi_\mu(k), \\ &\geq m_a \epsilon_V Pr\{\mu(Q) < \mu^*\}, \end{aligned}$$

where m_a is the tangent of the angle made by the line passing through $(a_i = a_\lambda, c(a_i))$ and $(a_{i-1}, c(a_{i-1}))$ with the line $l(\mu)$. Define $q_{\mu^*} = \inf\{q : \mu(q) \geq \mu^*\}$. As γ is admissible, we have that $Pr\{\mu(Q) < \mu^*\} = Pr\{Q < q_{\mu^*}\}$. Hence we have

$$\begin{aligned} Pr\{Q < q_{\mu^*}\} &\leq \frac{V}{m_a \epsilon_V}, \\ \text{and } \pi(q_{\mu^*} - 1) &\leq \frac{V}{m_a \epsilon_V}. \end{aligned}$$

We now choose $\epsilon_V = \epsilon$, a positive constant. To find a lower bound on $\bar{Q}(\gamma)$, in the following, we intend to find the largest \bar{q} such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$. But we note that $Pr\{Q < q_{\mu^*}\} \leq \frac{V}{m_a \epsilon}$ and for any $q < q_{\mu^*}$, $\pi(q) \leq \frac{V}{m_a \epsilon}$. Therefore, $\pi(q_{\mu^*}) \leq \pi(q_{\mu^*} - 1) \frac{\lambda}{\mu^*} \leq \frac{\lambda V}{m_a \epsilon \mu^*}$. Let \bar{q}_1 be the largest integer such that

$$\sum_{q=q_{\mu^*}}^{\bar{q}} \pi(q) \leq \frac{1}{2} - \frac{V}{m_a \epsilon},$$

then $\bar{q}_1 \leq \bar{q}$. Proceeding as for problem INTERVAL- μ CHOICE-1, we obtain a lower bound \bar{q}_2 on \bar{q}_1 by using an upper bound for $\pi(q)$. We note that if $q \geq q_{\mu^*}$ we have that $\pi(q-1)\lambda \geq \pi(q)\mu^*$. By induction, we obtain that for $q \in \{q_{\mu^*}, \dots\}$

$$\pi(q) \leq \pi(q_{\mu^*} - 1) \left(\frac{\lambda}{\mu^*} \right)^{q - q_{\mu^*} + 1},$$

and for any $q' \geq q_{\mu^*}$,
$$\sum_{q=q_{\mu^*}}^{q'} \pi(q) \leq \pi(q_{\mu^*} - 1) \sum_{m=1}^{q' - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*} \right)^m.$$

Using the above upper bound on $\sum_{q=q_{\mu^*}}^{q'} \pi(q)$, and $\pi(q_{\mu^*} - 1) \leq \frac{V}{m_a \epsilon}$, we obtain the following lower bound \bar{q}_2 to \bar{q}_1 .

If \bar{q}_2 is the largest integer such that

$$\frac{V}{m_a \epsilon} \sum_{m=1}^{\bar{q}_2 - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*} \right)^m \leq \frac{1}{2} - \frac{V}{m_a \epsilon},$$

then $\sum_{q=0}^{\bar{q}_2} \pi(q) \leq \frac{1}{2}$ and $\bar{q}_2 \leq \bar{q}_1$. Hence, \bar{q}_2 is the largest integer such that

$$\left(\frac{\lambda}{a_\lambda - \epsilon} \right)^{\bar{q}_2 - q_{\mu^*} + 1} \leq 1 + \left(\frac{\lambda - a_\lambda + \epsilon}{\lambda} \right) \frac{\epsilon m_a}{V} \left(\frac{1}{2} - \frac{V}{m_a \epsilon} \right)$$

$$\bar{q}_2 \leq q_{\mu^*} - 1 + \log_{\frac{\lambda}{a_\lambda - \epsilon}} \left[1 + \left(\frac{\lambda - a_\lambda + \epsilon}{\lambda} \right) \frac{\epsilon m_a}{V} \left(\frac{1}{2} - \frac{V}{m_a \epsilon} \right) \right].$$

Since $q_{\mu^*} \geq 0$, \bar{q}_2 is at least

$$\left\lceil \log_{\frac{\lambda}{a_\lambda - \epsilon}} \left[1 + \left(\frac{\lambda - a_\lambda + \epsilon}{\lambda} \right) \frac{\epsilon m_a}{V} \left(\frac{1}{2} - \frac{V}{m_a \epsilon} \right) \right] - 1 \right\rceil.$$

Therefore,

$$\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{\bar{q}_2}{2} \geq \frac{1}{2} \left\lceil \log_{\frac{\lambda}{a_\lambda - \epsilon}} \left[1 + \left(\frac{\lambda - a_\lambda + \epsilon}{\lambda} \right) \frac{\epsilon m_a}{V} \left(\frac{1}{2} - \frac{V}{m_a \epsilon} \right) \right] - 2 \right\rceil.$$

Hence, for any sequence γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we obtain that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

Remark 3.2.7. We note that as $V \downarrow 0$, $\pi(0) \downarrow 0$. We also note that as $V \downarrow 0$, there exists a set of queue lengths, \mathcal{Q}_h , occurring with high probability $(1 - V)$, such that $\mu(q) \in [a_\lambda, b_\lambda]$, $\forall q \in \mathcal{Q}_h$. If $\pi(0) \rightarrow 0$, then $|\mathcal{Q}_h| \rightarrow \infty$. We note that for $q \notin \mathcal{Q}_h$, $\pi(q) = \mathcal{O}(V)$. Then from the birth death structure of $Q(t)$ we obtain that the smallest queue length $q_{min} \in \mathcal{Q}_h$ has $\pi(q_{min}) = \mathcal{O}(V)$. Since for $q \in \mathcal{Q}_h$, $\mu(q) \in [a_\lambda, b_\lambda]$ and the policies that we consider are monotone, the stationary probability distribution of the queue lengths in \mathcal{Q}_h can be observed to be geometrically increasing and then decreasing, which leads to the $\log\left(\frac{1}{V}\right)$ growth for the average queue length.

Remark 3.2.8. We note that the lower bounding technique in [44] can be used to obtain the $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ lower bound by considering a uniformized version of $Q(t)$ as we discussed in Remark 3.2.4. We again note that obtaining the stationary probability of the queue length, as in the above proof has its advantages, since it gives us additional insights into the form of the optimal policy.

Lemma 3.2.9. For INTERVAL- μ CHOICE-2-3, for any sequence of non-idling admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$.

Proof. Consider a policy γ in the sequence γ_k , with $V_k = V$. Let $\mu^* \triangleq \lambda - \epsilon_V = a_\lambda - \epsilon_V$. Since $c(\mu_k) \geq l(\mu_k)$, we have

$$\begin{aligned} V &\geq \sum_{\mu_k < \lambda} (c(\mu_k) - l(\mu_k))\pi_\mu(k) \geq \sum_{\mu_k < \mu^*} (c(\mu_k) - l(\mu_k))\pi_\mu(k) \\ &\geq \sum_{\mu_k < \mu^*} m_a(\lambda - \mu_k)\pi_\mu(k) \geq m_a\epsilon_V Pr\{\mu(Q) < \mu^*\}, \end{aligned}$$

where m_a is the tangent of the angle made by the line passing through $(a_i = a_\lambda, c(a_i))$ and $(a_{i-1}, c(a_{i-1}))$ with the line $l(\mu)$. Let $\epsilon_V \triangleq a_2V$, where a_2 is chosen so that $\alpha \triangleq \frac{2V}{m_a\epsilon_V} < 1$. We note that a_2 can be chosen such that α is arbitrarily close to zero. Define $q_{\mu^*} = \inf\{q : \mu(q) \geq \mu^*\}$. Since γ is admissible, we have that $Pr\{Q < q_{\mu^*}\} = Pr\{\mu(Q) < \mu^*\} \leq \frac{V}{m_a\epsilon_V}$. We have

$$\begin{aligned} Pr\{Q < q_{\mu^*}\} &\leq \frac{\alpha}{2}, \\ \text{and } \pi(q_{\mu^*} - 1) &\leq \frac{\alpha}{2}. \end{aligned}$$

To find a lower bound on $\bar{Q}(\gamma)$, in the following, we intend to find the largest \bar{q} , such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$. But we note that $Pr\{Q < q_{\mu^*}\} \leq \frac{\alpha}{2}$ and for any $q < q_{\mu^*}$, $\pi(q) \leq \frac{\alpha}{2}$. Therefore, $\pi(q_{\mu^*}) \leq \pi(q_{\mu^*} - 1) \frac{\lambda}{\mu^*} \leq \frac{\lambda\alpha}{2\mu^*}$. If \bar{q}_1 is the largest integer such that

$$\sum_{q=q_{\mu^*}}^{\bar{q}} \pi(q) \leq \frac{1}{2} - \frac{V}{m_a\epsilon_V},$$

then $\bar{q}_1 \leq \bar{q}$. Proceeding as for problem INTERVAL- μ CHOICE-1, we obtain a lower bound \bar{q}_2 on \bar{q}_1 by using an upper bound for $\pi(q)$. We note that if $q \geq q_{\mu^*}$ we have that $\pi(q-1)\lambda \geq \pi(q)\mu^*$. By induction, we obtain that for $q \in \{q_{\mu^*}, \dots\}$

$$\begin{aligned} \pi(q) &\leq \pi(q_{\mu^*} - 1) \left(\frac{\lambda}{\mu^*}\right)^{q - q_{\mu^*} + 1}, \\ \text{and for any } q' \geq q_{\mu^*}, \sum_{q=q_{\mu^*}}^{q'} \pi(q) &\leq \pi(q_{\mu^*} - 1) \sum_{m=1}^{q' - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*}\right)^m. \end{aligned}$$

Using the above upper bound on $\sum_{q=q_{\mu^*}}^{q'} \pi(q)$ we obtain the following lower bound \bar{q}_2 to \bar{q}_1 . If \bar{q}_2

is the largest integer such that

$$\pi(q_{\mu^*} - 1) \sum_{m=1}^{\bar{q}_2 - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*} \right)^m \leq \frac{1 - \alpha}{2},$$

then $\bar{q}_2 \leq \bar{q}$.

We note that

$$\begin{aligned} V &\geq \sum_{\mu_k < \mu^*} m_a (\lambda - \mu_k) \pi_\mu(k), \\ &= \sum_{q < q_{\mu^*}} m_a (\lambda - \mu(q)) \pi(q). \end{aligned}$$

Again, since $\pi(q)\mu(q) = \pi(q-1)\lambda$, it follows that

$$V \geq m_a \lambda \pi(q_{\mu^*} - 1).$$

Now if \bar{q}_3 is the largest integer such that

$$\frac{V}{m_a \lambda} \sum_{m=1}^{\bar{q}_3 - q_{\mu^*} + 1} \left(\frac{\lambda}{\mu^*} \right)^m \leq \frac{1 - \alpha}{2},$$

then $\bar{q}_3 \leq \bar{q}_2$. We have that \bar{q}_3 satisfies

$$\begin{aligned} \bar{q}_3 &\leq q_{\mu^*} - 1 + \log_{\frac{\lambda}{\mu^*}} \left[1 + \frac{\epsilon_V m_a \lambda}{\lambda V} \frac{1 - \alpha}{2} \right], \\ \bar{q}_3 &\leq q_{\mu^*} - 1 + \frac{\log \left[1 + \frac{\epsilon_V m_a \lambda}{\lambda V} \frac{1 - \alpha}{2} \right]}{-\log \left(1 - \frac{\epsilon_V}{\lambda} \right)}. \end{aligned}$$

Since $q_{\mu^*} \geq 0$, and $\epsilon_V = a_2 V$, we have that \bar{q}_3 is at least

$$\left\lfloor \frac{\log \left[1 + m_a a_2 \frac{1 - \alpha}{2} \right]}{-\log \left(1 - \frac{a_2 V}{\lambda} \right)} - 1 \right\rfloor.$$

So that

$$\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{\bar{q}_2}{2} \geq \frac{\bar{q}_3}{2} \geq \frac{1}{2} \left[\frac{\log \left[1 + m_a a_2 \frac{1 - \alpha}{2} \right]}{-\log \left(1 - \frac{a_2 V}{\lambda} \right)} - 2 \right].$$

Since $\log \left(1 - \frac{a_2 V}{\lambda} \right) = \Theta(V)$ as $V \downarrow 0$, we have that for any sequence γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$. \square

Remark 3.2.10. We note that as $V \downarrow 0$, $\pi(0) \rightarrow 0$. We also note that as $V \downarrow 0$, there exists a set of queue lengths, \mathcal{Q}_h , occurring with high probability $(1 - V)$, such that $\mu(q) \rightarrow \lambda, \forall q \in \mathcal{Q}_h$.

If $\pi(0) \rightarrow 0$, then $|\mathcal{Q}_h| \rightarrow \infty$. Furthermore, for each $q \in \mathcal{Q}_h$, $\pi(q) = \mathcal{O}(V)$, rather than $\mathcal{O}(\sqrt{V})$ as in Remark 3.2.3. We also note that the stationary probability for each $q \in \mathcal{Q}_h$ become equal as $V \downarrow 0$. Then the average queue length is $\Omega\left(\frac{1}{V}\right)$. We note that the difference from the behaviour in Remark 3.2.3 arises since $c(\mu)$ is piecewise linear rather than being strictly convex.

3.2.2 Asymptotic behaviour of the tradeoff curve

In this section, we obtain asymptotic upper bounds for the cases INTERVAL- μ CHOICE-1, INTERVAL- μ CHOICE-2-1, INTERVAL- μ CHOICE-2-2, and INTERVAL- μ CHOICE-2-3. With the asymptotic lower bounds which were derived in the previous section, these bounds provide an almost complete order characterization of the tradeoff curve $Q^*(c_c)$.

Lemma 3.2.11. For INTERVAL- μ CHOICE-1, there exists a sequence of admissible policies γ_k with a sequence $V_k \downarrow 0$, such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. We evaluate the average queue length $\bar{Q}(\gamma)$ and average service cost $\bar{C}(\gamma)$ for a policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \lambda - \epsilon_U, \text{ for } q \in \{1, \dots, q_1\}, \\ \mu(q) &= \lambda + \epsilon'_U, \text{ for } q \in \{q_1 + 1, \dots, 2q_1\}, \\ \mu(q) &= \lambda + K, \text{ for } q \in \{2q_1 + 1, \dots\}.\end{aligned}$$

Let $\epsilon_U = \sqrt{U}$, $\epsilon'_U = \frac{\lambda \epsilon_U}{\lambda - \epsilon_U}$, and K be some positive constant such that $\lambda + K \leq r_{max}$. We also let $q_1 \triangleq \left\lceil \log\left(\frac{\lambda}{\lambda - \epsilon_U}\right) \left(1 + \frac{\epsilon_U}{U\lambda}\right) \right\rceil$. The sequence of policies γ_k is obtained by choosing U from a sequence $U_k \downarrow 0$.

For γ , we have

$$\pi(q) = \begin{cases} \pi(0) \left(\frac{\lambda}{\lambda - \epsilon_U}\right)^q & \text{for } q \in \{1, \dots, q_1\}, \\ \pi(0) \left(\frac{\lambda}{\lambda - \epsilon_U}\right)^{q_1} \left(\frac{\lambda}{\lambda + \epsilon'_U}\right)^{q - q_1} & \text{for } q \in \{q_1 + 1, \dots, 2q_1\}, \\ \pi(0) \left(\frac{\lambda}{\lambda - \epsilon_U}\right)^{q_1} \left(\frac{\lambda}{\lambda + \epsilon'_U}\right)^{q_1} \left(\frac{\lambda}{\lambda + K}\right)^{q - 2q_1} = \pi(0) \left(\frac{\lambda}{\lambda + K}\right)^{q - 2q_1} & \text{for } q \in \{2q_1 + 1, \dots\}.\end{cases}$$

Now, since $\sum_{q=0}^{\infty} \pi(q) = 1$ we have that

$$\begin{aligned} \pi(0) \left[1 + \sum_{q=1}^{q_1} \left(\frac{\lambda}{\lambda - \epsilon_U} \right)^q + \left(\frac{\lambda}{\lambda - \epsilon_U} \right)^{q_1} \sum_{q=1}^{q_1} \left(\frac{\lambda}{\lambda + \epsilon'_U} \right)^q + \sum_{q=1}^{\infty} \left(\frac{\lambda}{\lambda + K} \right)^q \right] &= 1, \\ \pi(0) \left[1 + \frac{\lambda}{\epsilon_U} \left(\left(\frac{\lambda}{\lambda - \epsilon_U} \right)^{q_1} - 1 \right) + \left(\frac{\lambda}{\lambda - \epsilon_U} \right)^{q_1} \frac{\lambda}{\epsilon'_U} \left(1 - \left(\frac{\lambda}{\lambda + \epsilon'_U} \right)^{q_1} \right) + \frac{\lambda}{K} \right] &= 1, \\ \pi(0) \left[1 + \frac{\lambda}{\epsilon_U} \left(\left(\frac{\lambda}{\lambda - \epsilon_U} \right)^{q_1} - 1 \right) + \left(\frac{\lambda}{\lambda - \epsilon_U} \right)^{q_1} \frac{\lambda}{\epsilon'_U} - \frac{\lambda}{\epsilon'_U} + \frac{\lambda}{K} \right] &= 1. \end{aligned}$$

From the above expression, using the lower bound $\log_{\left(\frac{\lambda}{\lambda - \epsilon_U}\right)} \left(1 + \frac{\epsilon_U}{U\lambda} \right) - 1$ on q_1 , we have that

$$\begin{aligned} \pi(0) &\leq \frac{1}{1 + \frac{\lambda}{\epsilon_U} \left(\left(\frac{\lambda - \epsilon_U}{\lambda} \right) \left(1 + \frac{\epsilon_U}{U\lambda} \right) - 1 \right) + \frac{\lambda - \epsilon_U}{\lambda} \frac{\lambda}{\epsilon'_U} \left(1 + \frac{\epsilon_U}{U\lambda} \right) - \frac{\lambda}{\epsilon'_U} + \frac{\lambda}{K}}, \\ &= \frac{\epsilon'_U U \lambda}{(\epsilon'_U + \epsilon)(\lambda(1 - U) - \epsilon_U) + U \lambda \epsilon'_U \left(1 + \frac{\lambda}{K} \right)}, \\ &= \mathcal{O}(U) \end{aligned} \tag{3.10}$$

We now evaluate $\bar{C}(\gamma)$. We have that

$$\begin{aligned} \bar{C}(\gamma) &= \pi(0) \cdot 0 + \pi_{\mu}(\lambda - \epsilon_U) c(\lambda - \epsilon_U) + \pi_{\mu}(\lambda + \epsilon'_U) c(\lambda + \epsilon'_U) + \pi_{\mu}(\lambda + K) c(\lambda + K), \\ &= \pi_{\mu}(\lambda - \epsilon_U) \left(c(\lambda) - \epsilon_U \frac{dc(\lambda)}{d\mu} + \mathcal{O}(\epsilon_U^2) \right) + \pi_{\mu}(\lambda + \epsilon'_U) \left(c(\lambda) + \epsilon'_U \frac{dc(\lambda)}{d\mu} + \mathcal{O}((\epsilon'_U)^2) \right) \\ &\quad + \pi_{\mu}(\lambda + K) \left(c(\lambda) + K \frac{dc(\lambda)}{d\mu} + G(K) \right), \end{aligned}$$

where $\frac{dc(\lambda)}{d\mu} \triangleq \frac{dc(\mu)}{d\mu} \Big|_{\mu=\lambda}$, and $G(K) = c(\lambda + K) - c(\lambda) - K \frac{dc(\lambda)}{d\mu}$. Combining the $c(\lambda)$ terms, bounding $\pi_{\mu}(\lambda - \epsilon_U) + \pi_{\mu}(\lambda + \epsilon'_U) + \pi_{\mu}(\lambda + K)$ above by 1 and using $\epsilon_U = \epsilon'_U = \mathcal{O}(\sqrt{U})$, we obtain that

$$\bar{C}(\gamma) \leq c(\lambda) + \mathcal{O}(U) + (-\epsilon_U \pi_{\mu}(\lambda - \epsilon_U) + \epsilon'_U \pi_{\mu}(\lambda + \epsilon'_U) + K \pi_{\mu}(\lambda + K)) \frac{dc(\lambda)}{d\mu} + G(K) \pi_{\mu}(\lambda + K).$$

In the following, we show that $\bar{Q}(\gamma) < \infty$ and therefore γ is admissible. Then, since

$$\begin{aligned} \pi_{\mu}(\lambda - \epsilon_U)(\lambda - \epsilon_U) + \pi_{\mu}(\lambda + \epsilon'_U)(\lambda + \epsilon'_U) + \pi_{\mu}(\lambda + K)(\lambda + K) &= \lambda, \text{ therefore} \\ \pi_{\mu}(\lambda - \epsilon_U)(-\epsilon_U) + \pi_{\mu}(\lambda + \epsilon'_U)(\epsilon'_U) + \pi_{\mu}(\lambda + K)(K) &= \pi(0)\lambda. \end{aligned}$$

Now from (3.10) we have that $\pi(0) = \mathcal{O}(U)$. Furthermore, $\pi_{\mu}(\lambda + K) = \pi(0) \frac{\lambda}{K} = \mathcal{O}(U)$. Therefore

$$\bar{C}(\gamma) \leq c(\lambda) + \mathcal{O}(U) + \frac{dc(\lambda)}{d\mu} \mathcal{O}(U) + G(K) \mathcal{O}(U).$$

For $V \triangleq \bar{C}(\gamma) - c(\lambda)$, $V = \mathcal{O}(U)$.

In order to obtain an upper bound on $\bar{Q}(\gamma)$, we use Proposition 2.A.1 with $q_\epsilon = 2q_1 + 1$ and $\epsilon = K$, to obtain that

$$\begin{aligned}\bar{Q}(\gamma) &\leq \frac{(2q_1 + 1)(K + \lambda)}{K} + \frac{\lambda + r_{max}}{2K}, \\ \bar{Q}(\gamma) &= \mathcal{O}(q_1) = \mathcal{O}\left(\frac{1}{\sqrt{U}} \log\left(\frac{1}{U}\right)\right).\end{aligned}$$

Hence γ is admissible.

Corresponding to the sequence U_k , we have a sequence $V_k = \mathcal{O}(U_k)$. Therefore, we have a sequence of policies γ_k with $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{U_k}} \log\left(\frac{1}{U_k}\right)\right)$. Then $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma) - c(\lambda) = V_k$. \square

We note that the asymptotic upper bound above for the sequence γ_k does not match the asymptotic lower bound $\Omega\left(\frac{1}{\sqrt{V_k}}\right)$, which was derived in Lemma 3.2.2.

Lemma 3.2.12. For INTERVAL- μ CHOICE-2-1, there exists a sequence of admissible policies γ_k , with a sequence of $V_k \downarrow 0$, such that $\frac{\lambda}{b_\lambda - \lambda} - \bar{Q}(\gamma_k) = \Theta\left(V_k \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. Consider the policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= b_\lambda, \text{ for } q \in \{1, \dots, q_{b_\lambda}\}, \\ \mu(q) &= r_{max}, \text{ for } q \in \{q_{b_\lambda} + 1, \dots\}.\end{aligned}$$

The sequence γ_k is obtained by choosing $q_{b_\lambda} = k \in \mathbb{Z}_+$. The rest of the proof is similar to the proof of Lemma 2.3.8. \square

Remark 3.2.13. We note that the above asymptotic upper bound does not match the asymptotic lower bound in Lemma 3.2.5. Since $\mu(q) \in [0, r_{max}]$, one expects that perhaps another sequence of policies for which $\mu(q) = b_\lambda + \epsilon_V$, for $q \in \{1, \dots, q_{b_\lambda}\}$ and where ϵ_V is a sequence decreasing to zero, achieves a better asymptotic upper bound. We have found out that this is not the case, and the asymptotic upper bound is the same as the one above.

Lemma 3.2.14. For INTERVAL- μ CHOICE-2-2, there exists a sequence of admissible policies γ_k , with a sequence $V_k \downarrow 0$, such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma_k) - c(\lambda) = V_k$.

Proof. Consider a policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= a_\lambda, \text{ for } q \in \{1, \dots, q_1\}, \\ \mu(q) &= b_\lambda, \text{ for } q \in \{q_1 + 1, \dots\},\end{aligned}$$

where $q_1 \triangleq \left\lceil \log_{\left(\frac{\lambda}{a_\lambda}\right)} \left(1 + \frac{\lambda - a_\lambda}{\lambda} \frac{1}{U}\right) \right\rceil$, with $U > 0$. The sequence γ_k is obtained by choosing U from a sequence $U_k \downarrow 0$. The rest of the proof is similar to the proof of Lemma 2.3.12. Since $V_k = \mathcal{O}(U_k)$, we obtain $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

Using the asymptotic lower bound on $\bar{Q}(\gamma_k)$ from Lemma 3.2.6, and the asymptotic upper bound above, and proceeding as in the proof of Proposition 2.3.9, we obtain the following result.

Proposition 3.2.15. For INTERVAL- μ CHOICE-2-2, we have that the optimal tradeoff curve $Q^*(c_{c,k})$ is $\Theta\left(\log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$, for a sequence $c_{c,k} = \bar{C}(\gamma_k)$, where γ_k is the sequence of policies in Lemma 3.2.14.

Lemma 3.2.16. For INTERVAL- μ CHOICE-2-3, there exists a sequence of admissible policies γ_k , with a sequence $V_k \downarrow 0$, such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$ and $\bar{C}(\gamma) - c(\lambda) = V_k$.

Proof. Consider a policy γ defined as follows :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \lambda, \text{ for } q \in \{1, \dots, q_1\}, \\ \mu(q) &= \lambda + K, \text{ for } q \in \{q_1 + 1, \dots\}.\end{aligned}$$

We note that $\lambda = a_i$, for some $i > 1$. Here K is a constant such that $\lambda + K \leq a_{i+1}$. We define $q_1 = \left\lceil \frac{1}{U} \right\rceil$, with $U > 0$. The sequence γ_k is obtained by choosing U from a sequence $U_k \downarrow 0$. The rest of the proof is similar to that of Lemma 2.3.14. Since $V_k = \mathcal{O}(U_k)$, we obtain $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. \square

Using the asymptotic lower bound on $\bar{Q}(\gamma_k)$ from Lemma 3.2.9, and the asymptotic upper bound above, and proceeding as in the proof of Proposition 2.3.9, we obtain the following result

Proposition 3.2.17. For INTERVAL- μ CHOICE-2-3, we have that the optimal tradeoff curve $Q^*(c_{c,k})$ is $\Theta\left(\frac{1}{c_{c,k} - c(\lambda)}\right)$, for a sequence $c_{c,k} = \bar{C}(\gamma_k)$, where γ_k is the sequence of policies in Lemma 3.2.16.

3.2.3 Asymptotic characterization of order-optimal admissible policies

As in Chapter 2, it is possible to obtain an asymptotic characterization of any sequence of order-optimal admissible policies using the bounds $P_l\{\cdot\}$, $P_u\{\cdot\}$, $\pi_l(\cdot)$, and $\pi_u(\cdot)$ and the inequalities (2.31) and (2.32). We only consider the cases where the minimum average queue length increases to infinity in the asymptotic regime \mathfrak{R} . The characterization of order-optimal policies for INTERVAL- μ CHOICE-2-2 and INTERVAL- μ CHOICE-2-3 can be obtained using similar methods as FINITE- μ CHOICE-2 and FINITE- μ CHOICE-3. Therefore, in this section we discuss the asymptotic characterization of the optimal policy for INTERVAL- μ CHOICE-1 only. We first obtain $P_l\{A\}$ and $P_u\{A\}$, where $A \subseteq [0, r_{max}]$ is a set of service rates.

Let $0 \leq \delta \leq 1$ and $a_2 > 0$. Let $A \subseteq [0, \lambda - a_2 V^{\frac{1-\delta}{2}}] \cup [\lambda + a_2 V^{\frac{1-\delta}{2}}, r_{max}]$. From (3.1), we have that

$$Pr\{\mu(Q) \in A\} \leq P_u\{A\} \triangleq \frac{V^\delta}{a_1 a_2}, \quad (3.11)$$

if $\delta > 0$. Suppose $\delta = 0$, then for $a_2 > \frac{1}{a_1}$, we have that

$$Pr\{\mu(Q) \in A\} \leq P_u\{A\} \triangleq \frac{1}{a_1 a_2}. \quad (3.12)$$

We then note that $P_l\{A^c\} \triangleq 1 - \frac{V^\delta}{a_1 a_2}$ if $\delta > 0$ and $P_l\{A^c\} \triangleq 1 - \frac{1}{a_1 a_2}$ if $\delta = 0$ and $a_2 > \frac{1}{a_1}$.

The upper bound $\pi_u(q)$ that we use in the following is obtained as in (3.4). We now proceed as in Chapter 2 to obtain a lower bound $\pi_l(0)$ on $\pi(0)$. From Lemma 3.2.11, we have that for any sequence of non-idling order-optimal policies γ_k , $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log\left(\frac{1}{V_k}\right)\right)$. We use the above upper bound since we are not able to show that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}}\right)$. Therefore, as in Section 2.3.3, we have that

$$\begin{aligned} \pi_l(0) &= \Omega\left(\frac{\sqrt{V}}{\log\left(\frac{1}{V}\right)} \left(1 - \frac{\lambda}{r_u}\right)^{\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)}\right), \\ &= \Omega\left(\frac{\sqrt{V}}{\log\left(\frac{1}{V}\right)} \frac{1}{\left(\frac{1}{1 - \frac{\lambda}{r_u}}\right)^{\left(\frac{1}{\sqrt{V}} \log\left(\frac{1}{V}\right)\right)}}\right). \end{aligned} \quad (3.13)$$

With $k \triangleq \frac{1}{\log \frac{1}{\left(\frac{1-\lambda}{ru}\right)} e}$ we then have that

$$\pi_l(0) = \Omega \left(\frac{\sqrt{V}}{\log \left(\frac{1}{V}\right)} \frac{1}{\left(\frac{1-\lambda}{ru}\right)^{\log \frac{1}{\left(\frac{1}{V}\right)} \left[\left(\frac{1}{V}\right)\right]^{\left(\frac{k}{\sqrt{V}}\right)}}} \right) = \Omega \left(\frac{V^{\left(\frac{k}{\sqrt{V}} + \frac{1}{2}\right)}}{\log \left(\frac{1}{V}\right)} \right). \quad (3.14)$$

Furthermore, we have the following lower bound $\pi_l(q')$ on $\pi(q')$, where $q' = \max \{q : \mu(q) < \lambda\}$.

We note that for $\mu < \lambda$, there exists a $m_1 > 0$ such that $c(\mu) - l(\mu) \leq m_1(\lambda - \mu)$. Also for $\mu \geq \lambda$, there exists a $m_2 > 0$ such that $c(\mu) - l(\mu) \leq m_2(\mu - \lambda)$. Then we have that

$$V = \mathbb{E}[c(\mu(Q)) - l(\mu(Q))] \leq \sum_{q:\mu(q)<\lambda} m_1(\lambda - \mu(q)) \pi(q) + \sum_{q:\mu(q)\geq\lambda} m_2(\mu(q) - \lambda) \pi(q).$$

$$\begin{aligned} V &\leq m_1 \sum_{q \leq q'} (\lambda - \mu(q)) \pi(q) + m_2 \sum_{q > q'} (\mu(q) - \lambda) \pi(q), \\ &= m_1 \left[\lambda \pi(0) + \sum_{1 \leq q \leq q'} (\lambda \pi(q) - \lambda \pi(q-1)) \right] + m_2 \sum_{q > q'} (\pi(q-1) - \pi(q)) \lambda, \\ &\leq m_1 \lambda \pi(q') + m_2 \lambda \pi(q'), \text{ or,} \\ \pi(q') &\geq \pi_l(q') \triangleq \frac{V}{\lambda(m_1 + m_2)}. \end{aligned} \quad (3.15)$$

We note that q' could be zero, in which case the sum $\sum_{1 \leq q \leq q'} (\lambda \pi(q) - \lambda \pi(q-1))$ is defined to be zero.

Lemma 3.2.18. For any sequence of non-idling order-optimal admissible policies γ_k , with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, and $Q_A = \{q : \mu(q) \in A\}$ for a $A \subseteq [0, r_{max}]$, we have that

$$|Q_A| = \begin{cases} \mathcal{O} \left(\frac{1}{V_k} \log \left(\frac{1}{V_k} \right) \right), & \text{if } A = \left[0, \lambda - a_2 V_k^{\left(\frac{1-\delta}{2}\right)} \right], \\ \Omega \left(\frac{1}{\sqrt{V_k}} \right), & \text{if } \left[\lambda - a_2 V_k^{\left(\frac{1}{2}\right)}, \lambda + a_2 V_k^{\left(\frac{1}{2}\right)} \right] \subseteq A, \\ \mathcal{O} \left(\frac{1}{V_k} \right), & \text{if } A = \{\lambda\}, \end{cases}$$

where $0 \leq \delta \leq 1$, and $a_2 > \frac{1}{a_1}$ if $\delta = 0$.

Proof. Consider a particular policy γ in the above sequence with $\bar{C}(\gamma_k) - c(\lambda) = V$. Let $A = \left[0, \lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)} \right]$. Then, from (3.11) we have that $P_u \{A\} = \frac{V^\delta}{a_1 a_2}$. Consider a q such that

$\mu(q) \in A$. Then $\pi(q) \geq \pi_l(q)$, where

$$\pi_l(q) \triangleq \pi_l(0) \left(\frac{\lambda}{\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)}} \right)^q,$$

and $\pi_l(0)$ is given in (3.14). As in (2.32) if $q_1 = \max \{q : \mu(q) \in A\}$, then the smallest integer $q_{1,u}$ such that

$$\sum_{q=0}^{q_{1,u}} \pi_l(q) \geq \frac{V^\delta}{a_1 a_2},$$

is an upper bound on q_1 . We have that

$$\pi_l(0) \left(\left(\frac{\lambda}{\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)}} \right)^{q_{1,u}+1} - 1 \right) \geq \frac{V^{\left(\frac{1+\delta}{2}\right)}}{a_1 \left(\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)} \right)}$$

From (3.14), we have that for small enough V , there exists a $a_3 > 0$ such that

$$\pi_l(0) \geq \left(\frac{a_3 V^{\left(\frac{k}{\sqrt{V}} + \frac{1}{2}\right)}}{\log\left(\frac{1}{V}\right)} \right).$$

So if $q'_{1,u}$ is the smallest integer such that

$$\left(\left(\frac{\lambda}{\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)}} \right)^{q'_{1,u}+1} - 1 \right) \geq \frac{V^{\left(\frac{1+\delta}{2}\right)}}{a_1 \left(\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)} \right)} \left(\frac{\log\left(\frac{1}{V}\right)}{a_3 V^{\left(\frac{k}{\sqrt{V}} + \frac{1}{2}\right)}} \right),$$

then $q'_{1,u} \geq q_{1,u}$. Or we have that $q'_{1,u}$ is the smallest integer such that

$$q'_{1,u} \geq \frac{\log \left(\frac{V^{\left(\frac{1+\delta}{2}\right)}}{a_1 \left(\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)} \right)} \left(\frac{\log\left(\frac{1}{V}\right)}{a_3 V^{\left(\frac{k}{\sqrt{V}} + \frac{1}{2}\right)}} \right) \right)}{\log \left(\frac{\lambda}{\lambda - a_2 V^{\left(\frac{1-\delta}{2}\right)}} \right)}$$

As $V \downarrow 0$, we have that $q'_{1,u} = \mathcal{O} \left(\frac{1}{V} \log \left(\frac{1}{V} \right) \right)$. If $Q_A = \{q : \mu(q) \in A\}$, then $|Q_A| = \mathcal{O} \left(\frac{1}{V} \log \left(\frac{1}{V} \right) \right)$.

Now consider $A = \left[\lambda - a_2 V^{\left(\frac{1}{2}\right)}, \lambda + a_2 V^{\left(\frac{1}{2}\right)} \right]$, where $a_2 > \frac{1}{a_1}$. We have that $P_l(A) = 1 - \frac{1}{a_1 a_2}$. Let $q_1 = \max \left\{ q : \mu(q) < \lambda - a_2 V^{\left(\frac{1}{2}\right)} \right\}$. Then from (3.7) we have that $\pi(q_1) \leq \sqrt{\frac{V}{a_1 \lambda^2}}$. Then for every $q \in A$, we have that

$$\pi(q) \leq \pi_u(q) \triangleq \sqrt{\frac{V}{a_1 \lambda^2}} \left(\frac{\lambda}{\lambda - a_2 V^{\left(\frac{1}{2}\right)}} \right)^{q-q_1}. \quad (3.16)$$

Let $Q_A = \{q : \mu(q) \in A\}$. We now obtain a lower bound on $|Q_A|$. From (2.31), if $q_{l,a}$ is the largest integer such that

$$\sum_{q=0}^{q_{l,a}} \sqrt{\frac{V}{a_1 \lambda^2}} \left(\frac{\lambda}{\lambda - a_2 V^{(\frac{1}{2})}} \right)^q \leq 1 - \frac{1}{a_1 a_2}, \quad (3.17)$$

then $q_{l,a} \leq |Q_A|$. This is equivalent to finding the largest integer $q_{l,a}$ such that

$$\begin{aligned} \sqrt{\frac{V}{a_1 \lambda^2}} \left(\left(\frac{\lambda}{\lambda - a_2 V^{(\frac{1}{2})}} \right)^{q_{l,a}} - 1 \right) &\leq \left(1 - \frac{1}{a_1 a_2} \right) \frac{a_2 V^{\frac{1}{2}}}{\lambda - a_2 V^{\frac{1}{2}}}, \\ \left(\frac{\lambda}{\lambda - a_2 V^{(\frac{1}{2})}} \right)^{q_{l,a}} - 1 &\leq \left(1 - \frac{1}{a_1 a_2} \right) \frac{a_1 \sqrt{a_2 \lambda^2}}{\lambda - a_2 V^{\frac{1}{2}}}. \end{aligned}$$

Then as in the proof of Lemma 3.2.2 we have that $q_{l,a} = \Omega\left(\frac{1}{\sqrt{V}}\right)$.

Consider $Q_\lambda = \{q : \mu(q) = \lambda\}$. We note that $\pi(Q_\lambda) \leq 1$. We also note that for every $q \in Q_\lambda$, $\pi(q) \leq \pi_l(q) \triangleq \pi_l(q')$ defined in (3.15). As in (2.32), if $q_{u,\lambda}$ is the smallest integer such that

$$q_{u,\lambda} \pi_l(q') \leq 1,$$

then $q_{u,\lambda}$ is an upper bound on $|Q_\lambda|$. Hence we obtain that $|Q_\lambda| = \mathcal{O}\left(\frac{1}{\sqrt{V}}\right)$. \square

Remark 3.2.19. From the above proof, the intuition behind choosing the buffer partition to scale as $\Omega\left(\frac{1}{\sqrt{V}}\right)$ can be observed.

3.2.4 Tradeoff problems which are similar to INTERVAL- μ CHOICE

In this section, we consider tradeoff problems which are similar to INTERVAL- μ CHOICE, for which an asymptotic characterization can be obtained using the techniques presented above for INTERVAL- μ CHOICE. We note that for INTERVAL- μ CHOICE, we restricted to admissible policies γ for which $\lambda(q) = \lambda, \forall q \in \mathbb{Z}_+$, and λ was such that $u(\lambda) \geq u_c$. The counterpart INTERVAL- λ CHOICE of INTERVAL- μ CHOICE is one in which we restrict to admissible policies γ for which $\mu(q) = \mu, \forall q \in \{1, 2, \dots\}$ and $\lambda(q) \in [0, r_{a,max}]$.

We note that $\bar{C}(\gamma)$ for such a policy γ is $(1 - \pi(0))c(\mu)$, which depends on the policy, unlike INTERVAL- μ CHOICE where the choice of λ fixed $\bar{U}(\gamma)$ to be $u(\lambda)$. For INTERVAL- λ CHOICE, we restrict to admissible γ such that $\mu(q) = \mu, \forall q > 0$, where μ is such that $c(\mu) \leq c_c$, so that $\bar{C}(\gamma) \leq c_c$. The tradeoff problem INTERVAL- λ CHOICE is

$$\begin{aligned} \text{minimize } & \gamma \in \Gamma_a \quad \bar{Q}(\gamma) \\ \text{and } & \bar{U}(\gamma) \geq u_c. \end{aligned} \quad (3.18)$$

The optimal value of the above problem is denoted as $Q^*(u_c)$. We also note from Lemma 2.2.2, that the maximum value of $\bar{U}(\gamma)$ over all admissible γ is $u(\mu)$. We obtain an asymptotic characterization of $Q^*(u_c)$ in the asymptotic regime \mathfrak{R} , where $u_c \uparrow u(\mu)$.

We assume that μ is such that $r_{a,max} > \mu$. If $r_{a,max} = \mu$, then we note that only the non-admissible policy γ , with $\lambda(q) = r_{a,max}, \forall q$, can achieve $u(\mu)$. Intuitively, if $u_c \uparrow u(\mu)$, then since $\mathbb{E}\lambda(Q) = \mathbb{E}\mu(Q) = (1 - \pi(0))\mu \uparrow \mu$, we have that $\pi(0) \downarrow 0$. Then for any policy which is feasible as $u_c \uparrow u(\lambda)$, $\pi(0) \downarrow 0$. Thus, intuitively for problem (3.18) we do not have a case where $Q^*(u_c)$ increases only up to a finite value as $u_c \uparrow u(\mu)$ (unlike INTERVAL- μ CHOICE-2-1).

As for INTERVAL- μ CHOICE, we consider the following cases for INTERVAL- λ CHOICE:

INTERVAL- λ CHOICE-1: $u(\lambda)$ is strictly concave for $\lambda \in [0, r_{a,max}]$.

INTERVAL- λ CHOICE-2: $u(\lambda)$ is piecewise linear and concave. That is, (a) there exists a minimal partition of $[0, r_{a,max}]$ into intervals $\{[a_i, b_i], i \in \{1, \dots, P\}\}$ with $a_1 = 0$, $b_P = r_{a,max}$, and $b_i = a_{i+1}$ and (b) there are linear functions f_i such that $\forall \mu \in [a_i, b_i], f_i(\lambda) = u(\lambda)$. This is further subdivided into two cases:

1. $a_\mu \triangleq a_i < \mu < b_\mu \triangleq b_i$, for some $i \geq 1$.
2. $\mu = a_\mu \triangleq a_i$, for some $i > 1$.

Remark 3.2.20. We note that for the discrete time queueing model, we only consider the case where the utility function is linear, since for such models we are interested in the average throughput as the performance measure. We note that INTERVAL- λ CHOICE-2-1 encompasses the case of linear utility functions. For linear utility functions, we can motivate the choice of system parameters for INTERVAL- λ CHOICE-2-1 as done for the case of INTERVAL- μ CHOICE. INTERVAL- λ CHOICE-2-1 with linear utility function is a simplified model for a discrete time queueing model, where the service batch size is *fixed* (but if the queue length is less than this fixed batch size, then the service batch size is equal to the queue length) and there is randomized admission control. The fixed service batch size is modelled by the fixed μ , while the admission control is modelled by the choice of the arrival rate $\lambda(q)$ as a function of q . In light of the discussion in Section 3.1 we assume that $\lambda(q)$ takes values in a finite interval.

We now present an asymptotic characterization of $Q^*(u_c)$ in the regime $u_c \uparrow u(\mu)$. We note that for INTERVAL- λ CHOICE-1 as well as for INTERVAL- λ CHOICE-2 it is possible to show (see Lemma 3.2.28) that there exists a sequence of admissible policies γ_k such that $\bar{U}(\gamma_k) \uparrow u(\mu)$. We note that for every $u_c < u(\mu)$, $\forall \epsilon > 0$, there exists some feasible $\gamma \in \Gamma_a$ such that $\bar{Q}(\gamma) \leq Q^*(u_c) + \epsilon$. Such an admissible policy is called ϵ -optimal in the following.

We first present asymptotic lower bounds on $Q^*(u_c)$ in the regime $u_c \uparrow u(\mu)$. Asymptotic lower bounds for INTERVAL- λ CHOICE are obtained along similar lines as for INTERVAL- μ CHOICE. For

all cases, we first obtain upper bounds on the stationary probability of certain arrival rates (rather than service rates), which go to zero as $u_c \uparrow u(\mu)$. Then as before, these upper bounds on the stationary probability of certain arrival rates lead to constraints on the stationary probability of all queue lengths. Since the stationary probability of the queue length determines the average queue length, the constraints determine the behaviour of average queue length as $u_c \uparrow u(\mu)$.

For the asymptotic analysis of INTERVAL- λ CHOICE, we define a line $l(\lambda)$ whose definition is similar to that of $l(\mu)$ for INTERVAL- μ CHOICE. For INTERVAL- λ CHOICE-1, $l(\lambda)$ is defined as the tangent to $u(\lambda)$ at $\lambda = \mu$. For INTERVAL- λ CHOICE-2-1, $l(\lambda)$ is defined as the line through $(a_\mu, u(a_\mu))$ and $(b_\mu, u(b_\mu))$, while for INTERVAL- λ CHOICE-2-2, $l(\lambda)$ is any line through $(a_\mu, u(a_\mu))$ with slope m , such that $\frac{du(\lambda)}{d\lambda}^-|_{\lambda=\mu} < m < \frac{du(\lambda)}{d\lambda}^+|_{\lambda=\mu}$. We note that $l(\lambda) \geq u(\lambda)$ and $\mathbb{E}l(\lambda(Q)) = l(\mathbb{E}\lambda(Q))$. We also note that the function $l(\lambda) - u(\lambda)$ is a convex function.

We now present a result, which formalizes the intuition that $\pi(0) \downarrow 0$ as $u_c \uparrow u(\mu)$.

Lemma 3.2.21. For INTERVAL- λ CHOICE, for any sequence of admissible policies such that $u(\mu) - \bar{U}(\gamma_k) = V_k \downarrow 0$, we have that $\pi(0) = \mathcal{O}(V_k)$. Therefore, as $u_c \uparrow u(\mu)$, $\pi(0) \downarrow 0$, for any sequence of feasible policies for (3.18).

Proof. Consider a particular policy γ in the sequence with $V_k = V$. We note that $\bar{U}(\gamma) \leq u(\mathbb{E}\lambda(Q))$. Since γ is admissible, we have that $\bar{U}(\gamma) \leq u(\mathbb{E}\mu(Q))$. We then have that $u^{-1}(\bar{U}(\gamma)) \leq \mathbb{E}\mu(Q) = (1 - \pi(0))\mu$, since $u^{-1}(\cdot)$ exists if $u(\lambda)$ is concave and increasing in λ . Therefore, we have that

$$\pi(0) \leq 1 - \frac{u^{-1}(\bar{U}(\gamma))}{\mu} = 1 - \frac{u^{-1}(u(\mu) - V)}{\mu}.$$

We note that $u^{-1}(x) \geq l^{-1}(x)$, $x \in \mathbb{R}_+$, where $l^{-1}(\cdot)$ is the inverse function of $l(\lambda)$. Then we have

$$\pi(0) \leq 1 - \frac{l^{-1}(u(\mu) - V)}{\mu} = 1 - \frac{l^{-1}(u(\mu)) - mV}{\mu} = \frac{mV}{\mu},$$

since $u(\mu) = l(\mu)$ and where m is the slope of l^{-1} . Therefore, for the sequence γ_k , $\pi(0) = \mathcal{O}(V_k)$. For INTERVAL- λ CHOICE, as $u_c \uparrow u(\mu)$, for any sequence γ_k of feasible policies, $u(\mu) - \bar{U}(\gamma_k) \downarrow 0$ and hence $\pi(0) \downarrow 0$. \square

For any policy, the set of arrival rates $\{\lambda(q) : q \in \mathbb{Z}_+\}$ is countable and is denoted as $(\lambda_0, \lambda_1, \dots)$, with $\lambda_k < \lambda_{k+1}$. For an admissible policy, let $\pi_\lambda(k)$ denote the stationary probability of using an arrival rate λ_k , i.e., $\pi_\lambda(k) = Pr\{\lambda(Q) = \lambda_k\} = \sum_{\{q:\lambda(q)=\lambda_k\}} \pi(q)$.

We make the following assumption, which is similar to (C2):

U2: For INTERVAL- λ CHOICE-1, the second derivative of $u(\lambda)$ at $\lambda = \mu$ is non-zero.

Lemma 3.2.22. For INTERVAL- λ CHOICE-1, for any sequence of non-idling admissible policies γ_k such that $u(\mu) - \bar{U}(\gamma_k) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$. Therefore, $Q^*(u_c) = \Omega\left(\frac{1}{\sqrt{u(\mu) - u_c}}\right)$.

Proof. The proof follows that of Lemma 3.2.2, but with some minor differences. We again consider a particular policy in the sequence with $V_k = V$. Let $\lambda^* \triangleq \mu + \epsilon_V$, where $\epsilon_V > 0$ is a function of V to be chosen later. As in the proof of Lemma 3.2.2, we have that

$$V = \sum_{q=0}^{\infty} \pi(q) \left[(\mu - \lambda(q)) \frac{du(\lambda)}{d\lambda} \Big|_{\lambda=\mu} + G(\lambda(q) - \mu) \right],$$

where $G(x)$ is a strictly convex function of x . As $\lambda(q) - \mu$ is bounded, we again have that there exists a positive a_1 such that

$$V \geq \sum_{q=0}^{\infty} \pi(q) \left[(\mu - \lambda(q)) \frac{du(\lambda)}{d\lambda} \Big|_{\lambda=\mu} + a_1(\lambda(q) - \mu)^2 \right].$$

Since $\sum_{q=0}^{\infty} \pi(q) \leq \mu$, we have that

$$V \geq \sum_{q=0}^{\infty} \pi(q) a_1 (\lambda(q) - \mu)^2.$$

Let $q_{\lambda^*} \triangleq \inf \{q : \lambda(q) \leq \lambda^*\}$. We note that unlike q_{μ^*} in Lemma 3.2.2, q_{λ^*} could be 0. We proceed as in the proof of Lemma 3.2.2 by choosing $\epsilon_V = a_2 \sqrt{V}$. Then $\Pr \{Q < q_{\lambda^*}\} \leq \frac{1}{a_1 a_2}$ if $q_{\lambda^*} > 0$. As before, we choose a_2 such that $\Pr \{Q < q_{\lambda^*}\} \leq \frac{\alpha}{2}$, where α can be made arbitrarily close to zero. If $q_{\lambda^*} = 0$, then $\Pr \{Q < q_{\lambda^*}\} = 0 \leq \frac{1}{a_1 a_2}$.

As in the proof of Lemma 3.2.2 we find the largest \bar{q} such that $\Pr \{Q \leq \bar{q}\} \leq \frac{1}{2}$. We note that if $q \geq q_{\lambda^*}$, then $\pi(q-1)\lambda^* \geq \pi(q)\mu$. Then by induction we obtain that for any $q \geq q_{\lambda^*}$,

$$\sum_{q=q_{\lambda^*}}^q \pi(q) \leq \pi(q_{\lambda^*}) \sum_{m=0}^{q-q_{\lambda^*}} \left(\frac{\lambda^*}{\mu}\right)^m. \quad (3.19)$$

We note that this is similar to (3.4), except that we express the above upper bound in terms of $\pi(q_{\lambda^*})$ rather than $\pi(q_{\mu^*} - 1)$ in (3.4), since q_{λ^*} could be zero.

If $q_{\lambda^*} = 0$, then from Lemma 3.2.21 we have that $\pi(q_{\lambda^*}) = \pi(0) = \mathcal{O}(V)$. If $q_{\lambda^*} > 0$, we obtain an upper bound on $\pi(q_{\lambda^*})$, as in the proof of Lemma 3.2.2. We have that

$$\frac{V}{a_1} \geq \sum_{q < q_{\lambda^*}} \pi(q) (\lambda(q) - \mu)^2 \geq \left(\sum_{q < q_{\lambda^*} - 1} (\lambda(q) - \mu) \pi(q) \right)^2.$$

Since for $q > 0$, since $\pi(q)\lambda(q) = \mu\pi(q+1)$, we proceed as in the proof of Lemma 3.2.2 to obtain that

$$\begin{aligned}\frac{V}{a_1} &\geq (\mu\pi(q_{\lambda^*}) - \mu\pi(0))^2, \\ &= \mu^2\pi(q_{\lambda^*})^2 + \mu^2\pi(0)^2 - 2\mu^2\pi(q_{\lambda^*})\pi(0).\end{aligned}$$

Since $\pi(0) \geq 0$ and $\pi(0) = \mathcal{O}(V)$ from Lemma 3.2.21, we have that

$$\begin{aligned}\frac{V}{a_1} + 2\mu^2\pi(q_{\lambda^*})\pi(0) &\geq \mu^2\pi(q_{\lambda^*})^2, \\ \frac{V}{a_1} + 2\mu^2\mathcal{O}(V) &\geq \mu^2\pi(q_{\lambda^*})^2, \\ \text{or } \pi(q_{\lambda^*}) &= \mathcal{O}(\sqrt{V}).\end{aligned}$$

We note that for both $q_{\lambda^*} = 0$ or $q_{\lambda^*} > 0$, we have that $\pi(q_{\lambda^*}) = \mathcal{O}(\sqrt{V})$.

We now proceed as in the proof of Lemma 3.2.2, by using (3.19), to find the largest integer \bar{q} such that

$$\pi(q_{\lambda^*}) \sum_{m=0}^{q-q_{\lambda^*}} \left(\frac{\lambda^*}{\mu}\right)^m \leq \frac{1}{2} - \frac{\alpha}{2}.$$

The rest of the proof is similar to that of Lemma 3.2.2, and we obtain that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$. Then given a sequence of $u_{c,k} \uparrow u(\mu)$, we have that there exists a sequence of feasible γ_k such that $\bar{Q}(\gamma_k) \leq Q^*(u_{c,k}) + \epsilon$, for some $\epsilon > 0$. Therefore, $Q^*(u_{c,k}) = \Omega\left(\frac{1}{\sqrt{u(\mu) - u_{c,k}}}\right)$, since $u_{c,k} \leq \bar{U}(\gamma_k)$. \square

Lemma 3.2.23. For INTERVAL- λ CHOICE-2-1, if $b_i = b_\mu$ and $i < P$, then for any sequence of non-idling admissible policies γ_k such that $u(\mu) - \bar{U}(\gamma_k) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. Therefore $Q^*(u_c) = \Omega\left(\log\left(\frac{1}{u(\mu) - u_c}\right)\right)$.

Proof. The proof follows that of Lemma 3.2.6. We define $\lambda^* = b_\mu + \epsilon$, where $\epsilon > 0$. Let $q_{\lambda^*} = \inf\{q : \lambda(q) \leq \lambda^*\}$. We note that q_{λ^*} could be 0, unlike q_{μ^*} in Lemma 3.2.6.

If $q_{\lambda^*} = 0$, then we have that $\pi(q_{\lambda^*}) = \pi(0) = \mathcal{O}(V)$. If $q_{\lambda^*} > 0$, then we have that

$$\begin{aligned}V &= \sum_{\lambda_k > \lambda^*} (l(\lambda_k) - u(\lambda_k)) \pi_\lambda(k), \\ &\geq m_a \epsilon Pr\{\lambda(Q) > \lambda^*\},\end{aligned}$$

where m_a is the tangent of the angle made by the line passing through $(b_i = b_\mu, u(b_\mu))$ and $(b_{i+1}, u(b_{i+1}))$ with $l(\mu)$. Or we have that $Pr\{Q < q_{\lambda^*}\} \leq \frac{V}{m_a \epsilon}$ and $\pi(q_{\lambda^*} - 1) \leq \frac{V}{m_a \epsilon}$. Since

$\pi(q_{\lambda^*} - 1)\lambda(q_{\lambda^*} - 1) = \pi(q_{\lambda^*})\mu$ we have that $\pi(q_{\lambda^*}) \leq \pi(q_{\lambda^*} - 1)^{\frac{r_{a,max}}{\mu}}$. We note that therefore $\pi(q_{\lambda^*}) = \mathcal{O}(V)$ for both $q_{\lambda^*} = 0$ and $q_{\lambda^*} > 0$.

Now proceeding as in the proof of Lemma 3.2.6 we have that for any $q \geq q_{\lambda^*}$ (we express the bound in terms of $\pi(q_{\lambda^*})$)

$$\sum_{q=q_{\lambda^*}}^q \pi(q) \leq \pi(q_{\lambda^*}) \sum_{m=0}^{q-q_{\lambda^*}} \left(\frac{\lambda^*}{\mu}\right)^m. \quad (3.20)$$

We note that independently of whether q_{λ^*} is 0 or not, if we find the largest \bar{q} such that

$$\sum_{q=q_{\lambda^*}}^{\bar{q}} \pi(q) \leq \frac{1}{2} - \frac{V}{a_1 \epsilon},$$

then $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2}$. We now proceed as in the proof of Lemma 3.2.6, using the upper bound in (3.20) and $\pi(q_{\lambda^*}) = \mathcal{O}(V)$ to obtain that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. Now given a sequence of $u_{c,k} \uparrow u(\mu)$, we have that there exists a sequence of feasible γ_k such that $\bar{Q}(\gamma_k) \leq Q^*(u_{c,k}) + \epsilon$, for some $\epsilon > 0$. Therefore, $Q^*(u_{c,k}) = \Omega\left(\log\left(\frac{1}{u(\mu) - u_{c,k}}\right)\right)$, since $u_{c,k} \leq \bar{U}(\gamma_k)$. \square

Remark 3.2.24. We note that the above lemma is used to obtain an asymptotic lower bound on $Q^*(u_c)$. As far as this asymptotic lower bound is concerned, the above lemma can be used even when if $i = P$, where i is such that $b_i = b_\mu$. We consider INTERVAL- λ CHOICE-2-1 for a larger \mathcal{X}_λ defined as follows. We extend \mathcal{X}_λ to $\bar{\mathcal{X}}_\lambda = \mathcal{X}_\lambda \cup (b_P, b_P + \delta]$, for some $\delta > 0$. We also extend the definition of $u(\cdot)$ to $\bar{\mathcal{X}}_\lambda$, by choosing a piecewise linear function on $(b_P, b_P + \delta]$ which preserves the strictly increasing concave property of $u(\cdot)$. We denote $Q^*(u_c)$ when $\lambda(q) \in \bar{\mathcal{X}}_\lambda$, by $Q_e^*(u_c)$. Then we note that $Q_e^*(u_c) \leq Q^*(u_c)$. The asymptotic lower bound for $Q_e^*(u_c)$ follows from the above lemma, which then also holds for $Q^*(u_c)$.

Remark 3.2.25. We note that the above asymptotic lower bound holds even in the case where $a_\mu = 0$. For INTERVAL- μ CHOICE-2, we note that the case with $a_\lambda = 0$ corresponds to the case INTERVAL- μ CHOICE-2-1, for which $Q^*(c_c)$ only increased to a finite value.

Remark 3.2.26. We note that in many cases, for queueing models with a single queue, the utility constraint is on the average throughput. Then we have that $u(\lambda)$ is a line segment, with $a_\mu = 0$ and $b_\mu = r_{a,max}$. We note that the asymptotic $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ lower bound holds for $Q^*(u_c)$, from the discussion in Remark 3.2.24.

Lemma 3.2.27. For INTERVAL- λ CHOICE-2-2, for any sequence of non-idling admissible policies γ_k such that $u(\mu) - \bar{U}(\gamma_k) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$.

Proof. The proof follows that of Lemma 3.2.9. We choose $\lambda^* = \mu + \epsilon_V = a_\mu + \epsilon_V$. Let $q_{\lambda^*} = \inf\{q : \lambda(q) \leq \lambda^*\}$. We note that q_{λ^*} could be 0, unlike q_{μ^*} in Lemma 3.2.9.

If q_{λ^*} is 0, then we note that $\pi(q_{\lambda^*}) = \pi(0) = \mathcal{O}(V)$ from Lemma 3.2.21. If $q_{\lambda^*} > 0$, then we have that

$$\begin{aligned} V &= \sum_{\lambda_k > \lambda^*} (l(\lambda_k) - u(\lambda_k)) \pi_{\lambda}(k), \\ &\geq m_a \epsilon_V Pr \{ \lambda(Q) > \lambda^* \}, \end{aligned}$$

where m_a is the tangent of the angle made by the line passing through $(a_i = a_{\mu}, u(a_{\mu}))$ and $(a_{i+1}, u(a_{i+1}))$ with $l(\mu)$. Or we have that $Pr \{ Q < q_{\lambda^*} \} \leq \frac{V}{m_a \epsilon_V}$ and $\pi(q_{\lambda^*} - 1) \leq \frac{V}{m_a \epsilon_V}$. We also note that since $\pi(q_{\lambda^*} - 1) \lambda(q_{\lambda^*} - 1) = \pi(q_{\lambda^*}) \mu$ we have that $\pi(q_{\lambda^*}) \leq \pi(q_{\lambda^*} - 1) \frac{r_{a, max}}{\mu}$. Therefore $\pi(q_{\lambda^*}) \leq \frac{V}{m_a \epsilon_V} \frac{r_{a, max}}{\mu}$. We choose $\epsilon_V = a_2 V$, so that $\frac{r_{a, max}}{m_a \mu a_2} \leq \frac{\alpha}{2}$, where $\alpha \ll 1$.

Then, as in the proof of Lemma 3.2.9, if \bar{q} is the largest integer such that

$$\sum_{q=q_{\lambda^*}}^{\bar{q}} \pi(q) \leq \frac{1}{2} - \frac{\alpha}{2},$$

then $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2}$, independently of whether $q_{\lambda^*} = 0$ or not.

We note that for any $q \geq q_{\lambda^*}$ we have that

$$\sum_{q=q_{\lambda^*}}^q \pi(q) \leq \pi(q_{\lambda^*}) \sum_{m=0}^{q-q_{\lambda^*}} \left(\frac{\lambda^*}{\mu} \right)^m. \quad (3.21)$$

We also note that if $q_{\lambda^*} > 0$, then we have that

$$\begin{aligned} V &\geq \sum_{\lambda_k > \lambda^*} m_a (\lambda_k - \mu) \pi_{\lambda}(k), \\ &= m_a \sum_{q < q_{\lambda^*}} (\lambda(q) - \mu) \pi(q), \\ &= m_a \mu \pi(q_{\lambda^*}) - \mu \pi(0). \end{aligned}$$

Then since $\pi(0) = \mathcal{O}(V)$ we have that $\pi(q_{\lambda^*}) = \mathcal{O}(V)$. Thus independently of whether $q_{\lambda^*} = 0$ or not, we have that $\pi(q_{\lambda^*}) = \mathcal{O}(V)$.

Now proceeding as in the proof of Lemma 3.2.9, using the above upper bound on $\pi(q_{\lambda^*})$ in (3.21) we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$. Now given a sequence of $u_{c,k} \uparrow u(\mu)$, we have that there exists a sequence of feasible γ_k such that $\bar{Q}(\gamma_k) \leq Q^*(u_{c,k}) + \epsilon$, for some $\epsilon > 0$. Therefore, $Q^*(u_{c,k}) = \Omega\left(\frac{1}{u(\mu) - u_{c,k}}\right)$, since $u_{c,k} \leq \bar{U}(\gamma_k)$. \square

We note that as for INTERVAL- μ CHOICE, using policies with similar structure as in Lemmas 3.2.11, 3.2.14, and 3.2.16 it is possible to obtain an asymptotic upper bound for INTERVAL- λ CHOICE-1

and tight asymptotic upper bounds for INTERVAL- λ CHOICE-2-1 and INTERVAL- λ CHOICE-2-2. Here, we obtain a single asymptotic upper bound for INTERVAL- λ CHOICE-1, INTERVAL- λ CHOICE-2-1, and INTERVAL- λ CHOICE-2-2, which shows that $u(\mu) = \sup_{\gamma \in \Gamma_{a,M}} \bar{U}(\gamma)$.

Lemma 3.2.28. There exists a sequence of admissible policies γ_k such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$ and $u(\mu) - \bar{U}(\gamma_k) = V_k$.

Proof. The proof follows that of Lemma 3.2.16 (which follows from that of Lemma 2.3.14). Consider a policy γ defined as follows:

$$\lambda(q) = \begin{cases} \mu & \text{for } q \in \{0, \dots, q_1\}, \\ \mu - K & \text{for } q \in \{q_1 + 1, \dots\}, \end{cases}$$

where $q_1 = \frac{1}{U}$ for positive U . The sequence of policies γ_k is obtained by choosing U from a sequence $U_k \downarrow 0$.

We note that

$$\begin{aligned} \bar{U}(\gamma) &= Pr\{Q \leq q_1\} u(\mu) + Pr\{Q > q_1\} u(\mu - K), \\ &= u(\mu) - Pr\{Q > q_1\} (u(\mu) - u(\mu - K)). \end{aligned}$$

Or $u(\mu) - \bar{U}(\gamma) = Pr\{Q > q_1\} (u(\mu) - u(\mu - K))$.

We also note that $\bar{Q}(\gamma) = \mathcal{O}(q_1) = \mathcal{O}\left(\frac{1}{U}\right)$. Therefore we have that $Pr\{Q \leq q_1\} \mu + Pr\{Q > q_1\} (\mu - K) = (1 - \pi(0))\mu$. Or

$$\begin{aligned} (1 - Pr\{Q > q_1\}) \mu + Pr\{Q > q_1\} (\mu - K) &= (1 - \pi(0))\mu, \\ \mu + Pr\{Q > q_1\} (-K) &= \mu - \mu\pi(0), \text{ or} \\ Pr\{Q > q_1\} &= \frac{\mu\pi(0)}{K}. \end{aligned}$$

Following the proof of Lemma 2.3.14, we have that $\pi(0) = \mathcal{O}(U)$. Hence, we have that $Pr\{Q > q_1\} = \mathcal{O}(U)$. Therefore, for the sequence of policies γ_k , we have that $u(\mu) - \bar{U}(\gamma_k) = \mathcal{O}(U_k)$. If $V_k \triangleq u(\mu) - \bar{U}(\gamma_k)$, then we have that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. \square

Remark 3.2.29. Since the techniques used in the analysis of INTERVAL- μ CHOICE and INTERVAL- λ CHOICE are similar, we expect that asymptotic bounds on any sequence of order-optimal policies can be obtained for INTERVAL- λ CHOICE as for INTERVAL- μ CHOICE. We note that the role of $\mu(q)$ and $\lambda(q)$ are interchanged. For example, using a similar sequence of steps as in the proof of Lemma 3.2.18, it is possible to show that for a sequence of non-idling order-optimal admissible policies with $u(\mu) - \bar{U}(\gamma_k) = V_k \downarrow 0$, $|Q_A| = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$ for INTERVAL- λ CHOICE-1, where $Q_A = \left\{q : \lambda(q) \in [\mu - a_2 V^{\frac{1}{2}}, \mu + a_2 V^{\frac{1}{2}}]\right\}$ and $a_2 > 0$.

3.3 Analysis of INTERVAL- $\lambda\mu$ CHOICE

We recall that for INTERVAL- $\lambda\mu$ CHOICE we restrict to policies γ such that $\lambda(q) \in [r_{a,min}, r_{a,max}]$ and $\mu(q) \in [0, r_{max}]$, $\forall q \in \mathbb{Z}_+$. The tradeoff problem for INTERVAL- $\lambda\mu$ CHOICE is:

$$\begin{aligned} & \text{minimize}_{\gamma \in \Gamma_a} && \bar{Q}(\gamma) \\ & \text{such that} && \bar{C}(\gamma) \leq c_c, \\ & \text{and} && \bar{U}(\gamma) \geq u_c, \end{aligned} \tag{3.22}$$

whose optimal value is $Q^*(c_c, u_c)$. Although it is possible to consider various forms of the function $c(\mu)$ as in the case of FINITE- μ CHOICE and INTERVAL- μ CHOICE, here we obtain a complete analysis for the case $c(\mu)$ being a strictly convex function of $\mu \in [0, r_{max}]$ (with assumption C2) and $u(\lambda)$ being either a strictly concave (with assumption U2) or a piecewise linear function of $\lambda \in [r_{a,min}, r_{a,max}]$. The reason for this assumption is the motivating discrete time problem, described below. We then comment on the asymptotic bounds for other forms of $c(\mu)$ in the following discussion.

Remark 3.3.1. INTERVAL- $\lambda\mu$ CHOICE corresponds to the tradeoff problem for the following discrete time queueing model. Work arrives in a batch of random size, in every slot, into an infinite buffer queue. The queue length is the amount of unfinished work and evolves on \mathbb{R}_+ . The amount of work which is admitted into the queue can be controlled, as a function, possibly randomized, of the current queue length. This feature is modelled by the control of the arrival rate, $\lambda(q)$, in INTERVAL- $\lambda\mu$ CHOICE. The amount of work done by the server in each slot, or the service batch size, can also be chosen as a function, possibly randomized, of the current queue length and is assumed to be a non-negative real number. This feature of the discrete time queue is modelled by the control of the service rate, $\mu(q)$, in INTERVAL- $\lambda\mu$ CHOICE. We note that the drift in the discrete time queueing model is real valued. In light of the discussion in Section 3.1 we assume that $\lambda(q)$ and $\mu(q)$ take any non-negative real values, but in finite intervals. In each slot, assume that there is a utility accrued with admitting customers and a service cost incurred in serving them. These are modelled by the utility rate and service cost rate functions $u(\cdot)$ and $c(\cdot)$ in INTERVAL- $\lambda\mu$ CHOICE. We choose $c(\mu)$ to be strictly convex as for INTERVAL- μ CHOICE-1. Another motivating factor for considering $c(\cdot)$ to be strictly convex, is the need to explain the logarithmic $\Theta(\log(\frac{1}{V}))$ behaviour of the average queue length when admission control is allowed, noticed by Neely in [44], compared to the $\Omega(\frac{1}{\sqrt{V}})$ behaviour of the average queue length, with strictly convex $c(\cdot)$, when admission control is not allowed. We note that in [44], since there was only a constraint on the average throughput, the function $u(\cdot)$ need only be linear, while the following result is presented for the case of strictly convex or piecewise linear $u(\cdot)$.

Remark 3.3.2. In this analysis, we assume that $u_c \leq u(r_{a,max})$. If $u_c > u(r_{a,max})$, then there does not exist any feasible policies for (3.22) We note that if $u_c = u(r_{a,max})$, then policies which

satisfy this utility constraint need to have $\lambda(q) = r_{a,max}, \forall q$, in which case the problem is the same as that considered in INTERVAL- μ CHOICE-1 (if $r_{a,max} < r_{max}$). We also note that the restriction of analysis to admissible policies implicitly requires that $u^{-1}(u_c) < r_{max}$.

Lemma 3.3.3. For INTERVAL- $\lambda\mu$ CHOICE, with $r_{a,min} > 0$, the service cost $\overline{C}(\gamma)$ for any admissible policy γ is bounded below by $c(u^{-1}(u_c))$.

Proof. To find a lower bound on the service cost for any admissible policy γ , we consider the following equivalent formulation of TRADEOFF (2.4) and use a series of relaxations on the constraints. The minimum average service cost for a given average queue length constraint q_c and a utility constraint u_c is given by

$$\begin{aligned} & \min_{\gamma \in \Gamma_a} \quad \mathbb{E}_\pi c(\mu(Q)), \\ \text{such that} \quad & \mathbb{E}_\pi u(\lambda(Q)) \geq u_c, \text{ and } \mathbb{E}_\pi Q \leq q_c. \end{aligned}$$

We note that for every $q_c < \infty$, we have that $\mathbb{E}_\pi \lambda(Q) = \mathbb{E}_\pi \mu(Q)$. So the optimal value of the optimization problem above is bounded below by the optimal value of

$$\begin{aligned} & \min_{\gamma \in \Gamma_a} \quad \mathbb{E}_\pi c(\mu(Q)), \\ \text{such that} \quad & \mathbb{E}_\pi u(\lambda(Q)) \geq u_c, \text{ and } \mathbb{E}_\pi \lambda(Q) = \mathbb{E}_\pi \mu(Q). \end{aligned}$$

Since $u(\lambda)$ is concave in λ , we have that for every γ such that $\mathbb{E}_\pi u(\lambda(Q)) \geq u_c$, $u(\mathbb{E}_\pi \lambda(Q)) \geq u_c$. Therefore the optimal value of the problem above is bounded below by the optimal value of

$$\begin{aligned} & \min_{\pi} \quad \mathbb{E}_\pi c(\mu(Q)), \\ \text{such that} \quad & \mathbb{E}_\pi \lambda(Q) \geq u^{-1}(u_c), \text{ and } \mathbb{E}_\pi \lambda(Q) = \mathbb{E}_\pi \mu(Q), \end{aligned}$$

where we are considering all possible distributions π for Q . Now since $c(\mu)$ is convex in μ , we obtain that the optimal value of the above problem is $\geq c(u^{-1}(u_c))$. Therefore $\overline{C}(\gamma) \geq c(u^{-1}(u_c))$. \square

In the following, as in the case of FINITE- μ CHOICE and INTERVAL- μ CHOICE, we consider INTERVAL- $\lambda\mu$ CHOICE in the asymptotic regime \mathfrak{R} where the service cost constraint c_c approaches $c(u^{-1}(u_c))$, where u_c is kept fixed.

3.3.1 Asymptotic lower bound

In this section we find an asymptotic lower bound on $\overline{Q}(\gamma_k)$ for any sequence of non-idling admissible policies γ_k for which $\overline{U}(\gamma_k) \geq u_c$ and $\overline{C}(\gamma_k) \downarrow c(u^{-1}(u_c))$. Subsequently, in Lemma 3.3.7 we show that there exists a sequence of non-idling admissible policies γ_k for which $\overline{C}(\gamma_k)$ approaches $c(u^{-1}(u_c))$ arbitrarily closely.

We note that, as in the case INTERVAL- μ CHOICE, even though the service rate $\mu(q)$ can take any value in $[0, r_{max}]$ and the arrival rate $\lambda(q)$ can take any value in $[r_{a,min}, r_{a,max}]$, the sets of service rates and arrival rates used by γ are both countable as the queue length takes only integer values. We first present the asymptotic lower bound for the case when $r_{a,min} > 0$, for which the proof is very similar to the case INTERVAL- μ CHOICE-1, and then consider the case $r_{a,min} = 0$.

Lemma 3.3.4. For INTERVAL- $\lambda\mu$ CHOICE, for any sequence of non-idling admissible policies γ_k such that $\bar{C}(\gamma_k) - c(u^{-1}(u_c)) = V_k \downarrow 0$ and $\bar{U}(\gamma_k) \geq u_c$, we have that

$$\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right).$$

Proof. Consider a particular policy γ in the sequence γ_k with $V_k = V$. Since γ is admissible, we have that $\mathbb{E}_\pi \mu(Q) = \mathbb{E}_\pi \lambda(Q)$. From the concavity of $u(\lambda)$, we have that $\mathbb{E}_\pi \lambda(Q) \geq u^{-1}(u_c)$. Let $\mu^* = u^{-1}(u_c) - \epsilon_V$, where ϵ_V is a function of V to be chosen later. Define $q_{\mu^*} = \inf\{q : \mu(q) \geq \mu^*\}$. We note that $\forall q < q_{\mu^*}, \mu(q) < \mu^*$. As $\mu(q)$ is non-decreasing, we have that

$$Pr\{Q < q_{\mu^*}\} = Pr\{\mu(Q) < \mu^*\}.$$

Let the countable set of service rates be denoted by $\{\mu_0 = 0, \mu_1, \dots\}$, where $\mu_i < \mu_{i+1}$ and $\mu_i \in [0, r_{max}]$. Let $l(\mu)$ be the tangent line at $(u^{-1}(u_c), c(u^{-1}(u_c)))$ to the curve $c(\mu)$. Then $V = \sum_{q=0}^{\infty} [c(\mu(q)) - l(\mu(q))] \pi(q)$. From Proposition 3.A.1, we have a positive a_1 such that

$$V \geq a_1 \sum_{q=0}^{\infty} [\mu(q) - u^{-1}(u_c)]^2 \pi(q) \geq a_1 \sum_{q=0}^{q_{\mu^*}-1} [\mu(q) - u^{-1}(u_c)]^2 \pi(q).$$

Hence

$$Pr\{Q \leq q_{\mu^*} - 1\} \leq \frac{V}{a_1 \epsilon_V^2},$$

$$\text{and } \pi(q_{\mu^*} - 1) \leq \frac{V}{a_1 \epsilon_V^2}.$$

Now, since $Q(t)$ is a birth death process $\forall q$, we have that $\pi(q)\lambda(q) = \pi(q+1)\mu(q+1)$. For any $q \geq q_{\mu^*}$,

$$\pi(q+1) = \frac{\pi(q)\lambda(q)}{\mu(q+1)} \leq \frac{\pi(q)r_{a,max}}{\mu^*},$$

$$\pi(q) \leq \pi(q_{\mu^*} - 1) \left(\frac{r_{a,max}}{\mu^*}\right)^{q - q_{\mu^*} + 1}. \quad (3.23)$$

Let \bar{q} be the largest integer such that $\sum_{q=0}^{\bar{q}} \pi(q) \leq \frac{1}{2}$. We find a lower bound on \bar{q} as in the proof

of Lemma 2.3.6. We note that $Pr\{Q \geq q_{\mu^*}\} \geq 1 - \frac{V}{a_1 \epsilon_V^2}$. Let $\epsilon_V = \epsilon$, where $0 < \epsilon < u^{-1}(u_c)$. For V small, let \bar{q}_1 be the largest integer such that

$$\sum_{q=q_{\mu^*}}^{\bar{q}_1} \pi(q) \leq \frac{1}{2} - \frac{V}{a_1 \epsilon_V^2}.$$

Then $\bar{q}_1 \leq \bar{q}$. We find a lower bound on \bar{q}_1 by using the upper bound on $\pi(q)$ from (3.23). Let \bar{q}_2 be the largest integer such that

$$\pi(q_{\mu^*} - 1) \sum_{q=1}^{\bar{q}_2 - q_{\mu^*} + 1} \left(\frac{r_{a,max}}{\mu^*} \right)^q \leq \frac{1}{2} - \frac{V}{a_1 \epsilon_V^2}.$$

Then $\bar{q}_2 \leq \bar{q}_1$. After substituting for μ^* , we have that any \bar{q}_2 satisfying the above inequality is such that

$$\bar{q}_2 - q_{\mu^*} + 1 \leq \log \left(\frac{r_{a,max}}{u^{-1}(u_c) - \epsilon_V} \right) \left(1 + \frac{r_{a,max} - u^{-1}(u_c) + \epsilon_V}{r_{a,max}} \frac{1}{\pi(q_{\mu^*} - 1)} \left(\frac{1}{2} - \frac{V}{a_1 \epsilon_V^2} \right) \right).$$

Hence we obtain that \bar{q}_2 is at least

$$\log \left(\frac{r_{a,max}}{u^{-1}(u_c) - \epsilon_V} \right) \left(1 + \frac{r_{a,max} - u^{-1}(u_c) + \epsilon_V}{r_{a,max}} \frac{1}{\pi(q_{\mu^*} - 1)} \left(\frac{1}{2} - \frac{V}{a_1 \epsilon_V^2} \right) \right) - 2.$$

We note that $\frac{1}{\pi(q_{\mu^*} - 1)} \geq \frac{a_1 \epsilon^2}{V}$ and is the dominant term in the regime where $V \downarrow 0$. Since $\bar{q} \geq \bar{q}_1 \geq \bar{q}_2$ and $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2}$, we have that for any sequence of γ_k with $\bar{C}(\gamma_k) - u^{-1}(u_c) = V_k \downarrow 0$, $\bar{Q}(\gamma_k) = \Omega \left(\log \left(\frac{1}{V_k} \right) \right)$. \square

Remark 3.3.5. We note that in this proof, no use was made of the assumption that the sequence of policies satisfies the constraint $\mathbb{E}_\pi u(\lambda(Q)) \geq u_c$. The difficulty in problem INTERVAL- $\lambda\mu$ CHOICE, is to actually construct such a sequence of policies.

In the above proof, we note that there exists a set \mathcal{Q}_h of queue lengths occurring with high probability, such that $\mu(q) \rightarrow u^{-1}(u_c)$, for every $q \in \mathcal{Q}_h$ as $V \downarrow 0$. But for every $q \in \mathcal{Q}_h$, it is possible to ensure through arrival rate control that the service rate $\mu(q)$ is not equal to the arrival rate $\lambda(q)$, while for Lemma 3.2.2 $\mu(q) = \lambda$ for all queue lengths with high probability. Hence, for admissible policies, for $q \in \mathcal{Q}_h$, the stationary probability distribution geometrically grows and then decays which leads to the $\log \left(\frac{1}{V} \right)$ behaviour.

Remark 3.3.6. *The case when $r_{a,min} = 0$ or $\lambda(q) \in [0, r_{a,max}]$:*

We note that if $r_{a,min} = 0$, then the birth death process may not be irreducible on \mathbb{Z}_+ . Therefore, in this case, admissible policies are assumed to induce a single positive recurrent class including zero. In Lemma 3.3.4 we had assumed that $r_{a,min} > 0$. If $r_{a,min} = 0$, then for an admissible policy γ , there could exist q such that $\lambda(q) = 0$. Let $q' = \inf \{q : \lambda(q) = 0\}$. Note that q' is in general

dependent on the parameters c_c and u_c . If q' is infinity, then the same approach as in Lemma 3.3.4 holds. If q' is finite, then as γ is admissible, $\forall q > q'$, $\lambda(q) = 0$. Hence all states $q > q'$ are transient and in steady state we need only consider the CTMC evolving on $\{0, \dots, q'\}$. Note that $\sum_{q=0}^{q'} \pi(q) = 1$ and therefore \bar{q} in Lemma 3.3.4 is smaller than q' . Therefore, the same approach as in Lemma 3.3.4 holds even if q' is finite.

3.3.2 Asymptotic behaviour of the tradeoff curve

In this section, we construct a sequence of admissible policies γ_k which achieves the minimum average service cost $c(u^{-1}(u_c))$ arbitrarily closely with $\bar{Q}(\gamma_k)$ scaling at the optimal rate as in Lemma 3.3.4. However, we are able to obtain an asymptotic upper bound only for the case where $u(\lambda)$ is strictly concave or linear (and not piecewise linear).

Lemma 3.3.7. For INTERVAL- $\lambda\mu$ CHOICE, with $u(\lambda)$ strictly concave or linear, there exists a sequence of admissible policies γ_k with a corresponding sequence $V_k \downarrow 0$ such that

$$\begin{aligned}\bar{Q}(\gamma_k) &= \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right), \\ \bar{C}(\gamma_k) - c(u^{-1}(u_c)) &= V_k, \\ \bar{U}(\gamma_k) &\geq u_c.\end{aligned}$$

The construction of the sequence of admissible policies γ_k is motivated by the following intuition, that we have obtained from the lower bound in Lemma 3.3.4. The sequence of policies should be such that as $V_k \downarrow 0$, the service rate used, at a queue length occurring with high probability, should be close to $u^{-1}(u_c)$. But the arrival rate $\lambda(q)$ should not exactly equal $u^{-1}(u_c)$, for all queue lengths q which occur with high probability. Then it should be possible to have a stationary distribution which is geometrically growing and then decaying, leading to the required $\log\left(\frac{1}{V_k}\right)$ scaling of $\bar{Q}(\gamma_k)$.

Proof. Consider a policy γ of the following form :

$$\begin{aligned}\mu(0) &= 0, \\ \mu(q) &= \mu_1 = u^{-1}(u_c) - \epsilon_U, \text{ for } q \in \{1, \dots, q_1\}, \\ \mu(q) &= \mu_2 = u^{-1}(u_c) + \epsilon_U, \text{ for } q \in \{q_1 + 1, \dots\}; \\ \text{and } \lambda(q) &= \lambda_1, \text{ for } q \in \{0, \dots, q_1 - 1\}, \\ \lambda(q) &= u^{-1}(u_c), \text{ for } q \in \{q_1, \dots, q_1 + K\}, \\ \lambda(q) &= \lambda_2, \text{ for } q \in \{q_1 + K + 1, \dots\}.\end{aligned}$$

Let $\epsilon_U = U$, $\lambda_1 > u^{-1}(u_c) > \lambda_2$, $\lambda_1 > \mu_1$, $\lambda_2 < \mu_2$, and $q_1 = \left\lceil \log\left(\frac{\lambda_1}{\mu_1}\right) \left(1 + \frac{\lambda_1 - \mu_1}{\lambda_1} \frac{1}{U}\right) \right\rceil$. We will specify K , λ_1 , and λ_2 later. Let $\frac{dc(u^{-1}(u_c))}{d\mu} \triangleq \frac{dc(\mu)}{d\mu}|_{\mu=u^{-1}(u_c)}$. We now obtain $\bar{C}(\gamma)$.

$$\begin{aligned}
\bar{C}(\gamma) &= \pi(0) \cdot 0 + \pi_\mu(\mu_1)c(\mu_1) + \pi_\mu(\mu_2)c(\mu_2), \\
&= \pi_\mu(\mu_1) \left(c(u^{-1}(u_c)) + (-\epsilon_U) \frac{dc(u^{-1}(u_c))}{d\mu} + \mathcal{O}(\epsilon_U^2) \right) + \\
&\quad \pi_\mu(\mu_2) \left(c(u^{-1}(u_c)) + (\epsilon_U) \frac{dc(u^{-1}(u_c))}{d\mu} + \mathcal{O}(\epsilon_U^2) \right), \\
&\leq c(u^{-1}(u_c)) + \mathcal{O}(U^2) + (-\epsilon_U \pi_\mu(\mu_1) + \epsilon_U \pi_\mu(\mu_2)) \frac{dc(u^{-1}(u_c))}{d\mu}, \\
&\leq c(u^{-1}(u_c)) + \epsilon_U \frac{dc(u^{-1}(u_c))}{d\mu} + \mathcal{O}(U^2), \\
&\leq c(u^{-1}(u_c)) + \mathcal{O}(U),
\end{aligned}$$

where $\frac{dc(u^{-1}(u_c))}{d\mu} \triangleq \frac{dc(\mu)}{d\mu}|_{\mu=u^{-1}(u_c)}$. Let $V = \bar{C}(\gamma) - c(u^{-1}(u_c))$, then we have that $V = \mathcal{O}(U)$. For γ , the average utility is

$$\bar{U}(\gamma) = u(\lambda_1) \left(\sum_{q=0}^{q_1-1} \pi(q) \right) + u_c \sum_{q=q_1}^{q_1+K} \pi(q) + u(\lambda_2) \sum_{q=q_1+K+1}^{\infty} \pi(q).$$

Let $\lambda_1 = u^{-1}(u_c) + \epsilon$ and $\lambda_2 = u^{-1}(u_c) - \epsilon$, where ϵ is a small positive constant. Then for strictly concave and linear $u(\cdot)$ we have that

$$\bar{U}(\gamma) \geq u_c + \left[\epsilon \sum_{q=0}^{q_1-1} \pi(q) - \epsilon \sum_{q=q_1+K+1}^{\infty} \pi(q) \right] D(u(u^{-1}(u_c))) + \mathcal{O}(u(u^{-1}(u_c))),$$

where $D(u(u^{-1}(u_c)))$ and $\mathcal{O}(u(u^{-1}(u_c)))$ are defined as follows. If $u(\cdot)$ is a strictly concave function, then it is differentiable at $u^{-1}(u_c)$ and $D(u(u^{-1}(u_c))) \triangleq \frac{du(\lambda)}{d\lambda}|_{\lambda=u^{-1}(u_c)}$ and $\mathcal{O}(u(u^{-1}(u_c))) = \mathcal{O}(\epsilon^2)$, with the above inequality being an equality. If $u(\cdot)$ is linear then $D(u(u^{-1}(u_c))) \triangleq \frac{du(\lambda)}{d\lambda}|_{\lambda=u^{-1}(u_c)}$ and $\mathcal{O}(u(u^{-1}(u_c))) = 0$, with the above inequality being an equality.

In the following we show that $\sum_{q=q_1+K+1}^{\infty} \pi(q) \leq \sum_{q=0}^{q_1-1} \pi(q)$, in which case we have that $\bar{U}(\gamma) \geq u_c$ for sufficiently small ϵ (which is fixed and independent of V). We have that

$$\begin{aligned}
\pi(q) &= \pi(0) \left(\frac{\lambda_1}{\mu_1} \right)^q, & \text{for } q \in \{1, \dots, q_1\}, \\
\pi(q) &= \pi(0) \left(\frac{\lambda_1}{\mu_1} \right)^{q_1} \left(\frac{u^{-1}(u_c)}{\mu_2} \right)^{q-q_1}, & \text{for } q \in \{q_1+1, \dots, q_1+K\}, \\
\pi(q) &= \pi(0) \left(\frac{\lambda_1}{\mu_1} \right)^{q_1} \left(\frac{u^{-1}(u_c)}{\mu_2} \right)^K \left(\frac{u^{-1}(u_c)}{\mu_2} \right) \left(\frac{\lambda_2}{\mu_2} \right)^{q-q_1-K-1}, & \text{for } q \in \{q_1+K+1, \dots\}.
\end{aligned}$$

Therefore,

$$\begin{aligned}\sum_{q=0}^{q_1-1} \pi(q) &= \pi(0) + \pi(0) \sum_{q=1}^{q_1-1} \left(\frac{\lambda_1}{\mu_1}\right)^q, \\ &= \pi(0) + \pi(0) \left(\frac{\lambda_1}{\lambda_1 - \mu_1}\right) \left(\left(\frac{\lambda_1}{\mu_1}\right)^{q_1-1} - 1\right).\end{aligned}$$

And,

$$\begin{aligned}\sum_{q_1+K+1}^{\infty} \pi(q) &= \pi(0) \left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{u^{-1}(u_c)}{\mu_2}\right)^K \left(\frac{u^{-1}(u_c)}{\mu_2}\right) \sum_{q=0}^{\infty} \left(\frac{\lambda_2}{\mu_2}\right)^q, \\ &= \pi(0) \left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{u^{-1}(u_c)}{\mu_2}\right)^K \left(\frac{u^{-1}(u_c)}{\mu_2 - \lambda_2}\right).\end{aligned}$$

If

$$\left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{u^{-1}(u_c)}{\mu_2}\right)^K \left(\frac{u^{-1}(u_c)}{\mu_2 - \lambda_2}\right) \leq \left(\frac{\lambda_1}{\lambda_1 - \mu_1}\right) \left(\left(\frac{\lambda_1}{\mu_1}\right)^{q_1-1} - 1\right) + 1 \quad (3.24)$$

then $\sum_{q=q_1+K+1}^{\infty} \pi(q) \leq \sum_{q=0}^{q_1-1} \pi(q)$. We note that (3.24) can be simplified to the question

$$\begin{aligned}\frac{\mu_1}{\lambda_1 - \mu_1} &\stackrel{?}{\leq} \left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{\mu_1}{\lambda_1 - \mu_1} - \left(\frac{u^{-1}(u_c)}{\mu_2}\right)^K \left(\frac{u^{-1}(u_c)}{\mu_2 - \lambda_2}\right)\right), \\ \frac{\mu_1}{\lambda_1 - \mu_1} &\stackrel{?}{\leq} \left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{u^{-1}(u_c) - \epsilon_U}{\epsilon + \epsilon_U} - \left(\frac{u^{-1}(u_c)}{u^{-1}(u_c) + \epsilon_U}\right)^K \left(\frac{u^{-1}(u_c)}{\epsilon + \epsilon_U}\right)\right), \\ \frac{\mu_1}{\lambda_1 - \mu_1} &\stackrel{?}{\leq} \left(\frac{\lambda_1}{\mu_1}\right)^{q_1} \left(\frac{u^{-1}(u_c)}{\epsilon + \epsilon_U}\right) \left(1 - \frac{\epsilon_U}{u^{-1}(u_c)} - \left(1 + \frac{\epsilon_U}{u^{-1}(u_c)}\right)^{-K}\right).\end{aligned}$$

We use the lower bound on q_1 , obtained by removing the ceiling, to arrive at the following question

:

$$\frac{u^{-1}(u_c) - \epsilon_U}{\epsilon + \epsilon_U} \stackrel{?}{\leq} \left(1 + \frac{\epsilon + \epsilon_U}{u^{-1}(u_c) - \epsilon_U} \frac{1}{U}\right) \left(\frac{u^{-1}(u_c)}{\epsilon + \epsilon_U}\right) \left(1 - \frac{\epsilon_U}{u^{-1}(u_c)} - \left(1 + \frac{\epsilon_U}{u^{-1}(u_c)}\right)^{-K}\right)$$

For sufficiently small U , with $\epsilon_U = U$, we have that $\left(1 + \frac{\epsilon_U}{u^{-1}(u_c)}\right)^{-K} \leq 1 - \frac{K\epsilon_U}{2u^{-1}(u_c)}$. So, instead of the above question we can ask the stronger question

$$\frac{u^{-1}(u_c) - \epsilon_U}{\epsilon + \epsilon_U} \stackrel{?}{\leq} \left(1 + \frac{\epsilon + \epsilon_U}{u^{-1}(u_c) - \epsilon_U} \frac{1}{V}\right) \left(\frac{u^{-1}(u_c)}{\epsilon + \epsilon_U}\right) \left(\frac{\epsilon_U}{u^{-1}(u_c)} \left(\frac{K}{2} - 1\right)\right).$$

We choose $K > 2 \left(1 + \frac{(u^{-1}(u_c))^2}{\epsilon}\right)$. Then we can ask the even stronger questions

$$\begin{aligned} u^{-1}(u_c) - \epsilon_U &\stackrel{?}{\leq} \left(1 + \frac{\epsilon + \epsilon_U}{u^{-1}(u_c) - \epsilon_U} \frac{1}{V}\right) \left(\epsilon_U \frac{(u^{-1}(u_c))^2}{\epsilon}\right), \\ u^{-1}(u_c) - \epsilon_U &\stackrel{?}{\leq} \frac{\epsilon + \epsilon_U}{\epsilon} \frac{(u^{-1}(u_c))^2}{u^{-1}(u_c) - \epsilon_U} \end{aligned}$$

which indeed hold. Hence for sufficiently small V and ϵ , $\bar{U}(\gamma) \geq u_c$.

We now obtain $\bar{Q}(\gamma)$ using Proposition 2.A.1 with $q_\epsilon = q_1 + 1$. Hence,

$$\begin{aligned} \bar{Q}(\gamma) &= \frac{(q_1 + 1)(\epsilon_U + \epsilon + r_{a,max})}{\epsilon_U + \epsilon} + \frac{r_{max}}{2(\epsilon_U + \epsilon)}, \\ \bar{Q}(\gamma) &\leq \frac{(q_1 + 1)(\epsilon + r_{a,max})}{\epsilon} + \frac{r_{max}}{2\epsilon}. \end{aligned}$$

As $q_1 = \mathcal{O}(\log(\frac{1}{V}))$, we obtain that $\bar{Q}(\gamma) = \mathcal{O}(\log(\frac{1}{V}))$. We note that the policy γ is admissible. The sequence of policies is obtained by choosing $U_k = \frac{1}{k}$. We note that then we have a corresponding sequence $V_k = \mathcal{O}(U_k)$. Thus, $\bar{Q}(\gamma_k) = \mathcal{O}(\log(\frac{1}{V_k}))$, and we have that there exists a sequence of admissible policies γ_k with a corresponding sequence $V_k \downarrow 0$ such that

$$\begin{aligned} \bar{Q}(\gamma_k) &= \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right), \\ \bar{C}(\gamma_k) - c(u^{-1}(u_c)) &= V_k, \\ \bar{U}(\gamma_k) &\geq u_c. \end{aligned}$$

□

Remark 3.3.8. We note that the above proof also applies if $u(\lambda)$ is piecewise linear and $(u^{-1}(u_c), u_c)$ lies on a linear segment of the piecewise linear function $u(\lambda)$. However, the proof does not apply if $u(\lambda)$ is piecewise linear and $u^{-1}(u_c)$ is such that the slope of $u(\lambda)$ changes at $(u^{-1}(u_c), u_c)$.

Using the asymptotic lower bound from Lemma 3.3.4, the asymptotic upper bound above, and proceeding as in the proof of Proposition 2.3.9, we arrive at the following result.

Proposition 3.3.9. For INTERVAL- $\lambda\mu$ CHOICE, for strictly concave or linear $u(\lambda)$, we have that the optimal tradeoff curve $Q^*(c_{c,k}, u_c) = \Theta\left(\log\left(\frac{1}{c_{c,k} - c(u^{-1}(u_c))}\right)\right)$, for the sequence $c_{c,k} = \bar{C}(\gamma_k)$, for the sequence of policies γ_k in Lemma 3.3.7.

Remark 3.3.10. For INTERVAL- $\lambda\mu$ CHOICE, an admissible policy γ can be specified by the sets $Q_{\mu,\lambda} = \{q : \mu(q) = \mu, \lambda(q) = \lambda\}$ for all possible μ and λ . However, we are only able to obtain bounds on sets of the form $Q_A = \{q : \mu(q) \in A \subseteq [0, r_{max}]\}$ for an asymptotic characterization of a sequence of order-optimal admissible policies. These bounds can be derived using similar techniques as in Section 3.2.3. For example, if $r_{a,min} > 0$, then it can be shown that if $A \subseteq$

$[0, u^{-1}(u_c) - a_2 V^{\frac{1-\delta}{2}}] \cup [u^{-1}(u_c) + a_2 V^{\frac{1-\delta}{2}}, \infty)$, for $0 < \delta \leq 1$ and $a_2 > 0$, then $|Q_{Ac}| = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$, for a sequence of non-idling order-optimal admissible policies with $\bar{C}(\gamma_k) - c(u^{-1}(u_c)) = V_k \downarrow 0$.

3.3.3 Tradeoff problems which are similar to INTERVAL- $\lambda\mu$ CHOICE

In this section, we first consider other asymptotic regimes for (2.4), which are similar to INTERVAL- $\lambda\mu$ CHOICE. We note that for INTERVAL- $\lambda\mu$ CHOICE, the utility constraint u_c was kept fixed while $c_{c,k} \downarrow c(u^{-1}(u_c))$. A similar problem (SP1) is one in which c_c is fixed and $u_{c,k} \uparrow u(c^{-1}(c_c))$. Another problem scenario (SP2) is one in which both $c_{c,k}$ and $u_{c,k}$ vary such that (a) $c_{c,k} - c(u^{-1}(u_{c,k})) \downarrow 0$ or (b) $u(c^{-1}(c_{c,k})) - u_{c,k} \downarrow 0$.

We note that SP2(b) encompasses SP1 since the sequence $u_{c,k}$ can be chosen such that $u_{c,k} = u_c, \forall k \in \mathbb{Z}_+$. We now show that the asymptotic regime for SP2(b) is equivalent to that for SP2(a), i.e., $c_{c,k} - c(u^{-1}(u_{c,k})) \downarrow 0$. We note that for any $(u_{c,k})$ and $(c_{c,k})$, for which the problem (3.22) is feasible, and also such that $u(c^{-1}(c_{c,k})) - u_{c,k} \downarrow 0$, we have that $\forall \epsilon > 0, \exists K_\epsilon$ such that, $\forall k > K_\epsilon$, $u(c^{-1}(c_{c,k})) - \epsilon \leq u_{c,k} \leq u(c^{-1}(c_{c,k}))$ (since $u_{c,k} \leq u(c^{-1}(c_{c,k}))$ if the problem (3.22) is feasible). Then we have that $u^{-1}(u(c^{-1}(c_{c,k})) - \epsilon) \leq u^{-1}(u_{c,k}) \leq c^{-1}(c_{c,k})$. For every $c_{c,k}$, we define $l_{1,k}(\lambda)$ (as in Section 3.2.4) to be (i) the tangent to $u(\lambda)$ at $(c^{-1}(c_{c,k}), u(c^{-1}(c_{c,k})))$, if $u(\lambda)$ is strictly convex, (ii) the line passing through $(a_\mu, u(a_\mu))$ and $(b_\mu, u(b_\mu))$, if $u(\lambda)$ is piecewise linear and $c^{-1}(c_{c,k})$ lies on a linear segment, and (iii) any line through $(a_\mu, u(a_\mu))$ with slope m , such that $\frac{du(\lambda)}{d\lambda}^-|_{\lambda=\mu} < m < \frac{du(\lambda)}{d\lambda}^+|_{\lambda=\mu}$, if $u(\lambda)$ is piecewise linear and $c^{-1}(c_{c,k})$ is a corner point of $u(\lambda)$. We note that $l_{1,k}(c^{-1}(c_{c,k})) = u(c^{-1}(c_{c,k}))$ in all three cases. Then

$$u^{-1}(u(c^{-1}(c_{c,k})) - \epsilon) \geq l_{1,k}^{-1}(u(c^{-1}(c_{c,k})) - \epsilon) = l_{1,k}^{-1}(u(c^{-1}(c_{c,k}))) - m_{1,k}\epsilon,$$

where $m_{1,k}$ is the slope of $l_{1,k}$. Since $l_{1,k}^{-1}(u(c^{-1}(c_{c,k}))) = c^{-1}(c_{c,k})$, we have that

$$\begin{aligned} c^{-1}(c_{c,k}) - m_{1,k}\epsilon &\leq u^{-1}(u_{c,k}) \leq c^{-1}(c_{c,k}), \\ c(c^{-1}(c_{c,k}) - m_{1,k}\epsilon) &\leq c(u^{-1}(u_{c,k})) \leq c_{c,k}. \end{aligned}$$

Let $l_{2,k}(\mu)$ be the tangent to $c(\mu)$ at $(c^{-1}(c_{c,k}), c_{c,k})$. Then we have that

$$\begin{aligned} l_{2,k}(c^{-1}(c_{c,k}) - m_{1,k}\epsilon) &\leq c(u^{-1}(u_{c,k})) \leq c_{c,k}, \\ c_{c,k} - m_{2,k}m_{1,k}\epsilon &\leq c(u^{-1}(u_{c,k})) \leq c_{c,k}, \end{aligned}$$

where $m_{2,k}$ is the slope of $l_{2,k}$. We note that $\exists M_1, M_2 \in \mathbb{R}_+$ such that $m_{1,k} \leq M_1$ and $m_{2,k} \leq M_2$ for every k , since both $u(\lambda)$ and $c(\mu)$ are defined on bounded domains. Since the above inequality holds for every $\epsilon > 0$ and $k > K_\epsilon$, we have that $c_{c,k} - c(u^{-1}(u_{c,k})) \downarrow 0$.

We now have the following result, under the stronger assumption that $u(\cdot)$ is m -strongly concave,

with $m > 0$. The proof is similar to that of Lemma 3.3.4.

Lemma 3.3.11. For INTERVAL- $\lambda\mu$ CHOICE, for any sequence of non-idling admissible policies γ_k and a sequence $u_{c,k} > 0$ such that $\bar{C}(\gamma_k) - c(u^{-1}(u_{c,k})) = V_k \downarrow 0$ and $\bar{U}(\gamma_k) \geq u_{c,k}$, we have that

$$\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right).$$

Proof. The proof follows that of Lemma 3.3.4 closely. Hence, we only state the differences here. We define $\mu^* = u^{-1}(u_{c,k}) - \epsilon_V$ and $q_{\mu^*} = \inf\{q : \mu(q) \geq \mu^*\}$. We note that unlike in the proof of Lemma 3.3.4, we define a different tangent line $l_k(\mu)$ for every $u_{c,k}$. Let $l_k(\mu)$ be the tangent line to $c(\mu)$ at $(u^{-1}(u_{c,k}), c(u^{-1}(u_{c,k})))$. From Proposition 3.A.1, we have a positive $a_{1,k}$ such that

$$V_k \geq a_{1,k} \sum_{q=0}^{q_{\mu^*}-1} [\mu(q) - u^{-1}(u_{c,k})]^2 \pi(q).$$

We note that unlike the proof of Lemma 3.3.4, here $a_{1,k}$ depends on the sequence $u_{c,k}$. Let $a \triangleq \inf_k \{a_{1,k}\}$. Since $u(\cdot)$ is m -strongly convex, we have that $a \geq m > 0$. Then we have that

$$\begin{aligned} \Pr\{Q \leq q_{\mu^*} - 1\} &\leq \frac{V}{a\epsilon_V^2}, \\ \text{and } \pi(q_{\mu^*} - 1) &\leq \frac{V}{a\epsilon_V^2}. \end{aligned}$$

Then, we proceed as in the proof of Lemma 3.3.4 to obtain that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

We note that an asymptotic upper bound can be obtained by evaluating $\bar{Q}(\gamma_k)$, $\bar{C}(\gamma_k)$, and $\bar{U}(\gamma_k)$ for a sequence of policies γ_k as in Lemma 3.3.7, but with u_c now being the sequence $u_{c,k}$. Then we have the following result

Proposition 3.3.12. For INTERVAL- $\lambda\mu$ CHOICE, for strongly concave or linear $u(\lambda)$, we have that the optimal tradeoff curve $Q^*(c_{c,k}, u_{c,k}) = \Theta\left(\log\left(\frac{1}{c_{c,k} - c(u^{-1}(u_{c,k}))}\right)\right)$, for the sequence $c_{c,k} = \bar{C}(\gamma_k)$ and $u_{c,k} = \bar{U}(\gamma_k)$, for the sequence of policies γ_k as above.

Remark 3.3.13. We note that for INTERVAL- $\lambda\mu$ CHOICE we have considered the case where $c(\mu)$ is strictly convex and $u(\lambda)$ is either strictly concave or linear (also piecewise linear for the asymptotic lower bound in Lemma 3.3.4). Although we have not presented the analysis for other forms of $c(\mu)$, such as when $c(\mu)$ is piecewise linear, here we outline how the methods presented in Chapter 2 as well as this chapter can be used in obtaining asymptotic lower bounds in these cases, in the asymptotic regime where $c_{c,k} \downarrow c(u^{-1}(u_c))$. Suppose $c(\mu)$ is piecewise linear and u_c is fixed. We note that as in INTERVAL- μ CHOICE-2, we can define service rates μ_l and μ_u with respect to $u^{-1}(u_c)$ rather than λ . Then the asymptotic behaviour of $Q^*(c_c, u_c)$ depends upon whether (i) $\mu_l < u^{-1}(u_c) < \mu_u$ and $\mu_l = 0$ or (ii) otherwise. For case (i), we note that $Q^*(c_c, u_c)$ only

increases to a finite value, since we can fix $\lambda(q) = u^{-1}(u_c)$ and apply the analysis of INTERVAL- μ CHOICE-2-1. However, we do not have an asymptotic lower bound in this case. For case (ii), we can proceed as in the proof of Lemma 3.3.4, except that $\mu^* \triangleq \mu_l - \epsilon$, where $\epsilon > 0$, to obtain that $Q^*(c_c, u_c)$ is $\Omega(\log(\frac{1}{V}))$.

3.4 Conclusions

In this chapter, we have considered the asymptotic characterization of the tradeoff problem for the state dependent M/M/1 model, with model features chosen such that insights can be obtained for the tradeoff problem for discrete time queueing models also. From the analysis, we see that for INTERVAL- μ CHOICE and INTERVAL- $\lambda\mu$ CHOICE the constraint on the average service cost, in the regime \mathfrak{R} , leads to constraints on the stationary probability distribution of the service rate $\mu(Q)$ and therefore the stationary probability distribution of the queue length which in turn determines the asymptotic growth of the minimum average queue length as a function of the average service cost. The exact nature of these constraints and the behaviour of the stationary distribution of the queue length depends on the nature of $c(\mu)$ at $\mu = \lambda$ and the extent of freedom in the choice of $\lambda(q)$ and $\mu(q)$ at a queue length q .

We observe that if $c(\mu)$ is strictly convex at λ , then in the asymptotic regime \mathfrak{R} as $V = c_c - c(\lambda) \downarrow 0$, the stationary probability of any queue length q such that $\mu(q) \neq \lambda$ goes to zero. More precisely, as $V \downarrow 0$, $\mu(q)$ for $q \in \mathcal{Q}_h$ (the set of queue lengths with high probability) has to approach λ , the stationary probabilities for such queue lengths are equal and each is $\mathcal{O}(\sqrt{V})$. This leads to the $\frac{1}{\sqrt{V}}$ behaviour for strictly convex $c(\mu)$. We expect that this is the same phenomenon which gives rise to the Berry-Gallager asymptotic lower bound [7], but for admissible policies. In Chapters 4 and 5, we see that this particular behaviour of the stationary probability does carry over to the discrete time model. Suppose, however that it is possible to control the arrival rates $\lambda(q)$, as in the case of INTERVAL- $\lambda\mu$ CHOICE. Then even though $\mu(q)$ for queue lengths $q \in \mathcal{Q}_h$ have to approach a single value $u^{-1}(u_c)$, the drift for such q need not be zero, since $\lambda(q)$ can be chosen to be different from $u^{-1}(u_c)$ for such q . In fact it is possible to choose $\lambda(q)$ such that the drift is initially positive and then negative for the set of queue lengths occurring with high probability so that the stationary probability of the queue length has a geometrically increasing and decaying behaviour which leads to the $\log(\frac{1}{V})$ asymptotic growth for the minimum average queue length. We expect that this is the reason for the $\mathcal{O}(\log(\frac{1}{V}))$ behaviour observed by Neely [44], but for admissible policies. In Chapter 5, we shall obtain an asymptotic lower bound for the discrete time model considered in [44] using the above idea, for admissible policies.

When $c(\mu)$ is piecewise linear in μ , we note that either $(\lambda, c(\lambda))$ can be a corner point as in

INTERVAL- μ CHOICE-2-3 or $(\lambda, c(\lambda))$ can lie on a linear portion of $c(\mu)$ as in INTERVAL- μ CHOICE-2-1 or INTERVAL- μ CHOICE-2-2. For INTERVAL- μ CHOICE-2-3, as in the case of INTERVAL- μ CHOICE-1, we observe that as $V \downarrow 0$, $\mu(q)$ for queue lengths in \mathcal{Q}_h has to approach λ , we again observe that the drift (proportional to $\lambda - \mu(q)$) for such queue lengths approaches zero, and the stationary probabilities for such queue lengths are equal and each is $\mathcal{O}(V)$. This leads to the $\frac{1}{V}$ behaviour for INTERVAL- μ CHOICE-2-3.

For INTERVAL- μ CHOICE-2-1 and INTERVAL- μ CHOICE-2-2, we note that as $V \downarrow 0$ service rates μ such that $\mu_l \leq \mu \leq \mu_u$ could be used, so that the drift for the set of queue lengths occurring with high probability is not zero. For INTERVAL- μ CHOICE-2-2, it is possible to choose $\mu(q)$ so that the drift is initially increasing and then decreasing for the set of queue lengths occurring with high probability. Then the stationary probability has a geometrically increasing and decaying behaviour which leads to the $\log(\frac{1}{V})$ behaviour. We expect that this is the reason behind the $\mathcal{O}(\log(\frac{1}{V}))$ growth observed for the case of piecewise linear cost functions in [43]. For INTERVAL- μ CHOICE-2-1 we note that as $V \downarrow 0$ the service rate 0 could be used, which for non-idling admissible policies implies that the queue length 0 has positive stationary probability, even for $V = 0$. This intuitively implies that the average queue length does not increase to infinity in this case.

Asymptotic bounds on any sequence of non-idling order-optimal admissible policies have been presented in Section 3.2.3. We have also discussed other variants of the tradeoff problem, such as INTERVAL- λ CHOICE, where only the arrival rate can be controlled. Even though INTERVAL- λ CHOICE is very similar to INTERVAL- μ CHOICE we find that there are no cases for INTERVAL- λ CHOICE where the minimum average queue length increases to a finite value unlike INTERVAL- μ CHOICE. We observe that the asymptotic behaviour of the minimum average queue length as the utility constraint is made arbitrarily close to the maximum value of the utility, can be obtained using ideas which are similar to that of INTERVAL- μ CHOICE, which shows that the method of obtaining the asymptotic behaviour of the average queue length through its stationary distribution for monotone policies is sufficiently general and provides a unified method which explains other scenarios also.

We also note that through the analysis of the state dependent M/M/1 model, with proper choice of model features, we get insights as to how to construct asymptotic lower bounds for the discrete time models. In Chapter 4, we find bounds on the stationary probability distribution of the queue length for the discrete time queueing model in the regime \mathfrak{R} , which have the asymptotic behaviour as suggested by the above analysis. Thus these bounds lead to the *right* asymptotic behaviour of the average queue length in the regime \mathfrak{R} for the discrete time queueing models.

Appendices

3.A A lower bound for strictly convex functions

Proposition 3.A.1. Let $g(x) : [0, R] \rightarrow \mathbb{R}_+$ be a finite strictly convex function, such that: (i) $g(0) = 0$, (ii) $g'(0) \geq 0$, and (iii) $g''(0) > 0$. Then there exists a positive constant a such that $g(x) \geq ax^2, \forall x \in [0, R]$.

Proof. We note that at 0, $g(0) = a \cdot 0$. If $\exists a > 0$ such that for all $x \in (0, R]$, $g'(x) \geq 2ax$, then $g(x) \geq ax^2$. So we have to prove that for all $x \in (0, R]$, $g'(x) \geq 2ax$. In essence, we have to prove that it is possible to find a positive a such that $\forall x \in (0, R], \frac{g'(x)}{x} \geq a$. We note that if $\inf_{x \in (0, R]} \frac{g'(x)}{x} > 0$, then it is possible to find such an a .

Now we prove that $\inf_{x \in (0, R]} \frac{g'(x)}{x} > 0$. Since both $g'(x)$ and x are non-negative, $\frac{g'(x)}{x} \geq 0$. Suppose we assume that $\inf_{x \in (0, R]} \frac{g'(x)}{x} = 0$. Then we have that $\forall \epsilon > 0, \exists x$ such that $\frac{g'(x)}{x} \leq \epsilon$.

Consider a sequence $\epsilon_n \downarrow 0$, then there exists a sequence x_n such that $\frac{g'(x_n)}{x_n} \leq \epsilon_n$. Note that as $x_n \leq R$, we have that 1) $\lim_{n \rightarrow \infty} g'(x_n) = 0$ and 2) $\lim_{n \rightarrow \infty} \frac{g'(x_n)}{x_n} = 0$. The sequence x_n may not be convergent. So we consider the subsequence y_m such that $\lim_{m \rightarrow \infty} y_m = \limsup_{n \rightarrow \infty} x_n = y$. Note that $\lim_{m \rightarrow \infty} g'(y_m) = 0$ and $\lim_{m \rightarrow \infty} \frac{g'(y_m)}{y_m} = 0$ as y_m is a subsequence of x_n . Now there are two cases: 1) $y > 0$ and 2) $y = 0$. Let $y > 0$, then by continuity of $g'(x)$ we have that $g'(y) = 0$. But $g'(y) > 0$ for every $y > 0$ and we have a contradiction on the assumption that $\inf_{x \in (0, R]} \frac{g'(x)}{x} = 0$. Consider the second case when $y = 0$, then we have that $\lim_{y_m \rightarrow 0} \frac{g'(y_m)}{y_m} = 0$. But note that $\lim_{x \rightarrow 0} \frac{g'(x)}{x} = g''(0) > 0$. Thus we again have a contradiction. Hence $\inf_{x \in (0, R]} \frac{g'(x)}{x} > 0$. We choose $a = \frac{1}{4} \inf_{x \in (0, R]} \frac{g'(x)}{x}$. \square

3.B Application of the Berry-Gallager lower bounding technique [7] in Remark 3.2.4 for INTERVAL- μ CHOICE-1

For INTERVAL- μ CHOICE, since we assume that $\mu(q) \leq r_{max}$, we can obtain a discrete time process $Q_d[m], m \in \mathbb{Z}_+$, by uniformizing the CTMC $Q(t)$ at rate $r_u = \lambda + r_{max}$ as in Appendix 2.A. We now outline how the lower bounding technique in [7] can be applied to $Q_d[m]$ to obtain the $\Omega\left(\frac{1}{\sqrt{c_c - c(\lambda)}}\right)$ asymptotic lower bound for $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$. We first show that a slightly modified form of [7, Lemma 4.1] holds for the uniformized process $Q_d[m]$ under an admissible policy γ (we note that we are using admissibility as defined in Chapter 2 and not as in [7]). As in [7, Appendix A], we have that $Pr\{Q < \lceil 2\bar{Q}(\gamma) \rceil\} > \frac{1}{2}$. Let $q_p = \arg \max_{q \in \{0, \dots, \lceil 2\bar{Q}(\gamma) \rceil\}} \pi(q)$. Then we have that $\pi(q_p) \geq \frac{1}{2 \lceil 2\bar{Q}(\gamma) \rceil}$. Now we define $\hat{Q}_d[m] = \max(Q_d[m], q_p)$. Then as in [7,

Appendix A] we can show that

$$\pi(q_p) \mathbb{E} \left[\widehat{Q}_d[m+1] - \widehat{Q}_d[m] | Q_d[m] = q_p \right] + \sum_{q=q_p+1}^{\infty} \pi(q) (\lambda - \mu(q)) \leq 0.$$

We note that $\mathbb{E} \left[\widehat{Q}_d[m+1] - \widehat{Q}_d[m] | Q_d[m] = q_p \right] = \frac{\lambda}{r_u}$. Therefore we obtain that

$$\begin{aligned} \pi(q_p) \frac{\lambda}{r_u} + \sum_{q=q_p+1}^{\infty} \pi(q) (\lambda - \mu(q)) &\leq 0, \text{ or,} \\ \sum_{q=q_p+1}^{\infty} \pi(q) (\lambda - \mu(q)) &\leq -\frac{\lambda}{r_u 2 \lceil 2\overline{Q}(\gamma) \rceil}. \end{aligned}$$

Then to obtain the asymptotic lower bound on $Q^*(c_c)$ we use the above upper bound in step (49) in the proof of [7, Proposition 4.2].

On the tradeoff of average queue length and average service cost for discrete time single server queues

4.1 Introduction

In this chapter we consider the tradeoff between average queue length and average service cost for the discrete time single server queueing models introduced in Chapter 1. Such discrete time models are in some cases more appropriate for modelling resource allocation problems in communication networks than the continuous time models considered in the previous chapters. A context in which such a model and the following analysis may be appropriate is that of wireless networks with fading, where one of the issues, which has been studied by many researchers (see [7], [29], [10], and [77]), is the optimal tradeoff of the average power and average delay, when the service batch size is dynamically chosen as a function of the fade state and the queue length. The characterization of the tradeoff between average error rate and average delay, for a point-to-point noisy link, when the service batch size is dynamically chosen as a function of the queue length, is another resource allocation problem that motivates the model studied in this chapter. The tradeoff of average power and average delay is dealt with in more detail in Chapter 5, whereas the tradeoff of average error rate with average delay is considered in Chapter 6. In this chapter, we consider two simplified discrete time models, with no admission control (i.e. $A[m] = R[m], \forall m$) and a single environment state, to develop the basic techniques for the characterization of such tradeoffs. The glossary of notation that we use in this chapter is given in Table 4.1. We now summarize the methodology that is used for obtaining the asymptotic lower bounds on the tradeoffs.

Symbol	Description
m	slot index
$A[m]$	random number of arrivals in slot m (after admission control)
A_{max}	maximum number of arrivals in any slot
λ, σ^2	mean and variance of $A[1]$
$S[m]$	batch service size in slot m
S_{max}	maximum batch service size
ϵ_a	probability of $A[m]$ exceeding S_{max}
$Q[m]$	queue length at the start of $(m + 1)^{th}$ slot
$\sigma[m]$	history of queue evolution
γ	policy - $(S[1], S[2], \dots)$
Γ	set of all policies
Γ_s	set of all stationary policies
$\bar{Q}(\gamma)$	average queue length for a policy γ
$c(s)$	service cost for I-model; a function of integer valued batch service size s
$c_R(s)$	service cost for R-model; a function of real valued batch service size s
$\bar{C}(\gamma)$	average service cost for a policy γ
c_c	average service cost constraint
β_{c_c}	non-negative Lagrange multiplier
$c_{\beta_{c_c}}(q, s)$	single stage cost; defined as $q + \beta_{c_c}c(s)$
g_β^*	optimal average cost for unconstrained MDP
$J_\beta(q)$	relative value function for unconstrained MDP
γ_β^*	optimal stationary policy for unconstrained MDP
$s_\beta^*(q)$	action at state q for the optimal policy γ_β^*
$\mathcal{R}_{\gamma_\beta^*}$	recurrence class for γ_β^*
Γ_a	set of all admissible policies
$\bar{s}(q)$	average service rate at a queue length q
π	stationary distribution for a policy which is clear from the context
π_γ	stationary distribution for policy γ
$Q^*(c_c)$	minimum average queue length over the set Γ_a under constraint c_c
s_l	largest service batch size $\leq \lambda$ at which slope of $c(\cdot)$ changes
s_u	smallest service batch size $\geq \lambda$ at which slope of $c(\cdot)$ changes

Table 4.1: Notation used in this chapter.

4.1.1 Methodology

In this section, we summarize a scheme that is used in Chapters 4, 5, and 6 to obtain asymptotic lower bounds on the minimum average queue length in the regime \mathfrak{R} as a function of the average service cost constraint (Chapters 4 and 6) or average power constraint (Chapter 5). As in Chapters 2 and 3, we obtain a lower bound on the stationary mean queue length using an upper bound on the stationary probability distribution of the queue length under the assumption that the queue length process is a DTMC and ergodic. The average service cost (or average power) is then related to the stationary probability distribution.

Let $Q \sim \pi$, where π is the stationary probability distribution of the queue length. From Markov inequality, we have that $\mathbb{E}_\pi Q \geq \bar{q} Pr\{Q \geq \bar{q}\}$. Suppose $Pr_u\{Q < q\}$ is any upper bound on the stationary probability distribution $Pr\{Q < q\}$. If \bar{q} is the largest q such that $Pr_u\{Q < q\} \leq \alpha$, then we have that $\mathbb{E}_\pi Q \geq \bar{q}(1 - \alpha)$. For convenience, we choose $\alpha = \frac{1}{2}$.

The upper bounds $Pr_u\{Q < q\}$ were obtained for the state dependent M/M/1 models in Chapters 2 and 3 using the detailed balance equations and lower bounds on the service rate $\mu(q)$. In Chapters 4, 5, and 6 we obtain geometric bounds on $Pr_u\{Q < q\}$, the first (Lemma 4.3.5) of which has been obtained by assuming certain properties for the transition probability distribution while the second (Lemma 4.3.7) has been obtained by extending the results available in Bertsimas et al. ([8] and [9]).

We now illustrate this scheme for the case of integer-valued queue evolution through an example. We will obtain that for the queueing process, the upper bound $Pr_u\{.\}$ has a geometric form, i.e., $Pr_u\{Q < q\} = \pi(0)\rho^q$. For our tradeoff problem, it will turn out that a non-negative function $D(q)$ can be obtained, such that $\mathbb{E}_\pi D(Q)$ is the difference between the average service cost and $c(\lambda)$ (which can be defined similarly as for INTERVAL- μ CHOICE). Then if $\mathbb{E}_\pi D(Q) \leq V$ for a $V > 0$, then we have that $\pi(0)D(0) \leq V$, where $D(0) > 0$. Therefore, we have a further upper bound on the stationary probability distribution, $Pr_u\{Q < q\} = \frac{V}{D(0)}\rho^q$. Then, the largest \bar{q}_1 such that $Pr_u\{Q < q\} \leq \frac{1}{2}$ satisfies $\frac{V}{D(0)}\rho^{\bar{q}_1} \leq \frac{1}{2}$. Or we have that $\bar{q}_1 = \left\lfloor \log_\rho \left(\frac{D(0)}{2V} \right) \right\rfloor$. In the asymptotic regime \mathfrak{R} , as $V \downarrow 0$, we obtain that \bar{q}_1 is $\Omega(\log(\frac{1}{V}))$ and therefore so is the average queue length as a function of V . Variations of this basic method are used throughout Chapters 4, 5, and 6 to obtain the asymptotic lower bounds.

4.1.2 System model - Integer valued queue evolution

We assume time to be slotted, with slots indexed by $m \in \mathbb{Z}_+$. In each slot m , a random number of customers $A[m] \in \mathbb{Z}_+$ arrive into the system. The arrival process $(A[m], m \geq 1)$ is assumed to be IID with $A[1] \leq A_{max}$, batch arrival rate $\mathbb{E}A[1] = \lambda < \infty$, $\text{var}(A[1]) = \sigma^2 < \infty$. The customers arrive into an infinite buffer queue. In slot m , a batch of customers of size $S[m] \in \mathbb{Z}_+$

is served. The batch of $S[m]$ customers is removed from the queue at the end of the m^{th} slot just before the new batch of customers which arrive in the m^{th} slot, $A[m]$, is admitted. We assume that $S[m] \leq S_{max}$, where S_{max} is the maximum batch size that can be served. We also assume that

$$\mathbf{A1} : Pr \{A[1] > S_{max}\} > \epsilon_a > 0.$$

We note that the above assumption is similar to the assumptions made in the definition of admissible policies in [7] and [43]. Furthermore, we note that the above assumption is reasonable, since the maximum number of arrivals usually exceeds the maximum capacity S_{max} of service, but $\lambda < S_{max}$.

The number of customers in the queue at the start of the $(m + 1)^{\text{th}}$ slot is denoted by $Q[m]$. We assume that $Q[0] = q_0 \in \mathbb{Z}_+$ customers. The queue evolution for $m \geq 0$ is given by:

$$Q[m + 1] = Q[m] - S[m + 1] + A[m + 1], \quad (4.1)$$

where $S[m + 1] \leq \min(S_{max}, Q[m])$. We define a policy γ to be the sequence of batch sizes $(S[1], S[2], \dots)$ under which the queue evolves according to (4.1). The set of all policies is denoted by Γ . A policy γ is stationary if $S[m + 1] = S(Q[m])$, $\forall m$, where $S(q)$ is a randomized function of q . The set of all stationary policies is denoted as Γ_s .

We assume that a service cost of $c(s)$ is incurred when serving a batch of size s . For example, this cost could be the expected number of symbols in error for the transmission of a batch of message symbols. The function $c(s) : \{0, 1, \dots, S_{max}\} \rightarrow \mathbb{R}_+$ is assumed to satisfy the following properties:

$$\mathbf{C1} : c(0) = 0,$$

$$\mathbf{C2} : c(s) \text{ is non-decreasing and convex }^1 \text{ for } s \in \{0, \dots, S_{max} - 2\}.$$

We now define the performance measures that we are interested in: a) The worst case average queue length for a policy $\gamma \in \Gamma$ is

$$\bar{Q}(\gamma, q_0) = \limsup_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=0}^{M-1} Q[m] \middle| Q[0] = q_0 \right], \quad (4.2)$$

and b) the worst case average service cost for a policy $\gamma \in \Gamma$ is

$$\bar{C}(\gamma, q_0) = \limsup_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M c(S[m]) \middle| Q[0] = q_0 \right]. \quad (4.3)$$

¹If $S_{max} = 1$, then there is no tradeoff. If $S_{max} \geq 2$, then $c(s + 2) - c(s + 1) \geq c(s + 1) - c(s)$, $\forall s \in \{0, \dots, S_{max} - 2\}$.

4.1.3 System model - Real valued queue evolution

We state only the differences from the model discussed in Section 4.1.2. We assume that the arrival random variable $A[m] \in [0, A_{max}] \subset \mathbb{R}_+$, with mean $\lambda < \infty$ and variance $\sigma^2 < \infty$. We also assume that the service batch size $S[m] \in [0, S_{max}] \subset \mathbb{R}_+$. Similar to assumption A1, we make the assumption:

$$\mathbf{RA1} : Pr \{A[1] - S_{max} > \delta_a\} > \epsilon_a,$$

where both δ_a and ϵ_a are positive real numbers. Let the initial queue length be $q_0 \in \mathbb{R}_+$. We have that the queue length $Q[m] \in \mathbb{R}_+$ and the evolution of $Q[m]$ is given by (4.1) with $S[m] \leq \min(S_{max}, Q[m-1])$. We assume that there is a service cost $c_R(s)$ associated with the service of a batch of size s . The function $c_R(s) : [0, S_{max}] \rightarrow \mathbb{R}_+$ is assumed to satisfy the following properties:

$$\mathbf{RC1} : c_R(0) = 0,$$

$$\mathbf{RC2} : c_R(s) \text{ is strictly convex and increasing in } s, \text{ for } s \in [0, S_{max}].$$

In the following, the model with integer valued queue evolution is referred to as the I-model, while the model with real valued queue evolution is referred to as the R-model. We note that R-model with the strictly convex $c_R(s)$ cost function is usually used as an approximation to the I-model, which has $c(s)$ as the cost function.

We note that our I-model is a simplified version of the model studied by Goyal et al. [29], wherein there is an additional state variable which is used to model fading, the arrival process is Markov, and $S_{max} = \infty$. Our R-model is a simplified version of the model studied by Berry and Gallager [7], wherein there is an additional fade state variable, and $S_{max} = \infty$.

4.1.4 Overview

The tradeoff problem (1.7) for I-model and R-model is formulated as a constrained Markov decision problem in Section 4.2. We consider the I-model first. For I-model, for certain values of the cost constraint, we consider an equivalent unconstrained Markov decision problem (as in (1.8)) in Section 4.3.1. We also identify several properties that are possessed by any stationary deterministic optimal policy for this problem. We then define the set of admissible policies, which are policies possessing the above properties (the definition of admissible policies is similar to that in Chapter 2).

From Section 4.3.2 onwards, we consider the tradeoff problem for the set of admissible policies. In Section 4.3.3 we characterize the infimum of the average service cost over all possible admissible policies, which is equal to the minimum average service cost required for mean rate stability. We identify three cases, which are similar to the three subcases for INTERVAL- μ CHOICE-2, for which

the asymptotic behaviour of the minimum average queue length is characterized. The asymptotic behaviour of the minimum average queue length is obtained as for INTERVAL- μ CHOICE-2, by first obtaining upper bounds on the stationary probability of the queue length. One of these bounds is a state dependent extension of the geometric bounds on the stationary probability of discrete time Markov chains presented in [8] and [9]. In Section 4.3.4 we show that depending on the value of the arrival rate, the minimum average queue length either (i) increases only to a finite value, or (ii) increases as $\log(\frac{1}{V})$, or (iii) increases as $\frac{1}{V}$, when the average service cost is V more than the infimum of the average service costs for admissible policies. Asymptotic bounds on order-optimal policies are presented in Section 4.3.5. We obtain an asymptotic lower bound on the minimum average queue length for ergodic arrival processes in Section 4.3.7, when the average service cost is V more than the infimum of the average service costs for admissible policies.

For R-model, we present an asymptotic analysis in Section 4.1.3. We obtain that the minimum average queue length is $\Omega\left(\frac{1}{\sqrt{V}}\right)$ when the average service cost is V more than the infimum of the average service costs for admissible policies, for strictly convex $c_R(s)$. Then we consider the case where $c_R(s)$ is piecewise linear and show that the asymptotic behaviour of the minimum average queue length is similar to that of the I-model.

4.2 Problem formulation for I-model and R-model

The tradeoff problem is to obtain $Q^*(c_c, q_0)$, which is the optimal value of the optimization problem

$$\underset{\gamma \in \Gamma}{\text{minimize}} \bar{Q}(\gamma, q_0) \text{ such that } \bar{C}(\gamma, q_0) \leq c_c, \quad (4.4)$$

where $c_c \geq 0$ is the average service cost constraint. The tradeoff curve $Q^*(c_c, q_0)$ is non-increasing and convex in c_c (see [6]) for any q_0 . We note that if $\lambda > S_{max}$ then any feasible policy for (4.4) is optimal, as the average queue length for any such feasible policy is infinity. Hence, in the following we assume that $\lambda \leq S_{max}$.

4.2.1 A constrained Markov decision process formulation

The tradeoff problem (4.4) can be formulated as a constrained Markov decision problem (CMDP) [2]. The state space of the CMDP is the state space of the queue length, which is \mathbb{Z}_+ for the I-model and \mathbb{R}_+ for the R-model. The action spaces at each state q are the sets $\{0, \dots, \min(q, S_{max})\}$ for the I-model and $[0, \min(q, S_{max})]$ for the R-model, both of which are compact for every q . The probabilistic evolution of the state of the CMDP from stage m to $m + 1$ is given by (4.1). Associated with the CMDP there are two single stage costs: (i) the holding cost q at state q , and (ii) the service cost $c(s)$ when an action s is taken at state q .

We redefine $c(s)$ as the lower convex envelope² of $c(s), s \in \{0, \dots, S_{max}\}$. We note that the redefined cost function $c(s)$ ($c(s) : [0, S_{max}] \rightarrow \mathbb{R}_+$) is a piecewise linear convex function. From [28] and [27], it is possible to show that if $c_c > c(\lambda)$ then (4.4) has an optimal solution and there exists an optimal policy $\gamma \in \Gamma_s$.

In the following, we show that for some values of c_c , there exists a stationary deterministic optimal policy for (4.4). Consider the following MDP:

$$\underset{\gamma \in \Gamma}{\text{minimize}} \left[\bar{Q}(\gamma, q_0) + \beta_{c_c}(\bar{C}(\gamma, q_0) - c_c) \right]. \quad (4.5)$$

We note that the above MDP has a single stage cost of $c_{\beta_{c_c}}(q, s) = q + \beta_{c_c}c(s)$ in state q . From Ma et al. [35], it is known that if there exists a $\beta_{c_c} > 0$, such that any stationary deterministic optimal policy for (4.5), has an average service cost equal to the constraint c_c , then the same policy is optimal for the constrained problem (4.4). The factor β_{c_c} can be interpreted as a Lagrange multiplier. The set of values of c_c for which such β_{c_c} exist is denoted as \mathcal{O}^u , as in Chapter 2. We note that the properties of any stationary deterministic optimal policy, which can be obtained from (4.5), carry over to (4.4) if $c_c \in \mathcal{O}^u$.

We note that the development in Altman [2] which leads to Theorem 12.7, which shows that for every value of the constraint c_c , there exists a Lagrange multiplier for which there is a stationary deterministic policy which is optimal for both the unconstrained MDP and the CMDP, requires assumption (B1) [2, Chapter 11], which does not hold for our model.

In the next section, for the I-model we study (4.5) in detail. The properties of any stationary deterministic optimal policy are then used to motivate the definition of a class of admissible policies. The tradeoff problem (4.4) is then analysed for the class of admissible policies. We note that for $c_c \in \mathcal{O}^u$, there exists at least one optimal admissible policy.

4.3 Asymptotic bounds for I-model

4.3.1 An unconstrained MDP formulation

The unconstrained MDP (4.5), which is obtained via the above Lagrange multiplier relaxation, is studied in [29]. However, we note that the development in [29] does not lead to an average cost optimality equation (ACOE). Since the ACOE enables us to obtain some additional properties of the optimal policy, in the next section we use the results from Sennott [67] to show that there exists a stationary deterministic average cost optimal policy, which also satisfies an ACOE, for a single stage cost of $c_{\beta}(q, s)$ for our model.

²We note that the lower convex envelope can be interpreted as the solution: $c(s) = \text{minimize } \mathbb{E}c(X)$, such that $\mathbb{E}X = s$.

The state space of the unconstrained MDP (4.5) is the state space \mathbb{Z}_+ of the queue length, which is countable. The action space at each state q is the set $\{0, \dots, \min(q, S_{max})\}$ which is compact. The probabilistic evolution of the state of the MDP from stage m to $m + 1$ is given by (4.1). The single stage cost for the unconstrained MDP is $c_\beta(q, s) = q + \beta c(s)$, where $\beta \geq 0$. We also assume that:

A2 : $Pr \{A[1] = a\} > 0$, for all $a \in \{0, \dots, A_{max}\}$,

Lemma 4.3.1. If $\lambda < S_{max}$, and if assumptions A1 and A2 hold, then there exists a stationary deterministic optimal policy γ_β^* for the unconstrained MDP, with optimal average cost g_β^* satisfying the following ACOE:

$$g_\beta^* + J_\beta(q) = \min_{s \in \{0, \dots, \min(q, S_{max})\}} \left\{ c_\beta(q, s) + \mathbb{E}J_\beta(q - s + A[1]) \right\}, \forall q \geq 0,$$

with γ_β^* using a batch size $s_\beta^*(q)$ at queue length q , satisfying

$$s_\beta^*(q) = \arg \min_{s \in \{0, \dots, \min(q, S_{max})\}} \left\{ c_\beta(q, s) + \mathbb{E}J_\beta(q - s + A[1]) \right\}, \forall q \geq 0,$$

where $J_\beta(q)$ is the optimal relative value function.

The proof is given in Appendix 4.A. From now on, we assume $\lambda < S_{max}$ and that A1 and A2 hold.

The following property of any optimal policy γ_β^* can also be obtained.

O1 : any stationary deterministic optimal policy γ_β^* is such that $s_\beta^*(q)$ is non-decreasing in q .

The proof of the above property is similar to that of Theorem 3.2 (iii) of [29], and is therefore omitted. We now state some observations which are obtained from the above lemma and O1.

O2 : The optimal average cost g_β^* is independent of the initial state q_0 and is finite.

O3 : For any policy, from assumptions A1 and A2, we note that from state 0 it is possible to reach any other state q . From O1, we obtain that any stationary deterministic optimal policy has a single recurrence class $\mathcal{R}_{\gamma_\beta^*}$, of the form $\{q_m, \dots\}$, where $q_m = \min \{q : \exists q' > q, p_{q',q} > 0\}$, and $p_{q',q} = Pr \{Q[m+1] = q | Q[m] = q'\}$ for the optimal policy under consideration.

O4 : We note that $s_\beta^*(q_m) = 0$ by definition. From A2, we have that q_m is an aperiodic state, and therefore the class $\mathcal{R}_{\gamma_\beta^*}$ of the Markov chain under γ_β^* is aperiodic.

O5 : From Lemma 4.3.1, for γ_β^* , for $q \in \mathcal{R}_{\gamma_\beta^*}$, we have that

$$g_\beta^* + J_\beta(q) = c_\beta(q, s_\beta^*(q)) + \mathbb{E}J_\beta(q - s_\beta^*(q) + A[1]),$$

which verifies the drift condition (10.13) from [37], with the Lyapunov function $V(q) = J_\beta(q)$, for a Markov chain restricted to $\mathcal{R}_{\gamma_\beta^*}$. We also note that $c_\beta(q, s_\beta^*(q))$ is near-monotone [37] in q . Therefore using [37, Theorem 10.3], we obtain that any stationary deterministic optimal policy is c_β -regular for a Markov chain on $\mathcal{R}_{\gamma_\beta^*}$. Then the Markov chain under γ_β^* is positive recurrent on $\mathcal{R}_{\gamma_\beta^*}$ with an associated invariant distribution. Furthermore the expected total cost of first passage from any state $q \in \mathcal{R}_{\gamma_\beta^*}$ to another state $q' \in \mathcal{R}_{\gamma_\beta^*}$ is finite [37, Theorem 10.3].

Using property O5 of γ_β^* we prove the following lemma, which shows that any optimal policy is in fact non-idling. We note that if γ_β^* is non-idling, then $R_\beta = \mathbb{Z}_+$ and the Markov chain under γ_β^* is irreducible.

Lemma 4.3.2. Any stationary deterministic optimal policy γ_β^* is non-idling, i.e., $s_\beta^*(q) > 0$, for all $q \in \mathcal{R}_{\gamma_\beta^*}$ and $q > 0$.

The proof shows that if γ_β^* is such that there exists a positive $q_1 \in \mathcal{R}_{\gamma_\beta^*}$ with $s_\beta^*(q_1) = 0$, then it is possible to construct a history dependent non-stationary policy for which the average $c_\beta(q, s)$ cost is strictly less, contradicting the optimality of γ_β^* . The essential steps in the proof are: a) we consider a particular sample path of the queue evolution, for which it is assumed that at the start of a slot m , $Q[m-1] = q_1$ with $S[m] = 0$, b) we obtain a new non-stationary policy which advances the service of one customer³ from one of the succeeding slots to m while keeping the departure times of all other customers unchanged, and c) we show that for this new policy the total $c_\beta(q, s)$ cost decreases because i) the delay of the customer whose departure time was advanced has decreased and ii) convexity of the service cost function implies that the service cost at m , $c(1)$, is less than or equal to the decrease in service cost at the slot where the customer was being served under policy γ_β^* . Extension of the proof to the case of average cost, with the optimal policy being not irreducible, is more technical and is therefore presented in Appendix 4.B.

We recall that by solving the unconstrained MDP, we are able to get solutions to problem (4.4), with service cost constraint c_c , only if a Lagrange multiplier β exists such that the optimal policy for the unconstrained MDP with single stage cost $c_\beta(q, s)$ has an average service cost equal to c_c , i.e, if $c_c \in \mathcal{O}^u$. Thus in general, the properties O1, O2, O3, O4, and O5, as well as Lemma 4.3.2 may not hold for all values of c_c in the original problem (4.4).

4.3.2 The tradeoff problem

As noted in Section 4.2.1, we consider problem (4.4) for the set of randomized stationary policies Γ_s , since there exists an optimal stationary policy. A policy $\gamma \in \Gamma_s$, specifies the service batch size $S(q)$

³Since q_1 is positive at least one customer will be present in the system at slot m , who will be served in some slot $> m$.

at a queue length q . We note that $S(q)$ is a random variable with support on $\{0, \dots, \min(S_{max}, q)\}$. We further restrict the study of problem (4.4) to the set of stationary *admissible* policies, whose definition is motivated by the properties O1, O2, O3, O4, O5, and Lemma 4.3.2. We now define the notion of stability for a policy.

Stability: A policy $\gamma \in \Gamma_s$ is said to be stable if: a) the Markov chain $Q[m]$ under γ is positive recurrent with stationary distribution π_γ on the recurrence class corresponding to q_0 and b) $\bar{Q}(\gamma, q_0) < \infty$.

Admissibility: A policy γ is called admissible if:

G1 : it is stable,

G2 : it induces an aperiodic, irreducible Markov chain $Q[m]$,

G3 : the average service rate at a queue length q , $\mathbb{E}S(q)$ is non-decreasing in q .

The set of all admissible policies is Γ_a . We note that the above properties of admissible policies are motivated by the observations about stationary deterministic policies made in Section 4.3.1. Property G1 is motivated by O2 and O5. Properties O3, O4 and Lemma 4.3.2 motivates property G2⁴, while G3 is motivated by O1. We note that the above definition of admissibility differs in the addition of property G3, from the definitions of admissible policies which were used by Berry and Gallager [7, Section IV] and Neely [43, Section III]. We also note that assumptions A1 and A2 have been used to motivate G2, through the properties O3 and O4, but as in Berry and Gallager [7], we could assume G2.

Remark 4.3.3. We now compare our definition of an admissible policy with that of Berry and Gallager [7]. In [7], it is required that a sequence of admissible policies form an ergodic Markov chain, i.e., an aperiodic, irreducible and positive recurrent Markov chain. Our admissible policies are also assumed to satisfy the same properties. In [7], it is required that a sequence of admissible policies γ_k are such that $\bar{Q}(\gamma_k) < \infty$ and $\lim_{k \rightarrow \infty} \bar{Q}(\gamma_k) = \infty$. We also assume that $\bar{Q}(\gamma_k) < \infty$. However, we do not assume that $\lim_{k \rightarrow \infty} \bar{Q}(\gamma_k) = \infty$. In fact, we shall see that for a particular case (Case 1), the optimal sequence of policies γ_k^* is such that $\lim_{k \rightarrow \infty} \bar{Q}(\gamma_k^*) < \infty$. The third property in [7], that admissible policies are assumed to satisfy is similar to our assumption A1. We note that the additional property G3 can be used to obtain additional insights about any stationary deterministic optimal policy. Furthermore, for $c_c \in \mathcal{O}^u$, there exists an admissible optimal policy.

From G1 and G2, the average queue length as well as the average service cost are independent of the initial queue length since the Markov chain ($Q[m]$) is aperiodic, irreducible and stable [37].

⁴We note that the development of asymptotic lower bounds also holds under an assumption weaker than irreducibility. We can assume that under the policy γ , (i) there is only a single positive recurrent class R_γ and (ii) the expected cumulative queue length and expected cumulative service cost starting from any state q_0 until R_γ is hit is finite. We note under assumption A2 and G3, R_γ is a contiguous set.

The average queue length and average service cost for $\gamma \in \Gamma_a$ are therefore denoted by $\overline{Q}(\gamma)$ and $\overline{C}(\gamma)$. We note that $\overline{Q}(\gamma) = \mathbb{E}_{\pi_\gamma} Q$ and $\overline{C}(\gamma) = \mathbb{E}_{\pi_\gamma} c(S(Q))$, where Q denotes the stationary queue length. Since, in general, for every c_c , we do not know if Lemma 4.3.2 holds for the optimal constrained policy, admissible policies are not required to be non-idling.

Objective : Our objective is to obtain the optimal tradeoff curve $Q^*(c_c)$ while restricting our attention to the class of admissible policies Γ_a , where $Q^*(c_c)$ is the optimal value of the following optimization problem

$$\text{TRADEOFF : } \underset{\gamma \in \Gamma_a}{\text{minimize}} \overline{Q}(\gamma) \text{ such that } \overline{C}(\gamma) \leq c_c.$$

We note that the TRADEOFF problem can be formulated for a larger class of policies, which are obtained by time sharing or mixing of policies $\gamma_k \in \Gamma_a$. Let $Q_M^*(c_c)$ denote the optimal tradeoff curve, when we consider time shared policies also. We note that the tradeoff curve $Q_M^*(c_c)$ which is obtained from time sharing is the lower convex envelope of the points $(c_c, Q^*(c_c))$. Since the asymptotic behaviour of $Q_M^*(c_c)$ can be obtained from that of $Q^*(c_c)$, as in Chapter 2, in the following we analyse $Q^*(c_c)$ only.

We note that given any $\epsilon > 0$, and for any c_c such that TRADEOFF is feasible, by definition there exists an admissible policy γ such that $\overline{Q}(\gamma) \leq Q^*(c_c) + \epsilon$ and $\overline{C}(\gamma) \leq c_c$. Such a feasible admissible policy is called ϵ -optimal in the following.

We note that for any $\gamma \in \Gamma_a$, we have that $\mathbb{E}_{\pi_\gamma} S(Q) = \lambda$. We recall that $c(s)$ was redefined as the piecewise linear lower convex envelope of $c(s)$, $s \in \{0, \dots, S_{max}\}$. Then, from Jensen's inequality, we have that for any policy $\gamma \in \Gamma_a$, $\overline{C}(\gamma) = \mathbb{E}_{\pi_\gamma} c(S(Q)) \geq c(\lambda)$. Therefore, $\inf_{\gamma \in \Gamma_a} \overline{C}(\gamma) \geq c(\lambda)$. We note that TRADEOFF does not have any feasible solutions if c_c is less than $c(\lambda)$.

We also note that $c(\lambda)$ is also the minimum average service cost which has to be expended for mean rate stability. In the following we show that $c(\lambda) = \inf_{\gamma \in \Gamma_a} \overline{C}(\gamma)$. We obtain an asymptotic characterization of $Q^*(c_c)$ in the asymptotic regime \Re as $c_c \downarrow c(\lambda)$ in the next section.

4.3.3 Asymptotic analysis of TRADEOFF - Preliminaries

The ideas used in the analysis of TRADEOFF are the same as those used for the analysis of INTERVAL- μ CHOICE-2 in Chapter 3. However, bounds on the stationary probability of the queue length for the DTMC ($Q[m]$) have to be developed, in order to relate the average service cost to the average queue length. As in Chapter 3, we identify three different cases based on the nature of the function $c(s)$ at $s = \lambda$. The cases are defined in terms of the quantities s_l and s_u defined as

follows:

$$s_u = \begin{cases} \min \{s : s \in \{\lceil \lambda \rceil, \dots, S_{max} - 1\}, c(s+1) - c(s) > c(s) - c(s-1)\} & \text{if this set is non-empty} \\ S_{max} & \text{otherwise.} \end{cases}$$

$$s_l = \begin{cases} \max \{s : s \in \{1, \dots, \lfloor \lambda \rfloor\}, c(s+1) - c(s) > c(s) - c(s-1)\} & \text{if this set is non-empty} \\ 0 & \text{otherwise.} \end{cases}$$

We note that s_l and s_u are analogous to the service rates μ_l and μ_u defined in chapter 3. The three cases that we consider are:

Case 1 : $s_l = 0, s_l < \lambda < s_u$

Case 2 : $s_l > 0, s_l < \lambda < s_u$

Case 3 : $s_l = \lambda = s_u$

We now show that $c(\lambda)$ can in fact be approached arbitrarily closely.

Lemma 4.3.4. There exists a sequence of policies $\gamma_{\epsilon_k} \in \Gamma_a$ such that $c(\lambda)$ can be approached arbitrarily closely, i.e., $\lim_{k \rightarrow \infty} \bar{C}(\gamma_{\epsilon_k}) = c(\lambda)$. Therefore $c(\lambda) = \inf_{\gamma \in \Gamma_a} \bar{C}(\gamma)$.

The proof is given in Appendix 4.C. We characterize the tradeoff curve $Q^*(c_c)$ in the asymptotic regime where the cost constraint c_c approaches $c(\lambda)$.

Similar to the definition of the line $l(\mu)$ in Chapter 3, here we define the line $l(s) : [0, S_{max}] \rightarrow \mathbb{R}_+$ as follows:

1. If $s_l < \lambda < s_u$, then $l(s)$ is the line through $(s_l, c(s_l))$ and $(s_u, c(s_u))$.
2. If $s_l = \lambda = s_u$, then $l(s)$ is a line through $(\lambda, c(\lambda))$ with slope m chosen such that $c(\lambda) - c(\lambda - 1) < m < c(\lambda + 1) - c(\lambda)$.

The different cases along with the line $l(s)$ are illustrated in Figure 4.1. We note that $\mathbb{E}_{\pi_\gamma} [c(S) - c(\lambda)] = \mathbb{E}_{\pi_\gamma} [c(S) - l(S)]$.

For a particular policy γ , if there is no source for confusion we use $\pi(q)$ to denote the stationary probability of queue length being q . The stationary probability of using a particular batch size $s \in \{0, \dots, S_{max}\}$ is denoted by $\pi_s(s)$. We now present two results which are used in the asymptotic characterization of $Q^*(c_c)$.

Lemma 4.3.5. For $\gamma \in \Gamma_a$, for some positive $\epsilon < s_l$, $q_s = \inf \{q : \mathbb{E}S(q) \geq s_l - \epsilon\}$, $\rho_d = \left(\frac{s_l - \epsilon}{S_{max}}\right) Pr \{A[1] = 0\}$, and $\rho = 1 + \frac{1}{\rho_d}$, if $Pr \{Q < q_s\} \left(1 + \frac{\rho}{\rho_d}\right) < \frac{1}{2}$, we have that

$$\bar{Q}(\gamma) \geq \frac{1}{2} \left[\log_\rho \left[\frac{1}{2Pr \{Q < q_s\}} \right] - 1 \right],$$




Figure 4.1: Illustration of the relationship between λ , s_l , and s_u along with the minimum average cost $c(\lambda)$ and the line $l(s)$ for the three cases

The proof is given in Appendix 4.D. The above result is useful in obtaining an asymptotic lower bound to $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$, since in Cases 2 and 3, as $c_c \downarrow c(\lambda)$, for any sequence of feasible policies for TRADEOFF, $Pr\{Q < q_s\} \downarrow 0$.

Remark 4.3.6. We note that a similar asymptotic lower bound has been derived in [44, Theorem 2] (where admission control is allowed) and in [72, Theorem 2], where the assumption G3 has not been used. Although, the above result has been derived independently, we note that underlying all the three derivations, there is the idea of bounding the probability of an event by a particular sequence of transitions for a Markov chain, i.e., a sequence of transitions in which the state of the Markov chain becomes successively smaller. Furthermore, in our proof, using assumption G3, we obtain geometric bounds on the stationary probability of any queue length, which is not available in [44] as well as [72].

We note that in Chapter 3, since the queue length process $Q(t)$ was a birth-death process, bounds on the stationary probability of the queue length could be obtained relatively easily. However, in this chapter, bounds on the stationary probability of discrete time Markov chains (DTMC) are required. In the following lemma, we present three bounds on the stationary probability of the queue length, one of which has been obtained by Bertsimas et al. [8] and [9] and the other two are state dependent extensions of the geometric bounds on the stationary probability of discrete time Markov chains presented in [8] and [9].

Lemma 4.3.7. Let $(Q[m], m \geq 0)$ be as in (4.1), for an admissible policy γ . Let ϵ_a be as defined in assumption A1. Then,

TAIL-PROB [8] : Suppose $\forall q \geq 0$, $\mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d$, where d is positive. Then for any finite q_1 and $k \geq 1$ we have

$$Pr\{Q \geq q_1 + k\} \geq \left(\frac{\epsilon_a}{\epsilon_a + d}\right)^k Pr\{Q \geq q_1\}.$$

TAIL-PROB-STATE-DEP-1 : Suppose there exists a q_d such that

$$\forall q \in \{0, \dots, q_d\}, \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d,$$

where d is positive. Then for any $q_1, k \geq 0$ such that $0 \leq q_1 + k \leq q_d$, we have

$$\begin{aligned} \Pr\{Q \geq q_1 + k\} &\geq \left(\frac{\epsilon_a}{\epsilon_a + d}\right)^k \Pr\{Q \geq q_1\} \\ &\quad + \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + d}\right)^k\right) \left[\Pr\{Q \geq q_d + 1\} \right. \\ &\quad \left. + \frac{1}{d} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \right]. \end{aligned}$$

TAIL-PROB-STATE-DEP-2 : Suppose there exists a q_d such that

$$\forall q \in \{0, \dots, q_d\}, \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d_1,$$

$$\text{and } \forall q \in \{q_d + 1, \dots\}, \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d_2,$$

where $d_2 > d_1 > 0$. Then for any $q_1, k \geq 0$ such that $0 \leq q_1 + k \leq q_d$, we have

$$\Pr\{Q \geq q_1 + k\} \geq \left(\frac{\epsilon_a}{\epsilon_a + d_1}\right)^k \Pr\{Q \geq q_1\} - \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + d_1}\right)^k\right) \frac{d_2 - d_1}{d_1} \Pr\{Q \geq q_d + 1\}.$$

The proof is presented in Appendix 4.E. We now present the asymptotic characterization of $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$.

4.3.4 Asymptotic characterization of $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$

We first consider Case 2, the proof of the following asymptotic lower bound has already been briefly discussed in Section 4.1.1. We use the geometric bound based lower bound on $\bar{Q}(\gamma)$ from Lemma 4.3.5 to obtain the asymptotic lower bound for Case 2.

Lemma 4.3.8. For Case 2, given any sequence of admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

Proof. Let us consider a particular policy γ in the sequence γ_k with $V_k = V$. From the definition of $l(s)$ we have that $\mathbb{E}_\pi \mathbb{E}_{S|Q} [c(S(Q)) - l(S(Q))] = V$. From the convexity of $c(s)$ and the linearity of $l(s)$ we have

$$\sum_{q=0}^{\infty} \pi(q) [c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))] \leq V.$$

Now as $c(s) \geq l(s)$, we have that

$$\sum_{q=0}^{q_s-1} \pi(q) (c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))) \leq V,$$

where q_s , as in Lemma 4.3.5, is $\inf \{q : \mathbb{E}S(q) \geq s_l - \epsilon\}$ for a positive $\epsilon < s_l$. We note that $q_s \geq 1$. We note that for $q < q_s$, $c(\mathbb{E}S(q)) - l(\mathbb{E}S(q)) \geq m\epsilon$, where m is the tangent of the angle made by the line passing through $(s_l - 1, c(s_l - 1))$ and $(s_l, c(s_l))$ with $l(s)$. Therefore we obtain that

$$\Pr \{Q < q_s\} \leq \frac{V}{m\epsilon}.$$

Using the above bound on $\Pr \{Q < q_s\}$, we have that for sufficiently small V , $\Pr \{Q < q_s\} < \frac{1}{2}$. Then, from Lemma 4.3.5 we have that

$$\bar{Q}(\gamma) \geq \frac{1}{2} \left[\log_{\rho} \left[\frac{m\epsilon}{2V} \right] - 1 \right],$$

where the upper bound $\frac{V}{m\epsilon}$ on $\Pr \{Q < q_s\}$ is used. Therefore for the sequence of policies γ_k with $V_k \downarrow 0$, we have that

$$\bar{Q}(\gamma_k) = \Omega \left(\log \left(\frac{1}{V_k} \right) \right).$$

□

Remark 4.3.9. For Case 2, as for INTERVAL- μ CHOICE-2-2, we have a set of queue lengths \mathcal{Q}_h , which occur with high probability, such that $\mathbb{E}S(q) \in [s_l, s_u], \forall q \in \mathcal{Q}_h$. Let the drift in state q be $\Delta(q) \triangleq \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q]$. We note that, intuitively the drift $\Delta(q)$ for $q \in \mathcal{Q}_h$ is increasing and then decreasing for in \mathcal{Q}_h . Then we expect that the stationary distribution of the queue length is geometrically increasing and then decreasing as for INTERVAL- μ CHOICE-2-2. However, unlike INTERVAL- μ CHOICE-2-2, for the discrete time model we are only able to obtain a geometrically increasing upper bound. This bound suffices to obtain the $\Omega \left(\log \left(\frac{1}{V_k} \right) \right)$ asymptotic lower bound in the above lemma.

We note that the model that is considered here is a simplified version of Neely's [43] model - there is only one queue evolving on \mathbb{Z}_+ and we have the fade state taking only a single value. The tradeoff optimal control algorithm (TOCA) of [43] achieves the logarithmic tradeoff stated in the above lemma but for the problem (4.4), since TOCA is not admissible. Therefore, we propose a sequence of admissible policies that achieves the asymptotic logarithmic tradeoff of $Q^*(c_c)$ in Lemma 4.3.8.

Lemma 4.3.10. Let a policy γ be defined as follows. At a queue length q , γ serves a batch size $\min(q, \tilde{s}(q))$, where

$$\tilde{s}(q) = \begin{cases} s_l, & \text{for } 0 \leq q < q_v, \\ s_u, & \text{for } q_v \leq q. \end{cases}$$

where $q_v > 0$. We obtain a sequence of policies γ_k , by choosing $q_v = \log\left(\frac{1}{V_k}\right)$, where $V_k < 1$ is a sequence decreasing to zero. Then for Case 2, γ_k is a sequence of admissible policies, such that $\overline{C}(\gamma_k) - c(\lambda) = \mathcal{O}(V_k)$ and $\overline{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$.

The proof of this lemma is given in Appendix 4.G. We that the structure of the sequence of policies γ_k is similar to that in Lemma 2.3.12. We note that this proof is motivated by and borrows ideas from the derivation of the asymptotic upper bound for the sequence of TOCA policies in [43, Corollary 2]. This leads to the following asymptotic characterization for Case 2.

Proposition 4.3.11. For Case 2, the optimal tradeoff curve $Q^*(c_{c,k}) = \Theta\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right)$ as $c_{c,k} \downarrow c(\lambda)$, for the sequence $c_{c,k} = \overline{C}(\gamma_k)$, where γ_k is the sequence of policies in Lemma 4.3.10.

Proof. For the sequence $c_{c,k} = \overline{C}(\gamma_k)$, we have that $Q^*(c_{c,k}) \leq \overline{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right)$. For $\epsilon > 0$, consider any sequence of feasible ϵ -optimal admissible policies γ'_k for the sequence $c_{c,k}$. We have that $\overline{Q}(\gamma'_k) = \Omega\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right)$ and $\overline{Q}(\gamma'_k) \leq Q^*(c_{c,k}) + \epsilon$. Therefore, $Q^*(c_{c,k}) = \Omega\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right)$. Hence, for $c_{c,k} = \overline{C}(\gamma_k)$, we have that $Q^*(c_{c,k}) = \Theta\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right)$. \square

Remark 4.3.12. TOCA algorithm : We note that the set of all available power values (denoted by Π in [43]) can be chosen such that the corresponding rates are $\{0, \dots, S_{max}\}$. As required in [43], the set is compact. The TOCA algorithm is parametrized by positive numbers w, ϵ, \tilde{q} , and β . The algorithm chooses at each slot $m \geq 1$, the batch size s_{TOCA} such that

$$s_{TOCA}[m] = \min\left(\arg\min_{s \in \{0, \dots, S_{max}\}} \left\{ \beta c(s) - W[m]s \right\}, Q[m-1]\right),$$

where

$$W[m] = \mathbb{I}\{Q[m-1] \geq \tilde{q}\} \left[we^{w(Q[m-1]-\tilde{q})} + 2X[m-1] \right] + \mathbb{I}\{Q[m-1] < \tilde{q}\} \left[-we^{w(\tilde{q}-Q[m-1])} + 2X[m-1] \right].$$

We note that $s_{TOCA}[m] = 0$ if $W[m] \leq 0$. The sequence $X[m], m \geq 0$ is obtained from a *virtual* queue which evolves according to

$$X[m+1] = \max(X[m] - s_{TOCA}[m+1] + \epsilon \mathbb{I}\{Q[m] < \tilde{q}\}, 0) + A[m+1] + \epsilon \mathbb{I}\{Q[m] \geq \tilde{q}\}.$$

As in [43], let $\delta_{max} = \max(A_{max}, S_{max})$. Let $0 < \epsilon < \min(\lambda - s_l, s_u - \lambda)$, $w = \frac{\epsilon}{\delta_{max}^2} e^{\frac{-\epsilon}{\delta_{max}}}$, and $\tilde{q} = \frac{2}{w} \log(\beta)$. A sequence of policies γ_k is generated by choosing a sequence $\beta_k = \frac{1}{V_k}$, for a sequence $V_k \downarrow 0$. Then from [43, Corollary 2], we have that

$$\overline{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right), \overline{C}(\gamma_k) = c(\lambda) + \mathcal{O}(V_k).$$

Therefore, we obtain that for the sequence of policies γ_k , $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{\bar{C}(\gamma_k) - c(\lambda)}\right)\right)$. We note that $W[m]$ is a non-decreasing function of q , where $Q[m-1] = q$. Since $s_{TOCA}[m]$ is a non-decreasing function of $W[m]$, we have that $s_{TOCA}[m]$ is a non-decreasing function of $Q[m-1]$. However, we note that $s_{TOCA}[m]$ is stationary only with respect to a state which includes an additional state variable $X[m-1]$ and hence is not admissible. Therefore, the above bound is an upper bound to the optimal value of (4.4) and not TRADEOFF.

If for any subsequence $c_{c,k}$ of \mathcal{O}^u such that there exists a constant $0 < m \leq 1$ and a subsequence $c_{TOCA,k}$ of $\bar{C}(\gamma_k)$ such that $c_{TOCA,k} \leq c_{c,k}$ and $c_{TOCA,k} - c(\lambda) \geq m(c_{c,k} - c(\lambda))$, then we have that $Q^*(c_{c,k}) \leq Q(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{c_{TOCA,k} - c(\lambda)}\right)\right) = \mathcal{O}\left(\log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right)$.

We now obtain an asymptotic characterization for Case 3. The method used is the same as that summarized in Section 4.1.1, except that the geometric bound used is obtained from Lemma 4.3.7(TAIL-PROB-STATE-DEP-1). Furthermore, we will see that in the asymptotic regime \mathfrak{R} , the geometric bound reduces to a constant bound, which leads to the specific form of asymptotic lower bound for Case 3.

Lemma 4.3.13. For Case 3, given any sequence of admissible policies γ_k with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$.

Proof. Let us consider a particular policy γ in the sequence γ_k with $V_k = V$. We define $q_d = \sup\{q : \mathbb{E}S(q) \leq \lambda + \epsilon_V\}$, where $\epsilon_V > 0$ will be chosen later. We note that as $\mathbb{E}S(0) = 0$, the above set is non-empty. Suppose we assume that q_d is finite.

We note that by the admissibility of γ , $\forall q \in \{0, \dots, q_d\}$, $\mathbb{E}S(q) \leq \lambda + \epsilon_V$. Hence, using $d = \epsilon_V$, we have from Lemma 4.3.7(TAIL-PROB-STATE-DEP-1), for a $\bar{q} < q_d$:

$$\Pr\{Q \geq \bar{q} + 1\} \geq \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}+1} + \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}+1}\right) \left[\Pr\{Q \geq q_d + 1\} + \frac{1}{\epsilon_V} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \right].$$

Or

$$\Pr\{Q \leq \bar{q}\} \leq 1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}+1} - \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}+1}\right) \left(\frac{1}{\epsilon_V} \sum_{q=q_d+1}^{\infty} \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \pi(q)\right) \quad (4.6)$$

as $\Pr\{Q \geq q_d + 1\} \geq 0$. For brevity, let $D_t \triangleq -\left(\frac{1}{\epsilon_V} \sum_{q=q_d+1}^{\infty} \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \pi(q)\right)$. We note that D_t is positive, as for $q \geq q_d + 1$, $\mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] < -\epsilon_V$. Consider

the expression for D_t . We have that for $q \geq q_d + 1$,

$$\mathbb{E}[Q[m] - Q[m+1] | Q[m] = q] = \mathbb{E}S(q) - \lambda$$

is positive as $\mathbb{E}S(q) > \lambda + \epsilon_V$. We note that by definition, $c(s)$ is piecewise linear. Let m be the tangent of the angle between (i) the line passing through $(\lambda + 1, c(\lambda + 1))$ and $(\lambda, c(\lambda))$, and (ii) $l(s)$. Then $m \sum_{q=q_d+1}^{\infty} \pi(q) (\mathbb{E}S(q) - \lambda) \leq \sum_{q=q_d+1}^{\infty} \pi(q) [c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))]$. Furthermore from the convexity of $c(\cdot)$, linearity of $l(\cdot)$, and as $c(s) - l(s) \geq 0$, we have that

$$\sum_{q=q_d+1}^{\infty} \pi(q) [c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))] \leq \mathbb{E}_{\pi} [c(\mathbb{E}S(Q)) - l(\mathbb{E}S(Q))] \leq \mathbb{E} [c(S(Q)) - l(S(Q))] = V.$$

Therefore

$$D_t \leq \frac{V}{m\epsilon_V}.$$

Now, as in the proof of Lemma 4.3.5, we find a lower bound $\frac{\bar{q}}{2}$ on $\bar{Q}(\gamma)$ by finding the largest \bar{q} such that $Pr\{Q \leq \bar{q}\} \leq \frac{1}{2}$. A lower bound \bar{q}_1 to \bar{q} can be obtained by using the upper bound (4.6) on $Pr\{Q \leq \bar{q}\}$. Let \bar{q}_1 be the largest integer, if one exists, such that

$$1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}_1+1} - \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}_1+1}\right) \left(\frac{1}{\epsilon_V} \sum_{q=q_d+1}^{\infty} \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \pi(q)\right) \leq \frac{1}{2}.$$

Then $\bar{q}_1 \leq \bar{q}$. Then we have to find \bar{q}_1 such that

$$\begin{aligned} 1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}_1+1} + \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}_1+1}\right) D_t &\leq \frac{1}{2}, \\ \text{or } \frac{1 + 2D_t}{2 + 2D_t} &\leq \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{\bar{q}_1+1}, \\ \text{or } \left(1 + \frac{\epsilon_V}{\epsilon_a}\right)^{\bar{q}_1+1} &\leq \frac{2 + 2D_t}{1 + 2D_t} \end{aligned}$$

We note that if $q_d = \infty$, then $D_t = 0$. However, \bar{q}_1 satisfying the above inequality for finite q_d is a lower bound for \bar{q}_1 for $q_d = \infty$. Hence, we proceed with finding the above \bar{q}_1 . Let \bar{q}_2 be the largest integer such that

$$\left(1 + \frac{\epsilon_V}{\epsilon_a}\right)^{\bar{q}_2+1} \leq \frac{2}{1 + 2D_t}. \quad (4.7)$$

Then $\bar{q}_2 \leq \bar{q}_1$. From (4.7) and the upper bound $\frac{V}{m\epsilon_V}$ on D_t , if \bar{q}_3 is the largest integer such that

$$\left(1 + \frac{\epsilon_V}{\epsilon_a}\right)^{\bar{q}_3+1} \leq \frac{2}{1 + 2\frac{V}{m\epsilon_V}},$$

then $\bar{q}_3 \leq \bar{q}_2$. Or, we have that \bar{q}_3 is the largest integer such that

$$\bar{q}_3 + 1 \leq \log_{\left(1 + \frac{\epsilon_V}{\epsilon_a}\right)} \left(\frac{2}{1 + \frac{2V}{m\epsilon_V}} \right).$$

We note that, as $V \downarrow 0$, if $\frac{V}{\epsilon_V} \rightarrow \infty$, then the bound will be negative. We choose $\epsilon_V = aV$, where $a > \frac{2}{m}$. Then we obtain that

$$\bar{q}_3 \leq \log_{\left(1 + \frac{aV}{\epsilon_a}\right)} \left(\frac{2}{1 + \frac{2}{ma}} \right) - 1,$$

where the RHS is positive as $V \downarrow 0$. Therefore the maximum \bar{q}_3 is at least

$$\left\lfloor \log_{\left(1 + \frac{aV}{\epsilon_a}\right)} \left(\frac{2}{1 + \frac{2}{ma}} \right) - 1 \right\rfloor.$$

Since $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2} \geq \frac{\bar{q}_2}{2} \geq \frac{\bar{q}_3}{2}$ and $\log\left(1 + \frac{aV}{\epsilon_a}\right) = \Theta(V)$, we have that for the sequence of policies γ_k , $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$. \square

Remark 4.3.14. As for INTERVAL- μ CHOICE-2-3, we note that there is a set of queue lengths \mathcal{Q}_h , which occur with high probability. Let the drift in state q be $\Delta(q) \triangleq \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q]$. Then \mathcal{Q}_h is the set of queue lengths such that $\Delta(q) \rightarrow 0$, for $q \in \mathcal{Q}_h$ as $c_c \downarrow c(\lambda)$. Intuitively, we expect that the stationary probabilities of all queue lengths in \mathcal{Q}_h are equal. With ϵ_V chosen to be aV as in the above proof, we have that

$$\begin{aligned} Pr\{Q \leq \bar{q}\} &\leq \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + aV}\right)^{\bar{q}+1}\right) \times \text{a constant}, \\ &\approx \bar{q} \frac{aV}{\epsilon_a} \times \text{another constant}. \end{aligned}$$

This suggests that the constant stationary probability for queue lengths in \mathcal{Q}_h is $\mathcal{O}(V)$, and therefore we obtain the $\Omega\left(\frac{1}{V_k}\right)$ asymptotic lower bound in the above lemma.

Remark 4.3.15. We note that the sequence of randomized policies γ_{ϵ_k} in the proof of Lemma 4.3.4 is such that γ_{ϵ_k} is admissible and $\bar{C}(\gamma_{\epsilon_k}) - c(\lambda) = \epsilon_k$ and $\bar{Q}(\gamma_{\epsilon_k}) = \mathcal{O}\left(\frac{1}{\epsilon_k}\right)$ as $\epsilon_k \downarrow 0$.

This leads to the following asymptotic characterization of case 3. The proof is similar to that of Proposition 4.3.11.

Proposition 4.3.16. For Case 3, the optimal tradeoff curve $Q^*(c_{c,k}) = \Theta\left(\frac{1}{c_{c,k} - c(\lambda)}\right)$ as $c_{c,k} \downarrow c(\lambda)$, for the sequence $c_{c,k} = \bar{C}(\gamma_{\epsilon_k})$.

Case 1: We note that in Case 1, $Q^*(c_c)$ does not grow to infinity as c_c approaches $c(\lambda)$. In fact, the policy γ_u , which serves $S[m+1] = \min(Q[m], s_u)$, has the finite minimum average

queue length over all admissible policies which achieve an average service cost of $c(\lambda)$. First of all, we note that any admissible policy γ which has $\bar{C}(\gamma) = c(\lambda)$, will have $Pr\{S(q) > s_u\} = 0, \forall q$ and $Pr\{S[m] > s_u\} = 0, \forall m \geq 1$. For a given realization of the arrival process and the randomization of the batch sizes, let $q^*[m]$ and $q[m]$ be the evolution of the queue process under γ_u and γ respectively. Then we note that $q^*[m] \leq q[m], \forall m$, and therefore γ_u has the least average queue length over all policies which have their average service cost equal to $c(\lambda)$. Furthermore, $Q^*(c_c) \leq \bar{Q}(\gamma_u)$ for $c_c \geq c(\lambda)$.

For Case 1, we are only able to obtain a tight lower bound for a restricted case. We show that if $s_u = 1$, then for any sequence of non-idling deterministic $\gamma_k \in \Gamma_a$, for which $\bar{C}(\gamma_k) - c(\lambda) = V_k$, we have that $\bar{Q}(\gamma_k) = \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} + \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right)$. We note that if $s_u = 1$, then $\bar{Q}(\gamma_u) = \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2}$, from [19]. Thus as $V \downarrow 0$, we have that the asymptotic lower bound has $\bar{Q}(\gamma_u)$ as the limit point. Furthermore, we note that the asymptotic order matches with what that was derived for FINITE- μ CHOICE-1 in Chapter 2.

We present the lower bound on the average queue length in a series of steps. Consider a particular policy γ in the above sequence in the sequence γ_k , with $\bar{C}(\gamma) - c(\lambda) = V$.

Let $q_u \triangleq \sup\{q : s(q) \leq s_u\}$. Since we have restricted attention to non-idling deterministic γ and $s_u = 1$, we have that $s(q) = s_u = 1$ for $q \in \{1, \dots, q_u\}$. In the following lemma, we obtain an upper bound on $Pr\{Q > q_u\}$, which will be used to obtain an upper bound on q_u .

We note that from Assumption A1, $A_{max} > S_{max}$. Therefore starting from any queue length q , the queue length in the next slot is at most $q - A_{max}$ or $q + A_{max}$.

Lemma 4.3.17. For any non-idling deterministic $\gamma \in \Gamma_a$, if $q_u \geq A_{max}$, then for any $k \geq 0$,

$$Pr\{Q > q_u + k\} \leq \min\left(\rho_u^{\left\lceil \frac{q_u + k}{2A_{max}} \right\rceil}, 1\right),$$

where $\rho_u = \frac{\lambda}{s_u}$.

The proof is discussed in Appendix 4.H. The proof is very similar to that of Lemma 1 of Bertsimas [8] and [9] but with some slight modification. We note that $c(s) \leq c_1(s)$, where

$$c_1(s) = \begin{cases} c(s), & \text{for } s \in [0, s_u], \\ c(s_u) + \frac{c(S_{max}) - c(s_u)}{S_{max} - s_u} (s - s_u), & \text{for } s \in (s_u, S_{max}]. \end{cases}$$

We note that $c_1(s) = l(s)$ for $s \in [0, s_u]$. Then,

$$\begin{aligned} V &= \bar{C}(\gamma) - c(\lambda) \leq \mathbb{E}_\pi [c_1(s(Q)) - l(s(Q))], \text{ and,} \\ \mathbb{E}_\pi [c_1(s(Q)) - l(s(Q))] &= \sum_{q>q_u} \pi(q) [c_1(s(q)) - l(s(q))], \text{ and,} \\ \sum_{q>q_u} \pi(q) [c_1(s(q)) - l(s(q))] &\leq \sum_{q>q_u} \pi(q) [m(S_{max} - s_u)], \end{aligned}$$

where m is the tangent of the angle between $c_1(s)$ and $l(s)$ at s_u . Therefore, we have that

$$\begin{aligned} V &\leq m(S_{max} - s_u) \sum_{q>q_u} \pi(q), \text{ or,} \\ Pr \{Q > q_u\} &\geq \frac{V}{m(S_{max} - s_u)}. \end{aligned} \tag{4.8}$$

We now proceed to find an upper bound on q_u by combining the above lower bound on $Pr \{Q > q_u\}$ with the upper bound derived in Lemma 4.3.17.

We note that if $q_u < A_{max}$, then A_{max} is an upper bound on q_u . Suppose $q_u \geq A_{max}$, then from Lemma 4.3.17 and (4.8), we have that

$$\frac{V}{m(S_{max} - s_u)} \leq Pr \{Q > q_u\} \leq \rho_u^{\left\lceil \frac{q_u}{2A_{max}} \right\rceil} \leq \rho_u^{\frac{q_u}{2A_{max}}}.$$

Since $\rho_u < 1$, we have that

$$\begin{aligned} \left(\frac{1}{\rho_u} \right)^{\frac{q_u}{2A_{max}}} &\leq \frac{m(S_{max} - s_u)}{V}, \text{ or,} \\ q_u &\leq 2A_{max} \log_{\frac{1}{\rho_u}} \left[\frac{m(S_{max} - s_u)}{V} \right]. \end{aligned} \tag{4.9}$$

Lemma 4.3.18. The average queue length for the policy γ ,

$$\bar{Q}(\gamma) \geq \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \frac{S_{max} - s_u}{s_u - \lambda} \left(q_u Pr \{Q > q_u\} + \sum_{q=q_u}^{\infty} Pr \{Q > q\} \right).$$

Proof. We note that the policy γ is admissible. Squaring both sides of the evolution equation (4.1), taking expectations with respect to the stationary distribution, and simplifying we obtain that

$$2\mathbb{E}[Q(S - A)] = \mathbb{E}A^2 + \mathbb{E}S^2 - 2\mathbb{E}A\mathbb{E}S.$$

We note that $\mathbb{E}A^2 - \mathbb{E}A\mathbb{E}S = \sigma^2$. Since $s_u = 1$, we have that $\mathbb{E}S^2 = \sum_{s=1}^{S_{max}} \pi_s(s)s^2 \geq$

$$s_u \sum_{s=1}^{S_{max}} \pi_s(s) s = \lambda s_u.$$

$$\begin{aligned} 2\mathbb{E}Q(S - A) &\geq \sigma^2 + \lambda(s_u - \lambda), \text{ or,} \\ \sum_{q=0}^{\infty} \pi(q)q(\mathbb{E}S(q) - \lambda) &\geq \frac{\sigma^2}{2} + \lambda(s_u - \lambda), \end{aligned} \tag{4.10}$$

Or, we have that

$$\begin{aligned} \sum_{q=0}^{\infty} \pi(q)q(s_u - \lambda) + \sum_{q=q_u+1}^{\infty} \pi(q)q(S_{max} - s_u) &\geq \frac{\sigma^2}{2} + \lambda(s_u - \lambda), \\ (s_u - \lambda) \sum_{q=0}^{\infty} \pi(q)q &\geq \frac{\sigma^2}{2} + \lambda(s_u - \lambda) - (S_{max} - s_u) \sum_{q=q_u+1}^{\infty} \pi(q)q \end{aligned} \tag{4.11}$$

Simplifying the term $\sum_{q=q_u+1}^{\infty} \pi(q)q$, we have

$$\begin{aligned} \sum_{q=q_u+1}^{\infty} \pi(q)q &= \sum_{q=1}^{\infty} (q_u + q)\pi(q_u + q), \\ &= q_u \sum_{q=q_u+1}^{\infty} \pi(q) + \sum_{q=1}^{\infty} q\pi(q_u + q). \end{aligned} \tag{4.12}$$

Then from (4.11) we have that

$$\bar{Q}(\gamma) \geq \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \frac{S_{max} - s_u}{s_u - \lambda} \left(q_u Pr\{Q > q_u\} + \sum_{q=q_u}^{\infty} Pr\{Q > q\} \right).$$

□

We now use the upper bound on q_u from (4.9) and the upper bound on $Pr\{Q \geq q_u + k\}$ from Lemma 4.3.17, in the above lower bound on $\bar{Q}(\gamma)$ to obtain our final result.

Lemma 4.3.19. If $s_u = 1$, then for any sequence of non-idling, deterministic $\gamma_k \in \Gamma_a$ such that $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \mathcal{O}\left(V_k \log\left(\frac{1}{V_k}\right)\right)$.

Proof. We consider a γ in the sequence γ_k , with $\bar{C}(\gamma) - c(\lambda) = V$. Then, from Lemma 4.3.18 we have that

$$\bar{Q}(\gamma) \geq \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \frac{S_{max} - s_u}{s_u - \lambda} \left(q_u Pr\{Q > q_u\} + \sum_{q=q_u}^{\infty} Pr\{Q > q\} \right).$$

Consider $\sum_{q=q_u}^{\infty} Pr \{Q > q\}$. We have that

$$\sum_{q=q_u}^{\infty} Pr \{Q > q\} = \sum_{m=0}^{\infty} \sum_{k=0}^{2A_{max}-1} Pr \{Q > q_u + m2A_{max} + k\}. \quad (4.13)$$

From the proof of Lemma 4.3.17 we have that

$$\begin{aligned} Pr \{Q > q_u + m2A_{max} + k\} &\leq \rho_u^{\left\lceil \frac{m2A_{max}+k}{2A_{max}} \right\rceil} Pr \{Q > q_u\}, \\ &\leq \rho_u^{\left\lceil \frac{m2A_{max}}{2A_{max}} \right\rceil} Pr \{Q > q_u\}, \text{ since } \rho_u < 1, \\ &= Pr \{Q > q_u\} \rho_u^m. \end{aligned}$$

Substituting in (4.13), we have that

$$\begin{aligned} \sum_{q=q_u}^{\infty} Pr \{Q > q\} &\leq 2A_{max} Pr \{Q > q_u\} \sum_{m=0}^{\infty} \rho_u^m, \\ &= 2A_{max} Pr \{Q > q_u\} \frac{1}{1 - \rho_u}. \end{aligned} \quad (4.14)$$

Now we obtain an upper bound on $Pr \{Q > q_u\}$. We note that

$$\bar{C}(\gamma) - c(\lambda) = \sum_{q>q_u} \pi(q) [c(s(q)) - l(s(q))].$$

Then

$$m_u \sum_{q>q_u} \pi(q) [s(q) - s_u] \leq \bar{C}(\gamma) - c(\lambda) = V,$$

where m_u is the tangent of the angle made by the line through $(c(s_u + 1), s_u + 1)$ and $(c(s_u), s_u)$ with $l(s)$. We note that for $q > q_u$, $s(q) - s_u \geq 1$. Therefore

$$\sum_{q>q_u} \pi(q) = Pr \{Q > q_u\} \leq \frac{V}{m_u}. \quad (4.15)$$

Using the above upper bound in (4.14), we obtain that

$$\sum_{q=q_u}^{\infty} Pr \{Q > q\} \leq 2A_{max} \frac{V}{m_u(1 - \rho_u)}.$$

From (4.12), we also obtain that

$$q_u Pr \{Q > q_u\} \leq \frac{2A_{max}V}{m_u} \log_{\frac{1}{\rho_u}} \left[\frac{m(S_{max} - s_u)}{V} \right].$$

Then we have that

$$\bar{Q}(\gamma) \geq \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \frac{S_{max} - s_u}{s_u - \lambda} \left(\frac{2A_{max}V}{m_u} \log_{\frac{1}{\rho_u}} \left[\frac{m(S_{max} - s_u)}{V} \right] + 2A_{max} \frac{V}{m_u(1 - \rho_u)} \right).$$

Thus, for the sequence of policies γ_k , with $V_k \downarrow 0$, we obtain that

$$\bar{Q}(\gamma_k) = \frac{\sigma^2}{2(s_u - \lambda)} + \frac{\lambda}{2} - \mathcal{O} \left(V_k \log \left(\frac{1}{V_k} \right) \right).$$

□

Remark 4.3.20. We note that the above asymptotic lower bounds can be obtained for admissible policies, even for general holding costs. Suppose, the holding cost is $h(q)$ in state q , instead of the queue length q . We assume that $h(q)$ is a strictly increasing function of q . Then, instead of $\bar{Q}(\gamma)$ for an admissible policy, we are interested in the average holding cost $\bar{H}(\gamma) = \sum_{q=0}^{\infty} \pi(q)h(q)$. We assume that for admissible policies $\bar{H}(\gamma)$ is finite. We note that asymptotic lower bounds on $\bar{H}(\gamma)$ can be obtained quite easily, from the above results. Consider the random variable $h(Q)$ for a policy γ . Then we obtain a lower bound $\frac{\bar{h}}{2}$ on $\bar{H}(\gamma)$, where \bar{h} is the largest number such that $Pr \{h(Q) \leq \bar{h}\} \leq \frac{1}{2}$. If the inverse function h^{-1} of h exists, then we have that \bar{h} is the largest number such that $Pr \{Q \leq h^{-1}(\bar{h})\} \leq \frac{1}{2}$. We note that we have already obtained lower bounds \bar{q}_l to \bar{q} , where \bar{q} is the largest integer such that $Pr \{Q \leq \bar{q}\} \leq \frac{1}{2}$. Therefore, we obtain that $\bar{H}(\gamma) \geq \frac{h(\bar{q}_l)}{2}$.

4.3.5 Asymptotic characterization of admissible policies for TRADEOFF

We consider a sequence of $c_{c,k} \downarrow c(\lambda)$ for TRADEOFF. Let γ_k be any sequence of feasible policies for the sequence $c_{c,k}$. In this section, we obtain an asymptotic characterization of γ_k . Our approach is similar to that in Sections 2.3.3 and 3.2.3 for the state dependent M/M/1 model. However, we are unable to obtain asymptotic upper bounds. Since, only the asymptotic upper bounds depended on the order-optimality property, the bounds that we derive here hold for any sequence of feasible policies for the sequence $c_{c,k}$.

We first obtain the bounds $P_l \{A\}$ and $P_u \{A\}$ as in Section 3.2.3. We note that the elements of sets A are average service rates $\mathbb{E}S(q)$. Let $A \subseteq [0, s_l - \epsilon_V] \cup [s_u + \epsilon_V, S_{max}]$. Let $Q_A = \{q : \mathbb{E}S(q) \in A\}$. Proceeding as in the proof of Lemma 4.3.8, we have that $Pr \{Q \in Q_A\} \leq P_u(Q_A) = \frac{V}{m\epsilon_V}$. Then, we have that

$$P_u(Q_A) = \begin{cases} \frac{V^{1-\delta}}{ma}, & \text{if } \epsilon_V = aV^\delta, 0 \leq \delta < 1, \\ \frac{1}{ma}, & \text{if } \epsilon_V = aV. \end{cases} \quad (4.16)$$

Let $A = [s_l - \epsilon_V, s_u + \epsilon_V]$. Let Q_A be defined as before. Then using (4.16), we have that

$Pr \{Q \in Q_A\} \geq P_l(Q_A)$, where

$$P_l(Q_A) = \begin{cases} 1 - \frac{V^{1-\delta}}{ma}, & \text{if } \epsilon_V = aV^\delta, 0 \leq \delta < 1, \\ 1 - \frac{1}{ma}, & \text{if } \epsilon_V = aV. \end{cases} \quad (4.17)$$

We note that the above two bounds hold for cases 2 and 3.

Consider the sets $A_1 = [0, s_l - \epsilon_V]$, $A_2 = (s_l - \epsilon_V, \tilde{s}]$, $A_3 = (\tilde{s}, s_u + \epsilon_V]$, and $A_4 = [s_u + \epsilon_V, S_{max}]$. From (4.16) we have that

$$\begin{aligned} \pi(A_2) + \pi(A_3) &= 1 - \pi(A_1) - \pi(A_4) \geq 1 - \frac{V}{m\epsilon_V}, \text{ and} \\ (s_l - \epsilon_V)\pi(A_2) + \tilde{s}\pi(A_3) &\leq \lambda. \end{aligned}$$

We have that

$$\begin{aligned} \pi(A_3) &\leq \frac{\lambda - (s_l - \epsilon_V)\pi(A_2)}{\tilde{s}}, \text{ and,} \\ \pi(A_2) + \frac{\lambda - (s_l - \epsilon_V)\pi(A_2)}{\tilde{s}} &\geq 1 - \frac{V}{m\epsilon_V}. \end{aligned}$$

Therefore,

$$\begin{aligned} \pi(A_2) \left[\frac{\tilde{s} - (s_l - \epsilon_V)}{\tilde{s}} \right] &\geq \frac{\tilde{s} - \lambda}{\tilde{s}} - \frac{V}{m\epsilon_V}, \\ \pi(A_2) &\geq \frac{\tilde{s} - \lambda}{\tilde{s} - (s_l - \epsilon_V)} - \frac{\tilde{s}V}{m\epsilon_V(\tilde{s} - (s_l - \epsilon_V))}. \end{aligned}$$

The above lower bound can be non-negative only if $\tilde{s} > \lambda$. In the following, we set $\tilde{s} = \lambda + \epsilon'_V$, where $\lambda + \epsilon'_V$ is assumed to be less than or equal to $s_u + \epsilon_V$ and $\lambda + \epsilon_V$ for cases 2 and 3 respectively. Then we have that,

$$\pi(A_2) \geq \frac{\epsilon'_V}{\lambda - s_l + \epsilon'_V + \epsilon_V} - \frac{\tilde{s}V}{m\epsilon_V(\lambda - s_l + \epsilon'_V + \epsilon_V)}. \quad (4.18)$$

We have the following result.

Lemma 4.3.21. For any sequence of admissible policies γ_k , with $\overline{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, and $Q_A = \{q : \mathbb{E}S(q) \in A\}$ for a $A \subseteq [0, S_{max}]$, we have that

$$|Q_A| = \begin{cases} \Omega\left(\log\left(\frac{1}{V_k}\right)\right), & \text{for Case 2, if } A = [s_l - aV^\delta, s_u + aV^\delta], 0 \leq \delta < 1, a > 0, \\ \Omega\left(\log\left(\frac{1}{V_k}\right)\right), & \text{for Case 2, if } A = [s_l - aV^\delta, \lambda + \epsilon], 0 \leq \delta < 1, a > 0, \epsilon > 0, \\ \Omega\left(\frac{1}{V_k}\right), & \text{for Case 3, if } A = [0, \lambda + aV], a > 0. \end{cases}$$

Proof. Consider a particular policy γ in the above sequence with $\overline{C}(\gamma) - c(\lambda) = V$. Let us first

consider the case where $A = [s_l - \epsilon_V, s_u + \epsilon_V]$. From (4.17), we have that $P_l(Q_A) \geq 1 - \frac{V^{1-\delta}}{ma}$, by choosing $\epsilon_V = aV^\delta, 0 \leq \delta < 1$.

We proceed as follows for Case 2. Let $q_s \triangleq \sup \{q : \mathbb{E}S(q) < s_l - \epsilon_V\}$. Then, we have that $Pr \{Q \leq q_s\} \leq \frac{V^{1-\delta}}{ma}$. Suppose q_l is the largest integer such that

$$Pr \{Q \leq q_s\} \sum_{q=1}^{q_l} \frac{\rho^k}{\rho_d} \leq 1 - \frac{V^{1-\delta}}{ma},$$

where ρ and ρ_d are as in Lemma 4.3.5, then $q_l \leq |Q_{A_V}|$. We then obtain that $q_l = \Omega(\log(\frac{1}{V}))$.

Now we consider A of the form $[s_l - \epsilon_V, \tilde{s}]$, for $\tilde{s} = \lambda + \epsilon'_V \leq s_u + \epsilon_V$ as in (4.18). For this A , since the asymptotic lower bound on $|Q_A|$ is $\Omega(\frac{1}{V})$ and is obtained as in the previous case, we do not present the derivation here. For Case 2, we choose $\epsilon_V = aV^\delta, 0 \leq \delta < 1$, and $\epsilon'_V = \epsilon$, where $\epsilon > 0$ is such that $\lambda + \epsilon < s_u$. Then from (4.18), we have that

$$\pi(A) \geq \frac{\epsilon'}{\lambda - s_l + \epsilon' + aV^\delta} - \frac{\tilde{s}V^{1-\delta}}{m(\lambda - s_l + \epsilon + aV^\delta)}. \quad (4.19)$$

Now proceeding as for the case when $A = [s_l - \epsilon_V, s_u + \epsilon_V]$ above, we obtain that for $A = [s_l - \epsilon_V, \lambda + \epsilon]$, $|Q_A| = \Omega(\log(\frac{1}{V}))$.

We now consider Case 3. For Case 3, we consider A to be $[0, s_u + \epsilon_V]$ rather than $[s_l - \epsilon_V, s_u + \epsilon_V]$. We also choose ϵ_V to be aV . Then, we have that $P_l(Q_A) \geq 1 - \frac{1}{ma}$. Let $q_d \triangleq \sup \{q : \mathbb{E}S(q) \leq s_u + \epsilon_V\}$. We note that $Pr \{Q \leq q_d - 1\} = 1 - Pr \{Q \geq q_d\}$. From TAIL-PROB-STATE-DEP-2 Lemma 4.3.7, we have that

$$Pr \{Q \leq q_d - 1\} \leq \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{q_d}\right) - \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + \epsilon_V}\right)^{q_d}\right) \left[\frac{1}{\epsilon_V} \sum_{q=q_d+1} [\lambda - \mathbb{E}S(q)] \pi(q)\right].$$

Then proceeding as in the proof of Lemma 4.3.13, we have that $q_d = \Omega(\frac{1}{V})$. Therefore, $|Q_A| = \Omega(\frac{1}{V})$. \square

4.3.6 An asymptotic lower bound for the tradeoff problem (4.4)

In this section, using the asymptotic results for TRADEOFF, derived in Section 4.3.4, we derive lower bounds for the optimal value of (4.4) for a set of c_c , via the Lagrange dual of (4.4). We note that the same approach applies to the tradeoff problems in Chapter 5. The asymptotic results for (4.4) are derived only for cases 2 and 3.

For (4.4), from Section 4.2.1, we have that there exists a stationary optimal policy. Therefore, the

optimal value of (4.4) is equal to that of

$$\begin{aligned} & \underset{\gamma \in \Gamma_s}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && \bar{C}(\gamma) \leq c_c. \end{aligned} \quad (4.20)$$

However, in this chapter, the problem that we have considered is

$$\begin{aligned} & \underset{\gamma \in \Gamma_a}{\text{minimize}} && \bar{Q}(\gamma), \\ & \text{such that} && \bar{C}(\gamma) \leq c_c. \end{aligned} \quad (4.21)$$

We note that the optimal value of (4.20) is lower bounded by the optimal value of its Lagrange dual:

$$\max_{\beta \geq 0} \left[\min_{\gamma \in \Gamma_s} [\bar{Q}(\gamma) + \beta (\bar{C}(\gamma) - c_c)] \right]. \quad (4.22)$$

From Lemma 4.3.1, we have that for any $\beta \geq 0$, there exists an admissible policy γ_β^* which achieves the minimum for the problem $\min_{\gamma \in \Gamma_s} [\bar{Q}(\gamma) + \beta (\bar{C}(\gamma) - c_c)]$. In the following, we show that an asymptotic lower bound to the solution of (4.20) can be obtained using (4.22) from the asymptotic behaviour of the optimal solution of (4.21), for certain sequences of c_c as $c_c \downarrow c(\lambda)$.

We note that for any sequence $\beta \uparrow \infty$, it can be shown that $\bar{C}(\gamma_\beta^*) \downarrow c(\lambda)$. For the following analysis, we also consider the MDP:

$$\min_{\gamma \in \Gamma_s} [\bar{Q}(\gamma) + \beta (\bar{C}(\gamma) - c(\lambda))]. \quad (4.23)$$

We note that for a $\beta \geq 0$, γ_β^* is optimal for both (4.22) and (4.23). We have the following result.

Proposition 4.3.22. Suppose γ_k is any sequence of policies such that $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$ and

$$\bar{Q}(\gamma_k) = \begin{cases} \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right), & \text{for Case 2,} \\ \mathcal{O}\left(\frac{1}{V_k}\right), & \text{for Case 3.} \end{cases}$$

Let $c_{c,k}$ be any sequence such that

$$c_{c,k} = \begin{cases} \Theta(\bar{C}(\gamma_k) - c(\lambda)), & \text{for Case 2,} \\ \Theta\left((\bar{C}(\gamma_k) - c(\lambda))^2\right), & \text{for Case 3.} \end{cases}$$

Then, for the tradeoff problem (4.20), for the sequence $c_{c,k} \downarrow c(\lambda)$ we have that

$$Q^*(c_{c,k}) = \begin{cases} \Omega\left(\log\left(\frac{1}{c_{c,k}-c(\lambda)}\right)\right), & \text{for Case 2,} \\ \Omega\left(\frac{1}{c_{c,k}-c(\lambda)}\right), & \text{for Case 3.} \end{cases}$$

Proof. We consider Case 2 first. Let γ_k be any sequence of policies (admissible or otherwise), which is such that $\overline{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$ and $\overline{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$. From Lemma 4.3.10, we note that at least one such sequence exists. Let $\tilde{\beta}_k = \frac{1}{(\overline{C}(\gamma_k) - c(\lambda))}$. We show that $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda) = \mathcal{O}\left(\frac{1}{\tilde{\beta}_k^\delta}\right)$, where $0 \leq \delta < 1$. We proceed by contradiction. Suppose $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)$ is not $\mathcal{O}\left(\frac{1}{\tilde{\beta}_k^\delta}\right)$. Then $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)$ is $\omega\left(\frac{1}{\tilde{\beta}_k^\delta}\right)$. Therefore, the optimal value of (4.23), $\overline{Q}(\gamma_{\tilde{\beta}_k}^*) + \tilde{\beta}_k (\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)) = \omega\left(\tilde{\beta}_k^{1-\delta}\right)$. However, we note that the sequence of policies γ_k is such that $\overline{Q}(\gamma_k) + \tilde{\beta}_k (\overline{C}(\gamma_k) - c(\lambda)) = \mathcal{O}\left(\log\left(\tilde{\beta}_k\right)\right)$, which contradicts the optimality of the sequence $\gamma_{\tilde{\beta}_k}^*$. Therefore, $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda) = \mathcal{O}\left(\frac{1}{\tilde{\beta}_k^\delta}\right)$.

Consider a sequence of $c_{c,k} \downarrow c(\lambda)$ for (4.20). Suppose $c_{c,k}$ is such that $c_{c,k} - c(\lambda) = \Theta\left(\overline{C}(\gamma_k) - c(\lambda)\right)$. Then (4.22) can be bounded below as

$$\max_{\beta \geq 0} \min_{\gamma \in \Gamma_s} [\overline{Q}(\gamma) + \beta (\overline{C}(\gamma) - c_{c,k})] \geq \overline{Q}(\gamma_{\tilde{\beta}_k}^*) + \tilde{\beta}_k (\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)) - \tilde{\beta}_k (c_{c,k} - c(\lambda)).$$

We have that $\tilde{\beta}_k (\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)) \geq 0$. Since $c_{c,k} - c(\lambda) = \mathcal{O}\left(\overline{C}(\gamma_k) - c(\lambda)\right)$, we have that $\tilde{\beta}_k (c_{c,k} - c(\lambda)) = \mathcal{O}(1)$. Furthermore, since $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda) = \mathcal{O}\left(\frac{1}{\tilde{\beta}_k^\delta}\right)$, we have that $\overline{Q}(\gamma_{\tilde{\beta}_k}^*) = \Omega\left(\log\left(\tilde{\beta}_k\right)\right)$. Since $c_{c,k} - c(\lambda)$ is also $\Omega\left(\overline{C}(\gamma_k) - c(\lambda)\right)$ we have that

$$\overline{Q}(\gamma_{\tilde{\beta}_k}^*) + \tilde{\beta}_k (\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)) - \tilde{\beta}_k (c_{c,k} - c(\lambda)) = \Omega\left(\log\left(\frac{1}{c_{c,k} - c(\lambda)}\right)\right),$$

which provides an asymptotic lower bound for (4.20).

We now consider Case 3. Let γ_k be any sequence of policies, which is such that $\overline{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$ and $\overline{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. From Lemma 4.3.4, we note that at least one such sequence exists. Let $\tilde{\beta}_k = \frac{1}{(\overline{C}(\gamma_k) - c(\lambda))^2}$. We show that $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda) = \mathcal{O}\left(\frac{1}{\sqrt{\tilde{\beta}_k}}\right)$. We proceed by assuming that $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)$ is not $\mathcal{O}\left(\frac{1}{\sqrt{\tilde{\beta}_k}}\right)$. Then $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)$ is $\omega\left(\frac{1}{\sqrt{\tilde{\beta}_k}}\right)$. Then we have that the optimal value of (4.23), $\overline{Q}(\gamma_{\tilde{\beta}_k}^*) + \tilde{\beta}_k (\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda)) = \omega\left(\sqrt{\tilde{\beta}_k}\right)$. We note that the sequence of policies γ_k is such that $\overline{Q}(\gamma_k) + \tilde{\beta}_k (\overline{C}(\gamma_k) - c(\lambda)) = \mathcal{O}\left(\sqrt{\tilde{\beta}_k}\right)$, which contradicts the optimality of the sequence $\gamma_{\tilde{\beta}_k}^*$. Therefore, $\overline{C}(\gamma_{\tilde{\beta}_k}^*) - c(\lambda) = \mathcal{O}\left(\frac{1}{\sqrt{\tilde{\beta}_k}}\right)$.

Then, proceeding as in Case 2, for any sequence of $c_{c,k}$ such that $c_{c,k} - c(\lambda) = \Theta\left(\left(\overline{C}(\gamma_k) - c(\lambda)\right)^2\right)$, we have that for Case 3, the optimal value of (4.20) is $\Omega\left(\frac{1}{c_{c,k} - c(\lambda)}\right)$. \square

We note that since for $c'_{c,k} \in \mathcal{O}^u$, there exists an admissible optimal policy for (4.4), the asymptotic lower bounds obtained in Lemma 4.3.8 and Lemma 4.3.13 apply directly. The above Lagrange dual approach shows that the asymptotic lower bounds also apply to (4.4) for the sequences $c_{c,k}$ considered above. Since the sequence of admissible optimal policies for the sequence $c'_{c,k}$ satisfy the properties required for γ_k stated in the above proposition, we note that the set of $c_{c,k}$ for which the above lower bound holds also contains \mathcal{O}^u . However, we are unable to show that for any sequence $c_{c,k} \downarrow c(\lambda)$ the asymptotic lower bounds in the above proposition hold.

4.3.7 Asymptotic lower bounds for ergodic $(A[m], m \geq 1)$

We note that when the arrival process $(A[m], m \geq 1)$ is an ergodic batch arrival process, the optimal policy for the tradeoff problem (4.4) may not be stationary. But in this section, we consider the set of policies Γ_s , which are such that the batch size $S(q)$ used for service in slot m is a function only of the queue length $Q[m-1] = q$, and is independent of anything else. We note that $S(Q[m-1])$ could be a randomized function of $Q[m-1]$. The asymptotic lower bound presented here is significant, in that it complements the asymptotic upper bound obtained for Markov batch arrival processes in [45, Section 4.9] and [30].

We assume that $(A[m])$ is ergodic, so that almost surely

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M A[m] = \mathbb{E}A[1] = \lambda,$$

and $\lambda < S_{max}$, where S_{max} is the largest batch size which can be served, as defined before. We also assume that the arrival process $(A[m])$ is such that

NA1 : Let $\sigma[m-1] = (Q[0] = q_0, A[1] = a_1, Q[1] = q_1, A[2] = a_2, \dots, A[m-1] = a_{m-1}, Q[m-1] = q_{m-1})$. We assume that

$$\inf_{m \in \mathbb{Z}_+} \min_{\substack{\{a_1, \dots, a_{m-1}\} \\ \{q_1, \dots, q_{m-1}\}}} Pr \{A[m] = 0 | \sigma[m-1]\} = \nu_a > 0.$$

We restrict to policies $\gamma \in \Gamma_s$ for which the following limits exist

$$\lim_{m \rightarrow \infty} Pr \{Q[m] = q | Q[0] = q_0\} = \pi(q), \forall q \in \mathbb{Z}_+, \quad (4.24)$$

with $\sum_{q=0}^{\infty} \pi(q) = 1$. We note that for policies $\gamma \in \Gamma_s$ for which the above limits exist, also have

well defined $\pi_s(s), \forall s \in \{0, \dots, S_{max}\}$, where

$$\pi_s(s) = \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M Pr \{S[m] = s | Q[0] = q_0\}. \quad (4.25)$$

For such a policy γ , the average service cost is $\bar{C}(\gamma) = \sum_{s=0}^{S_{max}} \pi_s(s)c(s)$ and the average queue length is $\bar{Q}(\gamma)$ is defined as $\sum_{q=0}^{\infty} \pi(q)q$.

To obtain an asymptotic characterization of the tradeoff, we again restrict to a class Γ_a of admissible policies. However, we use a weaker definition⁵ of admissible policies, compared with the definition in Section 4.3.2. A policy $\gamma \in \Gamma_s$ is admissible if:

NG1 : the limits in (4.24) exist for γ ,

NG2 : γ is *mean rate stable* (see [45]), i.e., $\sum_{s=0}^{S_{max}} \pi_s(s)s = \lambda$,

NG3 : the average service rate $\mathbb{E}S(q)$ is a non-decreasing function of q for γ .

Then for any admissible policy γ , we have that $\bar{C}(\gamma) \geq c(\lambda)$, as before, by applying Jensen's inequality.

Let U_V be a random variable with support on $\{0, \dots, S_{max}\}$ and $\mathbb{E}U_V = \lambda + V$. Let $(U_V[m], m \geq 1)$ be an IID sequence with $U_V[m] \sim U_V$. Consider a particular policy γ_V for a $V > 0$, which chooses $S[m] = \min(Q[m-1], U_V[m])$. Using [34, Lemma 1] we have that the limit $\pi(q) = \lim_{m \rightarrow \infty} Pr \{Q[m] = q\}$ exists. We note that then the evolution of the queue can be written as

$$Q[m] = \max \left(Q[m-1] - U_V[m], 0 \right) + A[m], \text{ for } m \geq 1, \text{ with } Q[0] = q_0.$$

Let $Q'[1] = \max(q_0 - U_V[1], 0)$, and

$$Q'[m] = \max \left(Q'[m-1] + A[m-1] - U_V[m], 0 \right), \text{ for } m \geq 2.$$

We note that $Q[m] = Q'[m] + A[m], m \geq 1$. Since the sequence of random variables $\zeta[m] = A[m] - U_V[m+1], m \geq 1$ is ergodic and $\mathbb{E}\zeta[1] = \lambda - (\lambda + V) < 0$, from [12, Chapter 1, Theorem 7] we have that $\lim_{m \rightarrow \infty} Pr \{Q'[m] < \infty\} = 1$ and $\lim_{m \rightarrow \infty} Pr \{Q'[m] = q\} = \pi'(q)$ exists. Since $A[1] \leq A_{max}$ we have that $\lim_{m \rightarrow \infty} Pr \{Q[m] < \infty\} = 1$. Thus γ_V satisfies properties NG1 and NG2. We also note that by construction γ_V satisfies property NG3 and hence γ_V is admissible. We note that $\bar{C}(\gamma_V) \leq \mathbb{E}c(U_V[1])$. Now consider the sequence of admissible policies γ_V for a sequence

⁵The set of admissible policies in this section contains the set of admissible policies defined in the previous sections. This shows that the development of the $\log(\frac{1}{V})$ asymptotic lower bound can be obtained under weaker assumptions than what was assumed in the previous sections.

$V \downarrow 0$. Then it is always possible to choose ⁶ the distribution of U_V such that $\bar{C}(\gamma_V) = c(\lambda) + \mathcal{O}(V)$ as $V \downarrow 0$. Therefore, as before, $c(\lambda) = \inf_{\gamma \in \Gamma_a} \bar{C}(\gamma)$.

We note that $c(\lambda)$ is piecewise linear and three cases arise depending on the value of λ , as shown in Section 4.3.3. We define s_l , s_u , and the line $l(s)$ as in Section 4.3.3. For Case 1, we do not have a tight asymptotic lower bound. We now obtain an asymptotic lower bound which applies to Cases 2 and 3. We note that the bound is obtained by generalizing the proofs of Lemmas 4.3.5 and 4.3.8. We also note that this asymptotic lower bound holds even if $(A[m])$ is just stationary rather than ergodic. Ergodicity of $(A[m])$ was required for proving that $c(\lambda) = \inf_{\gamma \in \Gamma_a} \bar{C}(\gamma)$.

For some positive $\epsilon < s_l$, let $q_s \triangleq \inf \{q : \mathbb{E}S(q) \geq s_l - \epsilon\}$. We note that the proof of the following asymptotic lower bound also follows the methodology summarized in Section 4.1.1, except that the geometric bound on the stationary probability distribution for the ergodic process $Q[m]$ is obtained from the assumptions NA1 and NG3.

Lemma 4.3.23. For an ergodic arrival process $(A[m], m \geq 1)$, satisfying NA1, and for any sequence of admissible policies γ_k (satisfying NG1, NG2, and NG3), with $\bar{C}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$, for Cases 2 and 3.

Proof. For a particular policy γ in the above sequence with $V_k = V$, as in the proof of Lemma 4.3.8, we have that

$$\sum_q \pi(q) \left[c(\mathbb{E}S(q)) - l(\mathbb{E}S(q)) \right] \leq V,$$

since $c(s)$ is convex and $l(s)$ is linear. Therefore,

$$\sum_{q=0}^{q_s-1} \pi(q) \leq \frac{V}{m_1 \epsilon},$$

where m_1 is the tangent of the angle between (i) the line passing through $(s_l - 1, c(s_l - 1))$ and $(s_l, c(s_l))$ and (ii) $l(s)$.

As in the proof of Lemma 4.3.5, we have that $Pr\{S(q) > 0\} \geq \frac{s_l - \epsilon}{S_{max}}, \forall q \geq q_s$, since for $q \geq q_s$, $\mathbb{E}S(q) \geq s_l - \epsilon$. Now we relate the stationary probability $\pi(q), q \geq q_s$ to $\sum_{q=0}^{q_s-1} \pi(q)$. We have that for a $q \geq q_s$ and for every $m \geq 0$

$$Pr\{Q[m+1] < q | Q[0] = q_0\} = Pr\{Q[m] - S(Q[m]) + A[m+1] < q | Q[0] = q_0\},$$

⁶For small enough V , a distribution for U_V that gives mass to either (i) λ and $\lambda + 1$, if λ is an integer or (ii) $\lfloor \lambda \rfloor$ and $\lceil \lambda \rceil$, if λ is not an integer, can be chosen such that $\mathbb{E}U_V = \lambda + V$ and $\mathbb{E}c(U_V) = c(\lambda) + mV$, where m is the slope of the line joining $(\lambda, c(\lambda))$ and $(\lambda + 1, c(\lambda + 1))$ for (i) or the line joining $(\lfloor \lambda \rfloor, c(\lfloor \lambda \rfloor))$ and $(\lceil \lambda \rceil, c(\lceil \lambda \rceil))$ for (ii).

which can be written as

$$= \mathbb{E}_{S[1], Q[1], \dots, Q[m-1], S[m]} \left[\Pr \left\{ Q[m] - S(Q[m]) + A[m+1] < q \mid Q[0] = q_0, S[1], Q[1], \dots, Q[m-1], S[m] \right\} \right],$$

which is

$$\geq \mathbb{E}_{S[1], Q[1], \dots, Q[m-1], S[m]} \left[\Pr \{ Q[m] = q \mid Q[0] = q_0, S[1], Q[1], \dots, Q[m-1], S[m] \} \times \Pr \{ S(Q[m]) > 0 \mid Q[0] = q_0, S[1], Q[1], \dots, Q[m-1], S[m], Q[m] = q \} \times \Pr \{ A[m+1] = 0 \mid Q[0] = q_0, S[1], Q[1], \dots, Q[m-1], S[m], Q[m] = q, \{ S(Q[m]) > 0 \} \} \right].$$

We note that the batch size $S(Q[m])$ is chosen independently of the history of the queue length evolution. We also note that the evolution $(Q[0] = q_0, S[1], Q[1], \dots, Q[m-1], S[m], Q[m] = q)$ is equivalent to the evolution $\sigma[m] = (q_0, A[1], Q[1], A[2], Q[2], \dots, A[m], Q[m] = q)$. Furthermore, $A[m+1]$ is independent of the batch size $S[m+1] = S(Q[m])$ given $Q[m]$. Therefore, using NA1, the above lower bound can be written as

$$= \mathbb{E}_{S[1], Q[1], \dots, Q[m-1], S[m]} \left[\Pr \{ Q[m] = q \mid Q[0] = q_0, Q[1], S[1], \dots, Q[m-1], S[m] \} \times \Pr \{ S(Q[m]) > 0 \mid Q[m] = q \} \times \Pr \{ A[m+1] = 0 \mid \sigma[m] \} \right].$$

Then using the property NA1 and the above lower bound on $\Pr \{ S(q) > 0 \}$ for $q \geq q_s$ we have that $\Pr \{ Q[m+1] < q \mid Q[0] = q_0 \}$

$$\geq \frac{s_l - \epsilon}{S_{max}} \nu_a \mathbb{E}_{S[1], Q[1], \dots, Q[m-1], S[m]} \left[\Pr \{ Q[m] = q \mid Q[0] = q_0, Q[1], S[1], \dots, Q[m-1], S[m] \} \right],$$

$$= \Pr \{ Q[m] = q \mid Q[0] = q_0 \} \frac{s_l - \epsilon}{S_{max}} \nu_a.$$

Defining $\rho_d = \frac{s_l - \epsilon}{S_{max}} \nu_a$ and $\rho = 1 + \frac{1}{\rho_d}$, and proceeding as in the proof of Lemma 4.3.5, we have that for any non-negative k , $q = q_s + k$, and for every $m \geq 0$,

$$\Pr \{ Q[m] = q \mid Q[0] = q_0 \} \leq \Pr \{ Q[m+1] < q_s \mid Q[0] = q_0 \} \frac{\rho^k}{\rho_d}.$$

Therefore,

$$\frac{1}{M} \sum_{m=0}^{M-1} \Pr \{ Q[m] = q \mid Q[0] = q_0 \} \leq \frac{\rho^k}{\rho_d} \left[\frac{1}{M} \sum_{m=0}^{M-1} \Pr \{ Q[m] < q_s \mid Q[0] = q_0 \} + \frac{\Pr \{ Q[M] < q_s \mid Q[0] = q_0 \} - \Pr \{ Q[0] < q_s \mid Q[0] = q_0 \}}{M} \right].$$

Then as $M \rightarrow \infty$, since γ is admissible, we have that

$$\pi(q) \leq \frac{\rho^k}{\rho_d} \sum_{q=0}^{q_s-1} \pi(q).$$

Proceeding similarly as in the proof of Lemma 4.3.5 (from (4.31)), we can show that

$$\bar{Q}(\gamma) \geq \frac{1}{2} \left[\log_{\rho} \left[\frac{1}{2 \sum_{q=0}^{q_s-1} \pi(q)} \right] - 1 \right].$$

Since $\sum_{q=0}^{q_s-1} \pi(q) \leq \frac{V}{m_1 \epsilon}$, we have that for the sequence of policies γ_k , $\bar{Q}(\gamma_k) = \Omega \left(\log \left(\frac{1}{V_k} \right) \right)$. \square

We note that the above asymptotic lower bound is not tight for Case 3 for IID ($A[m]$) since we have a $\Omega \left(\frac{1}{V_k} \right)$ lower bound on $\bar{Q}(\gamma_k)$.

In [45, Theorem 4.12] and [30], it has been shown that if ($A[m], m \geq 1$) is Markov, i.e.,

$$Pr \{A[m+1] = a_{m+1} | A[m] = a_m, \dots, A[1] = a_1\} = Pr \{A[m+1] = a_{m+1} | A[m] = a_m\},$$

then for a sequence of Quadratic Lyapunov Algorithm (QLA) policies, parametrized by a sequence $V_k \downarrow 0$, the average queue length is $\mathcal{O} \left(\frac{1}{V_k} \right)$ for an average service cost V_k more than $c(\lambda)$.

We note that the QLA algorithm chooses a deterministic batch size $s(q)$ for service based on the current queue length only (unlike, say TOCA, for which the batch size is chosen as a function of other auxiliary variables also). Therefore, the QLA algorithm falls in the restricted class of admissible policies considered in this section. We note that the batch sizes are chosen deterministically, therefore we have that

$$\begin{aligned} Pr \{A[m+1] = a_{m+1} | A[m] = a_m, \dots, A[1] = a_1, Q[0] = q_0\} = \\ Pr \{A[m+1] = a_{m+1} | Q[m] = q_m, A[m] = a_m, Q[m-1] = q_{m-1}, \dots, A[1] = a_1, Q[0] = q_0\}, \end{aligned}$$

for the given policy, for ($q_n, 0 \leq n \leq m$), such that $q_{n+1} = q_n - s(q_n) + a_{n+1}$. Then we note that the above asymptotic lower bound applies, for the QLA algorithm, under the assumption

$$\begin{aligned} Pr \{A[m] = a_m | A[1] = a_1, A[2] = a_2, \dots, A[m-1] = a_{m-1}, Q[0] = q_0\} = \\ Pr \{A[m] = a_m | A[1] = a_1, A[2] = a_2, \dots, A[m-1] = a_{m-1}\}. \end{aligned}$$

4.4 Asymptotic bounds for R-model

As for the integer valued case, the tradeoff problem is to obtain $Q^*(c_c, q_0)$, which is the optimal value of

$$\underset{\gamma \in \Gamma}{\text{minimize}} \bar{Q}(\gamma, q_0) \text{ such that } \bar{C}(\gamma, q_0) \leq c_c,$$

where $c_c \geq 0$ is the average service cost constraint.

We recall that the above problem can again be formulated as a constrained Markov decision problem over the class of policies Γ . However, since the R-model is usually used as an approximation for the I-model, we analyse the tradeoff problem only for a restricted class of stationary admissible policies, where the definition of this class of admissible policies is motivated by the definition in Section 4.1.2. We assume that $\lambda < S_{max}$ as before.

Stability : A policy $\gamma \in \Gamma_s$ is said to be stable if : a) the Markov chain $Q[m]$ under γ is positive Harris recurrent with stationary distribution π_γ on the recurrence class corresponding to q_0 , and b) $\bar{Q}(\gamma, q_0) < \infty$.

Admissibility : In the following we restrict ourselves to the class of admissible policies Γ_a which is defined below. A policy γ is called admissible if:

RG1 : it is stable,

RG2 : it induces an aperiodic, irreducible Harris Markov chain $Q[m]$, and,

RG3 : the average service rate at a queue length q , $\mathbb{E}S(q)$ is non-decreasing in q .

We note that the properties RG1 and RG2 are similar to the properties of admissible policies used in Berry and Gallager [7]. The additional property RG3 is motivated by the monotonicity property derived in Agarwal et al. [1] as well as the monotonic non-decreasing property of any stationary deterministic optimal policy derived for the integer valued case in Section 4.1.2. We note that as in Section 4.1.2, for $\gamma \in \Gamma_a$, the average queue length and average service cost are independent of q_0 and are therefore denoted by $\bar{Q}(\gamma)$ and $\bar{C}(\gamma)$ respectively. We also note that $\bar{Q}(\gamma) = \mathbb{E}_{\pi_\gamma} Q$ and $\bar{C}(\gamma) = \mathbb{E}_{\pi_\gamma} c_R(S(Q))$.

We have that the TRADEOFF problem, for the R-model, is to obtain $Q^*(c_c)$ which is the optimal value of

$$\underset{\gamma \in \Gamma_a}{\text{minimize}} \bar{Q}(\gamma), \text{ such that } \bar{C}(\gamma) \leq c_c.$$

From the convexity of $c_R(s)$ and Jensen's inequality, we obtain that for any $\gamma \in \Gamma_a$, $\bar{C}(\gamma) \geq c_R(\lambda)$. Similar to the proof of Lemma 4.3.4, it can be shown that there exists a sequence of policies $\gamma_\epsilon \in \Gamma_a$

such that $\bar{C}(\gamma_\epsilon) \downarrow c_R(\lambda)$. Therefore, we obtain that $\inf_{\gamma \in \Gamma_a} \bar{C}(\gamma) = c_R(\lambda)$. In the next section, we obtain an asymptotic characterization of $Q^*(c_c)$ in the asymptotic regime \mathfrak{R} as $c_c \downarrow c_R(\lambda)$.

4.4.1 Asymptotic lower bound

We note that this problem can be considered as a special case of the tradeoff problem considered by Berry and Gallager [7], with the fade state taking only a single value. The contribution in this section is a step towards an alternative explanation for the Berry-Gallager lower bound, but with the extra property RG3 for an admissible policy. We note that asymptotic bounds for any order-optimal policy can be obtained under the additional assumption RG3. In Chapter 5, we present a lower bound for multiple fade states. As in [7] and as for the integer valued service case, we obtain lower bounds on the minimum average queue length for any sequence of admissible policies γ_k such that $\bar{C}(\gamma_k) - c_R(\lambda) = V_k \downarrow 0$. The asymptotic lower bound is obtained via a lower bound on the stationary probability for a Markov chain which evolves on \mathbb{R}_+ . The lower bound is similar to Lemma 4.3.7-TAIL-PROB-STATE-DEP-1 but for a Markov chain with state space \mathbb{R}_+ .

Lemma 4.4.1. Let $(Q[m])$ be the queue length evolution process for an admissible policy γ . Let ϵ_a and δ_a be as in assumption RA1. Suppose there exists a q_d such that

$$\forall q \in [0, q_d], \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d,$$

where d is positive. Then for any $q_1, k \geq 0, \Delta > 0, \delta > 0, \Delta + \delta < \delta_a$, and $0 \leq q_1 + k\Delta \leq q_d$, we have

$$Pr\{Q \geq q_1 + k\Delta\} \geq \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + d}\right)^k Pr\{Q \geq q_1\} + \left(1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + d}\right)^k\right) \left[Pr\{Q \geq q_d\} - \frac{1}{d} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q)\right].$$

The proof is presented in Appendix 4.I. We note that the development of the above lower bound on $Pr\{Q \geq q\}$ is an extension of the geometric lower bound on stationary probability for countable space DTMCs available in Bertsimas et al. [8] and [9], to the case of DTMCs on \mathbb{R}_+ with state dependent drift.

Using the above result, we derive the following asymptotic lower bound. Similar to assumption C2 in Chapter 3 and as in [7], we assume that the second derivative of $c_R(s)$ is positive at $s = \lambda$.

Proposition 4.4.2. For any sequence of policies $\gamma_k \in \Gamma_a$ with $\bar{C}(\gamma_k) - c_R(\lambda) = V_k \downarrow 0, \bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$.

Proof. Consider a particular policy γ in the sequence γ_k with $V_k = V$. Let $q_d = \sup\{q : \mathbb{E}S(q) \leq \lambda + \epsilon_V\}$, where ϵ_V will be chosen later. Suppose q_d is finite. From the admissibility of γ , we have that

$\forall q \in [0, q_d], \mathbb{E}S(q) \leq \lambda + \epsilon_V$. Using $d = \epsilon_V$ in Lemma 4.4.1, we have for a $\bar{q} = k\Delta \leq q_d, k \geq 0$,

$$Pr \{Q \geq \bar{q}\} \geq \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k + 1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k \left[Pr \{Q \geq q_d\} - \frac{1}{\epsilon_V} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right].$$

Or we have that

$$Pr \{Q < \bar{q}\} \leq \left(1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k \right) \left[1 + \frac{1}{\epsilon_V} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) - Pr \{Q \geq q_d\} \right]. \quad (4.26)$$

We note that for $q \in [q_d, \infty), \mathbb{E}S(q) - \lambda \geq \epsilon_V$. For brevity, we denote $\frac{1}{\epsilon_V} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q)$ by D_t . Now we note that for the policy $\gamma, \mathbb{E}_\pi c_R(S(Q)) - c_R(\lambda) = V$. Define $l(s)$ as the tangent to the curve $c_R(s)$ at $(\lambda, c_R(\lambda))$. Then we have that $\mathbb{E}_\pi \mathbb{E}_{S|Q} [c_R(S(Q)) - l(S(Q))] = V$. Now as $c_R(s)$ is convex and $l(s)$ is linear, using Jensen's inequality we have that $\mathbb{E}_\pi [c_R(\mathbb{E}S(Q)) - l(\mathbb{E}S(Q))] \leq V$. As in [7, step (41)], $c_R(s) - l(s) = G(s - \lambda)$ where $G(x)$ is a strictly convex function with $G(0) = 0, G'(0) = 0$, and $G''(0) > 0$. Thus we have that $\mathbb{E}_\pi G(\mathbb{E}S(Q) - \lambda) \leq V$. Using the sequence of steps (45), (46), (47), and (48) of Berry and Gallager [7], we obtain that

$$\left[\int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right]^2 \leq \frac{V}{a_1},$$

where $a_1 > 0$ is such that $G(x) \geq a_1 x^2$ (see Proposition 3.A.1), for $x \in [-\lambda, S_{max} - \lambda]$. We note that then $D_t \leq \frac{1}{\epsilon_V} \sqrt{\frac{V}{a_1}}$. Choosing $\epsilon_V = 4\sqrt{\frac{V}{a_1}}$ we obtain that $D_t \leq \frac{1}{4}$.

We note that $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2}$, where $\bar{q} = \sup \{q : Pr \{Q < q\} \leq \frac{1}{2}\}$. Using the upper bound (4.26) and non-negativity of $Pr \{Q \geq q_d\}$, if $\bar{q}_1 = k_1\Delta$ where k_1 is the largest integer such that

$$\left(1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^{k_1} \right) [1 + D_t] \leq \frac{1}{2},$$

then $\bar{q}_1 \leq \bar{q}$. Therefore, k_1 is such that

$$\begin{aligned} \frac{1 + 2D_t}{2 + 2D_t} &\leq \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^{k_1}, \\ \left(1 + \frac{\epsilon_V}{\delta\epsilon_a} \right)^{k_1} &\leq \frac{2 + 2D_t}{1 + 2D_t} \end{aligned}$$

Let k_2 be the largest integer such that

$$\left(1 + \frac{\epsilon_V}{\delta\epsilon_a} \right)^{k_2} \leq \frac{2}{1 + 2D_t}. \quad (4.27)$$

Then $k_2 \leq k_1$. We note that even if q_d is infinite, $k_2\Delta$ is a lower bound to \bar{q} , since D_t is 0 in that case. The rest of the proof holds irrespective of whether q_d is finite or infinite.

Then, from (4.27) and using the upper bound $\frac{1}{4}$ on D_t , if k_3 is the largest integer such that

$$\left(1 + \frac{\epsilon_V}{\delta\epsilon_a}\right)^{k_3} \leq \frac{2}{1 + \frac{1}{2}},$$

then $k_3 \leq k_2$. We obtain that k_3 is at least

$$\log_{\left(1 + \frac{\epsilon_V}{\delta\epsilon_a}\right)}\left(\frac{4}{3}\right) - 1.$$

Since $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\Delta k_1}{2} \geq \frac{\Delta k_2}{2} \geq \frac{\Delta k_3}{2}$, we have that $\bar{Q}(\gamma) \geq \frac{\Delta}{2} \left(\log_{\left(1 + \frac{\epsilon_V}{\delta\epsilon_a}\right)}\left(\frac{4}{3}\right) - 1\right)$. Since $\log\left(1 + \frac{\epsilon_V}{\delta\epsilon_a}\right) = \Theta\left(\sqrt{V}\right)$, we have that for the sequence γ_k as $V_k \downarrow 0$, $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$. \square

We note that the problem considered here is a special case of the Berry-Gallager problem (with a single fade state) with admissible policies. The upper bounds for the average queue length and average service cost for the TOCA policy from Neely [43], can be used to obtain asymptotic upper bounds for this problem. However, we note that these bounds hold only for the problem (4.4), since the sequence of TOCA policies is not admissible. Therefore, as in Lemma 4.3.10, we present a sequence of admissible policies which achieve the above asymptotic growth rate up to a logarithmic factor.

Lemma 4.4.3. Let a policy γ be defined as follows. At a queue length q , γ serves a batch size $\min(q, \tilde{s}(q))$, where

$$\tilde{s}(q) = \begin{cases} \lambda - \epsilon_V, & \text{for } 0 \leq q \leq q_v, \\ \lambda + \epsilon_V, & \text{for } q_v < q \leq 2q_v, \\ \lambda + \epsilon, & \text{for } 2q_v < q. \end{cases}$$

where $q_v > 0$ and $\lambda + \epsilon \leq S_{max}$. We obtain a sequence of policies γ_k by choosing ϵ_V and q_v from the sequence ϵ_{V_k} and q_{v_k} defined as follows. Let $\omega_k = \sqrt{V_k}$, where $V_k \downarrow 0$. Let $\epsilon_{V_k} = \omega_k A_{max}^2 e^{\omega A_{max}}$ and $q_{v_k} = \frac{1}{\omega_k} \log\left(\frac{1}{\epsilon_{V_k}^3}\right)$. Then we have that γ_k is a sequence of admissible policies, such that $\bar{C}(\gamma_k) - c(\lambda) = \mathcal{O}(V_k)$ and $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log\left(\frac{1}{V_k}\right)\right)$.

The proof of this lemma is given in Appendix 4.J and is motivated by and borrows ideas from the proof of the asymptotic upper bound for TOCA policies in [43].

The TOCA policy is the same as that in Remark 4.3.12 except that for every $m \geq 1$, we have that the service batch size $s_{TOCA}[m]$ is chosen as

$$s_{TOCA}[m] = \min\left(\arg \min_{s \in [0, S_{max}]} \left\{ \beta c(s) - W[m]s \right\}, Q[m-1]\right).$$

Then, from [43, Theorem 3 and Corollary 1], for the TOCA policy γ as above, for $\beta > S_{max}$,

$w = \frac{\epsilon}{\delta_{max}^2} e^{-\frac{\epsilon}{\delta_{max}}}$, $\epsilon = \frac{1}{\sqrt{\beta}}$, and $\tilde{q} = \frac{6}{w} \log\left(\frac{1}{\epsilon}\right)$, we have that

$$\begin{aligned}\bar{Q}(\gamma) &= \mathcal{O}\left(\sqrt{\beta} \log(\beta)\right), \\ \bar{C}(\gamma) &= c(\lambda) + \mathcal{O}\left(\frac{1}{\beta}\right).\end{aligned}$$

For a sequence of policies γ_k , generated by choosing $\beta_k = \frac{1}{V_k}$, for a sequence $V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log\left(\frac{1}{V_k}\right)\right)$ and $\bar{C}(\gamma) = c(\lambda) + \mathcal{O}(V_k)$. The above bound is an upper bound for the optimal solution of (4.4). Furthermore, we note that the above upper bound is an upper bound to $Q^*(c_{c,k})$ for a sequence $c_{c,k}$ as in Remark 4.3.12.

4.4.2 R-model with a piecewise linear cost function

We note that R-model with a strictly convex $c_R(s)$ is usually used as an approximation to I-model. Usually, the function $c_R(s)$ coincides with $c(s)$ for $s \in \{0, \dots, S_{max}\}$. But we find that there are differences in the asymptotic behaviour of $Q^*(c_c)$ for I-model and R-model. We note that $c_R(\lambda) \leq c(\lambda), \forall \lambda \in (0, S_{max})$ and $c_R(\lambda) < c(\lambda)$ for $\lambda \notin \{0, \dots, S_{max} - 1\}$. Furthermore, R-model suggests that $Q^*(c_c)$ increases to infinity for all $\lambda \in (0, S_{max})$ as $c_c \downarrow c_R(\lambda)$. However, for Case 1, we see that $Q^*(c(\lambda))$ is finite (note that $c_R(\lambda) < c(\lambda)$ in this case). For Case 3, with $c_R(\lambda) = c(\lambda)$ we have that $Q^*(c_c) = \Omega\left(\frac{1}{\sqrt{c_c - c(\lambda)}}\right)$ for the R-model, whereas for the I-model $Q^*(c_c) = \Omega\left(\frac{1}{c_c - c(\lambda)}\right)$. So R-model with a strictly convex $c_R(s)$ overestimates the behaviour of $Q^*(c_c)$ for Case 1 and underestimates $c(\lambda)$ and $Q^*(c_c)$ for Cases 2 and 3. In the following, we briefly outline a method to show that a better approximation for I-model, is R-model with a piecewise linear $c_R(s)$. The service cost function $c_R(s)$ is chosen as the lower convex envelope of the service cost function $c(s), s \in \{0, \dots, S_{max}\}$ for the I-model. With this choice of $c_R(s)$, the asymptotic behaviour of I-model and its approximation, R-model, is the same.

We consider Case 2 first. We define s_l, s_u , and the line $l(s)$ as in Section 4.3.3. Consider any sequence of admissible policies γ_k with $\bar{C}(\gamma_k) - c_R(\lambda) = V_k \downarrow 0$. Then we have that $\mathbb{E}c_R(\mathbb{E}S(Q)) - c_R(\lambda) \leq V_k$. For a particular policy γ in the sequence, $0 < \epsilon < s_l$, and $q_s \triangleq s_l - \epsilon$ we have that $Pr\{Q < q_s\} \leq \frac{V}{m\epsilon}$ as in the proof of Lemma 4.3.8. We assume that $Pr\{A[1] \leq \frac{\Delta}{2}\} = \epsilon'_a > 0$, for some $0 < \Delta < \lambda$. We note that for the R-model, the queue evolution is on \mathbb{R}_+ . We discretize \mathbb{R}_+ into a countable number of intervals $([0, \frac{\Delta}{2}], [\frac{\Delta}{2}, \Delta], \dots)$. Then the proof of Lemma 4.3.8 can be modified to show that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. A complete illustration of this proof technique is given in Lemma 5.7.2.

For Case 3, we proceed as in the proof of Proposition 4.4.2 by defining q_d to be $\sup\{q : \mathbb{E}S(q) \leq \lambda + \epsilon_V\}$, where ϵ_V is a function of V to be chosen in the following. We recall that $D_t \triangleq \frac{1}{\epsilon_V} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q)$.

For the policy γ we have that

$$\begin{aligned} \mathbb{E} [c_R(\mathbb{E}S(Q)) - l(\mathbb{E}S(Q))] &\leq V, \\ \int_{q_d}^{\infty} (c_R(\mathbb{E}S(q)) - l(\mathbb{E}S(q))) d\pi(q) &\leq V, \text{ or,} \\ \frac{1}{\epsilon_V} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) &\leq \frac{V}{m\epsilon_V}, \end{aligned}$$

where m is the tangent of angle made by the line passing through $(\lambda - 1, c_R(\lambda - 1))$ and $(\lambda, c_R(\lambda))$ with $l(s)$. Now we choose $\epsilon_V = \frac{4V}{m}$ to obtain that $D_t \leq \frac{1}{4}$. Then we proceed as in the proof of Proposition 4.4.2 to obtain that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{V_k}\right)$.

We note that, by construction R-model has the service cost function $c_R(\lambda) = c(\lambda)$. Furthermore, the asymptotic behaviour for R-model and I-model coincide for Cases 2 and 3. For Case 1, it can be shown that $\bar{Q}(\gamma_u)$ is finite for R-model, so that $Q^*(c_R(\lambda))$ is also finite. However, we do not have asymptotic lower bounds for this case.

4.5 Conclusions

In this chapter, we have obtained an asymptotic characterization of the tradeoff curve $Q^*(c_c)$ in the asymptotic regime \mathfrak{R} for a discrete time queueing model (I-model). This asymptotic characterization has been obtained using the insights obtained from the analysis of INTERVAL- μ CHOICE in Chapter 3. We also consider a real valued approximation (R-model) to I-model, and compare the asymptotic results which are obtained for R-model with that for I-model.

For I-model we observe that the cost function $c(s)$ as a function of the average service rate s is piecewise linear. Then as for INTERVAL- μ CHOICE-2-2, we have three cases. For Case 2, motivated by INTERVAL- μ CHOICE-2-2, we construct an upper bound to the stationary probability distribution for the queue length which is geometrically increasing, which leads to $\Theta\left(\log\left(\frac{1}{V}\right)\right)$ asymptotic growth for $Q^*(c_c)$ as $c_c \downarrow c(\lambda)$. We note that this geometric upper bound on the stationary probability distribution can be obtained in general, even for Cases 1 and 3. However, for Case 3, motivated by INTERVAL- μ CHOICE-2-3, we expect that there is a set of queue lengths with high probability for which the average drift $\mathbb{E}S(q) - \lambda \downarrow 0$ as $V \downarrow 0$. We then expect that the stationary probability of such queue lengths should be equal and $\mathcal{O}(V)$. This intuition leads us to a refined bound on the stationary probability of the queue length, obtained by extending the bounds available in Bertsimas et al. [9], from which we obtained the $\Theta\left(\frac{1}{V}\right)$ asymptotic growth for $Q^*(c_c)$ as $V \downarrow 0$.

We note that Case 1 is similar to INTERVAL- μ CHOICE-2-1, however we are unable to obtain asymptotic lower bounds except for the restricted case where $s_u = 1$ and for the set of non-idling admissible policies. A direct translation of the ideas from INTERVAL- μ CHOICE-2-1 is not possible,

since for the discrete time model we cannot obtain a dominating policy γ' as for INTERVAL- μ CHOICE-2-1.

We note that R-model is similar to INTERVAL- μ CHOICE-1, and as in Case 3 above, motivated by the observation that for INTERVAL- μ CHOICE-1, the stationary probability of queue lengths occurring with high probability should be equal and $\mathcal{O}(\sqrt{V})$ as $V \downarrow 0$, we obtain a $\Omega\left(\frac{1}{\sqrt{V}}\right)$ asymptotic lower bound on $Q^*(c_c)$. We comparing the asymptotic behaviour of $Q^*(c_c)$ between I-model and R-model in their respective asymptotic regimes in Section 4.4.2. We observe that R-model with a strictly convex $c_R(s)$ overestimates the behaviour of $Q^*(c_c)$ for Case 1 and underestimates $c(\lambda)$ and $Q^*(c_c)$ for Cases 2 and 3. Therefore, we conclude that a more appropriate real valued approximation to I-model should have a cost function $c_R(s)$ chosen as the piecewise linear lower convex envelope of $c(s)$.

Since in our approach, we obtain bounds on the stationary probability of the queue length, we are able to obtain asymptotic bounds on any sequence of order-optimal policies in the asymptotic regime \mathfrak{R} . These bounds provide intuition for the design of buffer-partitioning policies and are presented in Section 4.3.5. We also obtain that the minimum average queue length is $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ for Cases 2 and 3, when the arrival process $A[m]$ is ergodic.

In [45, Section 4.8], it is observed that the *drift plus penalty* algorithm idles for certain values of the queue length. However, in Lemma 4.3.2, we have obtained that any optimal policy γ_β^* should be non-idling. Hence, the drift plus penalty algorithm has to be modified to be non-idling, for the models considered in this chapter. So Lemma 4.3.2 can be thought of as providing theoretical motivation for the *place-holder* method in [45, Section 4.8].

In the rest of the thesis, we use the above results to obtain asymptotic characterizations of some resource tradeoff problems arising in point-to-point communication links.

Appendices

4.A Proof of Lemma 4.3.1

Our approach is to verify the hypotheses of the single Theorem in Sennott [67] by showing that the Assumptions (1), (2), and (3*) of Sennott [67] are satisfied. The assumptions in Sennott [67] are as follows (note that the notation is as in [67]):

1. For every state i and discount factor α , the optimal expected total discount cost $V_\alpha(i)$ is finite,
2. Let $h_\alpha(i) \triangleq V_\alpha(i) - V_\alpha(0)$. There exists a non-negative N such that $-N \leq h_\alpha(i)$, for all states i and discount factors α ,
3. There exists non-negative M_i , such that $h_\alpha(i) \leq M_i$, for every state i and discount factor α . Let the transition probability under action a , from state i to j be $P_{i,j}(a)$. Then for all i , $\sum_j P_{i,j}(a(i))M_j < \infty$ for an action $a(i)$ feasible in state i .

Assumption (3*) assumes that in addition, $\sum_j P_{i,j}(a)M_j < \infty$, for all a feasible in state i . Let $C(i, a)$ be the single stage cost at state i , when action a is taken. If Assumptions (1), (2) and (3*) hold, the Theorem [67] states that:

Theorem 4.A.1. There exists a constant g , which is independent of the state i , and a function $h(i)$ with $-N \leq h(i) \leq M_i$, such that

$$g + h(i) = \min_a \left\{ C(i, a) + \sum_j P_{i,j}(a)h(j) \right\}, i \geq 0.$$

A policy f that attains the minimum in the RHS of the above equation is average cost optimal, with optimal average cost g .

For proving that the Assumptions (1) and (3) are satisfied we use [67, Proposition 5(i)] which states that:

Proposition 4.A.2. Assume that the Markov decision process has a stationary policy f inducing an irreducible, ergodic Markov chain satisfying $\sum_i \pi_i C(i, f(i)) < \infty$ ([67, Proposition 4, Condition (i)]), where π_i is the stationary probability for state i under f . Then Assumptions (1) and (3) hold.

Consider the stationary deterministic policy γ_f which uses a batch size $s(q) = \min(q, S_{max})$ when the queue length is q . Then, from assumption A2 we have that from any $q > 0$ the state 0 can be reached, since $Pr \{A[1] = 0\} > 0$. From state 0, any state $q > 0$ can be reached, which follows

from assumptions A1 and A2. Therefore γ_f is irreducible. We note that state 0 is aperiodic, therefore the Markov chain under γ_f is also aperiodic.

We now verify the drift condition (10.13) in [37] by choosing $V(q) = \frac{q^2}{2(S_{max}-\lambda)}$ which is non-negative if $\lambda < S_{max}$. Let $p_{q,q'} \triangleq Pr \{Q[m+1] = q' | Q[m] = q, S[m+1] = s(q)\}$. We have that

$$\sum_{q'} p_{q,q'} (V(q') - V(q)) \leq -q + \frac{3S_{max}^2 + \lambda^2 + \sigma^2}{2(S_{max} - \lambda)}, \forall q \geq 0.$$

As the function $c(q) = q$ is near-monotone [37] and $\frac{3S_{max}^2 + \lambda^2 + \sigma^2}{2(S_{max} - \lambda)} < \infty$ we have from [37, Theorem 10.3] that the Markov chain under the policy γ_f is c -regular, implying that it is also positive recurrent with invariant distribution π . Then under the same policy we have that $\mathbb{E}_\pi c_\beta(q, s) < \infty$ as $c_\beta(q, s) \leq q + \beta c(S_{max})$. This implies that we have verified condition (i) of Proposition 4 of [67] and therefore from [67, Proposition 5] the first and third assumptions hold. We note that as $A[1] \leq A_{max}$, once Assumption (3) holds, Assumption (3*) is implied.

To verify Assumption (2) of [67] it is sufficient to show that the optimal discounted cost $V_\alpha(q)$ is non-decreasing in q for every discount factor $\alpha \in (0, 1)$. This proposition follows as a special case of [29, Lemma C.1] and therefore we claim that $V_\alpha(q)$ is non-decreasing in q without proof.

Then from [67, Theorem] there exists a stationary deterministic optimal policy γ_β^* with optimal average cost g_β^* satisfying the following ACOE:

$$g_\beta^* + J_\beta(q) = \min_{s \in \{0, \dots, \min(q, S_{max})\}} \{c_\beta(q, s) + \mathbb{E}J_\beta(q - s + A[1])\},$$

with the stationary optimal policy γ_β^* using a batch size $s_\beta^*(q)$ at queue length q satisfying

$$s_\beta^*(q) = \arg \min_{s \in \{0, \dots, \min(q, S_{max})\}} \{c_\beta(q, s) + \mathbb{E}J_\beta(q - s + A[1])\},$$

where $J_\beta(q)$ is the optimal relative value function.

4.B Proof of Lemma 4.3.2

Proof. The proof proceeds by contradiction. Let $(Q[m], m \geq 0)$ be the evolution of the queue process under any stationary deterministic optimal policy γ_β^* starting from initial state q_0 . We assume that γ_β^* is such that there exists a queue length $q_1 > 0$, $q_1 \in \mathcal{R}_\beta$, such that $s_\beta^*(q_1) = 0$. Then we present a perturbation to this policy, which leads to a history dependent policy $\tilde{\gamma}$, which has a smaller average queue length as well as average service cost, which contradicts the assumed optimality of γ_β^* . However, to compare $\bar{Q}(\tilde{\gamma})$ with $\bar{Q}(\gamma_\beta^*)$ and $\bar{C}(\tilde{\gamma})$ with $\bar{C}(\gamma_\beta^*)$, we identify a delayed renewal process (X_k) embedded in $(Q[m])$. For γ_β^* and $\tilde{\gamma}$, we associate different reward

processes with (X_k) and use the Renewal Reward theorem to obtain $\bar{Q}(\gamma_\beta^*)$, $\bar{C}(\gamma_\beta^*)$, $\bar{Q}(\tilde{\gamma})$, and $\bar{C}(\tilde{\gamma})$.

Consider the evolution $(Q[m])$. The first cycle X_1 of the renewal process (X_k) is defined as follows :

$$X_1 = \min\{m : Q[m] = q_1\}.$$

At slot X_1 , $Q[m]$ enters the recurrence class $\mathcal{R}_{\gamma_\beta^*}$ and does not leave $\mathcal{R}_{\gamma_\beta^*}$ again. We note that X_1 depends only on the initial state q_0 as $(Q[m])$ is Markov. If $T_q(q_1)$ is the random time taken to hit q_1 starting from q , then $X_1 = T_{q_0}(q_1)$. Furthermore for γ_β^* , $X_1 < \infty$.

Now we note that by definition $Q[X_1] = q_1$ and $S[X_1 + 1] = 0$. Let $T_{q_1}^S(1)$ be the smallest positive integer such that $S[X_1 + T_{q_1}^S(1) + 1] > 0$, i.e., there is service of at least one customer in the $(X_1 + T_{q_1}^S(1) + 1)^{th}$ slot. As $(Q[m])$ is Markov, the distribution of $T_{q_1}^S(1)$ given $Q[X_1] = q_1$, is independent of $(Q[m], m < X_1)$. Let $\mathcal{S}_1 \subset R_\beta$ be the set of states in which at least one customer is served. The distribution of $T_{q_1}^S(1)$ is the same as that of the smallest random time T^S , to hit \mathcal{S}_1 , starting from q_1 . Furthermore, we note that the queue length random variable $Q(1) = Q[X_1 + T_{q_1}^S(1)]$ is distributed as the state Q^S of the Markov chain at the random time T^S when \mathcal{S}_1 is hit, starting from q_1 . The second cycle X_2 is defined as

$$X_2 = T_{q_1}^S(1) + T_{Q(1)}(q_1).$$

Similarly, the k^{th} cycle

$$X_k = T_{q_1}^S(k-1) + T_{Q(k-1)}(q_1),$$

where $Q(k-1) = Q[X_{k-1} + T_{q_1}^S(k-1)]$. We note that $T_{q_1}^S(k), k \geq 1$ are all IID and have the same distribution as T^S . The random variables $Q(k), k \geq 1$ are all IID and have the same distribution as Q^S . Hence, the random variables $X_k, k \geq 1$ are independent and $X_k, k \geq 2$ are identically distributed. Thus (X_k) constitutes a delayed renewal process. From property O5, we have that $\mathbb{E}X_k < \infty$.

We now associate a queue cost and service cost process with the renewal process. In each cycle k , we define the queue cost as the cumulative expected queue length :

$$C_k^Q = \sum_{m=0}^{X_k-1} \mathbb{E}[Q[T_k + m] | Q[T_k] = q_1],$$

where $T_k = \sum_{j=1}^{k-1} X_j$. Using property O5, we have that $\forall k \geq 2$, $C_k^Q < \infty$ with probability 1, as the optimal policy is c_β -regular [37]. Furthermore in each cycle k , we define the service cost as the




Figure 4.2: Illustration of the queue evolution under policy γ_β^* and its perturbation $\tilde{\gamma}$. The first three renewal cycles X_1, X_2 , and X_3 are also shown.

cumulative expected service cost :

$$C_k^C = \sum_{m=1}^{X_k} \mathbb{E} [c(S[T_k + m]) | Q[T_k] = q_1].$$

Again using property O5, we have that $\forall k \geq 2, C_k^C < \infty$ with probability 1, as the optimal policy is c_β -regular.

We note that the both C_k^Q and C_k^C are dependent only on X_k . Then from the renewal reward theorem we have that

$$\begin{aligned} \bar{Q}(\gamma_\beta^*) &= \frac{\mathbb{E}C_2^Q}{\mathbb{E}X_2}, \\ \bar{S}(\gamma_\beta^*) &= \frac{\mathbb{E}C_2^C}{\mathbb{E}X_2}. \end{aligned}$$

In the following, we perturb γ_β^* to obtain the policy $\tilde{\gamma}$. The perturbation, as well as the renewal cycle embedded in $Q[m]$ are illustrated in Figure 4.2.

We note that, given $(S[1], S[2], \dots)$, $\bar{Q}(\gamma)$ and $\bar{C}(\gamma)$ do not depend on the order of service for customers. Therefore, we can assume that at least one customer, chosen to be served at $T_{k+1} + T_{q_1}^S(k)$ was present in the queue at T_{k+1} . Also, from the definition of $T_{q_1}^S(k)$, the system is idle in the slots $\{T_{k+1}, \dots, T_{k+1} + T_{q_1}^S(k)\}$. The perturbed policy $\tilde{\gamma}$ advances the service of one customer

served in the $(T_k + T_{q_1}^S(k) + 1)^{th}$ slot to the $(T_k + 1)^{th}$ slot, for every cycle k . The policy does not change the service batch size used at any other slot. To implement $\tilde{\gamma}$, the evolution of the queue under the unperturbed γ_β^* policy is simulated for the same arrival process, in order to ascertain the slots at which the service batch size is to be changed. The policy $\tilde{\gamma}$ serves one customer at the $(T_{k+1} + 1)^{th}$ slot, and serves one customer less at the $(T_{k+1} + T_{q_1}^S(k) + 1)^{th}$ slot. To obtain $\bar{Q}(\tilde{\gamma})$ and $\bar{C}(\tilde{\gamma})$, we use the same renewal process (X_k) , but associate different per-cycle queue and service cost costs. In cycle $k \geq 2$, the new per-cycle queue cost is the cumulative expected queue cost

$$\tilde{C}_k^Q = q_1 + \sum_{m=1}^{X_k-1} \mathbb{E}[Q[T_k + m] - 1 | Q[T_k] = q_1].$$

As $\tilde{C}_k^Q < C_k^Q$ using the renewal reward theorem we obtain that $\bar{Q}(\tilde{\gamma}) < \bar{Q}(\gamma^*)$.

For $\tilde{\gamma}$, the cumulative expected service cost for a cycle $k \geq 2$,

$$\tilde{C}_k^C = c(1) + \sum_{m=2, m \neq T_k + T_{q_1}^S(k-1) + 1}^{X_k} \mathbb{E}[c(S[T_k + m]) | Q[T_k] = q_1] + \mathbb{E}[c(S[T_k + T_{q_1}^S(k-1) + 1]) | Q[T_k] = q_1].$$

The difference $C_k^C - \tilde{C}_k^C$ is

$$c(0) + c(S[T_k + T_{q_1}^S(k-1) + 1]) - c(1) - c(S[T_k + T_{q_1}^S(k-1) + 1] - 1).$$

From convexity for any $s \geq 1$, $c(s) - c(s-1) \geq c(1) - c(0)$. Therefore, $\tilde{C}_k^C \leq C_k^C$ and, using the renewal reward theorem for $\tilde{\gamma}$, we obtain that $\bar{C}(\tilde{\gamma}) \leq \bar{C}(\gamma_\beta^*)$. Therefore γ_β^* cannot be optimal. Hence, $\forall q \in \mathcal{R}_{\gamma_\beta^*}$ such that $q > 0$, $s_\beta^*(q) > 0$. \square

4.C Proof of Lemma 4.3.4

Proof. Case 1 : We first consider the case when $s_l = 0$ and $s_l < \lambda < s_u$. We note that the stationary deterministic policy, $s(q) = \min(q, s_u)$ achieves the minimum average service cost $c(\lambda)$. Furthermore, the above policy has finite average queue length.

Case 2 : We now consider the case when $s_l \geq 1$ and $s_l < \lambda < s_u$. Let $\alpha \in (0, 1)$ be such that $\alpha s_l + (1 - \alpha) s_u = \lambda + \epsilon$, where $\epsilon > 0$ and $\lambda + \epsilon < s_u$. Let $S_{\alpha, m}$ be a sequence of IID random variables with distribution

$$S_{\alpha, m} \sim \begin{cases} s_l \text{ w.p. } \alpha, \\ s_u \text{ w.p. } 1 - \alpha. \end{cases}$$

Consider a stationary policy $S(Q[m]) = \min(Q[m], S_{\alpha, m})$. We note that as $c(\cdot)$ is monotonically

non-decreasing,

$$\mathbb{E}_\pi c(S(Q)) \leq \mathbb{E}_\pi c(S_{\alpha,1}) \leq c(\lambda + \epsilon) = c(\lambda) + m\epsilon,$$

where $m = \frac{c(s_u) - c(s_l)}{s_u - s_l}$. On squaring both sides of the evolution equation (4.1) and taking expectations with respect to the stationary distribution, we have that

$$\begin{aligned} \mathbb{E}_\pi Q[m+1]^2 &= \mathbb{E}_\pi [Q[m]^2 + S(Q[m])^2 + A[m+1]^2 - 2S(Q[m])Q[m] - 2S(Q[m])A[m+1] + 2Q[m]A[m+1]] \\ \text{which implies that} \quad &2\mathbb{E}_\pi [Q[m]S(Q[m]) - Q[m]A[m+1]] \leq \mathbb{E} [S_{\alpha,1}^2] + \mathbb{E} [A[1]^2]. \end{aligned}$$

Let $\Delta_m = S_{\alpha,m} - S(Q[m])$, we note that $\Delta_m \leq s_u$. Furthermore $\Delta_m = 0$ if $Q[m] > s_u$. Therefore $Q[m]S(Q[m]) = Q[m](S_{\alpha,m} - \Delta_m) \geq Q[m](S_{\alpha,m}) - s_u^2$. Hence we have that

$$\begin{aligned} 2\mathbb{E}_\pi [Q[m]S(Q[m]) - Q[m]A[m+1]] &\leq \mathbb{E} S_{\alpha,1}^2 + \mathbb{E} A[1]^2, \\ 2\mathbb{E}_\pi [Q[m](S_{\alpha,m} - A[m+1])] - 2s_u^2 &\leq \mathbb{E} S_{\alpha,1}^2 + \mathbb{E} A[1]^2, \\ \text{or as } m \rightarrow \infty, 2\epsilon \mathbb{E}_\pi Q &\leq 2s_u^2 + \mathbb{E} S_{\alpha,1}^2 + \mathbb{E} A[1]^2, \\ \mathbb{E}_\pi Q &\leq \frac{1}{2\epsilon} [2s_u^2 + \mathbb{E} S_{\alpha,1}^2 + \mathbb{E} A[1]^2]. \end{aligned}$$

Since $\mathbb{E}_\pi Q$ is thus finite, we note that the above policy is admissible. By choosing ϵ from a sequence of $\epsilon_k \downarrow 0$ as $k \uparrow \infty$, we obtain a sequence of policies $\gamma_{\epsilon_k} \in \Gamma_a$ such that $\overline{C}(\gamma_{\epsilon_k}) - c(\lambda) \leq m\epsilon_k$ and $\overline{Q}(\gamma_{\epsilon_k}) = \mathcal{O}\left(\frac{1}{\epsilon_k}\right)$. Therefore $\lim_{k \uparrow \infty} \overline{C}(\gamma_{\epsilon_k}) = c(\lambda)$.

Case 3 : We next consider the case when $1 \leq s_l = \lambda = s_u < S_{max}$. Let $\alpha \in (0, 1)$ be such that $\alpha\lambda + (1 - \alpha)(\lambda + 1) = \lambda + \epsilon$, where $0 < \epsilon < 1$. Again we let $S_{\alpha,m}$ be a sequence of IID random variables with distribution

$$S_{\alpha,m} \sim \begin{cases} \lambda \text{ w.p. } \alpha, \\ \lambda + 1 \text{ w.p. } 1 - \alpha. \end{cases} \quad (4.28)$$

Consider a stationary policy $S(Q[m]) = \min(Q[m], S_{\alpha,m})$. The rest of the proof for Case 3 is similar to that of Case 2, and we obtain that there exists a sequence of policies $\gamma_{\epsilon_k} \in \Gamma_a$ such that $\lim_{\epsilon_k \downarrow 0} \overline{C}(\gamma_{\epsilon_k}) = c(\lambda)$. We note that for a sequence of $\epsilon_k \downarrow 0$, we obtain a sequence of policies $\gamma_{\epsilon_k} \in \Gamma_a$ such that $\overline{C}(\gamma_{\epsilon_k}) - c(\lambda) \leq m\epsilon_k$ and $\overline{Q}(\gamma_{\epsilon_k}) = \mathcal{O}\left(\frac{1}{\epsilon_k}\right)$. For both cases 2 and 3, since $\overline{C}(\gamma_{\epsilon_k}) - c(\lambda) \leq m\epsilon_k$, by redefining ϵ_k we can obtain a sequence of policies γ_{ϵ_k} such that $\overline{C}(\gamma_{\epsilon_k}) - c(\lambda) = \epsilon_k$ and $\overline{Q}(\gamma_{\epsilon_k}) = \mathcal{O}\left(\frac{1}{\epsilon_k}\right)$. \square

4.D Proof of Lemma 4.3.5

Proof. For a policy $\gamma \in \Gamma_a$, let $q_s \triangleq \inf \{q : \mathbb{E}S(q) \geq s_l - \epsilon\}$ for a fixed positive $\epsilon < s_l$. Then we note that from assumption G3, for $q \geq q_s$, $\mathbb{E}S(q) \geq s_l - \epsilon$. In the following, for $q \geq q_s$, we find a lower bound on the probability that at least one customer is served, i.e., a lower bound on

$Pr \{S(q) > 0\}$, given $q \geq q_s$. We note that for $q \geq q_s$.

$$\begin{aligned}
\mathbb{E}S(q) &\geq s_l - \epsilon, \\
\sum_{s=0}^{S_{max}} s Pr \{S(q) = s\} &\geq s_l - \epsilon, \\
S_{max} Pr \{S(q) > 0\} &\geq s_l - \epsilon, \text{ or,} \\
Pr \{S(q) > 0\} &\geq \frac{s_l - \epsilon}{S_{max}}.
\end{aligned} \tag{4.29}$$

For the policy γ , we have that $\pi = \pi \mathbb{P}$ where \mathbb{P} is the transition probability matrix of the Markov chain under policy γ , i.e. $\mathbb{P}_{q_1, q_2} = Pr \{Q[m+1] = q_2 | Q[m] = q_1\}$. We note that for a $q \geq q_s$, from (4.29), we have that there is a positive probability of reaching a state less than q , starting from q in one step, i.e.

$$P(Q[m+1] < q | Q[m] = q) \geq \left(\frac{s_l - \epsilon}{S_{max}} \right) Pr \{A[1] = 0\}.$$

Let $\rho_d \triangleq \left(\frac{s_l - \epsilon}{S_{max}} \right) Pr \{A[1] = 0\}$. For $q \geq q_s$, from $\pi = \pi \mathbb{P}$, we have

$$\sum_{q'=0}^{q-1} \pi(q') \geq \pi(q) \sum_{q'=0}^{q-1} \mathbb{P}_{q, q'} \geq \pi(q) \rho_d.$$

For $q = q_s$, we have

$$Pr \{Q < q_s\} \geq \pi(q_s) \rho_d.$$

For $q = q_s + 1$, we have

$$\begin{aligned}
Pr \{Q < q_s\} + \pi(q_s) &\geq \pi(q_s + 1) \rho_d, \\
\text{or } Pr \{Q < q_s\} \left(1 + \frac{1}{\rho_d} \right) &\geq \pi(q_s + 1) \rho_d.
\end{aligned} \tag{4.30}$$

Proceeding similarly, we obtain that for $q = q_s + k, k \geq 0$

$$\pi(q) \leq Pr \{Q < q_s\} \frac{\left(1 + \frac{1}{\rho_d} \right)^k}{\rho_d} = Pr \{Q < q_s\} \frac{\rho^k}{\rho_d},$$

where $\rho = 1 + \frac{1}{\rho_d} > 1$.

Recall that in Chapter 3, we obtained a lower bound $\frac{\bar{q}}{2}$ on the average queue length for a policy γ , where \bar{q} was such that $Pr \{Q \leq \bar{q}\} \leq \frac{1}{2}$. We use the same idea here. Let $\bar{q} = \sup\{q : \sum_{q'=0}^q \pi(q') \leq$

$\frac{1}{2}$. Suppose $Pr\{Q < q_s\} \left(1 + \frac{\rho}{\rho_d}\right) < \frac{1}{2}$. Let \bar{q}_1 be the largest integer such that

$$Pr\{Q < q_s\} + Pr\{Q < q_s\} \sum_{q=q_s}^{\bar{q}_1} \frac{\rho^{q-q_s}}{\rho_d} \leq \frac{1}{2}. \quad (4.31)$$

Then $\bar{q}_1 \leq \bar{q}$. We note that (4.31) is equivalent to finding the largest \bar{q}_1 such that

$$Pr\{Q < q_s\} \left[1 + \frac{1}{\rho_d} \frac{\rho^{\bar{q}_1 - q_s + 1} - 1}{\rho - 1}\right] \leq \frac{1}{2},$$

or $\bar{q}_1 \leq \log_\rho \left[1 + \rho_d(\rho - 1) \left(\frac{1}{2Pr\{Q < q_s\}} - 1\right)\right]$.

Hence we obtain that the \bar{q}_1 is at least

$$\log_\rho \left[\frac{1}{2Pr\{Q < q_s\}}\right] - 1.$$

Since $\bar{Q}(\gamma) \geq \frac{\bar{q}}{2} \geq \frac{\bar{q}_1}{2}$, we have that

$$\bar{Q}(\gamma) \geq \frac{1}{2} \left[\log_\rho \left[\frac{1}{2Pr\{Q < q_s\}}\right] - 1\right].$$

□

4.E Proof of Lemma 4.3.7

TAIL-PROB :

Proof. The proof of TAIL-PROB follows from that of Lemma 2 and Theorem 3 (2) of Bertsimas et al. [8], and is presented here for completeness. Let us define $\widehat{Q}[m] = \max(q_1, Q[m])$ and $\widehat{Q} = \max(q_1, Q)$. Then we note that as q_1 is finite, $\mathbb{E}_\pi \widehat{Q} < \infty$. Hence, $\sum_{q=0}^{\infty} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \mid Q[m] = q \right] = 0$. We split this sum into three parts leading to :

$$0 = \sum_{q=0}^{q_1-1} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \mid Q[m] = q \right] \quad (4.32)$$

$$+ \pi(q_1) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \mid Q[m] = q_1 \right] \quad (4.33)$$

$$+ \sum_{q=q_1+1}^{\infty} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \mid Q[m] = q \right]. \quad (4.34)$$

Now, as in [8] we note that for $q \in \{0, \dots, q_1 - 1\}$, $\widehat{Q}[m] = q_1$ and $\widehat{Q}[m+1] \geq q_1$. Therefore (4.32) ≥ 0 . Also we note that for $q \in \{q_1 + 1, \dots\}$, we have that $\widehat{Q}[m] = Q[m]$ and $\widehat{Q}[m+1] \geq Q[m+1]$,

so that (4.33) is bounded below by

$$\sum_{q=q_1+1}^{\infty} \pi(q) \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d \sum_{q=q_1+1}^{\infty} \pi(q).$$

Now let us consider (4.33). We have that (4.33) =

$$\mathbb{E} \left[\widehat{Q}[m+1] \middle| Q[m] = q_1 \right] - q_1 \geq \mathbb{E} \max(A[m+1] - S(q_1), 0) \geq 1 \cdot \epsilon_a,$$

where assumption A2 is used. Substituting these lower bounds in (4.32) and (4.34) we have that

$$\begin{aligned} 0 &\geq \pi(q_1) \epsilon_a - d \sum_{q=q_1+1}^{\infty} \pi(q), \\ &= \epsilon_a \Pr \{Q \geq q_1\} - \epsilon_a \Pr \{Q \geq q_1 + 1\} - d \Pr \{Q \geq q_1 + 1\}. \end{aligned}$$

$$\text{Hence } \Pr \{Q \geq q_1 + 1\} \geq \frac{\epsilon_a}{\epsilon_a + d} \Pr \{Q \geq q_1\}.$$

By redefining $\widehat{Q}[m] = \max(q_1 + 1, Q[m])$ we obtain that

$$\Pr \{Q \geq q_1 + 2\} \geq \frac{\epsilon_a}{\epsilon_a + d} \Pr \{Q \geq q_1 + 1\}.$$

Induction leads to the following bound, for $k \geq 1$:

$$\Pr \{Q \geq q_1 + k\} \geq \left(\frac{\epsilon_a}{\epsilon_a + d} \right)^k \Pr \{Q \geq q_1\}.$$

□

TAIL-PROB-STATE-DEP-1 :

Proof. We note that in this case the lower bound on $\Pr \{Q \geq q_1 + k\}$ is obtained in terms of expected drift over the tail of the queue length. Let us again define $\widehat{Q}[m] = \max(q_1, Q[m])$ and $\widehat{Q} = \max(q_1, Q)$. Then we note that as q_1 is finite, $\mathbb{E}_\pi \widehat{Q} < \infty$. Hence, we have that $\sum_{q=0}^{\infty} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] = 0$. We again split this sum into three parts leading to :

$$0 = \sum_{q=0}^{q_1-1} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] \tag{4.35}$$

$$+ \pi(q_1) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q_1 \right] \tag{4.36}$$

$$+ \sum_{q=q_1+1}^{\infty} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right]. \tag{4.37}$$

As in the case of (4.32) and (4.33) we lower bound (4.35) by zero and (4.36) by $\pi(q_1) \epsilon_a$. Let us

consider the case $q_d > q_1$. Then (4.37) can be written as

$$\begin{aligned} & \sum_{q=q_1+1}^{q_d} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] | Q[m] = q \right] + \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] | Q[m] = q \right], \\ & \geq -d \sum_{q=q_1+1}^{q_d} \pi(q) + \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q], \end{aligned}$$

using the definition of q_1 and $\widehat{Q}[m]$. Using the lower bounds above on (4.35), (4.36), and (4.37) we obtain that

$$\begin{aligned} 0 & \geq \epsilon_a \pi(q_1) - d \sum_{q=q_1+1}^{q_d} \pi(q) + \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q], \\ & = \epsilon_a Pr \{Q \geq q_1\} - \epsilon_a Pr \{Q \geq q_1 + 1\} - d Pr \{Q \geq q_1 + 1\} + d Pr \{Q \geq q_d + 1\} \\ & \quad + \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q], \text{ or,} \end{aligned}$$

$$Pr \{Q \geq q_1 + 1\} \geq \frac{\epsilon_a}{\epsilon_a + d} Pr \{Q \geq q_1\} + \frac{d}{\epsilon_a + d} Pr \{Q \geq q_d + 1\} \quad (4.38)$$

$$+ \frac{1}{\epsilon_a + d} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q], \quad (4.39)$$

Redefining $\widehat{Q}[m] = \max(q_1 + 1, Q[m])$, where $q_1 + 1 < q_d$, we obtain that

$$\begin{aligned} Pr \{Q \geq q_1 + 2\} & \geq \frac{\epsilon_a}{\epsilon_a + d} Pr \{Q \geq q_1 + 1\} + \frac{d}{\epsilon_a + d} Pr \{Q \geq q_d + 1\} \\ & \quad + \frac{1}{\epsilon_a + d} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q]. \end{aligned}$$

Induction leads to the following bound for $q_1 + k \leq q_d$, $k \geq 1$,

$$\begin{aligned} Pr \{Q \geq q_1 + k\} & \geq \left(\frac{\epsilon_a}{\epsilon_a + d} \right)^k Pr \{Q \geq q_1\} \\ & \quad + \frac{1 - \left(\frac{\epsilon_a}{\epsilon_a + d} \right)^k}{\frac{d}{\epsilon_a + d}} \left[\frac{d}{\epsilon_a + d} Pr \{Q \geq q_d + 1\} + \right. \\ & \quad \left. \frac{1}{\epsilon_a + d} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q] \right]. \end{aligned}$$

We note that the bound also holds trivially for $k = 0$. Simplifying we obtain that for $q_1, k \geq 0$,

such that $q_1 + k \leq q_d$,

$$\begin{aligned} Pr \{Q \geq q_1 + k\} &\geq \left(\frac{\epsilon_a}{\epsilon_a + d} \right)^k Pr \{Q \geq q_1\} \\ &+ \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + d} \right)^k \right) \left[Pr \{Q \geq q_d + 1\} + \right. \\ &\left. \frac{1}{d} \sum_{q=q_d+1}^{\infty} \pi(q) \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \right]. \end{aligned}$$

□

TAIL-PROB-STATE-DEP-2 :

Proof. The derivation of this bound is very similar to that of TAIL-PROB-STATE-DEP-1. We follow the steps in the proof of TAIL-PROB-STATE-DEP-1 till (4.39) with the drift d replaced by d_1 . In this case, for $0 \leq q_1 < q_d$, (4.39) is further simplified to :

$$Pr \{Q \geq q_1 + 1\} \geq \frac{\epsilon_a}{\epsilon_a + d_1} Pr \{Q \geq q_1\} + \frac{d_1}{\epsilon_a + d_1} Pr \{Q \geq q_d + 1\} - \frac{d_2}{\epsilon_a + d_1} Pr \{Q \geq q_d + 1\},$$

since $\forall q \geq q_d + 1$, $\mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \geq -d_2$. Hence, we have that

$$Pr \{Q \geq q_1 + 1\} \geq \frac{\epsilon_a}{\epsilon_a + d_1} Pr \{Q \geq q_1\} - \frac{d_2 - d_1}{\epsilon_a + d_1} Pr \{Q \geq q_d + 1\}.$$

Again by induction as before, we obtain that for $q_1 + k \leq q_d$, $k \geq 1$,

$$Pr \{Q \geq q_1 + k\} \geq \left(\frac{\epsilon_a}{\epsilon_a + d_1} \right)^k Pr \{Q \geq q_1\} - \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + d_1} \right)^k \right) \frac{d_2 - d_1}{d_1} Pr \{Q \geq q_d + 1\}.$$

The above lower bound also holds trivially for $k = 0$. □

4.F An upper bound on $\bar{Q}(\gamma)$ for a policy γ

Let γ be such that at a queue length q , a batch size $\min(q, S(q))$ is served, where $S(q) \leq S_{max}$ is a random function of the current queue length. Furthermore, let γ be such that $\mathbb{E}S(q)$ is a monotonically non-decreasing function of q such that there exists a finite queue length q_1 such that $\mathbb{E}S(q_1) \geq \lambda + \epsilon$, where $\epsilon > 0$. Then we have the following upper bound on $\bar{Q}(\gamma)$.

Proposition 4.F.1. For a policy γ as above, we have that

$$\bar{Q}(\gamma) \leq q_1 \frac{\lambda + \epsilon}{\epsilon} + \frac{\mathbb{E}A[1]^2 + S_{max}^2}{\epsilon}.$$

Proof. The proof is very similar to that of 2.A.1. Let $L(q) = q^2$ be a Lyapunov function. Then the Lyapunov drift

$$\begin{aligned}\Delta(q) &= \mathbb{E}[L(Q[m+1]) - L(Q[m]) | Q[m] = q], \\ &\leq -2q(\mathbb{E}S(q) - \lambda) + \mathbb{E}A[1]^2 + S_{max}^2.\end{aligned}$$

We note that for $q \geq q_1$, since $\mathbb{E}S(q) \geq \lambda + \epsilon$, we have that

$$\Delta(q) \leq -2q(\epsilon) + \mathbb{E}A[1]^2 + S_{max}^2.$$

For $q < q_1$, we have that

$$\begin{aligned}\Delta(q) &\leq -2q(\epsilon) + 2q[\lambda + \epsilon - \mathbb{E}S(q)] + \mathbb{E}A[1]^2 + S_{max}^2, \\ &\leq -2q(\epsilon) + 2q_1(\lambda + \epsilon) + \mathbb{E}A[1]^2 + S_{max}^2.\end{aligned}$$

Hence, $\forall q$, we have that

$$\Delta(q) \leq -2q(\epsilon) + 2q_1(\lambda + \epsilon) + \mathbb{E}A[1]^2 + S_{max}^2.$$

Now applying [36, Theorem A.4.3] we have that

$$\bar{Q}(\gamma) \leq q_1 \frac{\lambda + \epsilon}{\epsilon} + \frac{\mathbb{E}A[1]^2 + S_{max}^2}{\epsilon}.$$

□

4.G Proof of Lemma 4.3.10

Let $L(q) \triangleq e^{\omega(q_v - q)}$ be a Lyapunov function. Since for the policy γ , the batch size $\tilde{S}(q)$ could be more than q , the queue evolution equation under γ is written as

$$Q[m+1] = \max(Q[m] - \tilde{S}(Q[m]), 0) + A[m+1].$$

The expected Lyapunov drift is

$$\Delta(q) \triangleq \mathbb{E}[L(Q[m+1]) - L(Q[m]) | Q[m] = q].$$

Since $(A[m])$ is assumed to be IID, as in [43, Lemma 5(a)] we have that

$$\Delta(q) \leq e^{\omega(q_v - q)} \left[\mathbb{E} e^{\omega(\tilde{S}(q) - A[1])} - 1 \right].$$

As in the proof of [43, Lemma 5(a)] we have that

$$\mathbb{E} e^{\omega(\tilde{S}(q) - A)} \leq 1 + \omega(\mathbb{E}\tilde{S}(q) - \lambda) + \frac{\omega^2 A_{max}^2}{2} e^{\omega A_{max}}.$$

Hence, we have that

$$\Delta(q) \leq \omega e^{\omega(q_v - q)} \left[(\mathbb{E}\tilde{S}(q) - \lambda) + K \right].$$

where $K = \frac{\omega A_{max}^2}{2} e^{\omega A_{max}}$.

Now by definition, the policy γ is such that

$$\tilde{S}(q) = \begin{cases} s_l, & \text{for } 0 \leq q < q_v, \\ s_u, & \text{for } q_v \leq q. \end{cases}$$

Then we have that for $q < q_v$

$$\Delta(q) \leq -\omega e^{\omega(q_v - q)} [\lambda - s_l - K].$$

And for $q \geq q_v$,

$$\begin{aligned} \Delta(q) &\leq \omega e^{\omega(q - q_v)} [(s_u - \lambda) + K], \\ &= -\omega e^{\omega(q_v - q)} [\lambda - s_l - K] + \omega e^{\omega(q_v - q)} [s_u - s_l], \\ &\leq -\omega e^{\omega(q_v - q)} [\lambda - s_l - K] + \omega [s_u - s_l], \end{aligned}$$

Hence, for all q we have that

$$\Delta(q) \leq -\omega e^{\omega(q_v - q)} [\lambda - s_l - K] + \omega [s_u - s_l].$$

We choose ω such that $K < \lambda - s_l$. Proceeding as in the proof of [43, Theorem 3(c)], we have that

$$\mathbb{E} e^{\omega(q_v - Q)} \leq \frac{[s_u - s_l]}{(\lambda - s_l - K)}.$$

Since $\mathbb{E}e^{\omega(q_v-Q)} \geq \mathbb{E}[e^{\omega(q_v-Q)}|Q < s_l] Pr\{Q < s_l\}$, we therefore have that

$$Pr\{Q < s_l\} \leq e^{-\omega q_v} \frac{e^{\omega s_l} [s_u - s_l]}{(\lambda - s_l - K)}. \quad (4.40)$$

Now we note that

$$\begin{aligned} \bar{C}(\gamma) &= \sum_{s < s_l} \pi_s(s)c(s) + \pi_s(s_l)c(s_l) + \pi_s(s_u)c(s_u), \\ &\leq Pr\{Q < s_l\} c(s_l - 1) + c(\lambda) (1 - \pi_s(s_l) - \pi_s(s_u)) + m [s_u \pi_s(s_u) + s_l \pi_s(s_l) - \lambda(\pi_s(s_u) + \pi_s(s_l))], \end{aligned}$$

where m is the slope of $l(s)$. Then, we have that

$$\bar{C}(\gamma) - c(\lambda) \leq [c(s_l - 1) + m\lambda - c(\lambda)] Pr\{Q < s_l\}.$$

Now consider the sequence of policies γ_k for which $q_v = \log\left(\frac{1}{V_k}\right)$ for a sequence $V_k < 1$ such that $V_k \downarrow 0$. Then we have that $\bar{C}(\gamma) - c(\lambda) = \mathcal{O}(V_k)$. Furthermore, from Proposition 4.F.1 we have that $\bar{Q}(\gamma) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$. We note that γ_k is also a sequence of admissible policies, since $s(q)$ is a non-decreasing function of q and $\bar{Q}(\gamma_k) < \infty$.

4.H Proof of Lemma 4.3.17

Proof. Define $\hat{Q}[m] \triangleq \max(q_u, Q[m])$ and $\hat{Q} = \max(q_u, Q)$ for a $q_u \geq 0$. Then for any finite q_u we have that $\mathbb{E}\hat{Q} < \infty$ and therefore

$$\sum_{q=0}^{\infty} \pi(q) \mathbb{E} \left[\hat{Q}[m+1] - \hat{Q}[m] \mid Q[m] = q \right] = 0.$$

We note that this can be written as

$$0 = \sum_{q \leq q_u - A_{max}} \pi(q) \mathbb{E} \left[\hat{Q}[m+1] - \hat{Q}[m] \mid Q[m] = q \right] \quad (4.41)$$

$$+ \sum_{q_u - A_{max} < q \leq q_u + A_{max}} \pi(q) \mathbb{E} \left[\hat{Q}[m+1] - \hat{Q}[m] \mid Q[m] = q \right] \quad (4.42)$$

$$+ \sum_{q_u + A_{max} < q} \pi(q) \mathbb{E} \left[\hat{Q}[m+1] - \hat{Q}[m] \mid Q[m] = q \right]. \quad (4.43)$$

Then we note that for (4.41) as $q \leq q_u - A_{max}$, $\hat{Q}[m] = q_u$ and $Q[m+1] \leq q_u$, so that $\hat{Q}[m+1] = q_u$. Therefore (4.41) = 0. Now consider (4.43). We note that for $q > q_u + A_{max}$,

both $\widehat{Q}[m]$ and $\widehat{Q}[m+1]$ are equal to $Q[m]$ and $Q[m+1]$ respectively. Therefore we have that

$$(4.43) \quad = \sum_{q_u + A_{max} < q} \pi(q) [\lambda - s(q)],$$

$$\leq -(s_u - \lambda) Pr \{Q > q_u + A_{max}\}.$$

We also note that for $q_u - A_{max} < q \leq q_u + A_{max}$, $\widehat{Q}[m+1] - \widehat{Q}[m] \leq A[m+1]$. Therefore,

$$(4.42) \leq \lambda Pr \{q_u - A_{max} < Q \leq q_u + A_{max}\}.$$

Using these upper bounds in (4.41), (4.42), and (4.43) we obtain that

$$0 \leq \lambda Pr \{q_u - A_{max} < Q \leq q_u + A_{max}\} - (s_u - \lambda) Pr \{Q > q_u + A_{max}\},$$

which can be written as

$$0 \leq \lambda Pr \{Q > q_u - A_{max}\} - s_u Pr \{Q > q_u + A_{max}\}, \text{ or,}$$

$$Pr \{Q > q_u + A_{max}\} \leq \frac{\lambda}{s_u} Pr \{Q > q_u - A_{max}\}. \quad (4.44)$$

Let $\rho_u \triangleq \frac{\lambda}{s_u}$. Using (4.44) and inducting we obtain that

$$Pr \{Q > q_u\} \leq \rho_u^{\left\lceil \frac{q_u}{2A_{max}} \right\rceil}, \quad (4.45)$$

for any $q_u \geq A_{max}$.

By redefining q_u to be $q_u + k$, $k \geq 0$, we can show that

$$Pr \{Q > q_u + k\} \leq \rho_u^{\left\lceil \frac{q_u + k}{2A_{max}} \right\rceil}.$$

□

4.1 Proof of Lemma 4.4.1

Proof. This proof is similar to that of the proof of Lemma 4.3.7-TAIL-PROB-STATE-DEP-1. We define $\widehat{Q}[m] = \max(q_1, Q[m])$ and $\widehat{Q} = \max(q_1, Q)$. We note that since the policy is admissible and q_1 is finite, $\mathbb{E}_\pi \widehat{Q} < \infty$. Therefore

$$\int_0^\infty \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] d\pi(q) = 0.$$

We again split the above integral into three parts which leads to

$$0 = \int_0^{q_1 - \Delta} \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] d\pi(q) \quad (4.46)$$

$$+ \int_{q_1 - \Delta}^{q_1} \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] d\pi(q) \quad (4.47)$$

$$+ \int_{q_1}^{\infty} \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] d\pi(q). \quad (4.48)$$

where $\Delta > 0$ will be chosen in the following. We note that for $q \in [0, q_1 - \Delta)$ we have that $\widehat{Q}[m] = q_1$ and $\widehat{Q}[m+1] \geq q_1$, so that (4.46) ≥ 0 . Now as $q_1 \leq q_d$ and $\widehat{Q}[m] \geq Q[m]$ we obtain that (4.48)

$$\geq -dPr\{q_1 \leq Q < q_d\} + \int_{q_d}^{\infty} \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q] d\pi(q)$$

To obtain a lower bound on (4.47) we note that for $q \leq q_1$, $\widehat{Q}[m] = q_1$ and $\widehat{Q}[m+1] \geq q_1$. So $\widehat{Q}[m+1] - \widehat{Q}[m] \geq 0$. Then as in [7, steps (34), (35), and (36)] we use Markov inequality to lower bound $\mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right]$, $q \in [q_1 - \Delta, q_1)$.

$$\begin{aligned} \mathbb{E} \left[\widehat{Q}[m+1] - \widehat{Q}[m] \middle| Q[m] = q \right] &\geq \delta Pr \left\{ \widehat{Q}[m+1] - \widehat{Q}[m] \geq \delta \middle| Q[m] = q \right\}, \\ &\geq \delta Pr \left\{ Q[m+1] - Q[m] \geq \delta + \Delta \middle| Q[m] = q \right\}, \\ &= \delta Pr \left\{ A[m+1] - S[m+1] \geq \delta + \Delta \middle| Q[m] = q \right\}, \\ &\geq \delta Pr \left\{ A[m+1] - S_{max} \geq \delta + \Delta \middle| Q[m] = q \right\}, \\ &\geq \delta \epsilon_a. \end{aligned}$$

We note that Δ and δ have to be chosen so that $\Delta + \delta < \delta_a$. Thus we obtain that (4.47) $\geq \delta \epsilon_a Pr \{q_1 - \Delta \leq Q < q_1\}$. Combining these bounds and using $\mathbb{E} [Q[m+1] - Q[m] | Q[m] = q] = \mathbb{E}S(q) - \lambda$, we obtain that

$$\begin{aligned} 0 &\geq \delta \epsilon_a Pr \{q_1 - \Delta \leq Q < q_1\} - dPr\{q_1 \leq Q < q_d\} + \int_{q_d}^{\infty} \mathbb{E} [Q[m+1] - Q[m] | Q[m] = q] d\pi(q), \\ Pr \{Q \geq q_1\} &\geq \frac{\delta \epsilon_a}{\delta \epsilon_a + \epsilon_V} Pr \{Q \geq q_1 - \Delta\} + \frac{1}{\delta \epsilon_a + d} \left[dPr \{Q \geq q_d\} - \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right] \end{aligned}$$

Similarly, if we define $\widehat{Q}[m] = \max(q_1 + \Delta, Q[m])$ and if $q_1 + \Delta \leq q_d$, we obtain that

$$Pr \{Q \geq q_1 + \Delta\} \geq \frac{\delta \epsilon_a}{\delta \epsilon_a + \epsilon_V} Pr \{Q \geq q_1\} + \frac{1}{\delta \epsilon_a + d} \left[dPr \{Q \geq q_d\} - \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right].$$

By induction, we obtain that if $k \geq 0$, and $q_1 + k\Delta \leq q_d$, then $Pr \{Q \geq q_1 + k\Delta\}$

$$\begin{aligned} &\geq \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k Pr \{Q \geq q_1\} + \frac{1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k}{d} \left[dPr \{Q \geq q_d\} - \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right], \\ &= \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k Pr \{Q \geq q_1\} + \left(1 - \left(\frac{\delta\epsilon_a}{\delta\epsilon_a + \epsilon_V} \right)^k \right) \left[Pr \{Q \geq q_d\} - \frac{1}{d} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda) d\pi(q) \right] \end{aligned}$$

□

4.J Proof of Lemma 4.4.3

Consider a particular policy γ in the sequence of policies γ_k . We note that the policy γ is a stationary deterministic policy. Then, we have that

$$\bar{C}(\gamma) = \int_0^{S_{max}} c_R(s) d\pi_s(s).$$

Since $c_R(s)$ is a non-decreasing function, we have that

$$\begin{aligned} \bar{C}(\gamma) &\leq Pr \{S < \lambda - \epsilon_V\} c(\lambda - \epsilon_V) + \pi_s(\lambda - \epsilon_V) c(\lambda - \epsilon_V) + \pi_s(\lambda + \epsilon_V) c(\lambda + \epsilon_V) + \pi_s(\lambda + \epsilon) c(\lambda + \epsilon). \\ &\leq Pr \{S < \lambda - \epsilon_V\} c(\lambda) + \pi_s(\lambda - \epsilon_V) \left(c(\lambda) - \epsilon_V \frac{dc(\lambda)}{ds} + \mathcal{O}(\epsilon_V^2) \right) \\ &\quad + \pi_s(\lambda + \epsilon_V) \left(c(\lambda) + \epsilon_V \frac{dc(\lambda)}{ds} + \mathcal{O}(\epsilon_V^2) \right) + \pi_s(\lambda + \epsilon) \left(c(\lambda) + \epsilon \frac{dc(\lambda)}{ds} + G(\epsilon) \right), \end{aligned}$$

where $G(\epsilon) = c(\lambda + \epsilon) - \left(c(\lambda) + \epsilon \frac{dc(\lambda)}{ds} \right)$.

Since there exists a finite queue length at which a service rate greater than λ is used, we have that

$$\begin{aligned} \int_0^{\lambda - \epsilon_V} s d\pi_s(s) + (\lambda - \epsilon_V) \pi_s(\lambda - \epsilon_V) + (\lambda + \epsilon_V) \pi_s(\lambda + \epsilon_V) + (\lambda + \epsilon) \pi_s(\lambda + \epsilon) &= \lambda, \text{ or,} \\ -\epsilon_V \pi_s(\lambda - \epsilon_V) + \epsilon_V \pi_s(\lambda + \epsilon_V) + \epsilon \pi_s(\lambda + \epsilon) &= \int_0^{\lambda - \epsilon_V} s d\pi_s(s) \leq \lambda Pr \{S < \lambda - \epsilon_V\}. \end{aligned}$$

Hence, we have that

$$\bar{C}(\gamma) \leq c(\lambda) + Pr \{S < \lambda - \epsilon_V\} \lambda \frac{dc(\lambda)}{ds} + G(\epsilon) \pi_s(\lambda + \epsilon) + \mathcal{O}(\epsilon_V^2).$$

We note that $\pi_s(\lambda + \epsilon) = Pr \{Q > 2q_v\}$ and $Pr \{S < \lambda - \epsilon_V\} = Pr \{Q < \lambda - \epsilon_V\}$. We proceed to find upper bounds on $Pr \{Q < \lambda - \epsilon_V\}$ and $Pr \{Q > 2q_v\}$.

We first obtain an upper bound on $Pr \{Q < \lambda - \epsilon_V\}$ as in the proof of Lemma 4.3.10. Using the

Lyapunov function $L(q) = e^{\omega(q_u - q)}$ we have that the Lyapunov drift is

$$\Delta(q) \leq \omega e^{\omega(q_v - q)} (\tilde{s}(q) - \lambda + K),$$

where $K = \frac{\omega A_{max}^2}{2} e^{\omega A_{max}}$. Then, for $q \leq q_v$, we have that

$$\Delta(q) \leq \omega e^{\omega(q_v - q)} (-\epsilon_V + K).$$

For $q_v < q \leq 2q_v$, we have that

$$\begin{aligned} \Delta(q) &\leq \omega e^{\omega(q_v - q)} (\epsilon_V + K), \\ &= \omega e^{\omega(q_v - q)} (-\epsilon_V + K) + 2\omega e^{\omega(q_v - q)} \epsilon_V, \\ &\leq \omega e^{\omega(q_v - q)} (-\epsilon_V + K) + 2\omega \epsilon_V. \end{aligned}$$

For $2q_v < q$, we have that

$$\begin{aligned} \Delta(q) &\leq \omega e^{\omega(q_v - q)} (\epsilon + K), \\ &= \omega e^{\omega(q_v - q)} (-\epsilon_V + K) + 2\omega e^{\omega(q_v - q)} (\epsilon + \epsilon_V), \\ &\leq \omega e^{\omega(q_v - q)} (-\epsilon_V + K) + 2\omega (\epsilon + \epsilon_V). \end{aligned}$$

Therefore, for every q , we have that

$$\Delta(q) \leq \omega e^{\omega(q_v - q)} (-\epsilon_V + K) + 2\omega (\epsilon + \epsilon_V).$$

Let $K = \epsilon_V/2$. Or we have that $\epsilon_V = \omega A_{max}^2 e^{\omega A_{max}}$. Then,

$$\Delta(q) \leq -\omega e^{\omega(q_v - q)} \frac{\epsilon_V}{2} + 2\omega (\epsilon + \epsilon_V).$$

Now as in the proof of Lemma 4.3.10, we have that

$$Pr \{Q < \lambda - \epsilon_V\} \leq \frac{2(\epsilon + \epsilon_V)}{\epsilon_V} e^{-\omega(q_v - \lambda + \epsilon_V)}.$$

Let us choose $q_v = \frac{1}{\omega} \log\left(\frac{1}{\epsilon_V}\right)$. Then we have that

$$Pr \{Q < \lambda - \epsilon_V\} = \mathcal{O}(\epsilon_V^2) = \mathcal{O}(\omega^2). \quad (4.49)$$

To obtain an upper bound on $Pr \{Q > 2q_v\}$, we proceed similarly but with a Lyapunov function $L(q) = e^{\omega(q - q_v)}$. The expected Lyapunov drift is

$$\Delta(q) \triangleq \mathbb{E} \left[e^{\omega(Q[m+1] - q_v)} - e^{\omega(Q[m] - q_v)} \mid Q[m] = q \right].$$

Then we have that

$$\Delta(q) = e^{-\omega q_v} \mathbb{E} \left[e^{\omega((q-\tilde{s}(q))^+ + A)} - e^{\omega q} \right].$$

As in [43], we have that

$$\Delta(q) \leq e^{-\omega q_v} \mathbb{E} \left[e^{\omega(q-\tilde{s}(q)+A)} - e^{\omega q} \right] + e^{-\omega q_v + \omega \delta_{max}},$$

where $\delta_{max} = A_{max} + S_{max}$. Or we have that

$$\begin{aligned} \Delta(q) &\leq e^{\omega(q-q_v)} \mathbb{E} \left[e^{-\omega(\tilde{s}(q)-A)} - 1 \right] + e^{-\omega q_v + \omega \delta_{max}}, \\ &\leq -\omega e^{\omega(q-q_v)} \mathbb{E} [(\tilde{s}(q) - A) - K] + e^{-\omega q_v + \omega \delta_{max}}, \\ &= -\omega e^{\omega(q-q_v)} [\tilde{s}(q) - \lambda - K] + e^{-\omega q_v + \omega \delta_{max}}, \end{aligned}$$

where $K = \frac{\omega A_{max}^2}{2} e^{\omega A_{max}}$.

For $q \leq q_v$, we have that

$$\begin{aligned} \Delta(q) &\leq -\omega e^{\omega(q-q_v)} [-\epsilon_V - K] + e^{-\omega q_v + \omega \delta_{max}}, \\ &\leq -\omega e^{\omega(q-q_v)} [\epsilon_V - K] + 2\omega \epsilon_V + e^{-\omega q_v + \omega \delta_{max}}. \end{aligned}$$

For $q > q_v$, we have that

$$\Delta(q) \leq -\omega e^{\omega(q-q_v)} [\epsilon_V - K] + e^{-\omega q_v + \omega \delta_{max}}.$$

So for all q , we have that

$$-\omega e^{\omega(q-q_v)} [\epsilon_V - K] + 2\omega \epsilon_V + e^{-\omega q_v + \omega \delta_{max}}.$$

Since $K = \epsilon_V/2$, we have that

$$\begin{aligned} \omega \frac{\epsilon_V}{2} \mathbb{E} \left[e^{\omega(q-q_v)} \right] &\leq 2\omega \epsilon_V + e^{-\omega q_v + \omega \delta_{max}}, \\ \mathbb{E} \left[e^{\omega(q-q_v)} \right] &\leq 4 + \frac{2e^{-\omega q_v + \omega \delta_{max}}}{\omega \epsilon_V}. \end{aligned}$$

Since $e^{\omega q_v} Pr \{Q > 2q_v\} \leq \mathbb{E} [e^{\omega(q-q_v)}]$, we have that

$$Pr \{Q > 2q_v\} \leq 4e^{-\omega q_v} + e^{-\omega q_v} \frac{2e^{-\omega q_v + \omega \delta_{max}}}{\omega \epsilon_V}.$$

Since $q_v = \frac{1}{\omega} \log \left(\frac{1}{\epsilon_V} \right)^3$. Then we have that

$$Pr \{Q > 2q_v\} \leq 4\epsilon_V^3 + \epsilon_V^4 A_{max}^2 2e^{\omega(\delta_{max} + A_{max})}.$$

Therefore, from the above upper bound and (4.49), we have that both $Pr \{Q < \lambda - \epsilon_V\}$ and $Pr \{Q > 2q_v\}$ are $\mathcal{O}(\epsilon_V^2)$. Since $\mathcal{O}(\epsilon_V) = \mathcal{O}(\omega)$, we have that both $Pr \{Q < \lambda - \epsilon_V\}$ and $Pr \{Q > 2q_v\}$ are $\mathcal{O}(\omega^2)$.

For the sequence of policies γ_k , we have that $\bar{C}(\gamma_k) - c(\lambda) = \mathcal{O}(\omega_k^2) = \mathcal{O}(V_k)$. Since $q_{v_k} = \mathcal{O}\left(\frac{1}{\omega_k} \log \left(\frac{1}{\epsilon_V^3}\right)\right)$, we have that $q_{v_k} = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log \left(\frac{1}{V_k}\right)\right)$. Therefore, using Proposition 4.F.1, with $q_1 = 2q_v$, we have that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log \left(\frac{1}{V_k}\right)\right)$. We note that the policy γ_k is admissible. Therefore, we have a sequence of admissible policies γ_k such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{\sqrt{V_k}} \log \left(\frac{1}{V_k}\right)\right)$. and $\bar{C}(\gamma_k) = \mathcal{O}(V_k)$.

On the tradeoff of average power and average delay for fading wireless links

5.1 Introduction

Minimizing the average power as well as the average delay is a major requirement in current wireless communication networks, which brings the problem of designing good scheduling and power control policies to the forefront. In this chapter, we consider the characterization of the optimal tradeoff between average power and average queue length for a fading point to point link, with and without admission control, and obtain bounds on the tradeoff of average power and average delay by applying Little's law. The models that we consider capture some of the important issues underlying the general problem for wireless networks: there is bursty arrival of traffic which can be subjected to admission control, the channel gain varies unpredictably, and the transmitter can dynamically change its transmission rate by varying the transmission power. The bounds that we derive are obtained using the methods discussed in Chapter 4, and are asymptotic in nature. However, unlike the models in Chapter 4, where there was only a single environment state and no admission control, here we consider models with multiple environment states as well as with admission control. We also consider the asymptotic characterization of the tradeoff for models with multiple queues in this chapter. The glossary of notation that we use in this chapter is given in Table 5.1.

Symbol	Description
m	slot index
$R[m]$	random number of arrivals in slot m (before admission control)
$A[m]$	random number of arrivals in slot m (after admission control)
A_{max}	maximum number of arrivals in any slot
λ, σ^2	mean and variance of $A[1]$
\mathcal{H}	set of fade states
π_H	distribution of fade state
$H[m]$	fade state in slot m
$Q[m]$	queue length at the start of $(m + 1)^{th}$ slot
$S[m]$	batch service size in slot m
S_{max}	maximum batch service size
$\sigma[m]$	history of queue evolution
γ	policy - $(S[1], S[2], \dots)$
Γ	set of all policies
Γ_s	set of all stationary policies
$\bar{Q}(\gamma)$	average queue length for a policy γ
$P(h, s)$	power expended as a function of fade state h and batch size s
$\bar{P}(\gamma)$	average power for a policy γ
P_c	average power constraint
$c(\lambda)$	minimum average power required for queue stability for I-model
$c_R(\lambda)$	minimum average power required for queue stability for R-model
Γ_a	set of all admissible policies
$\bar{s}(q)$	average service rate at a queue length q
$Q^*(P_c)$	minimum average queue length over Γ_a under constraint P_c
ϵ_a	probability of $A[m]$ exceeding S_{max}
$\bar{A}(\gamma)$	average throughput for a policy γ
π	stationary distribution for a policy which is clear from the context
π_γ	stationary distribution for policy γ

Table 5.1: Notation used in this chapter.

5.1.1 Methodology

As stated in Section 4.1.1, the scheme for obtaining asymptotic lower bounds in this chapter is very similar to that in Chapter 4. We briefly summarize the differences. We again obtain \bar{q} which is the largest q such that $Pr_u\{Q < q\} \leq \frac{1}{2}$. However, we note that geometric bounds for $Pr_u\{Q < q\}$ are obtained as functions of the average drift of the queue length, where the averaging is done over the slot fade state also. For the tradeoff problem in this chapter, again a non-negative function $D(q)$ will be obtained, where $\mathbb{E}_\pi D(Q)$ is the difference between the average power (rather than the average service cost) and $c(\lambda)$. For illustrating the method, we again consider the case of integer valued queue evolution and obtain a geometric upper bound $Pr_u\{Q < q\}$. The steps which were followed for obtaining the $\Omega(\log(\frac{1}{\bar{v}}))$ asymptotic lower bound for the example in Section 4.1.1 can then be directly applied to obtain an $\Omega(\log(\frac{1}{\bar{v}}))$ asymptotic lower bound for the tradeoff of average delay with average power.

5.1.2 System model - Integer valued queue length evolution

We consider a discrete time system with slots indexed by the positive integer m . We assume that there is no admission control, so that $A[m] = R[m], \forall m$. In each slot m , a random number of packets $A[m] \in \mathbb{Z}_+$, where each packet is of the same size, arrive into the transmitter queue. The arrival sequence $(A[m], m \geq 1)$ is assumed to be IID with $A[1] \leq A_{max}$, batch arrival rate $\mathbb{E}A[1] = \lambda < \infty$, $\text{var}(A[1]) = \sigma^2 < \infty$. The packets are assumed to arrive into an infinite buffer, in which they wait until they are transmitted over a point to point fading channel. The fade state is assumed to be constant in a slot. The fade state takes values in a finite set \mathcal{H} , with $\min\{\mathcal{H}\} > 0$, and the fade state process $(H[m], m \geq 1)$, is assumed to be IID, with $H[1] \sim \pi_H$. The expectation with respect to π_H is denoted by \mathbb{E}_{π_H} . The processes $(A[m])$ and $(H[m])$ are assumed to be independent of each other.

The number of customers in the queue at the start of the $(m+1)^{th}$ slot is denoted by $Q[m]$. The system is assumed to start with $Q[0] = q_0 \in \mathbb{Z}_+$ customers. At the end of slot m , a batch with $S[m] \in \mathbb{Z}_+$ packets is removed from the transmitter queue just before the $A[m]$ new packets which arrive in the m^{th} slot are admitted. We assume that $S[m] \leq \min(Q[m-1], S_{max})$, where S_{max} is the maximum batch size that can be served. The queue evolution sampled at the slots is given by:

$$Q[m+1] = Q[m] - S[m+1] + A[m+1]. \quad (5.1)$$

The evolution of the queue length is illustrated in Figure 5.1.

At the start of slot m , the history of the system is defined as:

$$\sigma[m] \triangleq (q_0, H[1], S[1], Q[1], H[2], S[2], Q[2], \dots, Q[m-2], H[m-1]).$$




Figure 5.1: Evolution of the queue length $Q[m]$; the batch size $S[m]$ is chosen as a randomized function of the fade state $H[m]$, the queue length $Q[m - 1]$, and the history of the process $\sigma[m]$.

At the beginning of slot m , the scheduler observes $H[m]$ and chooses a batch service size $S[m] \in \mathbb{Z}_+$ as a randomized function of the history $\sigma[m]$, the current queue length $Q[m - 1]$, and the current fade state $H[m]$.

We define a policy γ to be the sequence of service batch sizes $(S[1], S[2], \dots)$. The set of all policies is denoted by Γ . If γ is such that $S[m + 1] = S(Q[m], H[m + 1])$, then γ is a stationary policy. The set of all stationary policies is denoted as Γ_s . In this chapter, we restrict attention to Γ_s in light of the discussion in Section 4.2.1, which also holds for this model. We also note that since $H[m]$ is assumed to be IID, the process $(Q[m], m \geq 0)$ is a Markov chain, if $\gamma \in \Gamma_s$.

The transmitter expends $P(h, s)$ units of power when transmitting s bits, when the fade state is h . We note that $P(h, s)$ is a function of the fading gain h^2 , when the fade state is h . Motivated by many examples (see [7] and [77]) of $P(h, s)$, we assume that $\forall h \in \mathcal{H}$, $P(h, s)$ satisfies the following properties:

C1 : $P(h, 0) = 0$, and

C2 : $P(h, s)$ is non-decreasing and convex in s , for $s \in \{0, \dots, S_{max}\}$.

The average power for a policy γ is

$$\bar{P}(\gamma, q_0) \triangleq \limsup_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M P(H[m], S[m]) \middle| Q[0] = q_0 \right]. \quad (5.2)$$

The average queue length for a policy γ is

$$\bar{Q}(\gamma, q_0) \triangleq \limsup_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=0}^{M-1} Q[m] \middle| Q[0] = q_0 \right]. \quad (5.3)$$

We consider the optimal tradeoff of $\bar{P}(\gamma, q_0)$ with $\bar{Q}(\gamma, q_0)$ for this model. The optimal tradeoff between $\bar{P}(\gamma, q_0)$ and average delay can be obtained from this using Little's law and is discussed in the following. We note that I-model considered in Chapter 4 is a special case of this model, with a

single fade state.

5.1.3 System model - Real valued queue length evolution

We state only the differences from the model discussed in the previous section. We assume that for $m \geq 1$, $A[m] \in [0, A_{max}]$, $S[m] \in [0, S_{max}]$, and $q_0 \in \mathbb{R}_+$. Hence, the queue length $Q[m] \in \mathbb{R}_+$, $\forall m \geq 0$. The function $P(h, s)$ is assumed to satisfy the following properties:

RC1 : $P(h, 0) = 0$, for every $h \in \mathcal{H}$,

RC2 : $P(h, s)$ is an increasing, strictly convex function in s , for $s \in [0, S_{max}]$, for every $h \in \mathcal{H}$.

The average power and average queue length are as in (5.2) and (5.3) respectively.

The tradeoff of average power and average queue length, for such models without admission control, has been studied by Berry and Gallager [7], Neely [43], Goyal et al. [29], Bettesh and Shamaï [10], Biyikoglu et al. [77] as well as many others. We note that this model is similar to that considered by Berry and Gallager [7], except that in [7], $S_{max} = \infty$. We recall that Berry and Gallager obtain that any sequence of policies γ_k , for which $\bar{P}(\gamma_k)$ is at most V_k more than the above minimum power required for queue stability, has $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$, as $V_k \downarrow 0$. We also note this model is similar to R-model considered in Chapter 4, except that in Chapter 4 we considered the case with a single fade state.

In the following, the model in the previous section, where the queue length evolution is assumed to be on the non-negative integers, is called the I-model, while the model described here, where the queue length evolution is assumed to be on the non-negative real numbers, is called the R-model. We note that as in Chapter 4, R-model with fading and $P(h, s)$ being strictly convex is usually used as an approximation for I-model. When used as an approximation, $P(h, s)$ for R-model coincides with $P(h, s)$ for I-model for $s \in \{0, \dots, S_{max}\}$, $\forall h$.

5.1.4 An example

Throughout this chapter, to illustrate the results for I-model and R-model, we use the following example. The number of packets $A[1]$ which arrive in a slot is assumed to be distributed according to a Binomial(A_{max}, p) distribution, with arrival rate $r_a kb/s$. The rate of service r in kb/s is assumed to be $200 \log_{10}(1 + SNR)$, where SNR is the received signal to noise ratio. We assume that $SNR = \frac{h^2 P}{L}$, where h^2 is the fading gain, P is the transmit power, and L encompasses the loss due to attenuation as well as noise power.

We assume that the slots are of duration $2ms$. We also assume that each packet has a size of 100-bits. Then the arrival rate of packets in a slot is $\lambda = \frac{r_a}{50}$. We assume that if $h^2 = 1$ and

$P = 1W$ then $r = 50$. Therefore, if $P(h, r)$ is the transmit power as a function of the fade state and the rate, we have that $P(h, r) = \frac{1.28}{h^2} (10^{r/200} - 1)$. We note that in one slot, the number of bits served is $2r$. We assume that the transmitter, in each slot, can choose its transmission rate in the set $\{0, 50, 100\}$ kb/s . To fit this example to our model, we express the queue length in units of 100 bits. Then in each slot, we have a Binomial arrival process of 100-bit packets and service of s 100-bit packets, where $s \in \{0, 1, 2\}$. The transmit power as a function of h and s is $P(h, s) = \frac{1.28}{h^2} (10^{50s/200} - 1)$, $s \in \{0, 1, 2\}$. We note that the average queue length, as defined, is in units of 100 bits.

For R-model, we assume the same distribution for $A[1]$. The set of possible batch sizes is assumed to be $\in [0, 2]$. The transmit power as a function of h and s is assumed to be $P(h, s) = \frac{1.28}{h^2} (10^{50s/200} - 1)$ but for $s \in [0, 2]$.

The above example uses a similar model for the rate as a function of the SNR as the example in Section 2.1.3. However, in this chapter, we consider a slot level model, which models the system on a faster time scale compared to the continuous time model in Chapter 2. Consideration of the faster time scale is necessary since we are interested in the average delay advantage that can be achieved by scheduling the packets in accordance with the channel variations in each slot.

5.1.5 Overview

As in the previous chapter, the objective in this chapter is to obtain an asymptotic characterization of the minimum average queue length as the average power is a small V more than the minimum average power required for stability. We formulate the tradeoff problem for I-model and R-model in Section 5.2 for a set of admissible policies, whose definition is similar to that in Chapter 4. We also obtain the infimum of the average power over the set of admissible policies, which is also the minimum average power required for stability of the queue, and discuss its properties in the same section. The asymptotic analysis of the tradeoff problem is then carried out in Section 5.3. For I-model, we show that depending on the value of λ , three cases arise, which are similar to those for the I-model in Chapter 4. For the first case, through numerical experiments, we show that the minimum average queue length does not increase to infinity for an admissible policy which achieves the infimum of the average power over the set of admissible policies. For the second and third cases, as in Chapter 4, we show that the minimum average queue length increases as $\Theta(\log(\frac{1}{V}))$ and $\Theta(\frac{1}{V})$ as $V \downarrow 0$. For R-model, we show that the minimum average queue length is $\Omega(\frac{1}{\sqrt{V}})$. We note that this is the same as the Berry-Gallager lower bound, but the set of admissible policies that we consider is a subset of the set of admissible policies considered in [7]. For the example in Section 5.1.4, we provide some numerical results to illustrate the bounds in Section 5.5. As in Chapter 4, we also obtain an asymptotic $\log(\frac{1}{V})$ lower bound on the minimum average queue length for I-model, when $(A[m])$ and $(H[m])$ are assumed to be ergodic in Section 5.6.

We then consider queueing models, I-model-U and R-model-U, which are similar to I-model and R-model, but with admission control, in Section 5.7. The tradeoff problems for I-model-U and R-model-U, are formulated in Section 5.7.2 and its asymptotic analysis is carried out in Section 5.7.3. We show that the minimum average queue length increases as $\Theta\left(\log\left(\frac{1}{V}\right)\right)$ when the average service cost is V more than the minimum, when $V \downarrow 0$ and with a lower bound constraint on the utility of average throughput.

We consider a single hop network model, which is an extension of R-model, in Section 5.8.

5.2 Problem formulation for I-model and R-model

Our objective is to characterize the minimum average queue length for a given constraint P_c on the average transmit power. The following formulation is for both the I-model and the R-model. The tradeoff problem is

$$\underset{\gamma \in \Gamma}{\text{minimize}} \bar{Q}(\gamma, q_0), \text{ such that } \bar{P}(\gamma, q_0) \leq P_c. \quad (5.4)$$

We note that as in Section 4.2.1 it is possible to formulate a CMDP for the above problem. The state space of the CMDP is $\mathbb{Z}_+ \times \mathcal{H}$, the action space at each $(q, h) \in \mathbb{Z}_+ \times \mathcal{H}$ is the set of batch sizes, and the evolution of the process is as given in (5.1).

Similar to the quantities $c(\lambda)$ and $c_R(\lambda)$ defined for I-model and R-model in Chapter 4, it is possible to obtain quantities $c(\lambda)$ and $c_R(\lambda)$ for I-model and R-model discussed above. The quantities $c(\lambda)$ and $c_R(\lambda)$ can be interpreted as the minimum average power required for mean rate stability of the queue for I-model and R-model respectively. Then it is possible to show that if $P_c > c(\lambda)$ for I-model (or $P_c > c_R(\lambda)$ for R-model), then there exists an optimal stationary policy γ^* for the above problem with stationary distribution π^* . So in the following we consider the above problem for policies in Γ_s .

For I-model, as discussed in Section 4.2.1, from Ma et al. [35], if P_c is such that there exists a Lagrange multiplier $\beta_{P_c} \geq 0$ and the average cost optimal policy $\gamma_{\beta_{P_c}}$ for a MDP with single stage cost $q + \beta_{P_c}P(h, s)$ in state (q, h) has $\bar{P}(\gamma_{\beta_{P_c}}, q_0) = P_c$, then $\gamma_{\beta_{P_c}}$ is optimal for (5.4). The set of all such P_c is denoted as \mathcal{O}^u (similar to \mathcal{O}^u in Chapter 4). It can then be shown that for all $P_c \in \mathcal{O}^u$, there exist optimal stationary deterministic policies for (5.4). Furthermore, such optimal stationary deterministic policies are such that the service batch size $s(q, h)$ is monotonically non-decreasing in q for every h . Then, as in Chapter 4, we consider the above tradeoff problem only for policies γ in an admissible set Γ_a for all values of P_c , where the definition of admissible policies is motivated by the above monotonicity property of any stationary deterministic optimal policy for $P_c \in \mathcal{O}^u$. Since R-model is an approximation to the I-model, the definition of admissible policies for R-model is motivated by the monotonicity property of any stationary deterministic optimal policy for $P_c \in \mathcal{O}^u$ for I-model.

We now state the properties which are satisfied by admissible policies for both I-model and R-model. A policy γ is *stable*, if (i) the Markov chain $(Q[m], m \geq 0)$ under γ is irreducible, aperiodic, and positive Harris recurrent with stationary distribution π , and (ii) $\bar{Q}(\gamma, q_0) < \infty$. A policy γ is *admissible* if:

G1 : it is stable, and,

G2 : the average service rate in state q , $\bar{s}(q) \triangleq \mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} S(q, H)$ is non-decreasing in q ¹.

For an admissible policy γ , we have that the performance measures $\bar{Q}(\gamma, q_0)$ and $\bar{P}(\gamma, q_0)$ are independent of the initial queue state q_0 and exist as limits. Therefore, in the following, these performance measures are denoted by $\bar{Q}(\gamma)$ and $\bar{P}(\gamma)$ respectively.

So in the following we consider the problem TRADEOFF:

$$\underset{\gamma \in \Gamma_a}{\text{minimize}} \bar{Q}(\gamma), \text{ such that } \bar{P}(\gamma) \leq P_c. \quad (5.5)$$

The optimal value of TRADEOFF is denoted as $Q^*(P_c)$. We note that whenever $P_c \in \mathcal{O}^u$, since there exists an optimal admissible policy $Q^*(P_c)$ is the solution to (5.4).

For an admissible policy γ , we note that since the arrival rate is constant, from Little's law the average delay for γ is $\frac{\bar{Q}(\gamma)}{\lambda}$. The tradeoff of average delay with average power can be obtained as $\frac{Q^*(P_c)}{\lambda}$.

As in Chapter 4, it is possible to consider a larger class of policies $\Gamma_{a,M}$, which is obtained by mixing or time sharing of policies in Γ_a . Let $Q_M^*(P_c)$ denote the optimal value of the above problem, but with the minimization carried out over the set $\Gamma_{a,M}$. We note that the asymptotic behaviour for $Q_M^*(P_c)$ can be directly obtained from $Q^*(P_c)$. Therefore, in the following we consider the asymptotic characterization of $Q^*(P_c)$ only.

If P_c is such that the above problem is feasible, then by definition there exists a feasible admissible policy γ such that $\bar{Q}(\gamma) \leq Q^*(P_c) + \epsilon$. Such a policy is called ϵ -optimal in the following.

We note that for any admissible policy γ , $\bar{Q}(\gamma) = \mathbb{E}_{\pi} Q$ and $\bar{P}(\gamma) = \mathbb{E}_{\pi} \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} P(H, S(Q, H))$. Since Q and H are independent, we also have that $\bar{P}(\gamma) = \mathbb{E}_{\pi} \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} P(H, S(Q, H))$. For any $\gamma \in \Gamma_a$, we note that the average arrival rate λ has to be equal to the average service rate, i.e., $\lambda = \mathbb{E}_{\pi} \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} S(Q, H) = \mathbb{E}_{\pi} \bar{s}(Q)$. Therefore, for $\gamma \in \Gamma_a$, $\bar{P}(\gamma)$ is lower bounded by the

¹We note for a stationary deterministic policy $s(q, h)$ is non-decreasing in q for every h . Therefore, $\mathbb{E}_{\pi_H} s(q, H)$ is non-decreasing in q . Since, we are considering randomized policies, we assume that $\mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} S(q, H)$ is non-decreasing in q .

optimal value of

$$\begin{aligned} & \underset{\gamma \in \Gamma_a}{\text{minimize}} && \mathbb{E}_\pi \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} P(H, S(Q, H)), \\ & \text{such that} && \mathbb{E}_\pi \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} S(Q, H) = \lambda, \end{aligned} \quad (5.6)$$

since the only constraint is on the average service rate. Now we note that $\mathbb{E}_\pi \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} S(Q, H) = \mathbb{E}_{\pi_H} \mathbb{E}_{Q|H} \mathbb{E}_{S|Q,H} S(Q, H)$. Then, we have that

$$\mathbb{E}_{Q|H} \mathbb{E}_{S|Q,H} S(Q, H) = \int_q \int_s s \cdot dp_{s,q|H} \cdot d\pi(q) = \int_s \int_q s \cdot dp_{s,q|H} = \int_s s \int_q dp_{s,q|H},$$

we have that $\mathbb{E}_{\pi_H} \mathbb{E}_{Q|H} \mathbb{E}_{S|Q,H} S(Q, H) = \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S$ where the conditional distribution of S given H depends upon the policy. A similar procedure can be carried out on $\mathbb{E}_\pi \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} P(H, S(Q, H))$ which leads to $\mathbb{E}_\pi \mathbb{E}_{H|Q} \mathbb{E}_{S|Q,H} P(H, S(Q, H)) = \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S)$. Then the optimal value of (5.6) is bounded below by the optimal value of

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S), \\ & \text{such that} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S = \lambda, \end{aligned} \quad (5.7)$$

where we minimize over all possible conditional distributions for the batch size S given H , irrespective of the policy. For the I-model, we denote the optimal value of (5.7) by $c(\lambda)$, while for the R-model we denote the optimal value of (5.7) by $c_R(\lambda)$ (we note that $c(\lambda)$ and $c_R(\lambda)$ are the minimum powers required for mean rate stability for the I-model and R-model respectively, see [43]). We note that for the R-model, the conditional distribution of the batch size has support on $[0, S_{max}]$, while for the I-model the conditional distribution has support on $\{0, \dots, S_{max}\}$. Hence, $c_R(\lambda) \leq c(\lambda), \forall \lambda \in [0, S_{max}]$. We note that feasible solutions exist for the above problem only if $\lambda \leq S_{max}$.

We have that $\forall \gamma \in \Gamma_a, \bar{P}(\gamma) \geq c(\lambda)$ for the I-model, and $\bar{P}(\gamma) \geq c_R(\lambda)$ for the R-model. From [43, Theorem 1], we have that if $\lambda < S_{max}$, then for every $P_c > c(\lambda)$ (or every $P_c > c_R(\lambda)$ for the R-model), there exists an admissible policy γ_{P_c} , such that $\bar{P}(\gamma_{P_c}) \leq P_c$ and as $P_c \downarrow c(\lambda)$, $\bar{Q}(\gamma_{P_c})$ grows without bound. Since, for an arrival rate of λ , $c(\lambda)$ (or $c_R(\lambda)$ for the R-model) can be approached arbitrarily closely by admissible policies, $c(\lambda)$ (or $c_R(\lambda)$ for the R-model) is the infimum of the average power for admissible policies.

For the I-model, since properties (C1) and (C2) are assumed to hold, from [43, Section VII], we have that $c(\lambda)$ is a piecewise linear (a proof is given in Appendix 5.A), non-decreasing convex function, for $\lambda \in [0, S_{max}]$, with $c(0) = 0$. Again from [43], $c_R(\lambda)$ is a non-decreasing, strictly convex function of $\lambda \in [0, S_{max}]$, with $c_R(0) = 0$. For the example discussed in Section 5.1.4, the function $c(\lambda)$ and $c_R(\lambda)$ are illustrated in Figure 5.2, for different \mathcal{H} and π_H as well as for the R-model and the I-model. We observe the following: a) for cases (i) and (iii) in Figure 5.2, for the same mean

fading gain $\mathbb{E}_{\pi_H} H^2$, for $\lambda < 1.2$, having a fading gain larger than the mean with some positive probability leads to a smaller $c(\lambda)$; b) for case (iv) we see that $c_R(\lambda)$ for the R-model is strictly convex. The function $c_R(\lambda) \leq c(\lambda)$, and coincides with $c(\lambda)$ for $\lambda \in \{0.5, 1, 1.5\}$; c) for both (i) and (ii) the function $c(\lambda)$ is piecewise linear, but the λ -s at which the slope changes is different and depends on π_H . We consider another example of $c(\lambda)$ and $c_R(\lambda)$ for the example in Section 5.1.4 with $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$. In this case, $c_R(\lambda) < c(\lambda)$ for all $\lambda \notin \{0.4, 0.8, 1.4\}$. From these examples, we can conclude for the R-model a smaller minimum average power $c_R(\lambda)$ is sufficient for stability, compared with $c(\lambda)$ for the I-model.




Figure 5.2: The optimal value $c(\lambda)$ of problem (5.7) for I-model: i) with $H \in \{0.5, 1\}$ with $\pi_H(0.5) = 0.5$ and $S \in \{0, 1, 2\}$, ii) with $H \in \{0.5, 1\}$ with $\pi_H(0.5) = 0.7$ and $S \in \{0, 1, 2\}$, and iii) with $S \in \{0, 1, 2\}$ and fade state fixed at $0.7906 = \sqrt{0.5 \times 0.5^2 + 0.5 \times 1^2}$ with the same average fading gain as (i). For R-model: (iv) with $H \in \{0.5, 1\}$ with $\pi_H(0.5) = 0.5$ and $S \in [0, 2]$.

As shown in Figures 5.2 and 5.3, it is possible that $c_R(\lambda_i) = c(\lambda_i)$ for some $\lambda_i \in [0, S_{max}]$. We will see from the asymptotic analysis in the next section, that the asymptotic growth rate of minimum average queue length suggested by the R-model and the I-model for such λ_i are different. We note that, in general, it is not known for what λ_i , if any, $c_R(\lambda_i) = c(\lambda_i)$. However, if $|\mathcal{H}| = 1$, it is clear that $c_R(s) = c(s), \forall s \in \{0, \dots, S_{max}\}$.

5.3 Asymptotic bounds for I-model

In this section, we obtain an asymptotic characterization of $Q^*(P_c)$ for I-model and R-model as $P_c \downarrow c(\lambda)$ and $P_c \downarrow c_R(\lambda)$ respectively. For I-model, as in Chapter 4, we assume

A1 : $Pr \{A[1] > S_{max}\} > \epsilon_a > 0$.

A2 : $Pr \{A[1] = a\} > 0$, for all $a \in \{0, \dots, A_{max}\}$,




Figure 5.3: The optimal value of problem (5.7): $c(\lambda)$ for I-model and $c_R(\lambda)$ for R-model, with $H \in \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$; $c(1.7)$ is 107% of $c_R(1.7)$.

For R-model as in Chapter 4, we assume

RA1 : $Pr \{A[1] - S_{max} > \delta_a\} > \epsilon_a$, for positive δ_a and ϵ_a .

We note that part of this analysis was presented in [61]. We consider the I-model first. Consider any $\gamma \in \Gamma_a$. To obtain the asymptotic behaviour of $Q^*(P_c)$ as $P_c \downarrow c(\lambda)$, we ascertain the asymptotic behaviour of $\pi(q)$ as $P_c \downarrow c(\lambda)$. We note that $\pi(q)$ determines the average queue length and the average power given the policy. As we have seen in Chapter 4, it turns out that the asymptotic behaviour of $\pi(q)$, is determined by the average drift of the queue, $\lambda - \bar{s}(q)$, at a queue length q . So we proceed by relating the average power to the average drift $\lambda - \bar{s}(q)$. We note that the average power used when the queue length is q is $\mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} P(H, S(Q, H))$, which is bounded below by the optimal value of

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S), \\ & \text{such that} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S = \bar{s}(q), \end{aligned}$$

where we have considered all possible conditional distributions on the batch size with support on $\{0, \dots, S_{max}\}$, subject only to the constraint that the average service rate is $\bar{s}(q)$. The above optimization problem is the same as (5.7) except that the constraint is now $\bar{s}(q)$ instead of λ . Therefore, the average power used when the queue length is q is bounded below by $c(\bar{s}(q))$. We note that any feasible policy γ for TRADEOFF has $\bar{P}(\gamma) \leq P_c$. Then for that γ , $\mathbb{E}_{\pi} c(\bar{s}(Q)) \leq \bar{P}(\gamma) \leq P_c$. We also note that from the convexity of $c(s)$, $\mathbb{E}_{\pi} c(\bar{s}(Q)) \geq c(\lambda)$, since $\mathbb{E}_{\pi} \bar{s}(Q) = \lambda$. Now as $P_c \downarrow c(\lambda)$, for any sequence of feasible policies for TRADEOFF, $\mathbb{E} c(\bar{s}(Q)) \downarrow c(\lambda)$.




Figure 5.4: Illustration of the relationship between λ , s_l , and s_u along with the minimum average cost $c(\lambda)$ and the line $l(s)$ for the three cases

The behaviour of $Q^*(P_c)$ as $P_c \downarrow c(\lambda)$ is observed to depend on the relationship of λ with $c(s)$, $s \in [0, S_{max}]$. Since $c(s)$ is piecewise linear, we can define a sequence of intervals $[a_p, b_p]$, $p \in \{1, \dots, P\}$, with $a_{p+1} = b_p$, $a_1 = 0$, and $b_P = S_{max}$. The sequence of intervals is such that for $s \in [a_p, b_p]$, $c(s)$ is linear. The following three cases arise:

1. $a_1 = 0 < \lambda < b_1$,
2. $a_p < \lambda < b_p$, $p > 1$, and,
3. $\lambda = a_p$, $p > 1$.

We note that $c_R(a_p) = c(a_p)$ in Figures 5.2 and 5.3. For Case 1, through numerical examples, we illustrate that $Q^*(P_c)$ does not grow to infinity as $P_c \downarrow c(\lambda)$. We obtain an asymptotic lower bound for Case 2 which is the asymptotic lower bound to the *super-fast* $\log\left(\frac{1}{V}\right)$ upper bound observed for the sequence of policies constructed by Neely in [43, Corollary 2], but for admissible policies. We also obtain an asymptotic lower bound for Case 3 which is used to illustrate another difference between the asymptotic behaviours of $Q^*(P_c)$ in the R-model and I-model.

For Cases 1 and 2, let $s_l \triangleq a_p$ and $s_u \triangleq b_p$, while for Case 3 let $s_l = s_u \triangleq a_p$. An example is shown in Figure 5.4. We define the line $l(s) : [0, S_{max}] \rightarrow \mathbb{R}_+$ as follows:

1. If $s_l < \lambda < s_u$, then $l(s)$ is the line through $(s_l, c(s_l))$ and $(s_u, c(s_u))$.
2. If $s_l = \lambda = s_u = a_p$ for some $p > 1$, then $l(s)$ is a line through $(\lambda, c(\lambda))$ with slope m chosen such that $\frac{c(a_p) - c(a_{p-1})}{a_p - a_{p-1}} < m < \frac{c(a_{p+1}) - c(a_p)}{a_{p+1} - a_p}$.

We note that $\mathbb{E}l(\bar{s}(Q)) = c(\lambda)$. We now present asymptotic lower bounds for $Q^*(P_c)$, which follow directly from the asymptotic lower bounds in Lemmas 4.3.8 and 4.3.13.

Lemma 5.3.1. For Case 2, given any sequence of admissible policies γ_k with $\bar{P}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

²We note that unlike in Chapter 4, s_l and s_u need not be integers.

Proof. We first note that for the sequence γ_k , since $\bar{P}(\gamma_k) - c(\lambda) = V_k$, we have that $U_k \triangleq \mathbb{E}_\pi c(\bar{s}(Q)) - c(\lambda) \downarrow 0$, with $U_k \leq V_k$. The rest of the proof is very similar to that of Lemma 4.3.8. With $\bar{s}(q)$ replacing $\mathbb{E}S(q)$ and U_k replacing V_k (with an inequality) in the proof of Lemma 4.3.8, and m being redefined as the tangent of the angle made by the line passing through $(a_{p-1}, c(a_{p-1}))$ and $(a_p, c(a_p))$ with $l(s)$, and proceeding as in the proof of Lemma 4.3.8 we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{U_k}\right)\right) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$, since $U_k \leq V_k$. \square

We note that an asymptotic upper bound for Case 2 can be obtained from a sequence of TOCA policies from [43]. However, TOCA policies are not admissible. Therefore, we now present an asymptotic upper bound for a sequence of policies, which is similar to the sequence of buffer partitioning policies proposed by Berry and Gallager [7].

Lemma 5.3.2. Let a policy γ be defined as follows. When the queue length is q and fade state is h , γ serves a batch size $\min(q, \tilde{S}(q, h))$, where

$$\tilde{S}(q, h) \sim \begin{cases} p_l(h), & \text{for } 0 \leq q < q_v, \\ p_u(h), & \text{for } q_v \leq q. \end{cases}$$

where $q_v > 0$, $p_l(h)$ and $p_u(h)$, $\forall h \in \mathcal{H}$ are the optimizing distributions for (5.7) with the constraint being s_l and s_u respectively. We obtain a sequence of policies γ_k , by choosing $q_v = \log\left(\frac{1}{V_k}\right)$, where $V_k < 1$ is a sequence decreasing to zero. Then for Case 2, γ_k is a sequence of admissible policies, such that $\bar{P}(\gamma_k) - c(\lambda) = \mathcal{O}(V_k)$ and $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$.

The proof of this asymptotic upper bound is given in Appendix 5.B. We note that the proof is quite similar to that of Lemma 4.3.10.

We then have the following result.

Proposition 5.3.3. For Case 2, we have that $Q^*(P_{c,k}) = \Theta\left(\log\left(\frac{1}{P_{c,k} - c(\lambda)}\right)\right)$ where $P_{c,k} \downarrow c(\lambda)$ and $P_{c,k} = \bar{P}(\gamma_k)$ for the sequence of policies in Lemma 5.3.2.

The proof is very similar to that of Proposition 4.3.11 and is therefore omitted.

Remark 5.3.4. We now consider a sequence of policies generated from a modified form of the tradeoff optimal control algorithm (TOCA) [43]. We note that in [43], the control variable is the power allocation, and the service rate is obtained as a function of the power allocation and the current fade state. The set of all possible power allocations (II in [43]) is the same for all fade states. However, in our model the control variable is the batch size. Therefore, we modify TOCA such that the set of power allocations is a function of the fade state. For each fade state, the possible power allocations are such that the batch sizes takes all possible values in $\{0, \dots, S_{max}\}$.

The modified TOCA algorithm is again parametrized by positive numbers w, ϵ, \tilde{q} , and β as in [43]. The algorithm chooses at each slot $m \geq 1$, the batch size s_{TOCA} such that

$$s_{TOCA}[m] = \min \left(\arg \min_{s \in \{0, \dots, S_{max}\}} \left\{ \beta P(H[m], s) - W[m]s \right\}, Q[m-1] \right),$$

where

$$W[m] = \mathbb{I}\{Q[m-1] \geq \tilde{q}\} \left[w e^{w(Q[m-1] - \tilde{q})} + 2X[m-1] \right] + \mathbb{I}\{Q[m-1] < \tilde{q}\} \left[-w e^{w(\tilde{q} - Q[m-1])} + 2X[m-1] \right].$$

We note that $s_{TOCA}[m] = 0$ if $W[m] \leq 0$. The sequence $X[m], m \geq 0$ is obtained from a *virtual* queue which evolves according to

$$X[m+1] = \max(X[m] - s_{TOCA}[m+1] - \epsilon \mathbb{I}\{Q[m] < \tilde{q}\}, 0) + A[m+1] + \epsilon \mathbb{I}\{Q[m] \geq \tilde{q}\}.$$

As in [43], let $\delta_{max} = \max(A_{max}, S_{max})$. For our Case 2, let $0 < \epsilon < \min(\lambda - a_p, b_p - \lambda)$, $w = \frac{\epsilon}{\delta_{max}^2} e^{\frac{-\epsilon}{\delta_{max}}}$, and $\tilde{q} = \frac{2}{w} \log(\beta)$. A sequence of TOCA policies γ_k is generated by choosing a sequence $\beta_k = \frac{1}{V_k}$, for a sequence $V_k \downarrow 0$. The proof of [43, Corollary 2] extends to the above version of the TOCA algorithm, and we have that

$$\bar{Q}(\gamma_k) = \mathcal{O} \left(\log \left(\frac{1}{V_k} \right) \right), \bar{P}(\gamma_k) = c(\lambda) + \mathcal{O}(V_k).$$

Therefore, we obtain that for the sequence of policies γ_k , $\bar{Q}(\gamma_k) = \mathcal{O} \left(\log \left(\frac{1}{\bar{P}(\gamma_k) - c(\lambda)} \right) \right)$. We note that $W[m]$ is a non-decreasing function of q , where $Q[m-1] = q$. Since $s_{TOCA}[m]$ is a non-decreasing function of $W[m]$, we have that $s_{TOCA}[m]$ is a non-decreasing function of $Q[m-1]$. However, γ_k is not a sequence of admissible policies, since the policies γ_k depend on the auxiliary variable $X[m-1]$. We note that the above upper bound is an upper bound on the optimal value of (5.4). Furthermore, as in Remark 4.3.12, for any subsequence $P_{c,k}$ of \mathcal{O}^u such that there exists a constant $m \geq 1$ and a subsequence $P_{TOCA,k}$ of $\bar{P}(\gamma_k)$ such that $P_{TOCA,k} - c(\lambda) \leq m(P_{c,k} - c(\lambda))$, we have that $Q^*(P_{c,k}) \leq Q(\gamma_k) = \mathcal{O} \left(\log \left(\frac{1}{P_{TOCA,k} - c(\lambda)} \right) \right) = \mathcal{O} \left(\log \left(\frac{1}{P_{c,k} - c(\lambda)} \right) \right)$.

Lemma 5.3.5. For Case 3, given any sequence of admissible policies γ_k with $\bar{P}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega \left(\frac{1}{V_k} \right)$.

Proof. We again note that for the sequence γ_k , since $\bar{P}(\gamma_k) - c(\lambda) = V_k$, we have that $U_k \stackrel{\Delta}{=} \mathbb{E}_\pi c(\bar{s}(Q)) - c(\lambda) \downarrow 0$, with $U_k \leq V_k$. The rest of the proof is very similar to that of Lemma 4.3.13. With $\bar{s}(q)$ replacing $\mathbb{E}S(q)$ and U_k replacing V_k (with an inequality) in the proof of Lemma 4.3.13, and m being redefined as the tangent of the angle made by the line passing through $(a_p, c(a_p))$ and $(a_{p+1}, c(a_{p+1}))$ with $l(s)$, and proceeding as in the proof of Lemma 4.3.13 we have that $\bar{Q}(\gamma_k) = \Omega \left(\frac{1}{U_k} \right) = \Omega \left(\frac{1}{V_k} \right)$, since $V_k \leq U_k$. \square

Remark 5.3.6. The derivation of the relationship between the difference of the $\bar{P}(\gamma)$ and $c(\lambda)$ and the *average drift* defined as $\sum_{q=q_d+1}^{\infty} \mathbb{E}[Q[m+1] - Q[m] | Q[m] = q] \pi(q)$, in the proof of Lemma 4.3.13, is motivated by the approach in [7]. We note that in our proof, q_d can be chosen arbitrarily by the choice of ϵ_V and then ϵ_V can be chosen so as to obtain the tightest asymptotic lower bound. However, in [7], q_d cannot be chosen arbitrarily. In fact, q_d is chosen as the queue length which has the maximal stationary probability of all queue lengths in the set $\{0, \dots, \lceil 2\bar{Q}(\gamma) \rceil\}$ for a policy. The freedom in the choice of q_d enables us to derive the $\Omega\left(\frac{1}{V}\right)$ asymptotic lower bound.

Remark 5.3.7. As in the analysis of Case 3 in Chapter 4, it is possible to show that a sequence of randomized policies $\{\gamma_k\}$ achieves the asymptotic lower bound derived above for Case 3. The sequence of policies is parametrized by a sequence $V_k \downarrow 0$. For a particular V_k , the policy chooses

$$S[m] = \min(S'(H[m]), Q[m-1]),$$

where $S'(H[m])$ is independently chosen for each m and is distributed according to any conditional distribution of batch size given $H[m]$ which is optimal for (5.7) but with the rate constraint being $\lambda + V_k$. For a particular V_k , it can then be shown that $\bar{P}(\gamma_k) = c(\lambda + V_k)$ and $\bar{Q}(\gamma_k) \leq \frac{\sigma^2 + 2S_{max}^2 + \mathbb{E}_{\pi_H} \mathbb{E} S(H)^2 - \lambda^2}{2V_k}$. Then as $k \rightarrow \infty$ and $V_k \downarrow 0$, $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$ and $\bar{P}(\gamma_k) = c(\lambda) + \mathcal{O}(V_k)$.

We have the following result.

Proposition 5.3.8. For Case 3, we have that $Q^*(P_{c,k}) = \Theta\left(\frac{1}{P_{c,k} - c(\lambda)}\right)$ where $P_{c,k} \downarrow c(\lambda)$ and $P_{c,k} = \bar{P}(\gamma_k)$ for the sequence of policies in the above remark.

The proof is very similar to that of Proposition 4.3.11 and is therefore omitted.

5.4 Asymptotic bounds for R-model

We now consider the asymptotic behaviour of $Q^*(P_c)$ in the asymptotic regime \mathfrak{R} as $P_c \downarrow c_R(\lambda)$ for the R-model. Similar to the I-model, it can be shown that the average power used when the queue length is q is bounded below by $c_R(\bar{s}(q))$. We again note that any feasible policy γ for TRADEOFF has $\bar{P}(\gamma) \leq P_c$, and for that γ , $\mathbb{E}_{\pi} c_R(\bar{s}(Q)) \leq \bar{P}(\gamma) \leq P_c$. We also note that from the convexity of $c_R(s)$, $\mathbb{E}_{\pi} c_R(\bar{s}(Q)) \geq c_R(\lambda)$, since $\mathbb{E}_{\pi} \bar{s}(Q) = \lambda$. Now as $P_c \downarrow c_R(\lambda)$, for any sequence of feasible policies for TRADEOFF, $\mathbb{E} c_R(\bar{s}(Q)) \downarrow c_R(\lambda)$. We note that the following result is similar to the Berry-Gallager lower bound, but is derived with the extra assumption G2.

Lemma 5.4.1. For any sequence of admissible policies γ_k with $\bar{P}(\gamma_k) - c_R(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$. Therefore, $Q^*(P_c) = \Omega\left(\frac{1}{\sqrt{P_c - c(\lambda)}}\right)$ as $P_c \downarrow c_R(\lambda)$.

Proof. We note that for the sequence γ_k , since $\bar{P}(\gamma_k) - c_R(\lambda) = V_k$, we have that $U_k \triangleq \mathbb{E}_\pi c_R(\bar{s}(Q)) - c_R(\lambda) \downarrow 0$, with $U_k \leq V_k$. The rest of the proof is similar to that of Proposition 4.4.2. We follow all the steps in the proof of Proposition 4.4.2, with $\bar{s}(q)$ replacing $\mathbb{E}S(q)$ and U_k replacing V_k (with an inequality) in the proof of Proposition 4.4.2, to obtain that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{U_k}}\right) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$. Let γ'_k be a sequence of ϵ -optimal policies for TRADEOFF for the sequence $P_{c,k}$. Then we have that $\bar{P}(\gamma'_k) \downarrow c_R(\lambda)$ and $\bar{Q}(\gamma'_k) = \Omega\left(\frac{1}{\sqrt{P_{c,k} - c_R(\lambda)}}\right)$. Since γ'_k is ϵ -optimal, we have that $Q^*(P_{c,k}) \geq \bar{Q}(\gamma'_k) - \epsilon$. Therefore, $Q^*(P_{c,k}) = \Omega\left(\frac{1}{\sqrt{P_{c,k} - c_R(\lambda)}}\right)$ as $P_{c,k} \downarrow c_R(\lambda)$. \square

Remark 5.4.2. A sequence of admissible policies γ_k can be obtained which achieves the above asymptotic lower bound up to a logarithmic factor, as in Lemma 4.4.3. The proof of this upper bound is very similar to that of Lemma 5.3.2 and is therefore omitted. A particular policy γ in the sequence is defined as follows. The policy γ serves a batch size $\min(q, \tilde{S}(q, h))$ when the queue length is q and fade state is h , where

$$\tilde{S}(q, h) \sim \begin{cases} p_-(h), & \text{for } 0 \leq q < q_v, \\ p_+(h), & \text{for } q_v \leq q < 2q_v, \\ p_\epsilon(h), & \text{for } 2q_v \leq q, \end{cases}$$

where $q_v > 0$, $p_-(h)$, $p_+(h)$, and $p_\epsilon(h)$, $\forall h \in \mathcal{H}$ are the optimizing distributions for (5.7) with the constraint being $\lambda - \epsilon_V$, $\lambda + \epsilon_V$, and $\lambda + \epsilon$ respectively. We obtain a sequence of policies γ_k , by choosing ϵ_V and q_v from sequences ϵ_{V_k} and q_{v_k} defined as follows. Let $\omega_k = \sqrt{V_k}$, where $V_k \downarrow 0$. Let $\epsilon_{V_k} = \omega_k A_{max}^2 e^{\omega A_{max}}$ and $q_{v_k} = \frac{1}{\omega_k} \log\left(\frac{1}{\epsilon_{V_k}^3}\right)$.

Remark 5.4.3. As in Remark 5.3.4, for a sequence of policies γ_k generated by the modified form of TOCA, it is possible to show as in [43], that for $V_k < \frac{1}{S_{max}}$, $\beta_k = \frac{1}{V_k}$, $\epsilon_k = \frac{1}{\sqrt{\beta_k}}$, $w_k = \frac{\epsilon_k}{\delta_{max}^2} e^{-\frac{\epsilon_k}{\delta_{max}}}$, and $\tilde{q}_k = \frac{6}{w_k} \log\left(\frac{1}{\epsilon_k}\right)$,

$$\begin{aligned} \bar{Q}(\gamma) &= \mathcal{O}\left(\sqrt{\frac{1}{V_k}} \log\left(\frac{1}{V_k}\right)\right), \\ \bar{P}(\gamma) &= c(\lambda) + \mathcal{O}(V_k), \end{aligned}$$

as $V_k \downarrow 0$. We note that the above asymptotic upper bound applies only to problem (5.4), since the sequence of policies considered above is not admissible.

5.5 A numerical example

To illustrate the results obtained in the previous section, we plot the optimal tradeoff curve for the example in Section 5.1.4. We note that the results in this section apply only to I-model. We assume

that $A_{max} = 5$. We consider first a case where the fading gain $\mathcal{H} = \{0.5, 1\}$ with $\pi_H(0.5) = 0.5$. The tradeoff curves $Q^*(P_c)$ for $\lambda \in \{0.9, 1.0, 1.1\}$, are shown in Figure 5.5. We note that these arrival rates correspond to bit arrival rates of 45, 50, and 55 kb/s respectively. Each point in the tradeoff curves is obtained by numerical solution of a suitably truncated MDP with state being the current queue length and the current fade state, and single stage cost $q + \beta P(h, s)$, where $\beta > 0$ is a Lagrange multiplier. Each tradeoff curve is obtained by varying β . From Figure 5.2 we have that $c(0.9) = 1.20$, $c(1.0) = 1.384$, and $c(1.1) = 1.781$. From the asymptotic characterization of $Q^*(P_c)$, we have that for $\lambda = 1.0$, $Q^*(P_c)$ increases as $1/(P_c - 1.384)$, while for $\lambda = 0.9$ and 1.1, $Q^*(P_c)$ increases as $\log\left(\frac{1}{P_c - c(\lambda)}\right)$. In Figure 5.5, we observe that for the same average queue length, the difference between the average service cost and $c(\lambda)$ increases from $\lambda = 0.9$ to $\lambda = 1.0$ and then decreases. This difference is even more pronounced in Figure 5.6, where λ is increased from 0.78 to 0.80 and then to 0.82 for the example in Section 5.1.4 with $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$.




Figure 5.5: The optimal tradeoff $Q^*(P_c)$ for the system in Section 5.1.4 with $\mathcal{H} = \{0.5, 1\}$ and $\pi_H(0.5) = 0.5$, for $\lambda \in \{0.9, 1.0, 1.1\}$ with $c(0.9) = 1.2$, $c(1.0) = 1.384$, and $c(1.1) = 1.781$.




Figure 5.6: The optimal tradeoff $Q^*(P_c)$ for the system in Section 5.1.4 with $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$, for $\lambda \in \{0.78, 0.80, 0.82\}$ with $c(0.78) = 1.0717$, $c(0.80) = 1.1071$, and $c(0.82) = 3.0995$.

For the example in Section 5.1.4, with $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$, for $\lambda = 0.82$, we plot the optimal tradeoff curve $Q^*(P_c)$, the asymptotic lower bound from Lemma 5.3.1 (using Lemma 4.3.8), and the upper bound for the sequence of policies from Remark 5.3.4, in Figure 5.7. We note that the upper bound is obtained via simulation, for a sequence of policies for which \tilde{q} has been chosen to be $\frac{1}{20}$ of what is suggested in [43, Theorem 3]. This heuristic has been used in [43]. The analytical upper bound for the sequence of TOCA policies from [43, Theorem 3] is found to be very weak. We note that the asymptotic bounds, although tight in the order sense, are very weak. To illustrate the bounds for Case 3, for the system in Section 5.1.4, with $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$, for $\lambda = 0.8$, we plot the optimal tradeoff curve $Q^*(P_c)$, the asymptotic lower bound $\frac{\bar{q}_3}{2}$ from Lemma 5.3.5 (using Lemma 4.3.13), and the upper bound for the sequence of policies from Remark 5.3.7, in Figure 5.8. Again the upper bound is obtained via simulation.




Figure 5.7: Comparison of the optimal tradeoff $Q^*(P_c)$ with the lower bound from Lemma 4.3.8 and the upper bound from Remark 5.3.4 for the system in Section 5.1.4, for $\lambda = 0.82$, for $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$




Figure 5.8: Comparison of the optimal tradeoff $Q^*(P_c)$ with the lower bound from Lemma 4.3.13 and the upper bound from Remark 5.3.7 for the system in Section 5.1.4, for $\lambda = 0.8$, for $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$

The numerical results illustrate that the asymptotic bounds which are obtained using the methods in Chapter 4 are weak, although they are tight in the order sense for Cases 2 and 3.

We consider the case $\lambda = 0.2$ in Figure 5.9, which corresponds to Case 1 for the example in Section 5.1.4, for both $\mathcal{H} = \{0.5, 1\}$ and $\mathcal{H} = \{0.1, 1\}$. We note that $c(0.2)$ is 0.1992 for both $\mathcal{H} = \{0.5, 1\}$ and $\mathcal{H} = \{0.1, 1\}$. We observe that $Q^*(P_c)$ approaches a finite value in both cases. We note that, for Case 1, if $|\mathcal{H}| = 1$, it is possible to show that there exists a policy γ such that $\bar{P}(\gamma) = c(\lambda)$ and $\bar{Q}(\gamma) < \infty$.




Figure 5.9: The optimal tradeoff $Q^*(P_c)$ for the system in Section 5.1.4, for $\lambda = 0.2$, for two cases of \mathcal{H} ; $Q^*(P_c)$ approaches a finite value in both cases.

The exact nature of $\bar{Q}^*(P_c)$ for the different cases depends on the *shape* of $\pi(q)$ in the regime \mathfrak{R} . We first consider Cases 2 and 3 for I-model. Intuition for the behaviour of $\bar{Q}^*(P_c)$ for R-model is similar to that of Case 3 for I-model. In Figures 5.10 and 5.11 we illustrate the behaviour of the probability mass function (PMF) of the queue length for Cases 2 and 3. For each case, the PMF has been obtained by solving the global balance equations of the DTMC for the optimal policy for the truncated MDP used in Figure 5.6.

We now present our observations about the nature of $\pi(q)$ in the regime \mathfrak{R} using the above example. We note that the intuition about the shape of $\pi(q)$ in the regime \mathfrak{R} , which we discuss below, holds

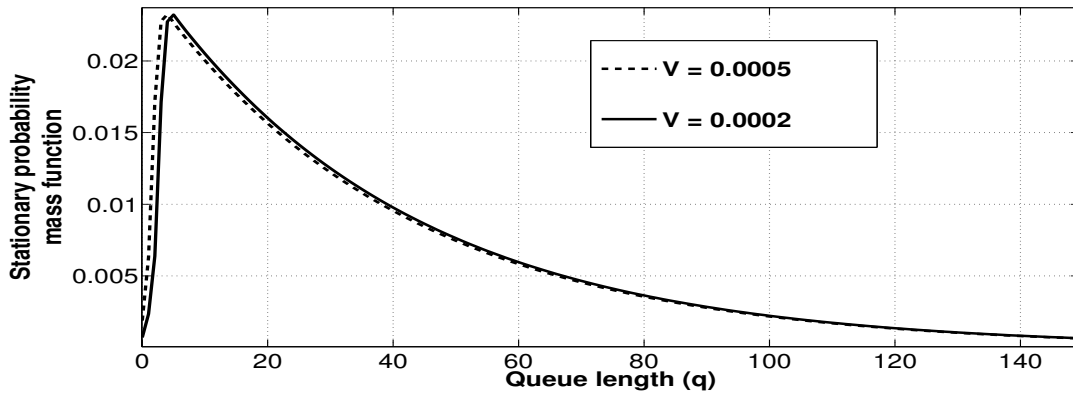


Figure 5.10: PMF of queue length for optimal policies in the regime \mathfrak{R} for Case 2 for the system in Section 5.1.4; $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$; $\lambda = 0.78$ and $c(0.78) = 1.0717$.

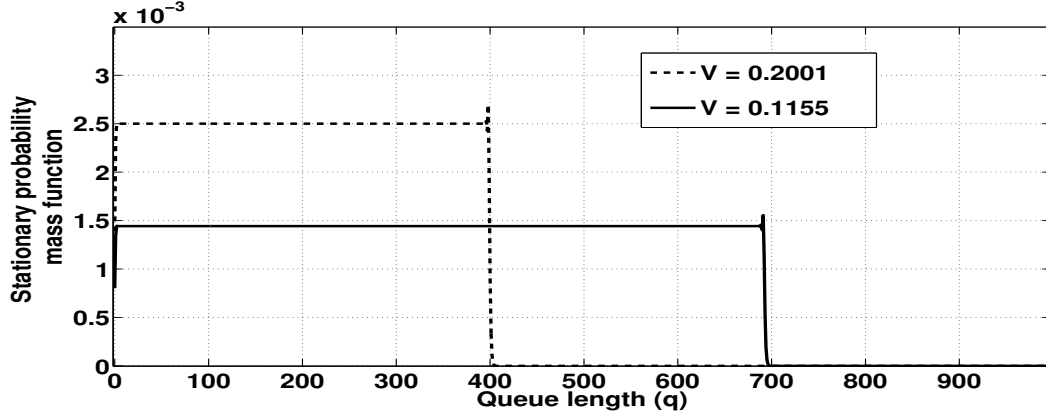


Figure 5.11: PMF of queue length for optimal policies in the regime \mathfrak{R} for Case 3 for the system in Section 5.1.4; $\mathcal{H} = \{0.1, 1\}$ and $\pi_H(0.1) = 0.6$; $\lambda = 0.8$ and $c(0.8) = 1.1071$.

for any sequence of feasible admissible policies, rather than just the sequence of optimal admissible policies.

For I-model, let $\mathcal{Q}_h = \{q : \mu(q) \in [s_l - \epsilon_V, s_u + \epsilon_V]\}$, where ϵ_V is $\omega(V)$ as $V \downarrow 0$. Then it can be shown that $Pr\{Q \in \mathcal{Q}_h\} \uparrow 1$ as $V \downarrow 0$, i.e., as $V \downarrow 0$, the service rates have to be chosen from the set $[s_l - \epsilon_V, s_u + \epsilon_V]$ or $[\lambda - \epsilon_V, \lambda + \epsilon_V]$ for Cases 2 and 3 respectively.

The expected drift of the queue length when the queue length is q is $\mathbb{E}[Q[m+1] - Q[m] | Q[m] = q]$. We note that for admissible policies the expected drift is a non-increasing function of q .

We consider Case 2 first. We note that since service rates can be chosen from the set $[s_l - \epsilon_V, s_u + \epsilon_V]$ as $V \downarrow 0$, intuitively the expected drift in the queue length for small queue lengths in \mathcal{Q}_h is positive, while for large queue lengths in \mathcal{Q}_h the expected drift is negative. From the geometric upper bound on $\pi(q)$, intuitively $\pi(q)$ grows geometrically for small queue lengths. Since the expected drift of the queue length for large queue lengths is negative, from the geometrically decreasing upper bound on $Pr\{Q > q-1\}$ and geometrically decreasing lower bound on $Pr\{Q > q\}$ from [9, Theorem 3] we have that $\pi(q)$ decreases geometrically to zero for large enough q . We note that this is the geometrically increasing and decreasing behaviour of $\pi(q)$ which is illustrated in Figure 5.10. We also note that this behaviour holds for any $V > 0$ and as illustrated in the example in above leads to $\log\left(\frac{1}{V}\right)$ behaviour of the average queue length for any sequence of admissible policies.

For Case 3, we note that for $q \in \mathcal{Q}_h$, as $V \downarrow 0$, the expected drift approaches 0. Suppose we define q_d as largest queue length in \mathcal{Q}_h and ϵ_V as aV where $a > 0$ is a constant. Then it can be shown that for any $q \in \mathcal{Q}_h$

$$Pr\{Q < q\} \leq \left(1 - \left(\frac{\epsilon_a}{\epsilon_a + d}\right)^q\right) \times \text{a constant}.$$

We note that the drift $d = \epsilon_V = aV$. Hence, $Pr\{Q < q\} \approx \mathcal{O}\left(q \frac{aV}{\epsilon_a}\right)$. Therefore, using the bounds on the stationary probability distribution, we are able to obtain the intuition that the stationary probability for any $q \in \mathcal{Q}_h$ is constant (this is illustrated in Figure 5.11) and $\mathcal{O}(V)$. Then we have

that the largest \bar{q} such that $Pr\{Q < \bar{q}\} \leq \frac{1}{2}$ is $\Omega\left(\frac{1}{\sqrt{V}}\right)$, and therefore so is the asymptotic lower bound on the minimum average queue length.

For R-model, we let $\mathcal{Q}_h = \{q : \mu(q) \in [\lambda - \epsilon_V, \lambda + \epsilon_V]\}$. The intuition for the shape of the distribution for R-model is similar to that for Case 3 above. With $\epsilon_V = a\sqrt{V}$ the stationary probability for any $q \in \mathcal{Q}_h$ is $\mathcal{O}\left(\sqrt{V}\right)$ and thus the asymptotic lower bound on \bar{q} is $\Omega\left(\frac{1}{\sqrt{V}}\right)$ and therefore so is the asymptotic lower bound on the minimum average queue length.

5.6 Asymptotic lower bounds for ergodic $(A[m])$ and $(H[m])$

In this section, we present an asymptotic lower bound for the optimal value of (5.4) when $(A[m])$ and $(H[m])$ are ergodic processes, for I-model.

We restrict to policies in Γ_s , which choose a batch service size $S[m] = S(Q[m-1], H[m])$ as a function of the current queue length and fade state and independently of anything else. We note that such policies may not be optimal. The development for ergodic $(A[m])$ and $(H[m])$ closely follows the development in Section 4.3.7.

Since $(A[m])$ is assumed to be ergodic, we have that almost surely

$$\lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M A[m] = \mathbb{E}A[1] = \lambda,$$

and $\lambda < S_{max}$, where S_{max} is the largest batch size which can be served, as defined before. We also assume that the arrival process $(A[m])$ is such that

NA1 : Let $\sigma[m-1] = (Q[0] = q_0, A[1] = a_1, Q[1] = q_1, A[2] = a_2, \dots, A[m-1] = a_{m-1}, Q[m-1] = q_{m-1})$. We assume that

$$\inf_{m \in \mathbb{Z}_+} \min_{\substack{\{a_1, \dots, a_{m-1}\} \\ \{q_1, \dots, q_{m-1}\}}} Pr\{A[m] = 0 | \sigma[m-1]\} = \nu_a > 0,$$

NA2 : $(A[m])$ is independent of $(H[m])$.

Since $(H[m])$ is assumed to be ergodic, we have that

$$Pr\{H[m] = h\} = \pi_H(h), \forall h \in \mathcal{H}, \forall m \geq 1,$$

with $\sum_{h \in \mathcal{H}} \pi_H(h) = 1$.

We restrict to policies $\gamma \in \Gamma_s$ for which the following limits exist

$$\lim_{m \rightarrow \infty} Pr\{Q[m-1] = q, H[m] = h\} = \pi_{Q,H}(q, h), \forall q \in \mathbb{Z}_+ \text{ and } h \in \mathcal{H}, \quad (5.8)$$

with $\sum_{q,h} \pi(q, h) = 1$. We note that for such a policy the following limits also exist

$$\begin{aligned} \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=0}^{M-1} Pr \{Q[m] = q\} &= \pi(q), \forall q \in \mathbb{Z}_+, \\ \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M Pr \{H[m] = h, S[m] = s\} &= \pi_{H,S}(h, s), \forall s \in \{0, \dots, S_{max}\}, \forall h \in \mathcal{H}, \\ \lim_{M \rightarrow \infty} \frac{1}{M} \sum_{m=1}^M Pr \{S[m] = s\} &= \pi_s(s), \forall s \in \{0, \dots, S_{max}\}. \end{aligned}$$

We again restrict to a set of admissible policies $\Gamma_a \subseteq \Gamma_s$, which are defined similarly as in Section 4.3.7. A policy $\gamma \in \Gamma_s$ is admissible if :

NG1 : the limits in (5.8) exist

NG2 : γ is mean rate stable, i.e., $\sum_{h \in \mathcal{H}} \sum_{s=0}^{S_{max}} \pi_{H,S}(h, s) s = \lambda$,

NG3 : the average service rate $\bar{s}(q) = \sum_h \pi_H(h) \mathbb{E}S(q, h)$ is a non-decreasing function of q for γ .

For any admissible policy γ , we have that $\bar{P}(\gamma) = \sum_{h \in \mathcal{H}} \sum_{s \in \{0, \dots, S_{max}\}} \pi_{H,S}(h, s) P(h, s)$ and $\bar{Q}(\gamma) \triangleq \sum_q \pi(q) q$.

For any admissible policy γ , we have that $\bar{P}(\gamma)$ is bounded below by the optimal value of

$$\begin{aligned} &\underset{\gamma \in \Gamma_a}{\text{minimize}} && \sum_{h \in \mathcal{H}} \sum_{s=0}^{S_{max}} \pi_{H,S}(h, s) P(h, s), \\ &\text{such that} && \sum_{h \in \mathcal{H}} \sum_{s=0}^{S_{max}} \pi_{H,S}(h, s) s = \lambda, \end{aligned}$$

where $\pi_{H,S}(h, s)$ is determined by the policy γ . The optimal value of the above problem is bounded below by the optimal value of

$$\begin{aligned} &\underset{p_{s|h}(s), h \in \mathcal{H}}{\text{minimize}} && \sum_{h \in \mathcal{H}} \pi_H(h) \sum_{s=0}^{S_{max}} p_{s|h}(s) P(h, s), \\ &\text{such that} && \sum_{h \in \mathcal{H}} \sum_{s=0}^{S_{max}} \pi_H(h) p_{s|h}(s) s = \lambda, \end{aligned}$$

which is $c(\lambda)$. We note that in the above optimization problem, the minimization is over all possible *conditional distributions* on the batch size. The conditional distribution $p_{s|h}(s)$ has the interpretation of the fraction of time a batch of size s is used when the fade state is h .

Let us denote the optimal conditional distribution that achieves $c(\lambda)$ in (5.7) by $p_{s|h}^*(\lambda), \forall h \in \mathcal{H}$. Similar to the approach in Section 4.3.7, we now show that there exists a sequence of admissible policies γ_V for a sequence $V \downarrow 0$ such that $\lim_{V \downarrow 0} \bar{P}(\gamma_V) = c(\lambda)$. We define random variables $U_V(h) \sim p_{s|h}^*(\lambda + V), \forall h \in \mathcal{H}$. For a particular $V > 0$, the policy γ_V , serves $S[m] = \min(Q[m-1], U_V[m])$, where for each m , $U_V[m]$ is an independent sample of $U_V(H[m])$. Then the queue

evolution under γ_V is

$$Q[m+1] = \max\left(Q[m] - U_V(H[m+1]), 0\right) + A[m+1], \text{ with } Q[0] = q_0, \forall m \geq 0.$$

Applying [34, Lemma 1] we have that the limit $\pi_{Q,H}(q, h) = \lim_{m \rightarrow \infty} Pr\{Q[m] = q, H[m+1] = h\}$ exists. With $Q'[1] = \max\left(q_0 - U_V(H[1]), 0\right)$, we can define a sequence $Q'[m]$ which evolves as

$$Q'[m+1] = \max\left(Q'[m] + A[m] - U_V(H[m+1]), 0\right), \forall m \geq 1.$$

We note that $Q[m] = Q'[m] + A[m], \forall m \geq 1$. Since the sequence $\zeta[m] = A[m] - U_V(H[m+1])$ is ergodic, $\mathbb{E}\zeta[1] = -V < 0$, and $A[1] < A_{max}$, we have that $\lim_{m \rightarrow \infty} Pr\{Q[m] < \infty\} = 1$.

We note that the sequence of policies γ_V is a sequence of admissible policies, with $\lim_{V \downarrow 0} \bar{P}(\gamma_V) = c(\lambda)$. Therefore $c(\lambda) = \inf_{\gamma \in \Gamma_a} \bar{P}(\gamma)$.

Now we obtain the following asymptotic lower bound, the proof of which is similar to that of Lemma 4.3.23. We note that the quantities s_l, s_u and the line $l(s)$ are defined as in Section 5.3.

Lemma 5.6.1. For an ergodic arrival process $(A[m], m \geq 1)$, satisfying NA1 and NA2, an ergodic fading process $(H[m], m \geq 1)$, and for any sequence of admissible policies γ_k (satisfying NG1, NG2, and NG3), with $\bar{P}(\gamma_k) - c(\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$, for cases 2 and 3.

Proof. For some positive $\epsilon < s_l$, let $q_s \triangleq \inf\{q : \mathbb{E}S(q, H) \geq s_l - \epsilon\}$. Then, as in the proof of Lemma 5.3.1, we have that $U_k \triangleq \mathbb{E}_{\pi_{\gamma_k}} c(\bar{s}(Q)) - c(\lambda) \leq V_k$. We consider a particular policy γ in the sequence with $\mathbb{E}_{\pi} c(\bar{s}(Q)) - c(\lambda) = U$. As in the proof of Lemma 4.3.23 with U in place of V , we can then show that $\sum_{q=0}^{q_s-1} \pi(q) \leq \frac{U}{m_1 \epsilon}$ and $Pr\{S(q, H) > 0\} \geq \frac{s_l - \epsilon}{S_{max}}, \forall q \geq q_s$, since for $q \geq q_s, \mathbb{E}S(q, H) \geq s_l - \epsilon$. Now we relate the stationary probability $\pi(q), q \geq q_s$ to $\sum_{q=0}^{q_s-1} \pi(q)$. We have that for a $q \geq q_s$ and for every $m \geq 0$

$$Pr\{Q[m+1] < q | Q[0] = q_0\} = Pr\{Q[m] - S(Q[m], H[m+1]) + A[m+1] < q | Q[0] = q_0\}.$$

Let $\sigma[m] = (S[1], Q[1], \dots, Q[m-1], S[m])$.

We have that $Pr\{Q[m] - S(Q[m], H[m+1]) + A[m+1] < q | Q[0] = q_0\}$

$$\begin{aligned} &= \mathbb{E}_{\sigma[m]} \left[Pr\{Q[m] - S(Q[m], H[m+1]) + A[m+1] < q | Q[0] = q_0, \sigma[m]\} \right], \\ &\geq \mathbb{E}_{\sigma[m]} \left[Pr\{Q[m] = q | Q[0] = q_0, \sigma[m]\} \times \right. \\ &\quad Pr\{S(Q[m], H[m+1]) > 0 | Q[0] = q_0, \sigma[m], Q[m] = q\} \times \\ &\quad \left. Pr\{A[m+1] = 0 | Q[0] = q_0, \sigma[m], Q[m] = q, \{S(Q[m], H[m+1]) > 0\}\} \right]. \end{aligned}$$

We note that since $\sigma[m]$ does not involve the history of the fade process we have that

$$Pr \{S(q, H[m+1]) > 0 | Q[0] = q_0, \sigma[m], Q[m] = q\} \geq \frac{sl - \epsilon}{S_{max}}.$$

Since $(Q[0] = q_0, \sigma[m], Q[m] = q, \{S(Q[m]) > 0\}) = (Q[0] = q_0, A[1], Q[1], \dots, A[m], Q[m] = q)$, using property NA1, we obtain the same lower bound as in the proof of Lemma 4.3.23 on $Pr \{Q[m+1] < q | Q[0] = q_0\}$. Following the rest of the steps in the proof of Lemma 4.3.23, we obtain that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{U_k}\right)\right)$. Since $U_k \leq V_k$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

We note that NA2 is not used in the asymptotic lower bound, while it was used in showing that a sequence of admissible policies exist for which $\bar{P}(\gamma_k) \downarrow c(\lambda)$.

The above asymptotic lower bound is weak for Case 3, since we have obtained an $\Omega\left(\frac{1}{V}\right)$ asymptotic lower bound when $(A[m])$ is IID. We note that the above asymptotic lower bound can be applied to the system considered by Huang and Neely [30]. We note that a similar asymptotic lower bound can be derived for the R-model also.

5.7 Queueing models with admission control

5.7.1 System model

We consider the optimal tradeoff of average queue length and average power for a fading point-to-point link, when the packets arriving to the link can be dropped, subject to a constraint on the utility of the time average throughput of the packets which are transmitted. We indicate only the differences from the models in Sections 5.1.2 and 5.1.3. In each slot m , a random number $R[m]$ of packets arrive into the system. We assume that $(R[m], m \geq 1)$ are IID, with $R[1] \leq A_{max}$, $\mathbb{E}R[1] = \lambda$, and $var(R[1]) = \sigma^2 < \infty$. We also denote the expectation with respect to the distribution of $R[1]$ as \mathbb{E}_R . At the end of slot m , $A[m] \leq R[m]$ packets are admitted into the infinite length transmitter queue, while $R[m] - A[m]$ packets are dropped. The fade state process, $(H[m], m \geq 1)$, is as before. We also assume that the arrival process $(R[m])$ is independent of $(H[m])$. At the beginning of slot $m \geq 1$, $Q[m-1]$ denotes the number of packets in the transmitter queue. The transmitter starts transmission of $S[m]$ packets at the start of slot m . We assume that just before the end of slot m , $A[m] \leq R[m]$ packets are admitted into the transmitter queue. We note that under our assumptions on $(S[m]), (A[m])$, and $q_0, Q[m] \in \mathbb{R}_+, \forall m \geq 0$. For this model, a policy γ for operation of the transmitter is the sequence of service and arrival batch sizes $(S[1], A[1], S[2], A[2], \dots)$. The set of all policies is denoted as Γ . If γ is such that $S[m+1] = S(Q[m], H[m+1])$ and $A[m+1] = A(Q[m], R[m+1], H[m+1])$, where $S(q, h)$ and $A(q, r, h)$ are randomized functions, then γ is a stationary policy. The set of stationary policies is denoted by Γ_s . Since $(R[m], H[m], m \geq 1)$ is assumed to be IID, we have that for a $\gamma \in \Gamma_s$,

$(Q[m], m \geq 0)$ is a Markov chain evolving on \mathbb{R}_+ .

If we assume that $A[m], R[m], q_0, S[m] \in \mathbb{Z}_+$, then the queue evolution $Q[m] \in \mathbb{Z}_+$ and the model is denoted as I-model-U. On the other hand, if $A[m], R[m], q_0, S[m] \in \mathbb{R}_+$, then the queue evolution $Q[m] \in \mathbb{R}_+$, and the model is denoted as R-model-U. Like R-model, R-model-U with a strictly convex $P(h, s)$ function is usually used as an approximation for I-model-U.

We define the average throughput of a policy $\gamma \in \Gamma_s$ as

$$\bar{A}(\gamma, q_0) = \liminf_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M A[m] \middle| Q[0] = q_0 \right]. \quad (5.9)$$

Let $U(a) : [0, A_{max}] \rightarrow \mathbb{R}_+$ be a strictly concave and increasing function of a , with $U(0) = 0$. The utility of transmitting the packets is $U(\bar{A}(\gamma, q_0))$, for a policy γ . The average power for a policy $\gamma \in \Gamma_s$ is $\bar{P}(\gamma, q_0)$ and the average queue length is $\bar{Q}(\gamma, q_0)$, as defined in (5.2) and (5.3) respectively.

The model considered by Neely [44] is the same as R-model-U. It is shown in [44] that there exists a sequence of policies $\gamma_k \in \Gamma_s$ with a corresponding sequence $V_k \downarrow 0$, such that $\bar{A}(\gamma_k, q_0) \geq \rho\lambda$ ($0 < \rho < 1$), $\bar{Q}(\gamma_k, q_0) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$, and $\bar{P}(\gamma_k, q_0)$ is at most V_k more than the minimum average power required for queue stability. It is also shown in [44] for $|\mathcal{H}| = 1$, that if γ_k is any sequence of policies, with $\bar{P}(\gamma_k, q_0)$ at most V_k more than the minimum average power required for queue stability and $\bar{A}(\gamma_k, q_0) \geq \rho\lambda$, then $\bar{Q}(\gamma_k, q_0) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$ as $V_k \downarrow 0$.

We consider the optimal tradeoff between $\bar{Q}(\gamma, q_0)$ and $\bar{P}(\gamma, q_0)$ subject to the average utility $U(\bar{A}(\gamma, q_0))$ being at least a positive $u_c < U(A_{max})$, for the class of stationary policies Γ_s , for I-model-U and R-model-U, in this chapter. The constraint $U(\bar{A}(\gamma, q_0)) \geq u_c$ is equivalent to having the constraint $\bar{A}(\gamma, q_0) \geq U^{-1}(u_c)$, where U^{-1} is the inverse function of U . We note that since the arrival rate is not the same for all $\gamma \in \Gamma_s$, minimization of the average queue length does not directly correspond to minimizing the average delay of the packets. Asymptotic bounds on the average delay can be derived using Little's law and are discussed in this chapter.

5.7.2 Problem formulation for I-model-U and R-model-U

The general tradeoff problem that we consider is

$$\underset{\gamma \in \Gamma}{\text{minimize}} \bar{Q}(\gamma, q_0) \text{ such that } \bar{P}(\gamma, q_0) \leq P_c \text{ and } \bar{A}(\gamma, q_0) \geq \rho\lambda,$$

where $0 < \rho < 1$. As in Section 4.2.1 we can show that if $P_c > c(\rho\lambda)$ for I-model-U or if $P_c > c_R(\rho\lambda)$ for R-model-U (which are the minimum average powers required for stability while supporting an arrival rate of $\rho\lambda$ rather than λ) then there exists an optimal stationary policy γ^* with stationary probability π^* . Therefore, we can restrict ourselves to the set of stationary policies.

As for I-model and R-model, we consider the above tradeoff problem for a set of admissible policies Γ_a . However, since there is admission control, we relax the irreducibility requirement (for I-model-U and R-model-U) as follows. For an admissible policy γ , the Markov chain $(Q[m], m \geq 0)$ has a single positive recurrent class \mathcal{R}_γ which contains 0. Furthermore, the cumulative expected queue cost as well as the cumulative expected power cost starting from any state q_0 until \mathcal{R}_γ is hit are finite.

We note that for a $\gamma \in \Gamma_a$, $\bar{Q}(\gamma, q_0) = \bar{Q}(\gamma)$, $\bar{P}(\gamma, q_0) = \bar{P}(\gamma)$, and $\bar{A}(\gamma, q_0) = \bar{A}(\gamma)$. A policy γ is defined to be admissible if: (i) the Markov process $(Q[m], m \geq 0)$ under γ is aperiodic and positive Harris recurrent on a single recurrence class \mathcal{R}_γ with stationary distribution π^3 , (ii) $\bar{Q}(\gamma) < \infty$, and (iii) $\bar{s}(q)$ is non-decreasing in q , where $\bar{s}(q) = \mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} S(q, H)$ is the average service rate at queue length q . We note that for a $\gamma \in \Gamma_a$,

$$\begin{aligned}\bar{Q}(\gamma) &= \mathbb{E}_\pi Q, \\ \bar{P}(\gamma) &= \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} P(H, S(Q, H)), \\ \bar{A}(\gamma) &= \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_R EA(Q, R, H).\end{aligned}$$

Let the average service rate be $\bar{S}(\gamma)$, then $\bar{S}(\gamma) = \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} S(Q, H)$.

The problem TRADEOFF that we consider is

$$\underset{\gamma \in \Gamma_a}{\text{minimize}} \bar{Q}(\gamma) \text{ such that } \bar{P}(\gamma) \leq P_c \text{ and } \bar{A}(\gamma) \geq \rho\lambda.$$

The optimal value of TRADEOFF is denoted as $Q^*(P_c, \rho)$. Suppose γ is feasible for TRADEOFF. Then

$$\bar{S}(\gamma) = \bar{A}(\gamma) \geq \rho\lambda.$$

Now we note that $\bar{P}(\gamma)$ is bounded below by the optimal value of

$$\begin{aligned}\underset{\gamma \in \Gamma_a}{\text{minimize}} & \quad \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} P(H, S(Q, H)), \\ \text{such that} & \quad \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} S(Q, H) \geq \rho\lambda.\end{aligned}\tag{5.10}$$

We note that $\mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} S(Q, H) = \mathbb{E}_{\pi_H} \mathbb{E}_\pi \mathbb{E}_{S|Q,H} S(Q, H)$. We have that $\mathbb{E}_{\pi_H} \mathbb{E}_\pi \mathbb{E}_{S|Q,H} S(Q, H) = \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S$ and $\mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} P(H, S(Q, H)) = \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S)$, where the conditional distribution of S given H depends on the policy γ . Then the optimal value of (5.10) is bounded below by

³We note that $\pi(A) = 0$ for any $A \notin \mathcal{R}_\gamma$

the optimal value of

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S), \\ & \text{such that} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S \geq \rho\lambda, \end{aligned} \tag{5.11}$$

where we minimize over all possible conditional distributions for S given h , irrespective of the policy γ . We note that for I-model-U, these distributions have support on $\{0, \dots, S_{max}\}$, whereas for R-model-U they have support on $[0, S_{max}]$.

We note that (5.11) has feasible solutions only if $\rho\lambda \leq S_{max}$. The optimal value of the above problem is $c(\rho\lambda)$ for I-model-U and $c_R(\rho\lambda)$ for R-model-U, since the constraint is satisfied with equality⁴. So, we have that for $\gamma \in \Gamma_a$, $\bar{P}(\gamma) \geq c(\rho\lambda)$ for I-model-U and $\bar{P}(\gamma) \geq c_R(\rho\lambda)$ for R-model-U. Thus, TRADEOFF has feasible solutions only if $P_c \geq c(\rho\lambda)$ for I-model-U and $P_c \geq c_R(\rho\lambda)$ for R-model-U.

We now show that $c(\rho\lambda)$ and $c_R(\rho\lambda)$ are both $\inf_{\{\gamma: \gamma \in \Gamma_a, \bar{A}(\gamma) \geq \rho\lambda\}} \bar{P}(\gamma)$ for I-model-U and R-model-U respectively. For I-model-U, we consider a sequence of policies γ_k , where for each γ_k , at each slot m , each customer in the batch $R[m]$ is admitted with probability ρ and dropped with probability $1 - \rho$. Then $\forall \gamma_k$ we have that $\bar{A}(\gamma_k) \geq \rho\lambda$. Now as for I-model, γ_k is such that $\bar{P}(\gamma_k) = c(\rho\lambda) + V_k$ and $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. Thus, if $\rho\lambda < S_{max}$, we have that there exists a sequence of admissible policies γ_k , such that $\bar{P}(\gamma_k) = c(\rho\lambda) + V_k$, $\bar{A}(\gamma_k) \geq \rho\lambda$, and $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$, for a sequence $V_k \downarrow 0$. For R-model-U, we choose $A[m] = \rho R[m]$, and then as for R-model, serve the customers using a sequence of policies such that $\bar{P}(\gamma_k) = c_R(\rho\lambda) + V_k$ and $\bar{Q}(\gamma_k) = \mathcal{O}\left(\frac{1}{V_k}\right)$. Hence, $c(\rho\lambda)$ and $c_R(\rho\lambda)$ are $\inf_{\{\gamma: \gamma \in \Gamma_a, \bar{A}(\gamma) \geq \rho\lambda\}} \bar{P}(\gamma)$ for I-model-U and R-model-U respectively.

In the following, we obtain an asymptotic characterization of $Q^*(P_c, \rho)$ in the asymptotic regimes \mathfrak{R} as $P_c \downarrow c(\rho\lambda)$ for I-model-U and $P_c \downarrow c_R(\rho\lambda)$ for R-model-U, under the assumption that $\rho\lambda < S_{max}$. We recall that $c(s)$ is a non-decreasing, piecewise linear, and convex function of $s \in [0, S_{max}]$, whereas $c_R(s)$ is a non-decreasing strictly convex function of $s \in [0, S_{max}]$, with $c(0)$ and $c_R(0)$ both being 0.

5.7.3 Asymptotic bounds

We first obtain an asymptotic lower bound for R-model-U. We then outline the derivation of the asymptotic lower bound for I-model-U, since it can be obtained using very similar techniques as for R-model-U and as in Section 5.3.

We assume that $R[1]$ satisfies the property:

⁴If the distribution which achieves the minimum in (5.11) is such that $\mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S > \rho\lambda$, then it is possible to show that there exists another distribution which has a strictly smaller $\mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S)$.

RA* : $Pr \{R[1] \leq \frac{\Delta}{2}\} = \epsilon'_a > 0$, for some Δ such that $0 < \Delta < \rho\lambda$.

Let γ be an admissible policy with $\bar{P}(\gamma) - c_R(\rho\lambda) = V$ and $\bar{A}(\gamma) \geq \rho\lambda$. For the policy γ , let $q_1 \triangleq \sup \{q : \bar{s}(q) \leq \rho\lambda - \epsilon\}$, for an ϵ chosen such that $0 < \epsilon < \rho\lambda - \Delta$. We note that the average power used when the queue length is q is $\mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} P(H, S(q, H))$, which is bounded below by the optimal value of

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S), \\ & \text{such that} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S \geq \bar{s}(q), \end{aligned} \quad (5.12)$$

where we minimize over all possible conditional distributions for S given h , irrespective of γ . We note that problem (5.12) is the same as (5.11) except that the constraint is $\bar{s}(q)$ instead of $\rho\lambda$. Therefore $\mathbb{E}_{\pi_H} \mathbb{E}_{S|q,H} P(H, S(q, H)) \geq c_R(\bar{s}(q))$. Since $\mathbb{E}_{\pi} \mathbb{E}_{\pi_H} \mathbb{E}_{S|Q,H} P(H, S(Q, H)) = \bar{P}(\gamma)$, we have that $\bar{P}(\gamma) \geq \mathbb{E}_{\pi} c_R(\bar{s}(Q))$. Therefore, $\mathbb{E}_{\pi} c_R(\bar{s}(Q)) - c_R(\rho\lambda) \leq \bar{P}(\gamma) - c_R(\rho\lambda) \leq V$. Since $c_R(\cdot)$ is a strictly convex and non-decreasing function, we assume that the second derivative of $c_R(s)$ is positive at $s = \rho\lambda$. As $\bar{S}(\gamma) = \mathbb{E}_{\pi} \bar{s}(Q) \geq \rho\lambda$, we have that

$$\mathbb{E}_{\pi} c_R(\bar{s}(Q)) = \mathbb{E}_{\pi} \left[c_R(\rho\lambda) + \left. \frac{dc_R(x)}{dx} \right|_{\rho\lambda} (\bar{s}(Q) - \rho\lambda) + G(\bar{s}(Q) - \rho\lambda) \right]$$

where $G(x)$ is a strictly convex function as in [7, eq (41)], with $G(0) = 0$ and $\frac{dG(x)}{dx}|_{x=0} = 0$. Since $\mathbb{E}_{\pi} \bar{s}(Q) \geq \rho\lambda$ we have that

$$\mathbb{E}_{\pi} c_R(\bar{s}(Q)) - c_R(\rho\lambda) \geq \mathbb{E}_{\pi} G(\bar{s}(Q) - \rho\lambda).$$

Thus, we have that $\mathbb{E}_{\pi} G(\bar{s}(Q) - \rho\lambda) \leq V$. Therefore, for q_1 as defined before,

$$\int_0^{q_1} G(\bar{s}(q) - \rho\lambda) d\pi(q) \leq V.$$

Since $G(\cdot)$ is strictly convex, for a positive a_1 we have that

$$\left[\int_0^{q_1} (\bar{s}(q) - \rho\lambda) d\pi(q) \right]^2 \leq \frac{V}{a_1},$$

where $a_1 > 0$. Since for $q < q_1$, $\bar{s}(q) \leq \rho\lambda - \epsilon$, we have that

$$Pr \{Q < q_1\} \leq \frac{V}{a_1 \epsilon^2}. \quad (5.13)$$

Let $S(q)$ be the random service batch size when in state q . We have that $Pr \{S(q) > s\} = \int_h Pr \{S(q, h) > s\} d\pi_H(h)$.

Lemma 5.7.1. For γ , for Δ as in RA*, with q_1 defined as above, we have that $\inf_{q > q_1} Pr \{S(q) > \Delta\} \geq$

$\delta_s > 0$, where $\delta_s = \frac{\rho\lambda - \epsilon - \Delta}{S_{max} - \Delta}$.

Proof. We note that $q_1 = \sup \{q : \bar{s}(q) \leq \rho\lambda - \epsilon\}$, for $0 < \epsilon < \rho\lambda - \Delta$, where Δ is as in RA*. Let $P(S(q))$ denote the distribution of $S(q)$. We note that by definition, $\forall q > q_1$,

$$\begin{aligned} \bar{s}(q) &> \rho\lambda - \epsilon, \text{ or,} \\ \int_0^{S_{max}} s dP(S(q)) &\geq \rho\lambda - \epsilon. \end{aligned} \quad (5.14)$$

Then,

$$\begin{aligned} \int_0^{\Delta} \Delta dP(S(q)) + \int_{\Delta}^{S_{max}} S_{max} dP(S(q)) &\geq \rho\lambda - \epsilon, \\ \Delta(1 - Pr\{S(q) > \Delta\}) + S_{max} Pr\{S(q) > \Delta\} &\geq \rho\lambda - \epsilon, \text{ or,} \\ Pr\{S(q) > \Delta\} &\geq \frac{\rho\lambda - \epsilon - \Delta}{S_{max} - \Delta}. \end{aligned}$$

Thus for any $q > q_1$, $Pr\{S(q) > \Delta\} \geq \delta_s > 0$, where $\delta_s = \frac{\rho\lambda - \epsilon - \Delta}{S_{max} - \Delta}$. \square

From the above result we have that $q_1 \geq \Delta$.

Lemma 5.7.2. For any sequence of admissible policies γ_k such that $\bar{A}(\gamma_k) \geq \rho\lambda$ and $\bar{P}(\gamma_k) - c_R(\rho\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

Proof. Let us consider a particular policy γ in the sequence γ_k with $V_k = V$. As γ is admissible we have that

$$Pr\{0 \leq Q < q_1\} = \int_0^{\infty} P(q, [0, q_1]) d\pi(q),$$

where $P(q, Q)$ is the transition kernel of the Markov chain. Hence, we have that

$$\begin{aligned} Pr\{0 \leq Q < q_1\} &\geq \int_{q_1}^{q_1 + \frac{\Delta}{2}} P(q, [0, q_1]) d\pi(q), \\ &\geq \epsilon'_a \delta_s Pr\left\{q_1 \leq Q < q_1 + \frac{\Delta}{2}\right\}, \end{aligned}$$

where we have used Lemma 5.7.1 and the property RA* to lower bound $Pr\{Q[m+1] < q_1 | Q[m] = q\}$ by $\epsilon'_a \delta_s$. Let $\rho_a = \epsilon'_a \delta_s$. Also, for any $q' > q$, let us denote $Pr\{q \leq Q < q'\}$ by $\pi[q, q']$.

Then we have obtained that

$$\pi[0, q_1] \geq \rho_a \pi\left[q_1, q_1 + \frac{\Delta}{2}\right).$$

Similarly, we have that

$$\pi\left[0, q_1 + \frac{\Delta}{2}\right] \geq \rho_a \pi\left[q_1 + \frac{\Delta}{2}, q_1 + 2\frac{\Delta}{2}\right),$$

which can be written as

$$\begin{aligned}\pi[0, q_1) + \pi\left[q_1, q_1 + \frac{\Delta}{2}\right) &\geq \rho_a \pi\left[q_1 + \frac{\Delta}{2}, q_1 + 2\frac{\Delta}{2}\right), \\ \pi[0, q_1) \left[1 + \frac{1}{\rho_a}\right] &\geq \rho_a \pi\left[q_1 + \frac{\Delta}{2}, q_1 + 2\frac{\Delta}{2}\right).\end{aligned}$$

By induction, for $m \geq 0$, we have that

$$\frac{\pi[0, q_1)}{\rho_a} \left(1 + \frac{1}{\rho_a}\right)^m \geq \pi\left[q_1 + m\frac{\Delta}{2}, q_1 + (m+1)\frac{\Delta}{2}\right).$$

Hence, we have that for $m \geq 0$,

$$\begin{aligned}\pi\left[q_1, q_1 + m\frac{\Delta}{2}\right) &= \begin{cases} \sum_{k=0}^{m-1} \pi\left[q_1 + k\frac{\Delta}{2}, q_1 + (k+1)\frac{\Delta}{2}\right), & \text{if } m > 0, \\ 0, & \text{if } m = 0, \end{cases} \\ &\leq \frac{\pi[0, q_1)}{\rho_a} \frac{\left(1 + \frac{1}{\rho_a}\right)^m - 1}{1 + \frac{1}{\rho_a} - 1} = \pi[0, q_1) \left[\left(1 + \frac{1}{\rho_a}\right)^m - 1\right]. \end{aligned} \quad (5.15)$$

Since $Pr\{Q < q_1\} \leq \frac{V}{a_1\epsilon^2}$ (from (5.13)), if m is the largest integer such that

$$\pi[0, q_1) + \pi\left[q_1, q_1 + m\frac{\Delta}{2}\right) \leq \frac{1}{2},$$

then $\bar{Q}(\gamma) \geq \frac{m\Delta}{4}$. Suppose m_1 is the largest integer such that

$$\pi[0, q_1) + \pi[0, q_1) \left[\left(1 + \frac{1}{\rho_a}\right)^{m_1} - 1\right] = \pi[0, q_1) \left(1 + \frac{1}{\rho_a}\right)^{m_1} \leq \frac{1}{2}. \quad (5.16)$$

Then $m_1 \leq m$. Using (5.13), if m_2 is the largest integer such that

$$\begin{aligned}\left(1 + \frac{1}{\rho_a}\right)^{m_2} &\leq \frac{a_1\epsilon^2}{2V}, \text{ or,} \\ m_2 &\leq \log_{\left(1 + \frac{1}{\rho_a}\right)}\left(\frac{a_1\epsilon^2}{2V}\right),\end{aligned}$$

then $m_2 \leq m_1$. We have that

$$m_2 = \left\lfloor \log_{\left(1 + \frac{1}{\rho_a}\right)}\left(\frac{a_1\epsilon^2}{2V}\right) \right\rfloor.$$

Since $\bar{Q}(\gamma) \geq \frac{m\Delta}{4} \geq \frac{m_1\Delta}{4} \geq \frac{m_2\Delta}{4}$, we obtain that

$$\bar{Q}(\gamma) \geq \frac{\Delta}{4} \left(\log_{\left(1 + \frac{1}{\rho_a}\right)}\left(\frac{a_1\epsilon^2}{2V}\right) - 1 \right).$$

So for the sequence of policies γ_k with $V_k \downarrow 0$ we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

We now outline the derivation of an asymptotic lower bound for I-model-U in the regime \mathfrak{R} . The analysis of I-model-U proceeds in a similar fashion as in Section 5.3; the piecewise linear function $c(\lambda)$ and the quantities $a_p, p \geq 1$ are similarly defined. The three cases which then arise are : (1) $0 < \rho\lambda < a_2$, (2), $a_p < \rho\lambda < a_{p+1}, p > 1$, and (3) $\rho\lambda = a_p, p > 1$. For Cases 2 and 3, proceeding similarly as in the proof of Lemma 5.3.1, it is possible to show that, for any sequence of admissible policies γ_k such that $\bar{A}(\gamma_k) \geq \rho\lambda$ and $\bar{P}(\gamma_k) - c(\rho\lambda) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. We note that we do not have any asymptotic results for Case 1, although numerically it can be shown that $Q^*(P_c) < \infty$ even if $P_c = c(\rho\lambda)$.

Remark 5.7.3. We note that a similar logarithmic asymptotic lower bound can be obtained for a different class of admissible policies. The difference is in the definition of the monotonicity property. For this new class of admissible policies, the average drift $\mathbb{E}S(q, H) - \mathbb{E}A(q, R, H)$ is assumed to be monotonically non-increasing in q .

Remark 5.7.4. We comment on an asymptotic upper bound for TRADEOFF, which is achieved by the sequence of Dynamic Packet Dropping (DPD) policies in [44]. A DPD policy is parametrized by the quantities β, ϵ, ω , and \bar{q} . The policy chooses a batch size $s_{DPD}[m]$ in each slot where

$$s_{DPD}[m] = \min \left(\arg \max_{s \in \{0, \dots, S_{max}\}} \left[s \left\{ X[m-1] - \omega e^{\omega(\bar{q} - Q[m-1])} \right\} - \beta P(H[m], s) \right], Q[m-1] \right),$$

where $(X[m])$ is a *virtual* queue which evolves according to

$$X[m] = \max(X[m-1] - s_{DPD}[m], 0) + (\rho + \epsilon)R[m].$$

For the DPD policy, through admission control, the queue length process $(Q[m], m \geq 0)$ evolves as

$$Q[m] = \min[\bar{q}, Q[m-1] - s_{DPD}[m] + R[m]],$$

We note that $A[m] = R[m]$ whenever $Q[m] \leq \bar{q}$, otherwise only that fraction of $R[m]$ is admitted so that $Q[m] = \bar{q}$. We note that $A(q, r, h)$ for this DPD policy, is a function only of the current queue length q and the current number of arrivals r .

From Theorem 1 [44], if $0 < \omega$ and $\omega e^{\omega S_{max}} \leq \frac{\lambda(1-\rho-\epsilon)}{\sigma^2}$, $\epsilon = \frac{1-\rho}{2\beta}$, $B = \frac{S_{max}^2 + (\rho+\epsilon)^2 A_{max}^2}{2} + 1$, $x \geq \frac{4S_{max}\epsilon\omega S_{max}B}{\lambda^2\omega(1-\rho-\epsilon)(1-\rho)}$, and $\bar{q} = \frac{\log(x\beta)}{\omega}$, then for the sequence of policies γ obtained by a sequence $\beta \uparrow \infty$, we have that $\bar{Q}(\gamma) = \mathcal{O}(\log(\beta))$, $\bar{P}(\gamma) = c_R(\rho\lambda) + \mathcal{O}\left(\frac{1}{\beta}\right)$, and $\bar{A}(\gamma) \geq \rho\lambda$.

We note that DPD policies are not stationary, since each policy depends on an auxiliary state $X[m]$. However, as noted in [44, Section III], using a sequence of admissible policies which are obtained from the admissible Positive-Drift Algorithm in [44] by choosing the parameter Q as $\log\left(\frac{1}{V}\right)$, where

$V \downarrow 0$, it can be shown that the above tradeoff is achievable. Therefore, the asymptotic lower bound derived in Lemma 5.7.2 is tight.

5.7.4 Discussion

Minimization of average delay: When average delay is the performance measure under consideration, then the problem that we are interested in is

$$\text{minimize}_{\gamma \in \Gamma_a} \frac{\bar{Q}(\gamma)}{\bar{A}(\gamma)} \text{ such that } \bar{P}(\gamma) \leq P_c \text{ and } \bar{A}(\gamma) \geq \rho\lambda.$$

Let the optimal value of the above problem be $D^*(P_c, \rho)$.

We note that, since $c(s)$ or $c_R(s)$ is a convex and non-decreasing function in $s \in [0, S_{max}]$, for any admissible policy γ , if γ is feasible for the above problem, we have that

$$\begin{aligned} \mathbb{E}_\pi c(\bar{s}(Q)) &\leq \bar{P}(\gamma) \leq P_c, \\ c(\mathbb{E}_\pi \bar{s}(Q)) &\leq P_c, \text{ or,} \\ \bar{S}(\gamma) = \mathbb{E}_\pi \bar{s}(Q) &\leq c^{-1}(P_c), \end{aligned}$$

where c^{-1} is the inverse function of c for I-model-U. Consider any sequence $P_{c,k} \downarrow c(\rho\lambda)$ as $k \uparrow \infty$. Since $\bar{A}(\gamma) = \bar{S}(\gamma)$, the objective function in the above optimization problem can be bounded above by $\frac{\bar{Q}(\gamma)}{\rho\lambda}$ and bounded below by $\frac{\bar{Q}(\gamma)}{c^{-1}(P_{c,1})}$. A similar bound can be obtained for R-model-U. Then, it follows that the asymptotic behaviour of $D^*(P_{c,k}, \rho)$ is the same as that of $Q^*(P_{c,k}, \rho)$ as $P_{c,k} \downarrow c(\rho\lambda)$ for I-model-U and $P_{c,k} \downarrow c_R(\rho\lambda)$ for R-model-U.

Relation to the asymptotic order optimal tradeoff in [42]: Neely [42] considers a system, with both admission control and service rate control, in which the arrival rate λ is larger than the maximum service rate S_{max} . The objective is to obtain a sequence of policies γ_k which achieve an order optimal minimum average queue length $\bar{Q}(\gamma_k)$ as the average utility $U(\bar{S}(\gamma_k))$ approaches the maximum utility value $u_{max} = U(S_{max})$. We note that there is no cost associated with the service of packets in [42]. It is shown that for any sequence of policies γ_k such that $u_{max} - U(\bar{S}(\gamma_k)) = V_k \downarrow 0$, $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. A sequence of policies γ_k such that $\bar{Q}(\gamma_k) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$ and $U(\bar{S}(\gamma_k)) = u_{max} - V_k$ is also obtained. We note that as the utility function is assumed to be strictly concave and increasing, the throughput value that maximizes the utility is S_{max} itself. For any sequence γ_k , if $\bar{U}(\gamma_k) \uparrow u_{max} = U(S_{max})$ it can be shown that the probability of using a service rate less than S_{max} decreases to zero. That is, with $q_1 \triangleq \sup\{q : \bar{s}(q) \leq S_{max} - \epsilon\}$, $Pr\{Q < q_1\} \downarrow 0$. Therefore, the proof of Lemma 5.7.2 can be applied to obtain an alternate proof for the asymptotic logarithmic lower bound on the average queue length obtained in [42], but for admissible policies.

Asymptotic bounds for a model with just admission control: We consider a queueing model in which there is no service batch size control and no service cost. The queue evolution is assumed to be as follows:

$$Q[m+1] = \max(Q[m] - s_b, 0) + A[m+1],$$

where $A[m+1] \leq R[m+1]$ and $s_b \leq S_{max}$ is a fixed batch size. The queue evolution is assumed to be on \mathbb{Z}_+ . We assume that $\mathbb{E}R[1] = \lambda > s_b$. We consider the case with a single fade state. We consider the tradeoff of average queue length and average throughput in the asymptotic regime where the average throughput approaches its maximum value s_b .

As in the previous sections, we consider this tradeoff problem for a class of admissible policies Γ_a . An admissible policy for this problem is a stationary policy which is stable (as defined in Section 5.7.2). However, we note that there is no service batch size control, and it is assumed that the average admitted rate at a queue length q , $\bar{a}(q) = \mathbb{E}A(q, R)$ is a non-increasing function of q . The tradeoff problem that we consider is:

$$\underset{\gamma \in \Gamma_a}{\text{minimize}} \bar{Q}(\gamma) \text{ such that } \bar{A}(\gamma) \geq t_c. \quad (5.17)$$

We note for any admissible policy γ , $\bar{A}(\gamma) \leq s_b$. Consider a sequence of policies γ_k defined as follows. A policy γ_k , in slot m , admits a packet from the batch of size $R[m]$ with probability ρ_k or rejects the packet with probability $1 - \rho_k$. Let the sequence $\rho_k \triangleq \frac{s_b - V_k}{\lambda}$, for a sequence $V_k \downarrow 0$. Since $\bar{A}(\gamma_k) = \rho_k \lambda < s_b$ we have that γ_k is a sequence of admissible policies with $\bar{A}(\gamma_k) \uparrow s_b$. Therefore, we have that $s_b = \sup_{\gamma \in \Gamma_a} \bar{A}(\gamma)$. We consider the tradeoff problem (5.17) in the asymptotic regime \mathfrak{R} where $t_c \uparrow s_b$.

We have the following asymptotic lower bound.

Lemma 5.7.5. For any sequence of admissible policies γ_k , with $s_b - \bar{A}(\gamma_k) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$.

Proof. Consider a policy γ in the sequence γ_k with $s_b - \bar{A}(\gamma) = V$. Since $\bar{S}(\gamma) = \bar{A}(\gamma)$, we have that

$$(s_b - 1)Pr\{s(Q) \neq s_b\} + s_b Pr\{s(Q) = s_b\} \geq s_b - V, \text{ or,}$$

$$Pr\{s(Q) \neq s_b\} = Pr\{s(Q) < s_b\} \leq V.$$

We note that $s(q) = \min(q, s_b)$.

We now proceed as in the proof of Lemma 4.3.8, but with the following changes. We define q_s to be s_b . Then, we have that $Pr\{Q < q_s\} = Pr\{s(Q) < s_b\} \leq V$.

Now we obtain a geometric upper bound on $\pi(q)$ for $q \geq q_s$ as in the proof of Lemma 4.3.5. We note that for $q \geq q_s$, we have that

$$Pr \{Q[m+1] < q | Q[m] = q\} \geq Pr \{A[1] = 0\},$$

since $s(q) = s_b$. Therefore, with ρ_d redefined to be just $Pr \{A[1] = 0\}$ in the proof of Lemma 4.3.5, we obtain that $\pi(q) \leq Pr \{Q < q_s\} \frac{\rho^k}{\rho_d}$ for $q = q_s + k$ and $k \geq 0$. Now proceeding as in the proof of Lemmas 4.3.5 and 4.3.8, we have that for the sequence γ_k , $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

We note that this model is the discrete time equivalent of INTERVAL- λ CHOICE. The problem considered here corresponds to INTERVAL- λ CHOICE-2-1.

5.8 Single hop networks

In this section, we illustrate how asymptotic lower bounds for the tradeoff of average power and total average queue length can be derived for a system with N source destination pairs communicating over single hop links. For example, this could be a N user multiple access or N user broadcast channel. The model and the associated tradeoff problem that we consider is motivated by Neely [43], who considered the problem of optimally trading off average power with average delay for a wireless downlink system, with no admission control.

5.8.1 System model

We first consider a model with real-valued queue evolution. We note that the model for the single hop network, with no admission control, is a straightforward extension of R-model in Section 5.1.3. We assume that there is an IID arrival process $(A_n[m], m \geq 1)$ to the queue for the n^{th} link, $n \in \{1, \dots, N\}$. We assume that $A_n[m] \leq A_{max}, \forall n$. The arrival processes to different links are assumed to be independent of each other. The n^{th} link is subjected to an IID fading process $(H_n[m], m \geq 1)$. The fading processes are assumed to be independent across links. The queue length at the start of the m^{th} slot for the n^{th} link is denoted as $Q_n[m-1]$. In the following, we use the notation \mathbf{X} to denote the vector (X_1, X_2, \dots, X_N) . The arrival rate vector is $\boldsymbol{\lambda} = (\mathbb{E}A_1[1], \mathbb{E}A_2[1], \dots, \mathbb{E}A_N[1])$. The fade state is assumed to take values in \mathcal{H} , with $\min\{\min_{n \in \{1, \dots, N\}} \{h_n^2\}, \mathbf{h} \in \mathcal{H}\} > 0$. The distribution of the fade state is denoted as $\pi_{\mathcal{H}}$ and the expectation with respect to this distribution as $\mathbb{E}_{\pi_{\mathcal{H}}}$. We assume that $|\mathcal{H}| < \infty$.

In each slot m , a service batch size vector $\mathbf{S}[m]$ is chosen as a randomized vector function $\mathbf{S}(\mathbf{Q}[m-1], \mathbf{H}[m])$ of the current queue length vector $\mathbf{Q}[m-1]$ and the current fade state vector $\mathbf{H}[m]$. We assume that $S_n[m] \leq S_{max}, \forall n, m$. A policy γ is the choice of the randomized function

$\mathbf{S}(\mathbf{q}, \mathbf{h}), \forall \mathbf{q}, \mathbf{h}$. The evolution of the N queues under γ is given by:

$$\mathbf{Q}[m+1] = \mathbf{Q}[m] - \mathbf{S}[m+1] + \mathbf{A}[m+1], \quad (5.18)$$

with $\mathbf{Q}[0] = \mathbf{q}_0$. We note that the process $(\mathbf{Q}[m])$ is a Markov process evolving on \mathbb{R}_+^N . The use of the service vector \mathbf{s} incurs a power $P(\mathbf{h}, \mathbf{s})$ when the fade state is \mathbf{h} . Similar to the properties (C1) and (C2), we assume that for every \mathbf{h} , $P(\mathbf{h}, \mathbf{0}) = 0$ and $P(\mathbf{h}, \mathbf{s})$ is a strictly convex function of $\mathbf{s}, \forall \mathbf{s} \in [0, S_{max}]^N$. These assumptions can be motivated by the properties of $P(\mathbf{h}, \mathbf{s})$ obtained in [6, Chapter 7]. The average power $\bar{P}(\gamma, \mathbf{q}_0)$ and total average queue length $\bar{Q}(\gamma, \mathbf{q}_0)$ are defined as:

$$\begin{aligned} \bar{P}(\gamma, \mathbf{q}_0) &\triangleq \lim_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=1}^M P(\mathbf{H}[m], \mathbf{S}[m]) \middle| \mathbf{Q}[0] = \mathbf{q}_0 \right], \\ \bar{Q}(\gamma, \mathbf{q}_0) &\triangleq \lim_{M \rightarrow \infty} \frac{1}{M} \mathbb{E} \left[\sum_{m=0}^{M-1} \sum_{n=1}^N Q_n[m] \middle| \mathbf{Q}[0] = \mathbf{q}_0 \right]. \end{aligned}$$

5.8.2 Problem formulation

We study the tradeoff of $\bar{P}(\gamma, \mathbf{q}_0)$ with $\bar{Q}(\gamma, \mathbf{q}_0)$ for a restricted set Γ_a of admissible policies. A policy γ is admissible if (i) it is stable with stationary distribution $\pi(\mathbf{q})$ (where stability is defined similarly as in Section 5.2), and (ii) instead of property G2, γ is such that

MG2 : the average service rate for queue n as a function of its queue length, $\bar{s}_n(q)$ is non-decreasing in q .

For example, for $N = 2$ and for the first queue, for an admissible policy γ we require that $\bar{s}_1(q_1) = \mathbb{E}_{Q_2|q_1} \mathbb{E}_{\pi_{\mathbf{H}}} \mathbb{E}_{S_1|q_1, Q_2, \mathbf{H}} S_1$ is non-decreasing in q_1 . We note that for a $\gamma \in \Gamma_a$, $\bar{P}(\gamma, \mathbf{q}_0) = \bar{P}(\gamma)$ and $\bar{Q}(\gamma, \mathbf{q}_0) = \bar{Q}(\gamma)$. As for the single link case, we study TRADEOFF for the single hop network, defined as:

$$\text{minimize}_{\gamma \in \Gamma_a} \bar{Q}(\gamma), \text{ such that } \bar{P}(\gamma) \leq P_c.$$

We define the function $c_R(\boldsymbol{\lambda})$ as the optimal value of

$$\begin{aligned} &\text{minimize} \quad \mathbb{E}_{\pi_{\mathbf{H}}} \mathbb{E}_{\mathbf{S}|\mathbf{H}} P(\mathbf{H}, \mathbf{S}), \\ &\text{such that} \quad \mathbb{E}_{\pi_{\mathbf{H}}} \mathbb{E}_{\mathbf{S}|\mathbf{H}} \mathbf{S} = \boldsymbol{\lambda}. \end{aligned} \quad (5.19)$$

We note that $c_R(\boldsymbol{\lambda})$ is similar to $c_R(\lambda)$ for the single link case. From [43], we have that $c_R(\boldsymbol{\lambda}) = \inf_{\gamma \in \Gamma_a} \bar{P}(\gamma)$, and $c_R(\boldsymbol{\lambda})$ is a strictly convex function of $\boldsymbol{\lambda}$, for $\boldsymbol{\lambda} \in [0, S_{max}]^N$, with $c_R(\mathbf{0}) = 0$. We now derive an asymptotic lower bound on the average queue length $\bar{Q}(\gamma_k)$ for any sequence of admissible policies γ_k for which $\bar{P}(\gamma_k) - c_R(\boldsymbol{\lambda}) \downarrow 0$. We note that the multiuser Berry-Gallager asymptotic lower bound [43, Theorem 2], is rederived in the following, with the extra assumption

MG2. The method of derivation illustrates how the asymptotic lower bounding technique can be extended to network scenarios. For ease of exposition, the asymptotic lower bound is derived for the case $N = 2$, but can be extended to any finite N .

5.8.3 Asymptotic lower bound

We assume that $A_n[1]$ satisfies

MA1 : $Pr \{A_n[1] - S_{max} > \delta_{n,a}\} > \epsilon_{n,a}$, for some positive $\delta_{n,a}$ and $\epsilon_{n,a}$, $\forall n$.

We first obtain a lower bound on the marginal stationary probability $Pr \{Q_1 \geq q\} = \int_{q_2=0}^{\infty} \int_{q_1=q}^{\infty} d\pi(q_1, q_2)$ of the two-dimensional Markov process, which is similar to Lemma 4.4.1. The lower bound on the marginal stationary probability is obtained for the first queue, but can be obtained for the second queue also by interchanging the indices of the two queues.

Lemma 5.8.1. Let $(Q[m], m \geq 0)$ be as in (5.18), evolving on \mathbb{R}_+^2 , with stationary probability π , for a $\gamma \in \Gamma_a$. Suppose there exists a $q_{1,d}$ such that

$$\forall q_1 \in [0, q_{1,d}], \mathbb{E}_{Q_2, \mathbf{H}} \left[\mathbb{E} \left[Q_1[m+1] - Q_1[m] \middle| Q_2[m], \mathbf{H}[m+1], Q_1[m] = q_1 \right] \right] \geq -d,$$

where d is positive. Then for any $\bar{q}_1, k \geq 0, \Delta > 0, \delta > 0, \Delta + \delta < \delta_{1,a}$, and $0 \leq \bar{q}_1 + k\Delta \leq q_{1,d}$,

$$\begin{aligned} Pr \{Q_1 \geq \bar{q}_1 + k\Delta\} &\geq \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + d} \right)^k Pr \{Q_1 \geq \bar{q}_1\} \\ &+ \left[1 - \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + d} \right)^k \right] \left[Pr \{Q_1 \geq q_{1,d}\} + \frac{1}{d} \int_{q_{1,d}}^{\infty} (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1) \right], \end{aligned}$$

where $\pi(q_1)$ is the marginal probability distribution of q_1 .

The proof is very similar to that of Lemma 4.4.1 and is presented in Appendix 5.C. Utilizing the above lower bound, we obtain the following result which is the extension of Lemma 5.4.1 to single hop networks. As for the single link case, we first express the average power $\bar{P}(\gamma)$ in terms of the function $c_R(s)$. We note that the average power used when the queue length vector is \mathbf{q} is bounded below by the optimal value of

$$\begin{aligned} &\text{minimize} && \mathbb{E}_{\pi_{\mathbf{H}}} \mathbb{E}_{\mathbf{S}|\mathbf{H}} P(\mathbf{H}, \mathbf{S}), \\ &\text{such that} && \mathbb{E}_{\pi_{\mathbf{H}}} \mathbb{E}_{\mathbf{S}|\mathbf{H}} \mathbf{S} = (\bar{s}_1(q_1), \bar{s}_2(q_2)). \end{aligned}$$

By definition, the optimal value of the above problem is $c_R(\bar{s}_1(q_1), \bar{s}_2(q_2))$. Then, we note that $\mathbb{E}_{\pi} c_R(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) \leq \bar{P}(\gamma)$. Furthermore, since $c_R(\cdot)$ is convex, $\mathbb{E}_{\pi} c_R(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) \geq c_R(\boldsymbol{\lambda})$.

Proposition 5.8.2. For any sequence of admissible policies γ_k with $\bar{P}(\gamma_k) - c_R(\lambda_1, \lambda_2) = V_k \downarrow 0$, we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$.

Proof. We note that for the sequence γ_k , since $\bar{P}(\gamma_k) - c_R(\lambda) = V_k$, we have that $\mathbb{E}_\pi c_R(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) - c_R(\lambda) = U_k \downarrow 0$, with $U_k \leq V_k$. Consider a particular policy γ in the sequence γ_k with $U_k = U$. Let $q_{1,d} = \sup\{q_1 : \bar{s}_1(q_1) \leq \lambda_1 + \epsilon_U\}$, where ϵ_U will be chosen later. From the admissibility of γ we have that $\forall q \leq q_{1,d}, \bar{s}_1(q_1) \leq \lambda_1 + \epsilon_U$. Assume that $q_d < \infty$. Using $d = \epsilon_U$ in Lemma 5.8.1, we have for a $0 < \Delta + \delta < \delta_a$, $k \geq 1$, and $\bar{q}_1 = k\Delta \leq q_{1,d}$

$$Pr\{Q_1 \geq \bar{q}_1\} \geq \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + \epsilon_U}\right)^k + \left[1 - \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + \epsilon_U}\right)^k\right] \left[Pr\{Q_1 \geq q_{1,d}\} + \frac{1}{\epsilon_U} \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1)\right].$$

Or we have that

$$Pr\{Q_1 < \bar{q}_1\} \leq \left(1 - \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + \epsilon_U}\right)^k\right) \left[1 - \frac{1}{\epsilon_U} \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1)\right],$$

where we have used the non-negativity of $Pr\{Q_1 \geq q_{1,d}\}$. Let $D_t \triangleq \frac{1}{\epsilon_U} \int_{q_{1,d}}^\infty (\bar{s}_1(q_1) - \lambda) d\pi(q)$. For γ we have that $\mathbb{E}_\pi c_R(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) - c_R(\lambda) = U$. Let $l(s_1, s_2)$ be the tangent plane to $c_R(s_1, s_2)$ at (λ_1, λ_2) . Then $U = \mathbb{E}_\pi [c_R(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) - l(\bar{s}_1(Q_1), \bar{s}_2(Q_2))]$, since $\mathbb{E}_\pi l(\bar{s}_1(Q_1), \bar{s}_2(Q_2)) = c_R(\lambda)$. Let $G(x_1, x_2) \triangleq c_R(x_1, x_2) - l(x_1, x_2)$. We note that $G(x_1, x_2)$ is strictly convex in x_1 and x_2 . Also $G(0, 0) = 0$ and $\frac{\partial G}{\partial x_1}(0, 0) = \frac{\partial G}{\partial x_2}(0, 0) = 0$. Proceeding as in steps (42)-(46) in [43, Appendix A] we obtain similarly as in the proof of Lemma 4.4.2 that

$$D_t \leq \frac{1}{\epsilon_U} \sqrt{\frac{U}{a_1}}, \quad (5.20)$$

for a positive a_1 . We choose $\epsilon_U = 4\sqrt{\frac{U}{a_1}}$. Let k_1 be the largest integer such that

$$\left(1 - \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + \epsilon_U}\right)^{k_1}\right) \left[1 - \frac{1}{\epsilon_U} \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1)\right] \leq \frac{1}{2}.$$

Then $Pr\{Q_1 < k_1\Delta\} \leq \frac{1}{2}$ and $\bar{Q}(\gamma) \geq \mathbb{E}_\pi Q_1 \geq \frac{k_1\Delta}{2}$. The same approach also holds if $q_d = \infty$. The rest of the proof is similar to that of Proposition 4.4.2 and we obtain that $k_1 = \Omega\left(\frac{1}{\sqrt{U}}\right)$, for small U . Therefore, $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{U_k}}\right) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$, since $U_k \leq V_k$. \square

We note that the proof of the above lemma illustrates how the lower bounding technique can be applied to a multiqueue case by considering each queue on its own, even though the service vector $\mathbf{S}(\mathbf{Q}, \mathbf{H})$ is chosen as a function of the queue length vector \mathbf{Q} . The proof depends on the upper bound on the marginal stationary probability of a particular queue, which can be obtained from the average drift for that queue, conditioned on its own queue length rather than on the queue length vector.

5.8.4 Discussion




Figure 5.12: Illustration of $c(\lambda_1, \lambda_2)$ as a function of the arrival rates λ_1 and λ_2 .

Model with integer valued queue evolution: We note that a single hop queueing network model with N users can be set up similarly as above, where $(A[m], m \geq 1)$, $(S[m], m \geq 1)$, and $(Q[m], m \geq 0)$ are assumed to evolve on \mathbb{Z}_+^N . Similar to the function $c(\lambda)$ defined for the single link case, a function $c(\lambda)$ (which is the counterpart of $c_R(\lambda)$) can be obtained as the optimal solution of the problem (5.19), but with the conditional distributions of service batch sizes having support on $\{0, 1, \dots, S_{max}\}$. We note that $c(\lambda)$ has a polyhedral structure [43, Section VII] (we recall that $c(\lambda)$ was piecewise linear). Then several cases may arise, e.g., for $N = 2$ we have the following cases, which are illustrated in Figure 5.12 : (1) λ lies on the interior of a face of $c(\lambda_1, \lambda_2)$ of type F1, which includes $\mathbf{0}$ (which corresponds to case 1 for the single link case), (2) λ lies on the interior of a face of $c(\lambda_1, \lambda_2)$ of type F2, not including $\mathbf{0}$ (which corresponds to Case 2 for the single link case), (3) λ lies on a vertex V (not $\mathbf{0}$) (which corresponds to Case 3 for the single link case), and (4) λ lies on an edge, which is (i) of type E2, having a projection on the (λ_1, λ_2) plane which is perpendicular to the λ_2 (or λ_1) axis or (ii) of type E1, having a projection on the (λ_1, λ_2) plane which is not perpendicular to either the λ_1 or λ_2 axes.

We again consider the TRADEOFF problem for this model in the asymptote of small V , where V is the difference between P_c and $c(\lambda)$. As illustrated in the proof of Proposition 5.8.2, the asymptotic lower bound for the total average queue length can be obtained by separate lower bounds on the average queue lengths of the individual queues. Therefore, we expect that the asymptotic lower bounds for Cases 2 and 3 can be obtained from straightforward extensions of Lemmas 5.3.1 and 5.3.5 respectively. As for the single-link case, we do not have any analytical results for Case 1.

For case 4(i) we expect that the total average queue length will grow as $\Omega\left(\frac{1}{V}\right)$ since the probability of the average service rate $\bar{s}_2(Q)$ being not equal to the arrival rate λ_2 should go to zero as $V \downarrow 0$. We note that in case 4(i) the average queue length of the first queue is expected to grow as

$\Omega(\log(\frac{1}{V}))$ since the probability of the average service rate $\bar{s}_1(Q)$ being greater than the arrival rate λ_1 does not go to zero in the asymptote of small V . We note that for the case of integer valued queue evolution, different queues may have different growth rates in the asymptote of small V .

For case 4(ii) we expect that the total average queue length will grow as $\Omega(\log(\frac{1}{V}))$, since the probability of the average service rate $\bar{s}_1(Q)$ and $\bar{s}_2(Q)$ being greater than λ_1 and λ_2 respectively is positive as the average power constraint approaches $c(\lambda_1, \lambda_2)$.

Model with admission control: We now comment on how the asymptotic lower bound can be derived for a N user model as above, but with admission control. We consider a model, which is a straightforward extension of R-model-U to the N user case.

We assume that the arrival rate vector into the system is $\mathbf{R}[m]$ in slot m , with $\mathbb{E}\mathbf{R}[1] = \boldsymbol{\lambda}$. For stationary policies, the number of packets admitted into the queue is a function $\mathbf{A}(\mathbf{Q}[m-1], \mathbf{R}[m], \mathbf{H}[m])$. For an admissible policy, defined as above, the average throughput is $\bar{\mathbf{A}}(\gamma) = \mathbb{E}_\pi \mathbb{E}_{\pi_H} \mathbb{E}_R \mathbf{A}(\mathbf{Q}, \mathbf{R}, \mathbf{H})$. The problem that we are interested in is:

$$\underset{\gamma \in \Gamma_a}{\text{minimize}} \bar{Q}(\gamma), \text{ such that } \bar{P}(\gamma) \leq P_c \text{ and } \bar{\mathbf{A}}(\gamma) \geq \rho \boldsymbol{\lambda},$$

for a $\rho < 1$. Similar to the single link case, it can be shown that $\inf_{\gamma \in \Gamma_a} \bar{P}(\gamma) = c_R(\rho \boldsymbol{\lambda})$. Then under the assumption that $\forall n \in \{1, \dots, N\}$, $Pr\{R_n[1] \leq \frac{\Delta_n}{2}\} = \epsilon'_{a,n} > 0$ for some Δ_n such that $0 < \Delta_n < \rho \lambda_n$ we can proceed as in the proof of Lemma 5.7.2 to prove that $\mathbb{E}_\pi Q_n = \Omega(\log(\frac{1}{V_k}))$, $\forall n$, for any sequence of admissible policies γ_k with $\bar{P}(\gamma_k) - c_R(\rho \boldsymbol{\lambda}) = V_k \downarrow 0$. Therefore, the total average queue length $\bar{Q}(\gamma_k)$ is also $\Omega(\log(\frac{1}{V_k}))$.

Similar results can be obtained for the case where the queue evolution is assumed to be integer valued, whenever $(\boldsymbol{\lambda}, c(\boldsymbol{\lambda}))$ does not lie on a face which contains $(\mathbf{0}, \mathbf{0})$.

5.9 Conclusions

We recall that R-model and R-model-U, both with strictly convex $P(h, s)$ functions, are usually used as approximations for I-model and I-model-U respectively. However, as in Chapter 4, we find that the asymptotic behaviours of $Q^*(P_c)$ in the asymptotic regimes \mathfrak{R} are different for the approximate models and the original models.

We note that the R-model suggests that a strictly smaller minimum average power $c_R(\boldsymbol{\lambda})$ is sufficient for stability, compared with $c(\boldsymbol{\lambda})$ for the I-model, for all $\boldsymbol{\lambda} \neq a_p, p > 1$. From Lemma 5.3.5, we observe that the asymptotic behaviour of the minimum average queue length is quite different for the I-model and R-model, for $\boldsymbol{\lambda} = a_p, p > 1$, for which $c_R(\boldsymbol{\lambda}) = c(\boldsymbol{\lambda})$. For such a $\boldsymbol{\lambda}$, for the

R-model, from Lemma 5.4.1, we have that the minimum average queue length is $\Omega\left(\frac{1}{\sqrt{V}}\right)$, if (a) the average power is V more than the minimum average power $c_R(\lambda)$ and (b) $P(h, s)$ is a strictly convex function of $s \in [0, S_{max}]$, for every $h \in \mathcal{H}$. In contrast, for the I-model, the asymptotic lower bound in Lemma 5.3.5 shows that the minimum average queue length is $\Omega\left(\frac{1}{V}\right)$. We note that the minimum average queue length for the R-model is always a lower bound to the minimum average queue length obtained from the I-model, for a given constraint on the average power, but the rate of increase of the minimum average queue length as the average power constraint is reduced is strictly smaller.

For Case 1, we find that the $Q^*(P_c)$ for R-model (as well as R-model-U) increases to infinity in the regime \mathfrak{R} , while for I-model we have numerically illustrated that $Q^*(c(\lambda)) < \infty$ (therefore, also for I-model-U). However, for R-model (as well as R-model-U), if we use a piecewise linear $P(h, s)$ function (as in Section 4.4.2), which is the lower convex envelope of the service cost function defined for the I-model (as well as I-model-U) on $\{0, \dots, S_{max}\}$, then the asymptotic behaviour of both the R-model (R-model-U) and I-model (I-model-U) matches. Therefore, a more appropriate approximation for I-model (or I-model-U), is a R-model (or R-model-U) with the above piecewise linear $P(h, s)$ function.

We note that Lemma 5.3.1 provides the asymptotic lower bound to the $\mathcal{O}\left(\log\left(\frac{1}{V}\right)\right)$ upper bound observed by Neely in [43, Corollary 2] for the I-model. This asymptotic lower bound was earlier shown only for a specific example ([43, Section VII-A]). We note that if $(A[m])$ and $(H[m])$ are ergodic sequences, independent of each other, then for Cases 2 and 3 for the I-model, we have obtained a $\Omega\left(\log\left(\frac{1}{V}\right)\right)$ asymptotic lower bound on the minimum average queue length in Lemma 5.6.1. As stated, for Case 1, if $|\mathcal{H}| > 1$, we do not have tight asymptotic upper or lower bounds on $Q^*(P_c)$, as $P_c \downarrow c(\lambda)$. Lemma 5.7.2 provides an asymptotic lower bound for R-model-U even if $|\mathcal{H}| > 1$, for admissible policies, while earlier an asymptotic lower bound was obtained only for the case $|\mathcal{H}| = 1$. We also illustrate how the asymptotic lower bound can be obtained for a N user single hop network and identify a case in which average queue lengths for different queues can have different asymptotic behaviours for integer valued queue evolution.

Appendices

5.A The function $c(\lambda)$ in Section 5.2

In this section, we show that the function $c(\lambda)$ defined as the optimal value of

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} P(H, S), \\ & \text{such that} && \mathbb{E}_{\pi_H} \mathbb{E}_{S|H} S = \lambda, \end{aligned}$$

is piecewise linear if $|\mathcal{H}|$ is finite. Let $p_{s|h}$ be the conditional distribution of the batch size s given the fade state h . Then the above problem can be written as

$$\begin{aligned} & \text{minimize} && \sum_h \pi_H(h) \sum_{s=0}^{S_{max}} p_{s|h} P(h, s), \\ & \text{such that} && \sum_h \pi_H(h) \sum_{s=0}^{S_{max}} p_{s|h} s = \lambda, \\ & && \sum_{s=0}^{S_{max}} p_{s|h} = 1, \forall h, \text{ and } p_{s|h} \geq 0, \forall s, h, \end{aligned}$$

which is a linear program in the variables $p_{s|h}$.

For ease of exposition, in the following we consider the case where $|\mathcal{H}| = 2$, but the approach holds for any finite $|\mathcal{H}|$. Let $\mathbf{s} = (0, 1, \dots, S_{max})$, $P(h_i, \mathbf{s}) = (P(h_i, 0), \dots, P(h_i, S_{max}))$, $p_{s|h} = (p_{0|h}, \dots, p_{S_{max}|h})$, and let $\mathbf{1}$ be a row vector of all ones and $\mathbf{0}$ a row vector of all zeros, both of size S_{max} . Then the above linear program can be written as

$$\begin{aligned} & \text{minimize} && [\pi_H(h_1)P(h_1, \mathbf{s}), \pi_H(h_2)P(h_2, \mathbf{s})] \begin{bmatrix} p_{\mathbf{s}|h_1}^t \\ p_{\mathbf{s}|h_2}^t \end{bmatrix} \\ & \text{such that} && \begin{bmatrix} \pi_H(h_1)\mathbf{s} & \pi_H(h_2)\mathbf{s} \\ \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \begin{bmatrix} p_{\mathbf{s}|h_1}^t \\ p_{\mathbf{s}|h_2}^t \end{bmatrix} = \begin{bmatrix} \lambda \\ 1 \\ 1 \end{bmatrix}, \\ & \text{and} && p_{s|h} \geq 0, \forall s, h. \end{aligned}$$

We note that the dual of this problem is

$$\begin{aligned} & \text{maximize} && [y_1, y_2, y_3] \begin{bmatrix} \lambda \\ 1 \\ 1 \end{bmatrix} \\ & \text{such that} && [y_1, y_2, y_3] \begin{bmatrix} \pi_H(h_1)\mathbf{s} & \pi_H(h_2)\mathbf{s} \\ \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \leq [\pi_H(h_1)P(h_1, \mathbf{s}), \pi_H(h_2)P(h_2, \mathbf{s})], \end{aligned}$$

where $y_1, y_2, y_3 \in \mathbb{R}$. Let $y_1^*(\lambda), y_2^*(\lambda)$, and $y_3^*(\lambda)$ be any optimizers for the dual for $\lambda \in [0, S_{max}]$. Since $c(\lambda)$ is convex in λ ([43]), it is differentiable at all $\lambda \in [0, S_{max}]$ except for $\lambda \in \mathcal{D}$, where \mathcal{D} is at most countable [20]. Now consider $c(\lambda)$ for a $\lambda \notin \mathcal{D}$. Then we have that $\frac{dc(\lambda)}{d\lambda}$ is well defined. Furthermore, since strong duality holds for the linear programs above, we also have that $\frac{dc(\lambda)}{d\lambda} = -y_1^*(\lambda)$ [14, Section 5.6]. Therefore, for $\lambda \notin \mathcal{D}$, $y_1^*(\lambda)$ is unique.

We note that the constraint set in the dual problem does not depend on λ and has finite number of vertices. Therefore, there are only finitely many ways in which $y_1^*(\lambda)$ can be unique. Hence, $\frac{dc(\lambda)}{d\lambda}$ can take only finitely many values. Since $c(\lambda)$ is also non-decreasing in λ , we have that $c(\lambda)$ is piecewise linear. We note that the above approach generalizes to any finite \mathcal{H} .

5.B Proof of Lemma 5.3.2

Let $L(q) \triangleq e^{\omega(q_v - q)}$ be a Lyapunov function. Since for the policy γ , the batch size $\tilde{S}(q)$ could be more than q , the queue evolution equation under γ is written as

$$Q[m+1] = \max(Q[m] - \tilde{S}(Q[m]), 0) + A[m+1].$$

The expected Lyapunov drift is

$$\Delta(q) \triangleq \mathbb{E}[L(Q[m+1]) - L(Q[m]) | Q[m] = q].$$

We note that the randomness in $\tilde{S}(q)$ arises from both the randomness in the fade state as well as the randomization of the batch size. The expectation of $\tilde{S}(q)$ is therefore with respect to this distribution. Proceeding as in the proof of Lemma 4.G we have that

$$\Delta(q) \leq \omega e^{\omega(q_v - q)} \left[(\mathbb{E}\tilde{S}(q) - \lambda) + K \right].$$

where $K = \frac{\omega A_{max}^2}{2} e^{\omega A_{max}}$.

Now by definition, the policy γ is such that

$$\mathbb{E}\tilde{S}(q) = \begin{cases} s_l, & \text{for } 0 \leq q < q_v, \\ s_u, & \text{for } q_v \leq q. \end{cases}$$

Then we have that for $q < q_v$

$$\Delta(q) \leq -\omega e^{\omega(q_v - q)} [\lambda - s_l - K].$$

And for $q \geq q_v$,

$$\begin{aligned}\Delta(q) &\leq \omega e^{\omega(q-q_v)} [(s_u - \lambda) + K], \\ &= -\omega e^{\omega(q_v-q)} [\lambda - s_l - K] + \omega e^{\omega(q_v-q)} [s_u - s_l], \\ &\leq -\omega e^{\omega(q_v-q)} [\lambda - s_l - K] + \omega [s_u - s_l],\end{aligned}$$

Hence, for all q we have that

$$\Delta(q) \leq -\omega e^{\omega(q_v-q)} [\lambda - s_l - K] + \omega [s_u - s_l].$$

We choose ω such that $K < \lambda - s_l$. Proceeding as in the proof of [43, Theorem 3(c)], we have that

$$\mathbb{E}e^{\omega(q_v-Q)} \leq \frac{[s_u - s_l]}{(\lambda - s_l - K)}.$$

Since $\mathbb{E}e^{\omega(q_v-Q)} \geq \mathbb{E}[e^{\omega(q_v-Q)} | Q < S_{max}] Pr\{Q < S_{max}\}$, we therefore have that

$$Pr\{Q < S_{max}\} \leq e^{-\omega q_v} \frac{e^{\omega S_{max}} [s_u - s_l]}{(\lambda - s_l - K)}. \quad (5.21)$$

Now we note that

$$\begin{aligned}\bar{P}(\gamma) &= Pr\{Q < S_{max}\} \mathbb{E}[P(H, S) | Q < S_{max}] + \\ &\quad Pr\{S_{max} \leq Q < q_v\} \mathbb{E}[P(H, S) | S_{max} < Q < q_v] + Pr\{q_v \leq Q\} \mathbb{E}[P(H, S) | q_v \leq Q].\end{aligned}$$

We note that for γ , for $S_{max} \leq q < q_v$, $\mathbb{E}P(H, S) = c(s_l)$ and for $q \geq q_v$, $\mathbb{E}P(H, S) = c(s_u)$. Furthermore for $q < S_{max}$, $\mathbb{E}P(H, S) \leq c(s_l)$. Hence, we have that

$$\bar{P}(\gamma) \leq Pr\{Q < S_{max}\} c(s_l) + Pr\{S_{max} \leq Q < q_v\} c(s_l) + Pr\{q_v \leq Q\} c(s_u).$$

Proceeding as in the proof of Lemma 4.3.10 we have that

$$\bar{P}(\gamma) - c(\lambda) \leq [c(s_l) + m\lambda - c(\lambda)] Pr\{Q < S_{max}\},$$

where m is the slope of $l(s)$.

Now consider the sequence of policies γ_k for which $q_v = \log\left(\frac{1}{V_k}\right)$ for a sequence $V_k < 1$ such that $V_k \downarrow 0$. Then we have that $\bar{P}(\gamma_k) - c(\lambda) = \mathcal{O}(V_k)$. Furthermore, from Proposition 4.F.1 we have that $\bar{Q}(\gamma) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$. We note that γ_k is also a sequence of admissible policies, since $s(q)$ is a non-decreasing function of q and $\bar{Q}(\gamma_k) < \infty$.

5.C Proof of Lemma 5.8.1

Proof. Define $\widehat{Q}_1[m] = \max(\bar{q}_1, Q_1[m])$ and $\widehat{Q}_1 = \max(\bar{q}_1, Q)$. As the policy is admissible and \bar{q}_1 is finite we have that $\mathbb{E}_\pi \widehat{Q} < \infty$. Therefore

$$\int_{q_1, q_2} \mathbb{E} \left[\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right] d\pi(q_1, q_2) = 0.$$

We split the integral over q_1 into three terms which leads to :

$$0 = \int_0^{\bar{q}_1 - \Delta} \int_0^\infty \mathbb{E} \left[\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right] d\pi(q_2|q_1) d\pi(q_1) \quad (5.22)$$

$$+ \int_{\bar{q}_1 - \Delta}^{\bar{q}_1} \int_0^\infty \mathbb{E} \left[\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right] d\pi(q_2|q_1) d\pi(q_1) \quad (5.23)$$

$$+ \int_{\bar{q}_1}^\infty \int_0^\infty \mathbb{E} \left[\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right] d\pi(q_2|q_1) d\pi(q_1) \quad (5.24)$$

We note that for $q_1 \leq \bar{q}_1 - \Delta$, (5.22) ≥ 0 . Consider (5.24), for which $\widehat{Q}_1[m] = Q_1[m]$ and $\widehat{Q}_1[m+1] \geq Q_1[m+1]$. Therefore

$$(5.24) \geq \int_{\bar{q}_1}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1).$$

Using the assumption MG3, the above integral can be further bounded below by

$$-dPr \{ \bar{q}_1 \leq Q_1 < q_{1,d} \} + \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1).$$

Consider (5.23), we have that $\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \geq 0$ for $q_1 \in [\bar{q}_1 - \Delta, \bar{q}_1]$. Hence as in the proof of Lemma 4.4.1 we use Markov inequality to lower bound (5.23).

$$\begin{aligned} \mathbb{E} \left[\widehat{Q}_1[m+1] - \widehat{Q}_1[m] \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right] &\geq \delta Pr \left\{ \widehat{Q}_1[m+1] - \widehat{Q}_1[m] \geq \delta \middle| Q_1[m] = q_1, Q_2[m] = q_2 \right\}, \\ &\geq \delta Pr \{ Q_1[m+1] - Q_1[m] \geq \delta + \Delta \middle| Q_1[m] = q_1, Q_2[m] = q_2 \}, \\ &\geq \delta \epsilon_{1,a}, \end{aligned}$$

where $\Delta > 0$ and $\delta > 0$ are chosen such that $\Delta + \delta < \delta_{1,a}$. Thus we obtain that

$$(5.23) \geq \delta \epsilon_{1,a} Pr \{ \bar{q}_1 - \Delta \leq Q_1 < \bar{q}_1 \}.$$

Combining the obtained lower bounds on (5.22), (5.23), and (5.24) we obtain that

$$\begin{aligned} 0 &\geq \delta \epsilon_{1,a} Pr \{ \bar{q}_1 - \Delta \leq Q_1 < \bar{q}_1 \} - dPr \{ \bar{q}_1 \leq Q_1 < q_{1,d} \} + \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1), \\ &= \delta \epsilon_{1,a} Pr \{ Q_1 \geq \bar{q}_1 - \Delta \} - (d + \delta \epsilon_a) Pr \{ Q_1 \geq \bar{q}_1 \} + dPr \{ Q_1 \geq q_{1,d} \} + \int_{q_{1,d}}^\infty (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1). \end{aligned}$$

Hence

$$Pr \{Q_1 \geq \bar{q}_1\} \geq \frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + d} Pr \{Q_1 \geq \bar{q}_1 - \Delta\} + \frac{1}{\delta\epsilon_{1,a} + d} \left[dPr \{Q_1 \geq q_{1,d}\} + \int_{q_{1,d}}^{\infty} (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1) \right].$$

By induction, as in the proof of Lemma 4.4.1, we obtain that if $k \geq 0$ and $\bar{q}_1 + k\Delta \leq q_{1,d}$, then

$$Pr \{Q_1 \geq \bar{q}_1 + k\Delta\} \geq \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + d} \right)^k Pr \{Q_1 \geq \bar{q}_1\} + \left[1 - \left(\frac{\delta\epsilon_{1,a}}{\delta\epsilon_{1,a} + d} \right)^k \right] \left[Pr \{Q_1 \geq q_{1,d}\} + \frac{1}{d} \int_{q_{1,d}}^{\infty} (\lambda_1 - \bar{s}_1(q_1)) d\pi(q_1) \right].$$

□

On the tradeoff of average error rate and average delay for point to point links

6.1 Introduction

In this chapter, we consider the transmission of a bursty information source over a noisy point to point link. Our objective is to transmit the randomly arriving source message symbols such that the message symbols are decoded *reliably* and with minimum *delay*. Reliability and delay performance are measured by the average symbol error rate and the average symbol delay respectively. Understanding the fundamental tradeoff between reliability and delay is significant due to the increasing use of cross layer scheduling, for allocation of highly constrained wireless resources, in modern high rate communication networks. We assume that the transmitter and receiver use a block code (such as a LDPC code) to reliably communicate the message symbols, as in many practical scenarios.

It is known [24, Chapter 24] that if the average information arrival rate into the system is less than the capacity of the channel, then arbitrarily low probability of error can be achieved using block coding with long codeword lengths and *finite* average delay. In this chapter, we first characterize how the minimum possible average delay grows when the average error rate is made arbitrarily small. However, there are cases where arbitrarily long codewords cannot be used. Then there is a positive infimum for achievable error rate. Then for this case, we characterize how the minimum possible average delay grows as the average error rate is made arbitrarily close to the above positive infimum of achievable error rates, using the techniques discussed in Chapter 4. We note that the service cost function in this chapter, which is the expected number of message symbols in error when a batch of message symbols is transmitted, turns out to be non-convex unlike the convex service cost

functions in the previous chapters. The notation that we use in this chapter is summarized in Tables 6.1 and 6.2.

6.1.1 System model

The system is assumed to evolve in discrete time units of slots. The slots are indexed by $n \in \{1, 2, \dots\}$. Each slot corresponds to one channel use. We note that these slots may be thought of as subslots in the slots for the discrete time models in Chapter 4 (this is why the slots are indexed by n rather than m in this chapter). The channel is assumed to be a memoryless channel, with input alphabet \mathcal{X} and output alphabet \mathcal{Y} . The transition probability function of the channel is denoted by $P_{Y|X}$.

We assume that there is no admission control throughout this chapter. The source generates a random number $A_s[n] \leq A_{max}$ of message symbols in each slot n . We assume that $A_s[n]$ is IID with $\mathbb{E}A_s[1] = \lambda$ and $var(A_s[1]) = \sigma^2$, both of which are finite. Each message symbol is generated independently and uniformly from a message alphabet \mathcal{M} of finite cardinality $|\mathcal{M}|$. The message symbols are assumed to enter the queue just before the slot boundary and reside in the transmitter buffer until they are encoded and transmitted. Transmission of a message symbol is assumed to require at least one slot. The transmitter buffer size is assumed to be infinite.

The transmitter is assumed to use random block coding [22, Chapter 5]. Each symbol of the transmitted channel codeword is picked IID from a distribution $\mathcal{Q}(\cdot)$ on \mathcal{X} . At a decision epoch, which occurs at, say the start of slot n , if the transmitter decides to transmit, it uses a random block codebook, with codewords generated as above. The codebook is characterized by two parameters, the number of message symbols (s) which are encoded and transmitted by the codeword and the length (τ) of the codeword. Starting from slot n , s message symbols are removed from the transmitter queue after τ slots. We assume that there is no path delay in the transmission of the codeword symbols. The receiver is assumed to decode the s message symbols jointly by using maximum likelihood decoding of the codeword. If the transmitter decides not to transmit, then the receiver is made aware of the idle state through the control channel and the transmitter idles for τ_0 slots. A decision epoch occurs after every transmission period or idle period. We assume that at the end of a slot n , events occur in the following order: a) if the service of a batch of symbols ends then the batch is removed from the queue, b) new arrivals in slot n are admitted into the queue, and c) the queue length state at the beginning of the next slot is obtained.

For the i^{th} message symbol, let $T_{a,i}$ denote the slot in which the symbol arrives into the transmitter queue. Let $T_{d,i}$ denote the slot in which the i^{th} symbol departs from the receiver. We note that $T_{d,i}$ is the slot in which transmission of the batch containing the i^{th} symbol finishes. The delay of the i^{th} symbol in the queue is $T_{d,i} - T_{a,i}$. Let E_i denote the event that the i^{th} symbol is in error.

A policy γ for operation of the transmitter consists of a sequence $((S[m], \mathcal{T}[m]), m \geq 1)$, where $m \in$

Symbol	Description
n	slot index
$A_s[n]$	random number of message symbol arrivals in slot n
A_{max}	maximum number of message symbol arrivals in a slot
λ, σ^2	mean and variance of $A_s[n]$
\mathcal{X}	input alphabet of a point to point channel
\mathcal{Y}	output alphabet of a point to point channel
$P_{Y X}$	channel transition probability function
\mathcal{Q}	distribution on the channel input symbol $\in \mathcal{X}$
s	a particular batch size
τ	a particular transmission duration
$T_{a,i}$	arrival time of i^{th} message
$T_{d,i}$	departure time of i^{th} message
E_i	event that the i^{th} symbol is in error
m	decision epoch index
$S[m]$	service batch size for the m^{th} transmission
S_{max}	maximum service batch size
$\mathcal{T}[m]$	duration of the m^{th} transmission
$\gamma[m]$	$= (S[m], \mathcal{T}[m])$
$Q[m]$	queue length at decision epoch m
$Q_s[n]$	queue length at start of slot n
$A[m]$	random number of message symbol arrivals between $(m-1)^{th}$ and m^{th} decision epochs
Γ	set of all policies
Γ_s	set of stationary policies
$P_e(\gamma)$	average error prob. for policy γ
$D(\gamma)$	average delay for policy γ
$P_{e,c}$	constraint on the probability of error of message symbols
D_c	constraint on the average delay of message symbols
\mathcal{M}	source message alphabet
$ \mathcal{M} $	cardinality of \mathcal{M}
N_c	fixed codeword length
$h(q, \tau)$	holding cost; is $q\tau + \frac{\tau(\tau-1)\lambda}{2}$
$\Gamma_{s,f}$	set of all policies with transmission time being a fixed parameter
Γ_{s,N_c}	set of all policies with transmission time being N_c

Table 6.1: Notation used in this chapter (part I).

Symbol	Description
$\bar{Q}(\gamma)$	average queue length for a policy γ
$c_s(s, \tau)$	expected number of messages in error for a random block code
$E_0(\rho, \mathcal{Q})$	Gallager's random coding error exponent
$P_e^*(D_c)$	minimum average error probability over Γ_s under delay constraint D_c
$P_{e,f}^*(D_c)$	minimum average error probability over $\Gamma_{s,f}$ under delay constraint D_c
$P_{e,N_c}^*(D_c)$	minimum average error probability over Γ_{s,N_c} under delay constraint D_c
$P_{N_c}(\lambda)$	minimum achievable average error probability for the system to be stable
R_c	cutoff rate of a point to point channel
C	capacity of a point to point channel
m_s	a particular message symbol, $m_s \in \mathcal{M}$
τ_0	duration of time the transmitter idles
$\tau(s)$	transmission time as a function of s

Table 6.2: Notation used in this chapter (part II).

$\{1, 2, 3, \dots\}$. If $S[m] > 0$, then $S[m]$ is the number of message symbols which start transmission at the $(m-1)^{th}$ decision epoch using a codeword of length $\mathcal{T}(m)$. If $S[m] = 0$, then the system idles for $\mathcal{T}[m] = \tau_0$ slots. Let $\gamma[m] = (S[m], \mathcal{T}[m])$. The evolution of the system is illustrated in Figure 6.1. For every m , $\gamma[m]$ can be a randomized function of (i) the history $(\gamma[1], \gamma[2], \dots, \gamma[m-1])$, (ii) the initial number of message symbols in the transmitter queue q_0 , and (iii) the arrival process up to the start of the $(m-1)^{th}$ decision epoch. The length of the transmitter queue at a decision epoch m is denoted by $Q[m]$. The queue length at the start of a slot n is denoted by $Q_s[n-1]$. The evolution of the system sampled at the decision epochs, for a policy γ , is given by the following equation:

$$Q[m+1] = Q[m] - S[m+1] + A[m+1], \quad (6.1)$$

where $S[m+1] \leq Q[m]$ and $A[m]$ is the random number of arrivals which have occurred in the period between the m^{th} and $(m+1)^{th}$ decision epochs. We note that $A[m] \sim \star_{\mathcal{T}[m]} A_s[1]$, the $\mathcal{T}[m]$ convolution of $A_s[1]$. Let Γ denote the set of all policies. The class of all stationary policies Γ_s is such that for any $\gamma \in \Gamma_s$, $S[m] = S(Q[m-1])$ and $\mathcal{T}[m] = \mathcal{T}(Q[m-1])$, where S and \mathcal{T} are functions (possibly randomized) of the queue length q . We note that if $\gamma \in \Gamma_s$, then $Q[m]$ is a Markov chain embedded in the random process $Q_s[n]$.

In the following, we consider two separate models, set up as follows.

R-model-A : $\forall m \geq 1, A[m] \in [0, A_{max}], S[m] \in \mathbb{R}_+$, and $q_0 \in \mathbb{R}_+$,

R-model-B : Same as the R-model-A, except that $S[m] \in [0, S_{max}]$, and $\mathcal{T}[m] = N_c, m \geq 1$.

We note that for R-model-A and R-model-B, the queue length evolution $(Q[m], m \geq 0)$ is on the




Figure 6.1: Evolution of the system for a block coding scheme

non-negative real numbers. Furthermore, for R-model-B, the maximum batch size is bounded by S_{max} and the decision epochs occur every N_c slots. Hence, for R-model-B, all codewords are of length N_c .

For R-model-A and R-model-B, for a $\gamma \in \Gamma$, we define the average delay as

$$D(\gamma) = \limsup_{I \rightarrow \infty} \frac{\sum_{i=1}^I (T_{d,i} - T_{a,i})}{I}. \quad (6.2)$$

We also define the average error rate as

$$P_e(\gamma) = \limsup_{I \rightarrow \infty} \frac{\sum_{i=1}^I \mathbb{I}_{E_i}}{I}, \quad (6.3)$$

where \mathbb{I}_{E_i} is the indicator function for the event E_i .

Our problem is then to

$$\begin{aligned} & \underset{\gamma \in \Gamma}{\text{minimize}} P_e(\gamma) & \text{or} & & \underset{\gamma \in \Gamma}{\text{minimize}} D(\gamma) \\ & \text{such that } D(\gamma) \leq D_c & & & \text{such that } P_e(\gamma) \leq P_{e,c}. \end{aligned}$$

for every $D_c \geq 1$ or $0 \leq P_{e,c} \leq 1$. For an average error rate constraint $P_{e,c} > \frac{|\mathcal{M}|-1}{|\mathcal{M}|}$, we have the following optimal solution: transmit all the arrivals in a slot in the succeeding slot, using a channel input symbol picked independently of the message symbols, and pick the message symbol estimates uniformly at the receiver (maximum likelihood decoding). In this chapter, we analyse the above problem, for the class of stationary policies Γ_s , in the asymptotic regime where (i) arbitrarily large reliability is required, that is, as $P_{e,c} \downarrow 0$ for R-model-A, and (ii) $P_{e,c}$ approaches the minimum

probability of error for N_c , when the codeword length τ is fixed to be N_c for R-model-B.

6.1.2 Related work

We note that for obtaining the average error rate of a policy, we would need to know how the error events E_i are related to the scheduling policy γ which chooses the parameters $(S[m], \mathcal{T}[m])$ for the random block codebook at every $m \geq 1$. It is intuitive that, for a stationary policy γ , apart from the noise introduced by the channel, E_i is a function only of the size s and the length τ of the codebook, from which the codeword used to transmit the i^{th} symbol is selected. Gallager [22, Chapter 5] has provided an upper bound on the average error probability of a random block code as a function of s and τ . The same reference also provides a lower bound on the average error probability of any block code as a function of s and τ . The above upper and lower bounds show that the average error probability of the random block code decays exponentially with the block length for codeword rates less than the capacity of the channel. The upper bound which is obtained as an ensemble average over random block codes (Gallager's random coding upper bound) and the lower bound (sphere packing bound) are found to coincide in the exponent for rates greater than the cutoff rate for the channel. So for large values of the block length and for codeword rates greater than the cutoff rate, the Gallager random coding upper bound can be used as a reasonable approximation for the average error probability of the block codeword. Recently Polyanskiy [49] has obtained upper and lower bounds as well as an analytically tractable approximation for the average error probability of block codewords as a function of the rate and block length of the codebook. The approximation has been observed to be tight even for small values of the block length. We provide a detailed review of the results from Gallager [22] in Section 6.2.1 wherein the Gallager random coding upper bound is used to approximate the average number of symbols in error, as a function of s and τ , when random block codewords are used, as in our model.

Now we provide a survey of prior work on characterizing the tradeoff between reliability and delay for point to point channels. Javidi and Swamy [31] consider a point to point link with a random stationary arrival process of bits at rate λ into an infinite transmitter buffer. Every N_c slots, a batch of $s = rN_c$ bits (with zero padding, if required) is encoded and transmitted over the noisy channel using a block code of fixed rate r and length N_c . The transmission of a bit is said to fail if : 1) the bit is part of a codeword that is decoded in error or 2) the bit is decoded past a given delay deadline D . The authors study the probability of a bit transmission failure as a function of r, N_c , and D . The analysis is asymptotic in nature - in the regime of large delay deadline D . It is assumed that $b > 2$ and the codeword length $N_c = \frac{D}{b}$ as $D \uparrow \infty$. The constant b can be interpreted as a parameter controlling the division of the overall delay budget D between queueing delay and transmission delay. A large deviations result is used to show that the delay deadline violation probability of a bit decays exponentially with the deadline D , as $D \uparrow \infty$. Using Gallager's random coding upper bound, it is then shown that bit transmission failure probability also decays

exponentially with the exponential rate being the minimum of $(1/b)$ times the Gallager exponent, and the delay exponent. The results obtained are in a similar vein as ours, as the exponential decay of the error rate is shown with respect to the delay deadline. However, we consider average delay as the metric and as we shall see, we also obtain the best decay rate over the class of stationary policies, using codewords of fixed length. We also show that scaling the codeword length linearly with average delay is optimal, with the scaling factor being $2/3$.

Musy and Telatar [40] consider a continuous time queueing model with a Poisson message symbol arrival process, for a point-to-point link with random block coding. The block coding scheme is assumed to be such that the codeword length can be varied as a function of the number of message symbols which are jointly encoded in a codeword. The codeword length is chosen such that the block error probability for every transmission is at most a fixed constant. The authors obtain upper and lower bounds on the minimum average delay for a fixed upper bound on the constant block error probability. A joint scheduling-coding scheme for a point to point link, with ARQ, has been analysed in [71] by Swamy and Javidi. The dependence of average error rate on delay has been considered for non block coding schemes with Poisson arrivals in [80] by Yoon, and for a periodic source in [47] by Negi and Goel.

6.1.3 Overview

We formulate the tradeoff problem in Section 6.2. For the class of stationary policies, we express the average error rate and average delay, in terms of quantities which are analytically more tractable, in the same section. In Section 6.2.1 we present a discussion on error exponents for discrete memoryless channels. In this chapter, the Gallager random coding upper bound is used to obtain an upper bound for the expected number of message symbols in error in every transmission, which is then used to approximate the average error rate of stationary policies. We then consider the asymptotic behaviour of the minimum average error rate subject to a constraint on the average delay, for R-model-A in Section 6.3. We show that the minimum average error rate decays exponentially to zero, as the constraint on the average delay increases to infinity. The exponential decay rate is shown to be two-thirds of the Gallager random coding exponent. We then comment on the exponential decay rate for a queueing model in which the queue length evolution is assumed to be on the set of non-negative integers. We also consider the asymptotic behaviour of the average error rate for a class of policies with codeword length dependent on the batch size, in the same section.

The asymptotic behaviour of the minimum average delay subject to a constraint on the average error rate is then characterized for R-model-B in Section 6.4. This particular problem is similar to the TRADEOFF problem, analysed in Chapter 4, except that the service cost function is non-convex. For R-model-B, as all codewords are of finite length N_c , it is intuitive that the infimum of achievable average error rates, for any finite average delay, is bounded away from zero. The

non-convex nature of the service cost function leads to the result that depending on the value of the arrival rate λ , the minimum average delay is either $\Omega\left(\frac{1}{\sqrt{V}}\right)$ or $\Theta\left(\log\left(\frac{1}{V}\right)\right)$, when the average error rate constraint is V more than the above positive infimum of achievable average error rates. We then consider the tradeoff problem for a similar model, where the queue length evolution is on integers, for which the results from Chapter 4 directly applies.

6.2 Problem formulation

In the following we restrict ourselves to a special subset of Γ_s , which is the set of *admissible* policies. A policy $\gamma \in \Gamma_s$ is admissible if:

1. the embedded Markov chain (EMC) $Q[m]$ is positive Harris recurrent, with stationary distribution π_γ , and,
2. the embedded average $\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} h(Q, \mathcal{T}(Q)) < \infty$, where $h(q, \tau) \triangleq q\tau + \frac{\tau(\tau-1)\lambda}{2}$.

We note that the function $h(q, \tau)$ can be interpreted as the cumulative expected queue length in a deterministic transmission/idle period of duration τ , when q message symbols are present in the queue at the beginning of that period. Every admissible policy is stable, according to the definition in Section 4.1.3 of Chapter 4. To avoid unnecessary notation, and since in the following we consider only admissible policies¹, we redefine Γ_s as the set of all admissible stationary policies. We recall that $\mathcal{T}[m]$ is the duration of the transmission/idle period beginning at the $(m-1)^{th}$ decision epoch. The set of all $\gamma \in \Gamma_s$, for which all transmission and idle periods are of a fixed duration N_c , i.e., $\mathcal{T}[m] = \tau_0 = N_c$, is denoted as Γ_{s, N_c} . Let $\Gamma_{s, f} \triangleq \bigcup_{N_c} \Gamma_{s, N_c}$. We note that when operating with a policy $\gamma \in \Gamma_{s, f}$, a transmission duration $N_c(\gamma)$ is chosen at the first decision epoch, which is kept fixed at all the other decision epochs. We have that $\Gamma \supset \Gamma_s \supset \Gamma_{s, f} \supset \Gamma_{s, N_c}$. We now express $D(\gamma)$ and $P_e(\gamma)$ as averages of functions of the queue length and block coding parameters at the decision epochs when $\gamma \in \Gamma_s$.

The average queue length $\bar{Q}(\gamma)$ for a policy $\gamma \in \Gamma_s$ is defined as²

$$\bar{Q}(\gamma) = \lim_{N \rightarrow \infty} \frac{\sum_{n=0}^{N-1} Q_s[n]}{N}. \quad (6.4)$$

For $\gamma \in \Gamma_s$, we note that $(Q[m], \mathcal{T}[m+1])$ is a Markov renewal process, with the renewal instants corresponding to the decision epochs. In order to express the time average $\bar{Q}(\gamma)$ in terms of a function of $Q[m]$, we use the Markov renewal reward theorem [33, Theorem D.16]. We associate a reward $R[m]$ with the m^{th} renewal cycle. The reward $R[m]$ is the cumulative queue length over

¹We note that the requirement of monotonicity as in the previous chapters is imposed only for R-model-B

²We note that this is a sample path definition, which is different from the definition in Chapter 4. However, for $\gamma \in \Gamma_s$, the two definitions are equivalent.

the renewal cycle, with the queue length at the beginning of the cycle being $Q[m-1]$ and cycle length being $\mathcal{T}[m]$. Since $Q[m]$ is Markov and $\gamma \in \Gamma_s$ we have that, conditioned on $Q[m-1]$ and $Q[m'-1]$, $R[m]$ is independent of $R[m']$ for any $m' \neq m$. Let $T[m]$ be the number of slots up to the start of the m^{th} renewal cycle ($T[1] = 0$). We have that the expected reward $\mathbb{E}[R[m]|Q[m-1]]$, in cycle m , is

$$r[m] = \mathbb{E} \left[Q[m-1]\mathcal{T}[m] + \sum_{n=1}^{\mathcal{T}[m]-1} \sum_{n'=1}^n A_s[T[m] + n'] \middle| Q[m-1] \right].$$

Since $A_s[n]$ are assumed to be IID we have that

$$r[m] = \mathbb{E} \left[Q[m-1]\mathcal{T}[m] + \frac{\mathcal{T}[m](\mathcal{T}[m]-1)\lambda}{2} \middle| Q[m-1] \right].$$

We define the holding cost at a decision epoch, where the queue length is q and a deterministic transmission/idle time τ is chosen, as a function $h(q, \tau)$ defined as

$$h(q, \tau) = q\tau + \frac{\tau(\tau-1)\lambda}{2}.$$

We note that $r[m] = \mathbb{E}h(q, \mathcal{T}(q))$, if $Q[m-1] = q$, and the expectation is over the randomized choice of τ . We note that for any $\gamma \in \Gamma_s$, since $\mathbb{E}_{\pi_\gamma} \mathbb{E}h(Q, \mathcal{T}(Q)) < \infty$, we also have that $\mathbb{E}_{\pi_\gamma} \mathbb{E}\mathcal{T}(Q) < \infty$. Then, we have from [33, Theorem D.16] that

$$\lim_{N \rightarrow \infty} \frac{\sum_{n=0}^{N-1} Q_s[n]}{N} \stackrel{a.s.}{=} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} h(Q, \mathcal{T}(Q))}{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} \mathcal{T}(Q)}.$$

Hence, for $\gamma \in \Gamma_s$, using Little's law [78], we have

$$D(\gamma) = \frac{\bar{Q}(\gamma)}{\lambda} = \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} h(Q, \mathcal{T}(Q))}{\lambda \mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} \mathcal{T}(Q)}. \quad (6.5)$$

To express $P_e(\gamma)$ as a time average we use the generalized $H = \lambda G$ form of Little's law [78]. Define $F_i[n] = 0$ for $n \in \{0, \dots, T_{d,i} - 1\}$ and $F_i[n] = \mathbb{I}_{E_i}$ for $n \in \{T_{d,i}, \dots\}$. Let G_I denote the fraction of the first I message symbols that are in error, and H_N denote the time-rate of message symbols errors over the first N slots, i.e.,

$$G_I = \frac{1}{I} \sum_{i=1}^I \sum_{n=1}^{\infty} (F_i[n] - F_i[n-1]),$$

$$H_N = \frac{1}{N} \sum_{n=0}^{N-1} \sum_{\{i: T_{d,i}=n\}} \mathbb{I}_{E_i}.$$

We note that $\lim_{I \rightarrow \infty} G_I = P_e(\gamma)$. If $H = \lim_{N \rightarrow \infty} H_N$, then $H = \lambda P_e(\gamma)$ from [78, Section 6]. For the batch transmission scheme, the set $\{i : T_{d,i} = n\}$ is empty, except for those n such that $n = T[m] = \sum_{m'=1}^{m-1} \mathcal{T}[m']$. We also have that $|\{i : T_{d,i} = T[m]\}| = S[m-1]$. We note that since the channel is memoryless and the decoding of a message batch of size $S[m-1]$ is done based on the channel outputs received only in the transmission time $\mathcal{T}[m-1]$, given $S[m-1]$ and $\mathcal{T}[m-1]$, the events $\{E_i : T_{d,i} = T[m]\}$ are independent of any other error event of a message symbol transmitted in a period other than $m-1$. As in the previous case we associate a reward $R[m]$ with the m^{th} renewal cycle, which is the total number of message symbols in error during the m^{th} transmission period, starting with $Q[m-1]$ message symbols in the queue. The expected reward $\mathbb{E}[R[m]|Q[m-1]]$ in the m^{th} renewal cycle is

$$\begin{aligned} r[m] &= \mathbb{E} \left[\sum_{\{i: T_{d,i}=T[m+1]\}} I_{E_i} \middle| Q[m-1] \right], \\ &= \mathbb{E}_{S[m], \mathcal{T}[m] | Q[m-1]} \left[\sum_{\{i: T_{d,i}=T[m+1]\}} Pr \{E_i | S[m], \mathcal{T}[m]\} \middle| Q[m-1] \right] \end{aligned}$$

where $Pr \{E_i | S[m] = s, \mathcal{T}[m] = \tau\}$ is the probability of error of the i^{th} symbol transmitted in the m^{th} period using a block code with size $|\mathcal{M}|^s$ and of codeword length τ . We define the error cost at a decision epoch to be $c_s(s, \tau) = \sum_{i \in \mathcal{C}(s, \tau)} Pr \{E_i | S = s, \mathcal{T} = \tau\}$, where at that decision epoch, $\mathcal{C}(s, \tau)$ is the set of s symbols which have been jointly encoded into a codeword of length τ . We note that for $\gamma \in \Gamma_s$, $\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T}|Q} \mathcal{T}(Q) < \infty$. Applying the Markov renewal reward theorem [33, Theorem D.16] we obtain that

$$H \stackrel{a.s.}{=} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{S, \mathcal{T} | Q} c_s(S(Q), \mathcal{T}(Q))}{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)}.$$

And therefore,

$$P_e(\gamma) = \frac{1}{\lambda} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{S, \mathcal{T} | Q} c_s(S(Q), \mathcal{T}(Q))}{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)}. \quad (6.6)$$

In the next section we discuss some approximations for $c_s(s, \tau)$ which are used in further analysis of the tradeoff problem.

6.2.1 The error cost $c_s(s, \tau)$

We note that $c_s(s, \tau) = \sum_{i \in \mathcal{C}(s, \tau)} \mathbb{E} [\mathbb{I}_{E_i} | S=s, \mathcal{T}=\tau]$, is the expected number of message symbols in error when a random block code is used to transmit s message symbols in τ slots. In the following, the error cost $c_s(s, \tau)$ is approximated using the Gallager random coding upper bound, since we are interested in an asymptotic characterization of the tradeoff curve as the error rate constraint $P_{e,c} \downarrow 0$. It is intuitive that policies $\gamma \in \Gamma_{s,f}$, for which $P_e(\gamma) \leq P_{e,c}$, have

block lengths growing to infinity as $P_{e,c} \downarrow 0$. Therefore in the regime of $P_{e,c} \downarrow 0$, approximating $c_s(s, \tau)$ using the Gallager random coding upper bound is not unreasonable [23]³. Let $P(s, \tau) \triangleq \Pr \{ \exists i \in \{1, \dots, s\} \text{ s.t. } m_{s,i} \neq \hat{m}_{s,i} | S = s, \mathcal{T} = \tau \}$, where $m_{s,i}$ and $\hat{m}_{s,i}$ are the transmitted and decoded symbols respectively.

For a random block code of length τ transmitting s symbols, we have from [22, Theorem 5.6.2] that the codeword error probability $P(s, \tau)$ is

$$\leq \min \left(1, e^{-\tau(E_0(\rho, \mathcal{Q}) - \rho \frac{s}{\tau} \ln |\mathcal{M}|)} \right) = \min \left(1, e^{-\tau E_r(\rho, s, \tau, \mathcal{Q})} \right), \quad (6.7)$$

where $E_0(\rho, \mathcal{Q}) = -\ln \left(\int_{y \in \mathcal{Y}} \left(\int_{x \in \mathcal{X}} d\mathcal{Q}(x) (dP_{Y|X}(y|x))^{\frac{1}{1+\rho}} \right)^{1+\rho} \right)$, $\rho \in [0, 1]$ and \mathcal{Q} is the channel input distribution. The best upper bound is obtained by minimizing the above bound over the parameter $\rho \in [0, 1]$ and the distribution \mathcal{Q} . We note that this can be done by maximising the exponent $E_r(\rho, s, \tau, \mathcal{Q})$ with respect to ρ and \mathcal{Q} . Depending on whether we do not choose/choose to optimise over ρ and/or \mathcal{Q} there are four different exponents: (a) $E_r(\rho, s, \tau, \mathcal{Q}) = (E_0(\rho, \mathcal{Q}) - \rho \frac{s}{\tau} \ln |\mathcal{M}|)$, (b) $E_r(\rho, s, \tau) = \max_{\mathcal{Q}} (E_0(\rho, \mathcal{Q}) - \rho \frac{s}{\tau} \ln |\mathcal{M}|)$, (c) $E_r(s, \tau, \mathcal{Q}) = \max_{\rho \in [0, 1]} (E_0(\rho, \mathcal{Q}) - \rho \frac{s}{\tau} \ln |\mathcal{M}|)$, and (d) $E_r(s, \tau) = \max_{\rho \in [0, 1]} \max_{\mathcal{Q}} (E_0(\rho, \mathcal{Q}) - \rho \frac{s}{\tau} \ln |\mathcal{M}|)$.

If a symbol is in error then the decoded codeword is also in error, therefore $\Pr \{ E_i | S = s, \mathcal{T} = \tau \} \leq P(s, \tau)$. Using the union bound on $P(s, \tau)$ and the above inequality we also obtain that

$$P(s, \tau) \leq \sum_{i=1}^s \Pr \{ E_i | S = s, \mathcal{T} = \tau \} \leq sP(s, \tau). \quad (6.8)$$

In the following, we approximate $c_s(s, \tau)$ by assuming that all symbols are decoded incorrectly if the codeword is in error. Different upper bounds on the codeword error probability (from the different random coding exponents) can be used to obtain approximations for $c_s(s, \tau)$. The following approximations are possible: (a) $s \min(1, e^{-\tau E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|})$ where both ρ and \mathcal{Q} are fixed, (b) $s \min(1, e^{\min_{\mathcal{Q}} (-\tau E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|)})$, where ρ is fixed, (c) $s e^{\min_{\rho \in [0, 1]} (-\tau E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|)}$, where \mathcal{Q} is fixed, and (d) $s e^{\min_{\rho \in [0, 1], \mathcal{Q}} (-\tau E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|)}$. We note that approximation (a) is found to be analytically tractable as ρ and \mathcal{Q} are fixed and do not depend on s and τ . In (b) the optimal \mathcal{Q} is a function of the fixed ρ and the following analysis applies, as it would be similar to that of (a) with \mathcal{Q} being fixed at the optimal \mathcal{Q} for the fixed ρ . We note that the minimizing ρ and \mathcal{Q} are not known explicitly as a function of s and τ even for simple channel models [22, Chap 5]. However, for R-model-A, we shall see that the exponential decay, of the minimum average error rate with the constraint on the average delay, can be shown to be governed by the best exponent (d), even with the following assumption. For R-model-B, we comment on how the asymptotic results can be

³In [23] Gallager has shown that the random coding upper bound is tight for the probability of block error for an IID ensemble of codes, by showing that the lower bound on the ensemble error probability is at most $\mathcal{O} \left(\frac{1}{\tau} \ln \left(\frac{1}{\sqrt{\tau}} \right) \right)$ away from the random coding upper bound for large block length τ , for code rates less than the capacity of the channel $P_{Y|X}$.

extended to other approximations of $c_s(s, \tau)$. We make the following assumption :

C1 : The error cost or the expected number of symbols in error for a random block code of codeword length τ transmitting s symbols is $c_s(s, \tau) = s \min \left(1, e^{(-\tau E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|)} \right)$, where ρ and \mathcal{Q} are fixed and are such that $\frac{E_0(\rho, \mathcal{Q})}{\rho} > \lambda \ln |\mathcal{M}|$.

The tradeoff problem that we consider is defined separately for R-model-A and R-model-B in the following sections.

6.3 Asymptotic analysis for R-model-A

6.3.1 Problem Statement

We consider the tradeoff of average error rate with average delay for $\gamma \in \Gamma_s$. The following analysis has been presented in [58] and [56]. The TRADEOFF problem is:

$$\begin{aligned} & \text{minimize}_{\gamma \in \mathcal{S}_b} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{S}, \mathcal{T} | Q} c_s(S(Q), \mathcal{T}(Q))}{\lambda \mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)} \\ & \text{such that} \quad \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} h(Q, \mathcal{T}(Q))}{\lambda \mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)} \leq D_c, \end{aligned} \quad (6.9)$$

where we have used (6.5) and (6.6). If $\mathcal{S}_b = \Gamma_s$, then let $P_e^*(D_c)$ denote the optimal value of TRADEOFF. If $\mathcal{S}_b = \Gamma_{s,f}$, then the optimal value of TRADEOFF is denoted by $P_{e,f}^*(D_c)$, and if $\mathcal{S}_b = \Gamma_{s,N_c}$, then the optimal value of TRADEOFF is denoted by $P_{e,N_c}^*(D_c)$. We can also consider the equivalent problem EQT-TRADEOFF:

$$\begin{aligned} & \text{minimize}_{\gamma \in \mathcal{S}_b} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} h(Q, \mathcal{T}(Q))}{\lambda \mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)} \\ & \text{such that} \quad \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{S}, \mathcal{T} | Q} c_s(S(Q), \mathcal{T}(Q))}{\lambda \mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{T} | Q} \mathcal{T}(Q)} \leq P_{e,c}. \end{aligned} \quad (6.10)$$

If $\mathcal{S}_b = \Gamma_s$, then let $D^*(P_{e,c})$ denote the optimal value of TRADEOFF. If $\mathcal{S}_b = \Gamma_{s,f}$, then the optimal value of TRADEOFF is denoted by $D_f^*(P_{e,c})$, and if $\mathcal{S}_b = \Gamma_{s,N_c}$, then the optimal value of TRADEOFF is denoted by $D_{N_c}^*(P_{e,c})$. In the next section we discuss the asymptotic behaviour of $P_{e,f}^*(D_c)$ and $P_e^*(D_c)$ as $D_c \rightarrow \infty$. We note that as in Section 4.2.1, we can show that there exists an optimal policy for the above problem.

6.3.2 Asymptotic analysis

In the following lemma, we obtain $\inf_{D_c} P_{e,N_c}^*(D_c)$.

Lemma 6.3.1. We have that $\inf_{D_c} P_{e,N_c}^*(D_c) = P_{N_c}(\lambda)$, where

$$P_{N_c}(\lambda) = \begin{cases} e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)}, & \text{if } 0 \leq \lambda N_c \leq s', \\ \frac{1}{\lambda N_c} \left((\lambda N_c - s') + s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|} \right), & \text{otherwise,} \end{cases}$$

where s' is the unique solution of the equation $(1 + \rho s' \ln |\mathcal{M}|) e^{\rho s' \ln |\mathcal{M}|} = e^{N_c E_0(\rho, \mathcal{Q})}$.

Proof. We first show that $P_{N_c}(\lambda) \leq P_{e,N_c}^*(D_c)$, $\forall D_c$. We note that as we are considering policies $\gamma \in \Gamma_{s,N_c}$, the expected transmission time is always N_c and therefore (6.5) and (6.6) simplify to $\frac{\mathbb{E}_{\pi, \gamma} h(Q, N_c)}{\lambda N_c}$ and $\frac{\mathbb{E}_{\pi, \gamma} \mathbb{E}_S |Q| c_s(S, N_c)}{\lambda N_c}$ respectively. We consider the following optimization problem:

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && \frac{1}{\lambda N_c} \mathbb{E}_{\pi} \left[S \min \left(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right) \right], \\ & \text{such that} && \mathbb{E}_{\pi} S \geq \lambda N_c + \epsilon, \epsilon \geq 0, \end{aligned} \quad (6.11)$$

where π is any distribution for S with support on \mathbb{R}_+ , irrespective of the policy γ . With $\epsilon = 0$ in (6.11) leading to the constraint $\mathbb{E}_{\pi} S \geq \lambda N_c$, the optimal value of the minimization problem (6.11) is a lower bound to $P_{e,N_c}^*(D_c)$ for any finite D_c as the minimization is over all possible distributions of S and we impose only the constraint that the average service rate $\frac{\mathbb{E}_{\pi} S}{N_c}$ has to be greater than or equal to the arrival rate λ . From Appendix 6.A we obtain that

$$P_{e,N_c}^*(D_c) \geq \begin{cases} e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)} & \text{if } 0 \leq \lambda N_c \leq s', \\ \frac{1}{\lambda N_c} \left((\lambda N_c - s') + s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|} \right) & \text{otherwise,} \end{cases} \quad (6.12)$$

where s' is the unique solution of the equation $(1 + \rho s' \ln |\mathcal{M}|) e^{\rho s' \ln |\mathcal{M}|} = e^{N_c E_0(\rho, \mathcal{Q})}$, for every finite D_c . Let $P_{N_c}(\lambda)$ denote the RHS of (6.12). We note that $P_{N_c}(\lambda)$ is equal to the value of the lower convex envelope of $\frac{c_s(s, N_c)}{\lambda N_c}$ as a function of s , at $s = \lambda N_c$.

Now we show that for any $\epsilon > 0$, it is possible to construct a policy $\gamma \in \Gamma_{s,N_c}$ such that $P_e(\gamma) \leq P_{N_c}(\lambda) + \epsilon$. Then we have that $P_{e,N_c}^*(D(\gamma)) \leq P_{N_c}(\lambda) + \epsilon$ and therefore $P_{N_c}(\lambda) = \inf_{D_c} P_{e,N_c}^*(D_c)$. In the first step of the construction we consider the optimization problem (6.11) but with an $\epsilon > 0$. The optimal value of the problem (6.11) as a function of ϵ is denoted by $s(\epsilon)$. A feasible distribution for S that achieves $s(\epsilon) + \delta$ is called (δ, ϵ) -optimal for a $\delta > 0$. In Appendix 6.A it is shown that there exists a distribution with finite support that is (δ, ϵ) -optimal for any $\epsilon > 0$ and $\delta > 0$.

Now we proceed to the second step of the construction. For an $\epsilon > 0$ we use any (δ, ϵ) -optimal distribution to construct a policy as follows. Let $U[m]$ be a sequence of IID random variables distributed as the (δ, ϵ) -optimal distribution. The policy γ uses batch sizes $S[m] = \min(Q[m] - 1, U[m])$. We note that γ is stationary. Since $\epsilon > 0$, $\mathbb{E}U > \lambda N_c$. Therefore, the EMC $Q[m]$ under γ is positive recurrent. It can also be shown that the average delay is finite as $U[m]$ has finite support. Therefore, γ is admissible, i.e., $\gamma \in \Gamma_s$. Evaluating the average error cost for γ we have

that $P_e(\gamma) \leq$

$$\frac{\mathbb{E}_{c_s}(U, N_c)}{\lambda N_c} = \begin{cases} e^{\epsilon \rho \ln |\mathcal{M}|} e^{-N_c E_0(\rho, \mathcal{Q}) + (\rho \ln |\mathcal{M}|) \lambda N_c} + \epsilon e^{-N_c E_0(\rho, \mathcal{Q}) + (\rho \ln |\mathcal{M}|) (\lambda N_c + \epsilon)}, & \text{if } \lambda N_c \in [0, s' - \epsilon], \\ \frac{1}{\lambda N_c} \left(\lambda N_c - s' + s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|} \right) + \frac{\epsilon + \delta}{\lambda N_c}, & \text{otherwise.} \end{cases}$$

Let $\delta = \epsilon$. Therefore, $P_e(\gamma) \leq P_{N_c}(\lambda) + \mathcal{O}(\epsilon)$. Hence, by an appropriate choice of ϵ in (6.11), we can show that $\forall \epsilon > 0$, there exists $D_c = D(\gamma)$ such that $P_{e, N_c}^*(D_c) \leq P_e(\gamma) \leq P_{N_c}(\lambda) + \epsilon$. Hence, $P_{N_c}(\lambda) = \inf_{D_c} P_{e, N_c}^*(D_c)$. \square

Remark 6.3.2. Since any policy $\gamma \in \Gamma_{s, N_c}$ is feasible for TRADEOFF for any $D_c \geq D(\gamma)$, we also have that $P_{N_c}(\lambda) = \inf_{\gamma \in \Gamma_{s, N_c}} P_e(\gamma)$. We note that $P_{e, N_c}^*(D_c)$ is a non-increasing function of D_c . Therefore, $P_{N_c}(\lambda) = \lim_{D_c \rightarrow \infty} P_{e, N_c}^*(D_c)$. We note that $P_{N_c}(\lambda)$ is similar to $c(\lambda)$ in Chapter 4.

We now present a lemma which formalizes our intuition that $P_{N_c}(\lambda)$ approaches zero as the code-word length N_c approaches infinity.

Lemma 6.3.3. If $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, then $P_{N_c}(\lambda) = e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)}$ for large enough N_c , and therefore, $\lim_{N_c \rightarrow \infty} P_{N_c}(\lambda) = 0$.

Proof. From Lemma 6.3.1 we have that

$$P_{N_c}(\lambda) = \begin{cases} e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)} & \text{if } 0 \leq \lambda N_c \leq s', \\ \frac{1}{\lambda N_c} \left((\lambda N_c - s') + s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|} \right) & \text{otherwise,} \end{cases}$$

where s' solves the equation $(1 + \rho s' \ln |\mathcal{M}|) e^{\rho s' \ln |\mathcal{M}|} = e^{N_c E_0(\rho, \mathcal{Q})}$. We note that if we show that, $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$ implies that $\lambda N_c \leq s'$ for large enough N_c , then we have that $P_{N_c}(\lambda) = e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)}$ for large enough N_c , which leads to both results of the lemma. We have that

$$\begin{aligned} e^{-N_c E_0(\rho, \mathcal{Q})} \left[e^{\rho s' \ln |\mathcal{M}|} \{1 + \rho s' \ln |\mathcal{M}|\} \right] &= 1 \text{ or ,} \\ \rho s' \ln |\mathcal{M}| + \ln(1 + s' \rho \ln |\mathcal{M}|) &= N_c E_0(\rho, \mathcal{Q}). \end{aligned} \quad (6.13)$$

We note that the LHS of (6.13) is an increasing continuous function of s' with LHS = 0 at $s' = 0$. Thus for any N_c , there is a unique s' which solves the equation (6.13). Also as N_c increases this unique solution $s'(N_c)$, as a function of N_c , also increases, and as $N_c \uparrow \infty$, $s'(N_c) \uparrow \infty$. We multiply both sides of (6.13) by λ , divide by $s'(N_c)$ and $E_0(\rho, \mathcal{Q})$ to yield

$$\lambda \ln |\mathcal{M}| \frac{\rho}{E_0(\rho, \mathcal{Q})} + \frac{\lambda}{E_0(\rho, \mathcal{Q})} \frac{\ln(1 + s'(N_c) \rho \ln |\mathcal{M}|)}{s'(N_c)} = \frac{\lambda N_c}{s'(N_c)} \quad (6.14)$$

Since we have assumed that $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, $\delta \triangleq 1 - \lambda \ln |\mathcal{M}| \frac{\rho}{E_0(\rho, \mathcal{Q})} > 0$. For sufficiently large N_c , $s'(N_c)$ can be made sufficiently large, so that $\frac{\ln(1 + s'(N_c) \rho \ln |\mathcal{M}|)}{s'(N_c)} < \frac{\delta}{2} \frac{E_0(\rho, \mathcal{Q})}{\lambda}$ which implies that

$\frac{\lambda N_c}{s'(N_c)} \leq 1 - \frac{\delta}{2}$. Hence for large N_c , $\frac{\lambda N_c}{s'(N_c)} < 1$ and therefore $P_{N_c}(\lambda) = e^{-N_c(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)}$. Therefore, $P_{N_c}(\lambda) \downarrow 0$ as $N_c \uparrow \infty$ if $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$. \square

We now consider the asymptotic behaviour of $P_{e,f}^*(D_c)$, as $D_c \uparrow \infty$. We first show that if $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, then $P_{e,f}^*(D_c) \downarrow 0$ as $D_c \uparrow \infty$. This result is obtained by showing that for a sequence of policies $\mu_k \in \Gamma_{s,f}$, for which $D(\mu_k) \uparrow \infty$, $P_e(\mu_k) \downarrow 0$. Then $P_{e,f}^*(D(\mu_k)) \downarrow 0$. Furthermore, since $D_{\mu_k} \uparrow \infty$, $P_{e,f}^*(D_c) \downarrow 0$ for any sequence $D_c \uparrow \infty$. The sequence of policies $\mu_k \in \Gamma_{s,f}$ is parametrized by the codeword length $N_{c,k}$, where $N_{c,k}$ is an unbounded increasing sequence. A fixed rate of transmission r is chosen such that $\lambda \ln |\mathcal{M}| < r \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$. Let $R_k = rN_{c,k}$. The policy μ_k chooses $S[m] = \min(Q[m-1], R_k)$. We note that since $r > \lambda$, it can be shown that $\mu_k \in \Gamma_{s,f}, \forall k$.

Proposition 6.3.4. If $\lambda \ln |\mathcal{M}| < r \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, for the sequence of policies μ_k defined above, we have that $\lim_{k \rightarrow \infty} \frac{-\ln P_e(\mu_k)}{D(\mu_k)} = \frac{r-\lambda}{r-\lambda+r\lambda} (E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|)$, with $D(\mu_k) \uparrow \infty$. Therefore $P_e(\mu_k) \downarrow 0$.

Proof. For μ_k , we note that $P_e(\mu_k) \leq \frac{\mathbb{E}_{\pi_{\mu_k}} [S e^{-N_{c,k} E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|}]}{\lambda N_{c,k}}$. As $S \leq rN_{c,k}$ we have that

$$P_e(\mu_k) \leq \frac{r}{\lambda} \mathbb{E}_{\pi_{\mu_k}} \left[e^{-N_{c,k} E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right], \quad (6.15)$$

$$\leq \frac{r}{\lambda} \left(e^{E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|} \right)^{-N_{c,k}}. \quad (6.16)$$

From Appendix 6.D we have that

$$\mathbb{E}_{\mu_k} Q \leq \frac{\sigma^2}{2(r-\lambda)} + \frac{\lambda N_{c,k}}{2} + \frac{r \lambda N_{c,k}}{r-\lambda}.$$

We note that in our case, there is an extra time average holding cost of $\frac{\lambda(N_{c,k}-1)}{2}$ due to the customers waiting during a transmission period of $N_{c,k}$ slots. Therefore,

$$\begin{aligned} D(\mu_k) &\leq \frac{\sigma^2}{2\lambda(r-\lambda)} - \frac{1}{2} + N_{c,k} \left(1 + \frac{r\lambda}{r-\lambda} \right), \\ \frac{r-\lambda}{r-\lambda+r\lambda} \left(D(\mu_k) - \frac{\sigma^2}{2\lambda(r-\lambda)} + \frac{1}{2} \right) &\leq N_{c,k}. \end{aligned} \quad (6.17)$$

Since $D(\mu_k)$ is at least $N_{c,k}$ we have that $D(\mu_k) \uparrow \infty$ as $N_{c,k} \uparrow \infty$.

Substituting the lower bound on $N_{c,k}$ from (6.17) in (6.16), we have that

$$P_e(\mu_k) \leq \frac{r}{\lambda} \left(e^{E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|} \right)^{-\frac{r-\lambda}{r-\lambda+r\lambda} \left(D(\mu_k) - \frac{\sigma^2}{2\lambda(r-\lambda)} + \frac{1}{2} \right)}.$$

Hence $\lim_{k \rightarrow \infty} \frac{-\ln P_e(\mu_k)}{D(\mu_k)} = \frac{r-\lambda}{r-\lambda+r\lambda} (E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|)$, and $P_e(\mu_k) \downarrow 0$.

□

We now show that any sequence of policies $\gamma_k \in \Gamma_{s,f}$, for which $P_e(\gamma_k) \rightarrow 0$, has $N_c(\gamma_k) \rightarrow \infty$.

Lemma 6.3.5. For any sequence of policies $\gamma_k \in \Gamma_{s,f}$, $N_c(\gamma_k) \rightarrow \infty$ if $P_e(\gamma_k) \rightarrow 0$.

Proof. Assume that there exists a sequence of policies $\gamma_k \in \Gamma_{s,f}$ for which $P_e(\gamma_k) \rightarrow 0$, but $N_c(\gamma_k) \not\rightarrow \infty$. Therefore $\limsup_{k \rightarrow \infty} N_c(\gamma_k) = N_{c,l} < \infty$. From Lemma 6.3.1 we have that $P_e(\gamma_k) \geq P_{N_{c,k}}(\lambda), \forall k$. If $N_{c,l} < \infty$ we have that $\lim_{k \rightarrow \infty} P_e(\gamma_k) \geq P_{N_{c,l}}(\lambda) > 0$, which contradicts the assumption that $P_e(\gamma_k) \downarrow 0$. Therefore $N_c(\gamma_k)$ should necessarily grow to infinity for the sequence of policies γ_k . □

From the above result and since $P_{e,f}^*(D_c) \downarrow 0$ as $D_c \uparrow \infty$, we know that to obtain an asymptotic characterization of $P_{e,f}^*(D_c)$ as $D_c \uparrow \infty$, we need only consider sequences of policies $\gamma_k \in \Gamma_{s,f}$ with $N_c(\gamma_k) \uparrow \infty$. Now for any sequence $\gamma_k \in \Gamma_{s,f}$, with $N_c(\gamma_k) \rightarrow \infty$, we present an upper bound on the decay rate $\lim_{D_c \rightarrow \infty} -\frac{\ln(P_{e,f}^*(D_c))}{D_c}$.

Proposition 6.3.6. If $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, the exponential decay rate of $P_{e,f}^*(D_c)$ has the following upper bound:

$$\lim_{D_c \rightarrow \infty} \frac{-\ln(P_{e,f}^*(D_c))}{D_c} \leq \frac{2}{3} (E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|),$$

for a fixed ρ and \mathcal{Q} . The best upper bound on the decay rate is obtained by fixing ρ and \mathcal{Q} to be ρ^* and \mathcal{Q}^* respectively, where

$$(\rho^*, \mathcal{Q}^*) = \arg \max_{\rho \in [0,1], \mathcal{Q}} \frac{2}{3} (E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|).$$

Proof. From Proposition 6.3.4, we have that $P_{e,f}^*(D_c) \downarrow 0$ as $D_c \uparrow \infty$. From Lemma 6.3.5, we need only consider any sequence of $\gamma_k \in \Gamma_{s,f}$ such that $N_{c,k} = N_c(\gamma_k) \rightarrow \infty$. We have from Lemma 6.3.1 that $P_e(\gamma_k) \geq P_{N_{c,k}}(\lambda)$. Then from Lemma 6.3.3 for large enough k , as $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$ we have that $P_e(\gamma_k) \geq e^{-N_{c,k}(E_0(\rho) - \rho \lambda \ln |\mathcal{M}|)}$. From (6.5), for a fixed codeword length $N_{c,k}$ we have that $D(\gamma_k) = \frac{\mathbb{E}_\pi \mathcal{Q}}{\lambda} + \frac{N_{c,k} - 1}{2}$. For any γ_k , $\mathbb{E}_\pi \mathcal{Q} \geq \mathbb{E}_\pi \mathbb{E}_{S|Q} S = \lambda N_{c,k}$ as $Q \geq S$. Therefore $D(\gamma_k) \geq \frac{3N_{c,k}}{2} - \frac{1}{2}$ or $\frac{2}{3}(D(\gamma_k) + \frac{1}{2}) \geq N_{c,k}$. Hence $P_e(\gamma_k) \geq e^{-\frac{2}{3}(D(\gamma_k) + \frac{1}{2})(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)}$ or $\lim_{k \rightarrow \infty} \frac{-\ln(P_e(\gamma_k))}{D(\gamma_k)} \leq \frac{2}{3}(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)$. We note that the above upper bound on the exponential decay rate holds for any sequence of policies γ_k . For (6.9), we have by definition that for $D_c \geq 1$ and $\epsilon > 0$, there exists a γ such that $D(\gamma) \leq D_c$ and $P_e(\gamma) \leq P_{e,f}^*(D_c) + \epsilon$. For any sequence γ_k , choose ϵ from the sequence $\epsilon_k = e^{-e^{N_{c,k}}}$. Then, we have that $\lim_{D_c \rightarrow \infty} \frac{-\ln(P_{e,f}^*(D_c))}{D_c} \leq \frac{2}{3}(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)$. The best upper bound on decay rate is then obtained by choosing the fixed ρ and \mathcal{Q} to be ρ^* and \mathcal{Q}^* respectively. □

Now we consider the decay rate, $\lim_{D_c \uparrow \infty} -\frac{\ln P_e^*(D_c)}{D_c}$. Unlike the class of policies considered above, where the transmission duration was always fixed to be a parameter N_c , for $\gamma \in \Gamma_s$ we have that at each decision epoch m , the transmitter can choose both the batch size $S[m]$ and the transmission duration $\mathcal{T}[m]$. We obtain a lower bound to the exponential decay rate of $P_e^*(D_c)$ as $D_c \uparrow \infty$, by obtaining the exponential decay rate $\lim_{k \rightarrow \infty} -\frac{\ln P_e(e_k)}{D(e_k)}$, for a sequence of exhaustive service (EXH) policies e_k , with $D(e_k) \uparrow \infty$. Consider an EXH policy e . The policy e chooses $S[m] = Q[m-1]$ and $\mathcal{T}[m]$ such that the codeword error probability is at most $P_{e,b}$. Intuition suggests that the codeword length $\mathcal{T}[m]$ has to be at least a minimum value, say $\tau(S[m])$, which is a function of the batch size $S[m]$, to guarantee that the codeword error probability $P(S[m], \mathcal{T}[m])$ is at most $P_{e,b}$ for every m . The message alphabet size is $|\mathcal{M}|$, therefore the alphabet size of the batch of s message symbols is $|\mathcal{M}|^s$. The rate in nats is $R = \frac{s \ln |\mathcal{M}|}{\tau(s)}$. Consider a randomly generated codebook, in which each codeword symbol is chosen independently according to the distribution \mathcal{Q} on the input alphabet \mathcal{X} . The receiver is assumed to do maximum likelihood decoding of the joint message. From [22, Theorem 5.6.2], we have that the average probability of codeword error is bounded above as:

$$P(s, \tau) \leq e^{-\tau\{-\rho R + E_0(\rho, \mathcal{Q})\}},$$

$$\text{or } \ln P(s, \tau) \leq \rho s \ln |\mathcal{M}| - \tau E_0(\rho, \mathcal{Q}),$$

where $\rho \in [0, 1]$, $E_0(\rho, \mathcal{Q}) = -\ln \int_{y \in \mathcal{Y}} \left[\int_{x \in \mathcal{X}} \mathcal{Q}(x) P_{Y|X}(y|x) \left(\frac{1}{1+\rho}\right) dx \right]^{1+\rho} dy$. To guarantee the average error rate requirement, we constrain $P(s, \tau)$ to be $\leq P_{e,r}$. If

$$e^{-\tau\{-\rho R + E_0(\rho, \mathcal{Q})\}} \leq P_{e,r},$$

then $P(s, \tau) \leq P_{e,r}$. For every s , if τ is chosen as a function $\tau(s)$ to satisfy the above inequality, we have that

$$\frac{-\ln P_{e,r}}{E_0(\rho, \mathcal{Q})} + \frac{\rho s \ln |\mathcal{M}|}{E_0(\rho, \mathcal{Q})} \leq \tau(s).$$

Thus $\tau(s)$ has to be chosen as the smallest integer greater than or equal to $as + b$, where $a = \frac{\rho \ln |\mathcal{M}|}{E_0(\rho, \mathcal{Q})}$ and $b = \frac{-\ln P_{e,r}}{E_0(\rho, \mathcal{Q})}$. Therefore $\tau(s) = \lceil as + b \rceil$, $s > 0$, so that $P(s, \tau) \leq P_{e,r}$.

We assume that e chooses $\mathcal{T}[m] = \tau(S[m])$ if $S[m] > 0$ and $\mathcal{T}[m] = 1$ if $S[m] = 0$. Let the stationary distribution corresponding to policy e be π . Then for e , $P_e(e)$ is less than or equal to $\frac{\mathbb{E}_\pi S P_{e,b}}{\lambda \mathbb{E}_\pi \mathcal{T}}$. Since $\mathbb{E}_\pi S = \lambda \mathbb{E}_\pi \mathcal{T}$ we have that $P_e(e) \leq P_{e,b}$. For a given $P_{e,b}$ let the average delay for e be denoted by $D_{EXH}(P_{e,b})$. We have the following upper bound $D_u(P_{e,b})$ on the average delay $D_{EXH}(P_{e,b})$.

Proposition 6.3.7. $D_{EXH}(P_{e,b}) \leq D_u(P_{e,b})$, where

$$D_u(P_{e,b}) = \left(\frac{b+1}{b} \right) \left(\frac{a\sigma^2}{2(1-a\lambda)} + \frac{3(b+1)\lambda}{2(1-a\lambda)} + \frac{a\sigma^2}{2(1-a^2\lambda^2)} - \frac{\lambda}{2} \right).$$

The proof is given in Appendix 6.B. We now compare the above upper bound with the upper bound obtained by Musy and Telatar [40]. We note that the model considered by Musy and Telatar in [40] is a continuous time model with Poisson arrivals. In [40], if a batch size $s > 0$ is used, then the service time is $D + ks$. For simplicity, assume that D and k are integers. Let δ be such that $b\delta = D$ and $a\delta = k$, where a and b are integers. Then if we assume that each slot in our model is of δ duration, then the time taken for service of s customers in our model and the model in [40] is the same. Let the Poisson arrival rate of customers in Musy's model be λ_m . Now assume that the arrival process in our model is Bernoulli with an arrival probability (rate) of $\lambda = \lambda_m\delta$ in each slot. We note that if a and b are integers, then $\tau(s) = as + b$ for every $s > 0$. Then the upper bound in Proposition 6.3.7 simplifies to

$$\frac{a\sigma^2}{2(1-a\lambda)} + \frac{3b\lambda}{2(1-a\lambda)} + \frac{a\sigma^2}{2(1-a^2\lambda^2)} - \frac{\lambda}{2}. \quad (6.18)$$

We note that $\sigma^2 = \lambda - \lambda^2$ for a Bernoulli arrival process. Substituting $\lambda = \lambda_m\delta$, $a = \frac{k}{\delta}$ and $b = \frac{D}{\delta}$ and taking the limit as $\delta \downarrow 0$ (along a sequence such that a and b are integers) we obtain that the average queue length in the limit is

$$\frac{3D\lambda_m}{2(1-\lambda_mk)} + \frac{\lambda_mk(2 + \lambda_mk)}{2(1-\lambda_mk)}. \quad (6.19)$$

We note that this is the same as that obtained by Musy in Section 2.3.2 of his thesis [39].

The above upper bound leads to the following characterization of the exponential decay rate of the average error rate.

Proposition 6.3.8. The exponential decay rate : $\lim_{D_c \rightarrow \infty} \frac{-\ln P_e^*(D_c)}{D_c} \geq \frac{2}{3}(E_0(\rho^*, \mathcal{Q}^*) - \rho^* \lambda \ln |\mathcal{M}|)$.

Proof. Let $P_{e,u}(D_c) \triangleq \inf \{p \in [0, 1] : D_u(p) \leq D_c\}$. We note that if $D_u(p) \leq D_c$, then the e policy corresponding to p has $D(e) \leq D_c$ and therefore $P_e^*(D_c) \leq P_e(e) \leq p$. Therefore, $P_e^*(D_c) \leq P_{e,u}(D_c)$. We have that $P_{e,u}(D_c)$ is such that

$$\left(\frac{b_u + 1}{b_u \lambda}\right) \left(\frac{3(b_u + 1)\lambda}{2(1-a\lambda)} + \frac{a\sigma^2}{2(1-a\lambda)} + \frac{a\sigma^2}{2(1-a^2\lambda^2)} - \frac{\lambda}{2}\right) = D_c,$$

where $b_u = \frac{-\ln P_{e,u}(D_c)}{E_0(\rho, \mathcal{Q})}$. We obtain the following quadratic equation in b_u :

$$(b_u + 1)((b_u + 1)c_1 + c_2) = b_u \lambda D_c,$$

where $c_1 = \frac{3\lambda}{2(1-a\lambda)}$ and $c_2 = \frac{a\sigma^2(2+a\lambda)}{2(1-a^2\lambda^2)} - \frac{\lambda}{2}$. Or

$$\frac{-\ln P_{e,u}(D_c)}{E_0(\rho, \mathcal{Q})} = \frac{1}{2c_1} \left[\lambda D_c - 2c_1 - c_2 + \sqrt{(\lambda D_c - 2c_1 - c_2)^2 - 4c_1(c_1 + c_2)} \right],$$

since we want the largest exponent. Let $D_{c,k}$ be a sequence such that $D_{c,k} \uparrow \infty$. Then, we have a sequence $p_k = \inf \{p : D_u(p) \leq D_{c,k}\}$. As above, we obtain

$$\frac{-\ln p_k}{E_0(\rho, \mathcal{Q})} = \frac{1}{2c_1} \left[\lambda D_{c,k} - 2c_1 - c_2 + \lambda D_{c,k} \sqrt{\left(1 - \frac{2c_1 + c_2}{D_{c,k}}\right)^2 - \frac{4c_1(c_1 + c_2)}{D_{c,k}^2}} \right].$$

Therefore, we have that $\lim_{D_{c,k} \rightarrow \infty} \frac{-\ln p_k}{D_{c,k}} = \frac{\lambda E_0(\rho, \mathcal{Q})}{c_1} = \frac{2}{3}(E_0(\rho, \mathcal{Q}) - \rho \lambda \ln |\mathcal{M}|)$, after substituting for a in c_1 . Hence, we have that $\lim_{D_c \rightarrow \infty} \frac{-\ln P_e^*(D_c)}{D_c} \geq \frac{2}{3}(E_0(\rho^*, \mathcal{Q}^*) - \rho^* \lambda \ln |\mathcal{M}|)$, by fixing ρ and \mathcal{Q} to be ρ^* and \mathcal{Q}^* respectively. \square

We note that we do not have any upper bound on the exponential decay rate $\lim_{D_c \rightarrow \infty} \frac{-\ln P_e^*(D_c)}{D_c}$, over the class of policies Γ_s .

6.3.3 Integer valued queue evolution

We note that for R-model-A, the queue length process evolves on \mathbb{R}_+ . Suppose we consider the case, where the arrival process $A[m] \in \{0, 1, \dots, A_{max} \in \mathbb{Z}_+\}$, $q_0 \in \mathbb{Z}_+$, and $S[m] \in \mathbb{Z}_+$. Then the queue length process $(Q[m], m \geq 0)$ would evolve on \mathbb{Z}_+ . Let us denote this model as I-model-A. Intuitively, $P_e^*(D_c)$, $P_{e,f}^*(D_c)$, or $P_{e,N_c}^*(D_c)$ (or $D^*(P_{e,c})$, $D_f^*(P_{e,c})$, or $D_{N_c}^*(P_{e,c})$) for R-model-A (with $Pr \{A[m] = a\} > 0$ only for $a \in \{0, 1, \dots, A_{max}\}$) would be a lower bound to the same performance measures for I-model-A, as the set of feasible policies for R-model-A would always be a superset of the set of feasible policies for I-model-A for the optimization problem (6.9) (or (6.10)), from which $P_e^*(D_c)$, $P_{e,f}^*(D_c)$, or $P_{e,N_c}^*(D_c)$ (or $D^*(P_{e,c})$, $D_f^*(P_{e,c})$, or $D_{N_c}^*(P_{e,c})$) is obtained as the optimal value.

For I-model-A, we note that $P_{e,f}^*(D_c) \downarrow 0$ as $D_c \uparrow \infty$, since the sequence of policies μ_k can be restricted to use only non-negative integer valued $R_k = rN_{c,k}$ by an appropriate choice of the sequence $N_{c,k}$ and a rational number $r > \lambda$. Furthermore, in this case, the following tighter bound on the exponential decay rate can be obtained for the sequence μ_k .

Proposition 6.3.9. If $\lambda \ln |\mathcal{M}| < r \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q})}{\rho}$, for the sequence of policies μ_k defined above, we have that $\lim_{k \rightarrow \infty} \frac{-\ln P_e(\mu_k)}{D(\mu_k)} = \frac{2}{3}(E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|)$, with $D(\mu_k) \uparrow \infty$. Therefore $P_e(\mu_k) \downarrow 0$.

Proof. For μ_k , we note that $P_e(\mu_k) \leq \frac{\mathbb{E}_{\pi_{\mu_k}} [S e^{-N_{c,k} E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|}]}{\lambda N_c}$. As $S \leq r N_{c,k}$ we have that

$$P_e(\mu_k) \leq \frac{r}{\lambda} \mathbb{E}_{\pi_{\mu_k}} \left[e^{-N_{c,k} E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right], \quad (6.20)$$

$$\leq \frac{r}{\lambda} \left(e^{E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|} \right)^{-N_{c,k}}. \quad (6.21)$$

From Denteneer et al. [19, equation (9) and the upper bound in (12)] we have that

$$\mathbb{E}_{\pi_{\mu_k}} Q \leq \frac{\sigma^2}{2(r-\lambda)} + \frac{\lambda N_{c,k}}{2} + \frac{1}{2} \min(\lambda N_{c,k}, r N_{c,k} - 1).$$

For k large enough, as $N_{c,k} \rightarrow \infty$ and $\lambda < r$, we have that

$$\mathbb{E}_{\pi_{\mu_k}} Q \leq \frac{\sigma^2}{2(r-\lambda)} + \lambda N_{c,k}.$$

We note that in our case, there is an extra time average holding cost of $\frac{\lambda(N_{c,k}-1)}{2}$ due to the customers waiting during a transmission period of $N_{c,k}$ slots. Therefore an upper bound on the time average queue length is

$$\frac{\sigma^2}{2(r-\lambda)} + \frac{3\lambda N_{c,k}}{2} - \frac{\lambda}{2}.$$

Therefore

$$\begin{aligned} D(\mu_k) &\leq \frac{\sigma^2}{2\lambda(r-\lambda)} + \frac{3N_{c,k}}{2} - \frac{1}{2}, \\ \frac{2}{3} \left(D(\mu_k) - \frac{\sigma^2}{2\lambda(r-\lambda)} + \frac{1}{2} \right) &\leq N_{c,k}. \end{aligned} \quad (6.22)$$

Since $D(\mu_k)$ is at least $N_{c,k}$ we have that $D(\mu_k) \uparrow \infty$ as $N_{c,k} \uparrow \infty$. Substituting the lower bound on $N_{c,k}$ from (6.22) in (6.21), we have that

$$P_e(\mu_k) \leq \frac{r}{\lambda} \left(e^{E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|} \right)^{-\frac{2}{3} \left(D(\mu_k) - \frac{\sigma^2}{2\lambda(r-\lambda)} + \frac{1}{2} \right)}.$$

Hence $\lim_{k \rightarrow \infty} \frac{-\ln P_e(\mu_k)}{D(\mu_k)} = \frac{2}{3} (E_0(\rho, \mathcal{Q}) - \rho r \ln |\mathcal{M}|)$, and $P_e(\mu_k) \downarrow 0$. \square

Remark 6.3.10. We note that Proposition 6.3.9 provides a lower bound on the exponential decay rate of $\lim_{k \rightarrow \infty} \frac{-\ln(P_{e,f}^*(D_c))}{D_c}$, since for any $D_{c,k} \uparrow \infty$, we have a subsequence of μ_k , such that $D(\mu_k) \leq D_{c,k}$ and therefore $P_{e,f}^*(D_{c,k}) \leq P_e(\mu_k)$. We note that the capacity (C) of the discrete memoryless channel is given by $\frac{dE_0(\rho, \mathcal{Q}^*)}{d\rho} |_{\rho=0}$. Let us consider the case when $C - \delta \leq \lambda \ln |\mathcal{M}| < C$, where δ is a small positive constant. In the following, we show that the lower bound to the exponential decay rate of $P_{e,f}^*(D_c)$ achieved by the sequence of policies μ_k approximately matches with the upper bound on the exponential decay rate of $P_{e,f}^*(D_c)$, $\frac{2}{3} (E_0(\rho^*, \mathcal{Q}^*) - \rho \lambda \ln |\mathcal{M}|)$, obtained

in Proposition 6.3.6. We note that if $C - \delta \leq \lambda \ln |\mathcal{M}| < C$, for small positive δ , then any ρ satisfying $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho, \mathcal{Q}^*)}{\rho}$ is approximately zero. Furthermore, ρ^* is such that $\lambda \ln |\mathcal{M}| < \frac{E_0(\rho^*, \mathcal{Q}^*)}{\rho^*}$ and is therefore approximately zero. Then, from (6.20), since $\mathbb{E}_{\pi_{\mu_k}} e^{\rho^* S \ln |\mathcal{M}|} \approx e^{\rho^* \lambda N_{c,k} \ln |\mathcal{M}|}$, we have that $P_e(\mu_k) \lesssim \frac{r}{\lambda} (e^{E_0(\rho^*, \mathcal{Q}^*) - \rho^* \lambda \ln |\mathcal{M}|})^{-N_{c,k}}$, which yields an approximate lower bound on the exponential decay rate $\lim_{k \rightarrow \infty} \frac{-\ln P_e(\mu_k)}{D(\mu_k)} \approx \frac{2}{3} (E_0(\rho^*, \mathcal{Q}^*) - \rho^* \lambda \ln |\mathcal{M}|)$, which matches with the upper bound in Lemma 6.3.6.

Recall that $P_{N_c}(\lambda) = \inf_{D_c} P_{e, N_c}^*(D_c)$ for both R-model-A and I-model-A, where we have restricted to the set of policies Γ_{s, N_c} for both models. We have that Proposition 6.3.6 also holds for I-model-A. The proof of Proposition 6.3.6 holds for I-model-A as: (a) $P_{e, f}^*(D_c) \downarrow 0$ as $D_c \uparrow \infty$, from Proposition 6.3.9, and (b) $P_e(\gamma_k) \geq P_{N_{c,k}}(\lambda)$, which holds since $P_{N_c}(\lambda)$ for R-model-A is a lower bound to $P_{N_c}(\lambda)$ defined for I-model-A.

We also note that the upper bound $D_u(P_{e,c})$ holds for I-model-A under a EXH policy which serves only integer number of message symbols. Hence Proposition 6.3.7 and therefore Proposition 6.3.8 also holds for I-model-A.

Remark 6.3.11. In the next section, we present an asymptotic analysis for R-model-B. We recall that for R-model-B, the codeword length N_c is fixed. The asymptotic analysis of R-model-B is significant, since the tradeoff problem for R-model-B is a subproblem for R-model-A with the restriction to policies in $\Gamma_{s, f}$. In fact, we study how $D_{N_c}^*(P_{e,c})$ behaves as $P_{e,c} \downarrow P_{N_c}(\lambda)$.

6.4 Asymptotic analysis for R-model-B

6.4.1 Problem Statement

We state the tradeoff problem for R-model-B so that it is similar to the definition of the problem TRADEOFF in Chapter 4. The TRADEOFF problem for R-model-B is

$$\begin{aligned} & \text{minimize}_{\gamma \in \Gamma_s} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E} h(Q, N_c)}{\lambda N_c} \\ & \text{such that} \quad \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{\mathcal{S}|Q} c_s(S(Q), N_c)}{\lambda N_c} \leq P_{e,c}, \end{aligned}$$

where $h(q, N_c) = qN_c + \frac{\lambda N_c(N_c-1)}{2}$ as defined before, and $P_{e,c}$ is a constraint on the average error rate. As in Section 4.2.1, we can show that there exists a set \mathcal{O}^u of $P_{e,c}$ such that there exists a stationary deterministic optimal policy for $P_{e,c} \in \mathcal{O}^u$. This stationary deterministic policy is optimal for an unconstrained MDP with single stage cost $h(q, N_c) + \beta_{P_{e,c}} c_s(s, N_c)$, where $\beta_{P_{e,c}} \geq 0$ is a Lagrange multiplier.

However, we consider the above problem only for a subset $\Gamma_a \subset \Gamma_s$, which is the set of monotone

admissible policies. We redefine the set of admissible policies for R-model-B as follows. A policy $\gamma \in \Gamma_a$ if:

RG1 : $\gamma \in \Gamma_s$,

RG2 : it induces an aperiodic, irreducible Harris Markov chain $Q[m]$,

RG3 : the average service rate at a queue length q , $\mathbb{E}S(q)$ is non-decreasing in q .

We note the above properties are similar to those defined in Section 4.4 of Chapter 4.

We note that the function $c_s(s, N_c)$ is not convex. Consider a $P_{e,c} \in \mathcal{O}^u$ and the corresponding unconstrained MDP with Lagrange multiplier $\beta_{P_{e,c}}$. Existing proofs of the monotonicity property of the optimal policy for the unconstrained MDP, such as those in Goyal et al. [29], require that the function $c_s(s, N_c)$ be convex. We are not able to prove that the batch size is a monotonically non-decreasing function of the queue length for the optimal policy even if $P_{e,c} \in \mathcal{O}^u$, unlike in Chapter 4. However, we observe that the optimal policy prescribes a batch size which is monotonically increasing in the queue length in numerical solutions of the MDP. This is the only motivation for assuming RG3.

The TRADEOFF problem for R-model-B is to obtain $D_{N_c}^*(P_{e,c})$ which is the optimal value of

$$\begin{aligned} & \text{minimize}_{\gamma \in \Gamma_a} \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}h(Q, N_c)}{\lambda N_c} \\ & \text{such that} \quad \frac{\mathbb{E}_{\pi_\gamma} \mathbb{E}_{S|Q} c_s(S(Q), N_c)}{\lambda N_c} \leq P_{e,c}. \end{aligned}$$

We note that for every $P_{e,c}$ such that the above problem is feasible, for every $\epsilon > 0$, by definition there is a feasible admissible policy γ such that $D(\gamma) \leq D_{N_c}^*(P_{e,c}) + \epsilon$. We call such an admissible policy ϵ -optimal for $P_{e,c}$.

For R-model-B, suppose that $S_{max} > \frac{N_c E_0(\rho, \mathcal{Q})}{\rho \ln |\mathcal{M}|}$. As in Chapter 4, we define $c(s) : [0, S_{max}] \rightarrow \mathbb{R}_+$ as the lower convex envelope of $\{(s, c_s(s, N_c)), s \in [0, S_{max}]\}$. Then there exists a $s' < S_{max}$, which satisfies the following equation:

$$S_{max} e^{N_c(E_0(\rho, \mathcal{Q}))} = e^{\rho s' \ln |\mathcal{M}|} [S_{max} + \rho \ln |\mathcal{M}| s' (S_{max} - s')].$$

We note that the tangent drawn from $(S_{max}, c_s(S_{max}, N_c) = S_{max})$ touches the $c_s(s, N_c)$ curve at $(s', c_s(s', N_c))$. We also note that the definition of s' is similar to that for R-model-A, except that for R-model-B, the slope of $c_s(s, N_c)$ at $s = s'$ is not one. Furthermore s' for R-model-B is always greater than or equal to s' for R-model-A. Figure 6.2 shows an example. We note that the lower convex envelope $c(s) = c_s(s, N_c)$, for $s \in [0, s']$ and for $s \in (s', S_{max})$, $c(s)$ is the tangent line segment drawn from $(s', c_s(s', N_c))$ to (S_{max}, S_{max}) .




Figure 6.2: Illustration of the error cost function $c_s(s, N_c)$ and the lower convex envelope $c(s)$. The lower convex envelope coincides with $c_s(s, N_c)$ for all $s \in [0, s']$.

For R-model-B, the infimum of all achievable average probabilities of error, $\inf_{\gamma \in \Gamma_a} P_e(\gamma)$, is again denoted as $P_{N_c}(\lambda)$. Similar to the proof of Lemma 6.3.1, we can show that $P_{N_c}(\lambda) = \frac{c(\lambda)}{\lambda N_c}$. In the following, we obtain an asymptotic lower bound to $D_{N_c}^*(P_{e,c})$ as $P_{e,c} \downarrow P_{N_c}(\lambda)$.

6.4.2 Asymptotic analysis

We assume that the following properties hold for $A[1]$.

RA1 : $Pr \{A[1] - S_{max} > \delta_a\} > \epsilon_a$,

RA2 : $Pr \{A[1] \leq \frac{\Delta_a}{2}\} = \epsilon'_a > 0$, for some $0 < \Delta_a < s'$.

We note that RA1 is the same as that defined in Section 4.4 of Chapter 4. We also note that the property RA2, is similar to the property A2 assumed in Section 4.1.2 of Chapter 4, in that there is a positive probability of the number of arrivals being in an interval including zero.

We note that the service cost function $c(s)$, as defined above, satisfies the following properties, which are similar to the properties RC1 and RC2 defined in Section 4.4 of Chapter 4.

RC1 : $c(0) = 0$, and,

RC2 : $c(s)$ is strictly convex for $s \in [0, s']$ and linear for $s \in [s', S_{max}]$.

The asymptotic behaviour of $D_{N_c}^*(P_{e,c})$ as $P_{e,c} \downarrow P_{N_c}(\lambda)$ is different depending on whether $\lambda \leq \frac{s'}{N_c}$ or $\lambda > \frac{s'}{N_c}$. We characterize the asymptotic behaviour in both of these cases in the following.

Lemma 6.4.1. For $\lambda \leq \frac{s'}{N_c}$, and for any sequence of admissible policies γ_k with $P_e(\gamma_k) - P_{N_c}(\lambda) = V_k \downarrow 0$ we have that

$$\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{V_k}}\right).$$

Proof. For an admissible policy γ , define $\bar{C}(\gamma) = \mathbb{E}_{\pi_\gamma} c(S(Q))$. We note that as $c_s(s, N_c) \geq c(s)$, $P_e(\gamma_k) = \frac{\mathbb{E}_{\pi_\gamma} c_s(S(Q), N_c)}{\lambda N_c} \geq \frac{\bar{C}(\gamma_k)}{\lambda N_c}$. Then we have that there exists a sequence U_k such that $U_k = \bar{C}(\gamma_k) - c(\lambda) \leq \lambda N_c (P_e(\gamma_k) - P_{N_c}(\lambda)) = \lambda N_c V_k \downarrow 0$. Let us consider a particular policy γ in the sequence γ_k with $U_k = U$ and $\pi_{\gamma_k} = \pi$. To obtain a lower bound on $\bar{Q}(\gamma)$ we proceed, as in the proof of Proposition 4.4.2, with $q_d = \sup \{q : \mathbb{E}S(q) \leq \lambda N_c + \epsilon_U\}$ where ϵ_U will be chosen in the following. We follow all the steps in the proof of Proposition 4.4.2, with q_d as above and ϵ_U in place of ϵ_V , up to the step where $\bar{Q}(\gamma)$ is bounded below by $\frac{k_2 \Delta}{2}$, where the k_2 is the largest integer satisfying the inequality (4.27). That is, we have that $\bar{Q}(\gamma) \geq \frac{k_2 \Delta}{2}$, where k_2 is the largest integer, such that

$$\left(1 + \frac{\epsilon_U}{\delta \epsilon_a}\right)^{k_2} \leq \frac{2}{1 + 2D_t},$$

where $D_t = \frac{1}{\epsilon_U} \int_{q_d}^{\infty} (\mathbb{E}S(q) - \lambda N_c) d\pi(q)$. We note that in Proposition 4.4.2, a lower bound on k_2 was obtained via the upper bound $\frac{1}{4}$ on D_t . We note that, unlike in Proposition 4.4.2, the function $c(s)$ is not strictly convex for $s \in [0, S_{max}]$. So we derive an upper bound on D_t in a slightly different way. Define $q_\lambda = \sup \{q : \mathbb{E}S(q) \leq \lambda N_c\}$. Let $l(s)$ be the tangent line to the curve $c(s)$ at $(\lambda, c(\lambda))$. We have that $\mathbb{E}_\pi \{c(S(Q)) - l(S(Q))\} = U$. That is,

$$\int_0^{\infty} \mathbb{E} \{c(S(q)) - l(S(q))\} d\pi(q) = U.$$

As $c(s) \geq l(s)$ we have that

$$\int_0^{q_\lambda} \mathbb{E} \{c(S(q)) - l(S(q))\} d\pi(q) \leq U.$$

As $c(s)$ is convex and $l(s)$ is linear we have that

$$\int_0^{q_\lambda} \{c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))\} d\pi(q) \leq U.$$

We note that for $\lambda \leq \frac{s'}{N_c}$, there exists an $a_1 > 0$ such that $c(s) - l(s) \geq a_1(s - \lambda N_c)^2$ for $s \leq \lambda N_c$. Furthermore as for $q \leq q_\lambda$, $\mathbb{E}S(q) \leq \lambda N_c$, we have that

$$\int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c)^2 d\pi(q) \leq \frac{U}{a_1},$$

which can be written as

$$\int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c)^2 d\pi(q) + \int_{q_\lambda}^{\infty} 0 d\pi(q) \leq \frac{U}{a_1}.$$

By Jensen's inequality we then have that

$$\left(\int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) + \int_{q_\lambda}^\infty 0 d\pi(q) \right)^2 \leq \frac{U}{a_1}.$$

Hence we obtain that

$$\int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) \geq -\sqrt{\frac{U}{a_1}}. \quad (6.23)$$

Now we note that for an admissible policy γ ,

$$\begin{aligned} \int_0^\infty (\mathbb{E}S(q) - \lambda N_c) d\pi(q) &= 0, \text{ or,} \\ \int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) + \int_{q_\lambda}^{q_d} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) + \int_{q_d}^\infty (\mathbb{E}S(q) - \lambda N_c) d\pi(q) &= 0. \end{aligned}$$

For $q > q_\lambda$, $\mathbb{E}S(q) > \lambda N_c$, so that $\int_{q_\lambda}^{q_d} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) > 0$, which implies that

$$\begin{aligned} \int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) + \int_{q_d}^\infty (\mathbb{E}S(q) - \lambda N_c) d\pi(q) &\leq 0, \\ \int_{q_d}^\infty (\mathbb{E}S(q) - \lambda N_c) d\pi(q) &\leq -\int_0^{q_\lambda} (\mathbb{E}S(q) - \lambda N_c) d\pi(q) \leq \sqrt{\frac{U}{a_1}}, \end{aligned}$$

from (6.23). Therefore we obtain that $D_t \leq \frac{1}{\epsilon_U} \sqrt{\frac{U}{a_1}}$. We choose $\epsilon_U = 4\sqrt{\frac{U}{a_1}}$, and proceed as in the proof of Proposition 4.4.2 to obtain that $\bar{Q}(\gamma) \geq \frac{\Delta}{2} \left(\log_{\left(1 + \frac{\epsilon_U}{\delta \epsilon_a}\right)} \left(\frac{4}{3}\right) - 1 \right)$. Therefore, for the sequence of policies γ_k , we have that $\bar{Q}(\gamma_k) = \Omega\left(\frac{1}{\sqrt{U_k}}\right) = \Omega\left(\frac{1}{\sqrt{V_k}}\right)$ as $U_k \leq \lambda N_c V_k$. \square

Remark 6.4.2. We note that if S_{max} is such that $S_{max} \leq \frac{N_c E_0(\rho, \mathcal{Q})}{\rho \ln |\mathcal{M}|}$, then for all $\lambda < \frac{S_{max}}{N_c}$, the above asymptotic lower bound would hold.

Proposition 6.4.3. For $\lambda \leq \frac{s'}{N_c}$, as $P_{e,c} \downarrow P_{N_c}(\lambda)$, we have that

$$D^*(P_{e,c}) = \Omega\left(\frac{1}{\sqrt{P_{e,c} - P_{N_c}(\lambda)}}\right).$$

Proof. Consider the sequence of policies γ_k^* which are ϵ -optimal for TRADEOFF for a sequence p_k of $P_{e,c}$ such that $p_k \downarrow P_{N_c}(\lambda)$ for some $\epsilon > 0$. Since $P_e(\gamma_k^*) \leq p_k$, we have that from Lemma 6.4.1 that $\bar{Q}(\gamma_k^*) = \Omega\left(\frac{1}{\sqrt{p_k - P_{N_c}(\lambda)}}\right)$. Since $D^*(p_k) \geq \frac{1}{\lambda N_c} [\bar{Q}(\gamma_k^*) + \frac{\lambda}{2} N_c (N_c - 1)] - \epsilon$, we obtain that $D^*(p_k) = \Omega\left(\frac{1}{\sqrt{p_k - P_{N_c}(\lambda)}}\right)$. \square

Remark 6.4.4. We note that an asymptotic upper bound, which is tight upto a logarithmic factor, can be obtained for the above case, using a sequence of policies as in Lemma 4.4.3. We note that

for this sequence of policies, it is possible to choose batch sizes s such that $c_s(s, N_c) = c(s)$.

In the following, we obtain an asymptotic lower bound to $D_{N_c}^*(P_{e,c})$ as $P_{e,c} \downarrow P_{N_c}(\lambda)$ for $\lambda > \frac{s'}{N_c}$. This lower bound is obtained by extending Lemma 4.3.8 to the case when the state space of the Markov chain is the set of non-negative real numbers. As in Lemma 4.3.5 we first obtain a lower bound on the queue length as a function of the stationary probability of the queue length being in a certain set. Then we relate the stationary probability of the queue length being in the above set to the average error rate.

Let $q_{s'} \triangleq \inf \{q : \mathbb{E}S(q) \geq s' - \epsilon\}$, where ϵ is a small positive constant, chosen such that $\Delta_a < s' - \epsilon$, where Δ_a is as in RA2. We note that for an admissible policy γ , $\mathbb{E}_\pi \mathbb{E}S(Q) = \lambda N_c$, and therefore there would exist finite q , for which $\mathbb{E}S(q) \geq \lambda N_c$. Thus, $q_{s'}$ is finite. The following lemma shows that for a queue length $q \geq q_{s'}$, there is a positive minimum probability of serving at least a certain number of customers.

Lemma 6.4.5. For an admissible policy γ , with $q_{s'}$ defined as above, for Δ_a as in RA2, we have that $\inf_{q \geq q_{s'}} Pr \{S(q) > \Delta_a\} \geq \delta_s > 0$, where $\delta_s = \frac{s' - \epsilon - \Delta_a}{S_{max} - \Delta_a}$.

The proof of this lemma is similar to that of Lemma 5.7.1 and is presented in Appendix 6.C.

Lemma 6.4.6. For $\lambda > \frac{s'}{N_c}$, and for any sequence of admissible policies γ_k with $P_e(\gamma_k) - P_{N_c}(\lambda) = V_k \downarrow 0$, we have that

$$\bar{Q}(\gamma_k) = \Omega \left(\log \left(\frac{1}{V_k} \right) \right).$$

Proof. For an admissible policy γ , define $\bar{C}(\gamma) = \mathbb{E}_{\pi_\gamma} c(S(Q))$. We note that as $c_s(s, N_c) \geq c(s)$, we have that there exists a sequence U_k such that $U_k = \bar{C}(\gamma_k) - c(\lambda) \leq \lambda N_c (P_e(\gamma_k) - P_{N_c}(\lambda)) = \lambda N_c V_k \downarrow 0$. We consider a particular policy γ in the sequence γ_k with $U_k = U$ and $\pi_{\gamma_k} = \pi$. As in the proof of Lemma 5.7.2, since γ is admissible, we have that if m_1 is the largest integer such that

$$\pi [0, q_{s'}) + \pi [0, q_{s'}) \left[\left(1 + \frac{1}{\rho} \right)^{m_1} - 1 \right] = \pi [0, q_{s'}) \left(1 + \frac{1}{\rho} \right)^{m_1} \leq \frac{1}{2}, \quad (6.24)$$

then $\bar{Q}(\gamma) \geq \frac{m_1 \Delta_a}{4}$. Let us define the line $l(s)$ as the line passing through $(s', c(s'))$ and $(S_{max}, c(S_{max}))$.

For the policy γ we have that

$$\int_0^\infty \mathbb{E} \{c(S(q)) - l(S(q))\} d\pi(q) = U.$$

As $c(s)$ is convex and $l(s)$ is linear we have that

$$\int_0^\infty \{c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))\} d\pi(q) \leq U.$$

Also as $c(s) \geq l(s)$,

$$\int_0^{q_{s'}} \{c(\mathbb{E}S(q)) - l(\mathbb{E}S(q))\} d\pi(q) \leq U.$$

We note that for $q < q_{s'}$ we have that $\mathbb{E}S(q) < s' - \epsilon$ and there exists $a_1 > 0$ such that $c(\mathbb{E}S(q)) - l(\mathbb{E}S(q)) \geq a_1(\mathbb{E}S(q) - s')^2$. Hence we obtain that

$$\int_0^{q_{s'}} (\mathbb{E}S(q) - s')^2 \leq \frac{U}{a_1},$$

and as $q < q_{s'}$, $(\mathbb{E}S(q) - s')^2 > \epsilon^2$. Therefore we obtain that

$$\int_0^{q_{s'}} \epsilon^2 d\pi(q) \leq \frac{U}{a_1}.$$

And hence $\pi[0, q_{s'}] \leq \frac{U}{a_1 \epsilon^2}$.

From (6.24), if m_2 is the largest integer such that

$$\begin{aligned} \left(1 + \frac{1}{\rho}\right)^{m_2} &\leq \frac{a_1 \epsilon^2}{2U}, \text{ or,} \\ m_2 &\leq \log_{\left(1 + \frac{1}{\rho}\right)} \left(\frac{a_1 \epsilon^2}{2U}\right), \end{aligned}$$

then $m_2 \leq m_1$. We note that m_2 is at least

$$\left\lceil \log_{\left(1 + \frac{1}{\epsilon}\right)} \left(\frac{a_1 \epsilon^2}{2U}\right) \right\rceil.$$

Since $\bar{Q}(\gamma) \geq \frac{m\Delta_a}{4} \geq \frac{m_1\Delta_a}{4} \geq \frac{m_2\Delta_a}{4}$, we obtain that

$$\bar{Q}(\gamma) \geq \frac{\Delta_a}{4} \left(\log_{\left(1 + \frac{1}{\epsilon}\right)} \left(\frac{a_1 \epsilon^2}{2U}\right) - 1 \right).$$

So for the sequence of policies γ_k with $U_k \downarrow 0$ we have that $\bar{Q}(\gamma_k) = \Omega\left(\log\left(\frac{1}{U_k}\right)\right) = \Omega\left(\log\left(\frac{1}{V_k}\right)\right)$. \square

Proposition 6.4.7. For $\lambda > \frac{s'}{N_c}$, as $P_{e,c} \downarrow P_{N_c}(\lambda)$, we have that

$$D^*(P_{e,c}) = \Omega\left(\log\left(\frac{1}{P_{e,c} - P_{N_c}(\lambda)}\right)\right).$$

The proof of the above result is similar to that of Proposition 6.4.3.

Remark 6.4.8. An asymptotic upper bound to $D_{N_c}^*(P_{e,c})$, as $P_{e,c} \downarrow P_{N_c}(\lambda)$ can be obtained from Lemma 4.3.10. We note that in this case, $s_l = s'$ and $s_u = S_{max}$. Then, as in Lemma 4.3.10 we

can show that there exists a sequence of admissible policies γ_k , for which $D^*(P_{e,c}) = \mathcal{O}\left(\log\left(\frac{1}{V_k}\right)\right)$ and $P_e(\gamma_k) - P_{N_c}(\lambda) = V_k \downarrow 0$. We note that $P_e(\gamma_k)$ can be computed as in Lemma 4.3.10 since $c_s(s', N_c) = c(s')$ and $c_s(S_{max}, N_c) = c(S_{max})$.

Remark 6.4.9. We can set up a queueing model I-model-B, which is similar to R-model-B, except that $q_0 \in \mathbb{Z}_+$, $A[m] \in \{0, 1, \dots, A_{max} \in \mathbb{Z}_+\}$, and $S[m] \in \{0, 1, \dots, S_{max}\}$, $m \geq 1$. Therefore, for I-model-B the queue length evolution is on \mathbb{Z}_+ . We note that queueing model I-model-B is analogous to I-model-A. Asymptotic lower bounds to $D_{N_c}^*(P_{e,c})$ can be obtained as in Section 4.3.3 of Chapter 4, with $c(s)$ defined as the piecewise linear lower convex envelope of $\{(s, c_s(s, N_c)), s \in \{0, \dots, S_{max}\}\}$. Asymptotic upper bounds to $D_{N_c}^*(P_{e,c})$ can be obtained as in Remark 6.4.8.

Remark 6.4.10. For R-model-B, we note that asymptotic lower bounds can be derived for other approximations for $c_s(s, N_c)$, if the lower convex envelopes for such approximations have the same form as $c(s)$ above. We note that the asymptotic nature of the bounds only depended on: (a) $c(s)$ being strictly convex in $[0, s']$, and (b) $c(s)$ being linear in $(s', S_{max}]$. In Appendix 6.E we consider some examples for the approximation $\tilde{c}_s(s, N_c)$ instead of $c_s(s, N_c)$ for the error cost, where

$$\tilde{c}_s(s, N_c) = \min_{\rho \in [0,1]} s e^{-N_c E_0(\rho, Q) + \rho s \ln |\mathcal{M}|},$$

and illustrate that the lower convex envelope $\tilde{c}(s)$ of $\tilde{c}_s(s, N_c)$ as a function of s has a similar form as $c(s)$.

We note that instead of using Gallager's random coding upper bound we could also use Polyanskiy's normal approximation for the codeword error probability to derive an approximation for $c_s(s, N_c)$. From [49], if $0 \leq s \leq \frac{CN_c}{\log_2 |\mathcal{M}|}$, then we have that the codeword error probability $P_{e,b}$ satisfies the following approximation:

$$N_c \approx \left(\frac{\mathbb{Q}^{-1}(P_{e,b})}{1 - \frac{s \log_2 |\mathcal{M}|}{N_c C}} \right)^2 \frac{V}{C^2},$$

where \mathbb{Q} is the Gaussian Q function, C is the channel capacity (in bits/channel use), and V is the channel dispersion. Then, we have that

$$P_{e,b} \approx \mathbb{Q} \left(\sqrt{\frac{N_c C^2}{V}} \left(1 - \frac{s \log_2 |\mathcal{M}|}{N_c C} \right) \right),$$

and the approximation

$$\tilde{c}_s(s, N_c) = s \mathbb{Q} \left(\sqrt{\frac{N_c C^2}{V}} \left(1 - \frac{s \log_2 |\mathcal{M}|}{N_c C} \right) \right).$$

We note that if $s > \frac{CN_c}{\log_2 |\mathcal{M}|}$, then we have the approximation $\tilde{c}_s(s, N_c) = s$.

It can be shown that if $\frac{S_{max} \log_2 |\mathcal{M}|}{N_c} \leq C$, then $\tilde{c}_s(s, N_c)$ is strictly convex in s . Then, as in Proposition 6.4.3, we have that for any $\lambda < \frac{C}{N_c}$, as $P_{e,c} \downarrow P_{N_c}(\lambda)$, we have that $D^*(P_{e,c}) = \Omega\left(\frac{1}{\sqrt{P_{e,c} - P_{N_c}(\lambda)}}\right)$.

6.5 Conclusions

We have shown that for R-model-A (from Proposition 6.3.6) as well as I-model-A (the discussion in Section 6.3.3), the exponential decay of average error rate with average delay is at most two-thirds of the Gallager random coding exponent, when we restrict attention to the set $\Gamma_{s,f}$ of policies. We also note that for I-model-A, for information arrival rate $\lambda \ln |\mathcal{M}|$ approaching the capacity of the channel, a sequence of policies $\in \Gamma_{s,f}$, that uses a fixed service rate, approximately achieves the best exponential decay rate of two-thirds of the Gallager random coding exponent. We note that the $\frac{2}{3}$ factor arises because of the fundamental limitation of block codes; message symbols arriving in a transmission period of duration N_c have an average waiting period of $\frac{N_c}{2}$ until the succeeding transmission period and have to wait for at least an additional N_c slots before leaving the queue. Therefore, the average delay is at least $\frac{3}{2}N_c$. So naturally the question arises whether *streaming codes* have a better exponential decay rate. In [57], we show that a sequence of fixed rate *randomly time varying* streaming codes, with increasing *constraint lengths* achieve an exponential decay rate of average error rate with average delay which is equal to the Gallager random coding exponent.

We then considered the exponential decay rate for a sequence of exhaustive service policies, which ensure a constant block error probability per transmission by varying the codeword length. We obtain that for the above sequence of policies, the exponential decay rate is at least two-thirds of the Gallager random coding exponent. The performance of such policies has been studied in detail in [53], [54], [55], [59], and [60]. We note that the above exponential decay rate provides a lower bound to the exponential decay rate achievable by a sequence of policies in Γ_s . However, we do not have an upper bound on the exponential decay rate over any sequence of policies in Γ_s .

The analysis of R-model-B illustrates the application of the lower bounding technique in Chapter 4 to cases where the service cost function is not convex. We observe that the asymptotic lower bound depends on the nature of the lower convex envelope of the service cost function at λ . Since the lower convex envelope $c(s)$ is convex, the analysis and observations from Chapter 4 apply. We note that the form of $c(s)$ considered in this chapter is not strictly convex or piecewise linear, but piecewise convex. The asymptotic upper bounds show that a sequence of admissible policies using only the service rates s for which the service cost $c_s(s, N_c) = c(s)$ is order optimal (only for the $\log(\frac{1}{V})$ case). We note that this is reminiscent of Crbill's exclusion principle⁴ for the control of M/M/1

⁴which states that for a M/M/1 queue with controllable service rates, the stationary optimal policy that minimizes the time average of the single stage cost $Q(t) + \beta c_s(\mu(Q(t)))$ ($\beta > 0$), uses only service rates which are such that $c_s(s) = c(s)$, where $c(\cdot)$ is the lower convex envelope of the service cost function $c_s(\cdot)$.

queues with non-convex service costs [25], but optimality in our case is only in the asymptotic order sense. From this analysis, we can conclude that an asymptotic characterization of the minimum average queue length in the asymptotic regime \mathfrak{R} for admissible policies, can be obtained from the techniques in Chapter 4 by considering the lower convex envelope of the service cost function, even if the service cost function is not convex.

Appendices

6.A Optimization problem (6.11)

Consider the optimization problem (6.11) :

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && \frac{1}{\lambda N_c} \mathbb{E}_{\pi} \left[S \min \left(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right) \right] \\ & \text{such that} && \mathbb{E}_{\pi} S \geq \lambda N_c + \epsilon, \end{aligned} \quad (6.25)$$

where π is any distribution for S and $\epsilon \geq 0$. We denote the optimal value of the above problem by $p_e(\epsilon)$. For any distribution π , we have that the point

$$\left(\mathbb{E}_{\pi} S, \frac{1}{\lambda N_c} \mathbb{E}_{\pi} \left[S \min \left(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right) \right] \right),$$


lies in the convex hull of the set of points

$$\left(s, \frac{1}{\lambda N_c} \left[s \min \left(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|} \right) \right] \right), s \in \mathbb{R}_+.$$

With the constraint $\mathbb{E}_{\pi} S \geq \lambda N_c + \epsilon$, it is clear that there would exist some distribution π' such that the point $(\mathbb{E}_{\pi'} S, \frac{1}{\lambda N_c} \mathbb{E}_{\pi'} [S \min(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|})])$, with $\mathbb{E}_{\pi'} S = \lambda N_c + \epsilon$, lies on the lower convex envelope of $\frac{c_s(s, N_c)}{\lambda N_c}$ (the curve AC, as shown in Figure 6.3(a)). Therefore, π' is optimal, and we have that

$$p_e(\epsilon) = \begin{cases} \frac{\lambda N_c + \epsilon}{\lambda N_c} e^{-N_c E_0(\rho, \mathcal{Q}) + \rho(\lambda N_c + \epsilon) \ln |\mathcal{M}|} & \text{if } 0 \leq (\lambda N_c + \epsilon) \leq s', \\ \frac{1}{\lambda N_c} \left(((\lambda N_c + \epsilon) - s') + s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|} \right) & \text{otherwise,} \end{cases} \quad (6.26)$$

where s' is such that $\frac{d\{s e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s \ln |\mathcal{M}|}\}}{ds} \Big|_{s=s'} = 1$.



(a) Illustration of the convex hull (the region between and including the curves AB and AC) and the optimal values $p_e(\epsilon_1)$ and $p_e(\epsilon_2)$. $Pr\{S = \lambda N_c + \epsilon_1\} = 1$, and (ii) for ϵ_2 , the distribution puts mass at s' and s_1 .

Now we obtain distributions which are (δ, ϵ) -optimal, where $\delta > 0$. A distribution d is said to be (δ, ϵ) -optimal for (6.25) if

$$\begin{aligned}\mathbb{E}_d S &= \lambda N_c + \epsilon, \text{ and,} \\ \frac{1}{\lambda N_c} \mathbb{E}_d \left[S \min \left(1, e^{-N_c E_0(\rho, \mathcal{Q}) + \rho S \ln |\mathcal{M}|} \right) \right] &= p_e(\epsilon) + \delta.\end{aligned}$$

We note that if : (i) $\lambda N_c + \epsilon \leq s'$ then the distribution that gives probability 1 to $\lambda N_c + \epsilon$ is (δ, ϵ) -optimal $\forall \delta \geq 0$, (ii) $\lambda N_c + \epsilon > s'$ and $\delta \geq \delta_b = s' - s' e^{-N_c E_0(\rho, \mathcal{Q}) + \rho s' \ln |\mathcal{M}|}$ then a distribution that gives probability one to $\lambda N_c + \epsilon$ is δ -optimal, and (iii) $\lambda N_c + \epsilon > s'$ and $\delta < \delta_b$, then a distribution that gives mass to two points s' and s_1 , characterized in the following lemma, is (δ, ϵ) -optimal.

Lemma 6.A.1. If $\lambda N_c + \epsilon > s'$ and $\delta < \delta_b$, then the distribution d defined as follows is (δ, ϵ) -optimal. Let $s_1 = \frac{s'}{\delta} [(\lambda N_c + \epsilon - s') + \delta] - \frac{c_s(s', N_c)(\lambda N_c + \epsilon - s')}{\delta}$.

$$\begin{aligned}Pr \{S = s'\} &= \frac{s_1 - (\lambda N_c + \epsilon)}{s_1 - s'}, \\ Pr \{S = s_1\} &= \frac{\lambda N_c + \epsilon - s'}{s_1 - s'}.\end{aligned}$$

We provide an outline of the proof. We note that if the distribution needs to achieve only $p_e(\epsilon) + \delta$, then we can choose the points $p_1 = (s_1, c_s(s_1, N_c))$ (as shown in Figure 6.3(b)) and $p_2 = (s', c_s(s', N_c))$. The point $(\lambda N_c + \epsilon, p_e(\epsilon) + \delta)$ lies on the line joining these two points and is therefore a convex combination of the points p_1 and p_2 . Solving for s_1 and the convex combination leads to the proof of the above lemma.

We note that the (δ, ϵ) -optimal distribution has finite support for any $\delta > 0$ and $\epsilon > 0$.

6.B Proof of Proposition 6.3.7

Proof. For the policy e , under stationary conditions, we have that $Q = s(Q)$ and $\tau(s) = \lceil as + b \rceil, \forall s$. Suppose $\mathbb{E}_{\pi_e} Q < \infty$ for e , then we have that $\mathbb{E}_{\pi_e} Q = \lambda \mathbb{E}_{\pi_e} \tau(Q)$. For brevity, in this proof we use $\mathbb{E}[\cdot]$ to denote $\mathbb{E}_{\pi_e}[\cdot]$. Since $as + b \leq \tau(s) \leq as + b + 1$, we obtain that $\mathbb{E}Q \geq \frac{b\lambda}{1-a\lambda}$ and $\mathbb{E}Q \leq \frac{(b+1)\lambda}{1-a\lambda}$. Also $\frac{b}{1-a\lambda} \leq \mathbb{E}\tau(Q) \leq \frac{b+1}{1-a\lambda}$. We also have that

$$\begin{aligned}\mathbb{E}Q^2 &= \sigma^2 \mathbb{E}\tau(Q) + \lambda^2 \mathbb{E}\tau(Q)^2, \\ &\leq \sigma^2 (a\mathbb{E}Q + b + 1) + \lambda^2 (a^2 \mathbb{E}Q^2 + (b+1)^2 + 2a(b+1)\mathbb{E}Q) \text{ or,} \\ \mathbb{E}Q^2 &\leq \frac{(b+1)\sigma^2 + (b+1)^2\lambda^2 + \frac{(b+1)\lambda}{1-a\lambda}(a\sigma^2 + 2a(b+1)\lambda^2)}{(1-a^2\lambda^2)}.\end{aligned}$$

For the policy e , $\mathbb{E}Q \leq \frac{\mathbb{E}Q(aQ+b+1) + \frac{\mathbb{E}(aQ+b)(aQ+b+1)\lambda}{2}}{\mathbb{E}\tau(Q)}$. Substituting the upper bounds for $\mathbb{E}Q^2$, and $\mathbb{E}Q$, and the lower bound for $\mathbb{E}\tau(Q)$ in the above expression leads to the upper bound on the

average queue length for the EXH policy. \square

6.C Proof of Lemma 6.4.5

Proof. We note that $q_{s'} = \inf \{q : \mathbb{E}S(q) \geq s' - \epsilon\}$, for ϵ such that $0 < \epsilon < s' - \Delta$, for Δ in RA2. We note that by definition, $\forall q \geq q_{s'}$,

$$\begin{aligned} \mathbb{E}S(q) &\geq s' - \epsilon, \\ \int_0^{S_{max}} S(q).dP(S(q)) &\geq s' - \epsilon. \end{aligned} \quad (6.27)$$

Then,

$$\begin{aligned} \int_0^{\Delta} \Delta dP(S(q)) + \int_{\Delta}^{S_{max}} S_{max} dPS(q) &\geq s' - \epsilon, \\ \Delta(1 - Pr\{S(q) > \Delta\}) + S_{max}Pr\{S(q) > \Delta\} &\geq s' - \epsilon, \text{ or,} \\ Pr\{S(q) > \Delta\} &\geq \frac{s' - \epsilon - \Delta}{S_{max} - \Delta}. \end{aligned}$$

Thus for any $q \geq q_{s'}$, $Pr\{S(q) > \Delta\} \geq \delta_s > 0$, where $\delta_s = \frac{s' - \epsilon - \Delta}{S_{max} - \Delta}$. \square

6.D Upper bound for average queue length for the policy μ in Proposition 6.3.4

For the policy μ with parameters N_c and r , we have that a batch of size $s(q) = \min(q, rN_c)$ is served when the queue length is q . We note that if the arrival distribution is such that $A_{max} < r$, so that $A[1] < rN_c$, then the average queue length $\mathbb{E}_{\mu}Q = \lambda N_c$.

The following upper bound, holds for any arrival distribution with $\mathbb{E}A^2 = \mathbb{E}(A[1]^2) < \infty$. From (6.1) we have that

$$\begin{aligned} \mathbb{E}\{Q[m+1]^2 - Q[m]^2 | Q[m] = q\} &= -2q(s(q) - \lambda N_c) + \mathbb{E}A^2 + s(q)^2 - 2s(q)\lambda N_c, \\ &= -2qN_c(r - \lambda) + 2q(rN_c - s(q)) + \mathbb{E}A^2 + s(q)^2 - 2s(q)\lambda N_c. \end{aligned}$$

Now, by taking expectations with respect to the stationary distribution π of μ , we have that

$$\begin{aligned} \mathbb{E}_{\pi}Q &= \frac{1}{2N_c(r - \lambda)} [2\mathbb{E}_{\pi}[Q(rN_c - s(Q))] + \mathbb{E}A^2 + \mathbb{E}_{\pi}s(Q)^2 - 2\lambda^2 N_c^2], \\ &= \frac{1}{2N_c(r - \lambda)} [\sigma^2 N_c + 2\mathbb{E}_{\pi}[Q(rN_c - s(Q))] + \mathbb{E}_{\pi}s(Q)^2 - \lambda^2 N_c^2]. \end{aligned} \quad (6.28)$$

We now simplify and provide an upper bound for $2\mathbb{E}_{\pi}[Q(rN_c - s(Q))] + \mathbb{E}_{\pi}s(Q)^2 - \lambda^2 N_c^2$.

We denote the stationary probability $Pr\{s(Q) = rN_c\} = Pr\{Q > rN_c\}$ by p_r . Then we have that $2\mathbb{E}_\pi [Q(rN_c - s(Q))] + \mathbb{E}_\pi s(Q)^2 - \lambda^2 N_c^2$

$$\begin{aligned}
&= 2 \int_{q < rN_c} q(rN_c - q) d\pi(q) + \int_{q < rN_c} q^2 d\pi(q) + p_r r^2 N_c^2 - \lambda^2 N_c^2, \\
&= 2 \int_{q < rN_c} qrN_c d\pi(q) - \int_{q < rN_c} q^2 d\pi(q) + p_r r^2 N_c^2 - \lambda^2 N_c^2, \\
&= 2rN_c [\lambda N_c - p_r r N_c] - \int_{q < rN_c} q^2 d\pi(q) + p_r r^2 N_c^2 - \lambda^2 N_c^2, \\
&= 2r\lambda N_c^2 - \int_{q < rN_c} q^2 d\pi(q) - p_r r^2 N_c^2 - \lambda^2 N_c^2, \\
&= \lambda N_c^2 (r - \lambda) + r\lambda N_c^2 - p_r r^2 N_c^2 - \int_{q < rN_c} q^2 d\pi(q). \tag{6.29}
\end{aligned}$$

Therefore,

$$2\mathbb{E}_\pi [Q(rN_c - s(Q))] + \mathbb{E}_\pi s(Q)^2 - \lambda^2 N_c^2 \leq \lambda N_c^2 (r - \lambda) + r\lambda N_c^2. \tag{6.30}$$

Substituting in (6.28), we obtain that

$$\mathbb{E}_\pi Q \leq \frac{\sigma^2}{2(r - \lambda)} + \frac{\lambda N_c}{2} + \frac{r\lambda N_c}{2(r - \lambda)}.$$

We now present another upper bound on $2\mathbb{E}_\pi [Q(rN_c - s(Q))] + \mathbb{E}_\pi s(Q)^2 - \lambda^2 N_c^2$, obtained by lower bounding p_r in (6.29). The lower bound on p_r is obtained under the assumption that $Pr\{A[1] \geq rN_c\} > 0$. Let $p_l = 1 - p_r$. We note that p_l is the fraction of time a batch size less than rN_c is used. We note that the evolution of $(Q[m])$ can be divided into cycles of random duration. Each cycle comprises of two periods, where each period is also of random duration. The first period starts in slot m , if $Q[m - 1] \geq rN_c$ and $Q[m] < rN_c$. The duration of the first period is distributed according to a Geometric distribution with mean $\frac{1}{Pr\{A[1] \geq rN_c\}}$. Following the first period, we have the second period which starts in a slot m such that $Q[m - 1] < rN_c$ and $Q[m] \geq rN_c$. We note that throughout the second period the queue length is greater than or equal to rN_c and a batch service of rN_c occurs in each slot.

Let Q_p be the random queue length at the start of the second period. We note that $Q_p \sim A[1]$ conditioned on $\{A[1] \geq rN_c\}$. Let $Q_p = q_p$. Since $(Q[m])$ is Markov, the duration of the second period $T_p(q_p)$ is then

$$T_p(q_p) = \min \left\{ t : q_p + \sum_{m=1}^t [A[m] - rN_c] < rN_c \right\}.$$

Then, applying Wald's lemma for q_p , and taking expectations over the distribution of Q_p , we have

that

$$\mathbb{E}T_p(Q_p) \geq \frac{\mathbb{E}_{[A[1]|A[1] \geq rN_c]}[Q_p - rN_c]}{rN_c - \lambda N_c}.$$

We note that

$$\begin{aligned} p_l &= \frac{\frac{1}{Pr\{A[1] \geq rN_c\}}}{\frac{1}{Pr\{A[1] \geq rN_c\}} + \mathbb{E}T_p(Q_p)}, \\ &\leq \frac{\frac{1}{Pr\{A[1] \geq rN_c\}}}{\frac{1}{Pr\{A[1] \geq rN_c\}} + \frac{\mathbb{E}_{[A[1]|A[1] \geq rN_c]}[Q_p - rN_c]}{rN_c - \lambda N_c}}. \end{aligned}$$

Since $\mathbb{E}_{[A[1]|A[1] \geq rN_c]}Q_p = \frac{\int_{a \geq rN_c} a dP_{A[1]}(a)}{Pr\{A[1] \geq rN_c\}}$, we have that

$$\begin{aligned} p_l &\leq \frac{1}{1 + \frac{\left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}{rN_c - \lambda N_c}}, \\ &= \frac{N_c(r - \lambda)}{N_c(r - \lambda) + \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}. \end{aligned}$$

Therefore,

$$p_r \geq \frac{\left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}{N_c(r - \lambda) + \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}.$$

Substituting the above lower bound in the RHS of (6.29) and using $\int_{q < rN_c} q^2 d\pi(q) \geq 0$, we have that $2\lambda N_c^2(r - \lambda) + \lambda N_c^2 \left(\lambda - \frac{p_r r^2}{\lambda} \right)$

$$\leq 2\lambda N_c^2(r - \lambda) + N_c^2 \left(\lambda^2 - \frac{\left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}{N_c(r - \lambda) + \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]} r^2 \right).$$

The above expression can then be simplified to

$$2\lambda N_c^2(r - \lambda) + (r - \lambda)N_c^2 \frac{\left(\lambda^2 N_c - \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right] (r + \lambda) \right)}{N_c(r - \lambda) + \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}. \quad (6.31)$$

Therefore, $2\mathbb{E}_\pi [Q(rN_c - s(Q))] + \mathbb{E}_\pi s(Q)^2 - \lambda^2 N_c^2$

$$\leq 2\lambda N_c^2(r - \lambda) + (r - \lambda)N_c^2 \frac{\left(\lambda^2 N_c - \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right] (r + \lambda) \right)}{N_c(r - \lambda) + \left[\int_{a \geq rN_c} (a - rN_c) dP_{A[1]}(a) \right]}. \quad (6.32)$$




Figure 6.3: Illustration of $\tilde{c}_s(s, N_c)$ and its lower convex envelope, for $N_c \in \{100, 150, 200\}$, for the parameters in Example 1

Substituting in (6.28), we obtain that

$$\mathbb{E}_\pi Q \leq \frac{\sigma^2}{2(r - \lambda)} + \lambda N_c + \frac{N_c}{2} \frac{\left(\lambda^2 N_c - \left[\int_{a \geq r N_c} (a - r N_c) dP_{A[1]}(a) \right] (r + \lambda) \right)}{N_c(r - \lambda) + \left[\int_{a \geq r N_c} (a - r N_c) dP_{A[1]}(a) \right]}.$$

6.E Examples for $\tilde{c}_s(s, N_c)$ in Remark 6.4.10

In this section we provide representative numerical examples which illustrate that the lower convex envelope $\tilde{c}(s)$ of $\tilde{c}_s(s, N_c)$ has a similar form as $c(s)$. The transition probability matrices and the input distribution Q have been generated randomly for these examples.

Example 1 : For this example we take $|\mathcal{M}| = 2$, and a DMC with $|\mathcal{X}| = 5$ and $|\mathcal{Y}| = 10$. The transition probability matrix $P_{Y|X}$ is :

$$\begin{bmatrix} 0.0459 & 0.2101 & 0.1339 & 0.1138 & 0.1094 & 0.1212 & 0.0348 & 0.1859 & 0.0309 & 0.0140 \\ 0.0108 & 0.2152 & 0.0646 & 0.0593 & 0.0210 & 0.0341 & 0.2328 & 0.1601 & 0.0612 & 0.1409 \\ 0.1251 & 0.1038 & 0.0403 & 0.1645 & 0.1146 & 0.1375 & 0.0872 & 0.0727 & 0.0657 & 0.0886 \\ 0.0728 & 0.1252 & 0.1291 & 0.1038 & 0.0062 & 0.1191 & 0.1417 & 0.1202 & 0.1211 & 0.0607 \\ 0.1000 & 0.1678 & 0.0060 & 0.1042 & 0.0149 & 0.1505 & 0.1352 & 0.0174 & 0.1673 & 0.1368 \end{bmatrix}$$

The input distribution Q is $(0.2018, 0.2551, 0.2515, 0.0943, 0.1972)$. Then the numerically obtained $\tilde{c}_s(s, N_c)$ for three values of N_c is shown in Figure 6.3.




Figure 6.4: Illustration of $\tilde{c}_s(s, N_c)$ and its lower convex envelope, for $N_c \in \{100, 150, 200\}$, for the parameters in Example 2

Example 2 : For this example we take $|\mathcal{M}| = 20$, and a DMC with $|\mathcal{X}| = 5$ and $|\mathcal{Y}| = 5$. The transition probability matrix $P_{Y|X}$ is :

$$\begin{bmatrix} 0.0649 & 0.2915 & 0.1642 & 0.1986 & 0.2808 \\ 0.1557 & 0.2722 & 0.0509 & 0.1063 & 0.4148 \\ 0.1883 & 0.1538 & 0.2409 & 0.1187 & 0.2984 \\ 0.4160 & 0.0809 & 0.0959 & 0.2615 & 0.1457 \\ 0.0860 & 0.2355 & 0.2115 & 0.1459 & 0.3212 \end{bmatrix}$$

The input distribution Q is $(0.0739, 0.0141, 0.2423, 0.3808, 0.2889)$. Then the numerically obtained $\tilde{c}_s(s, N_c)$ for three values of N_c is shown in Figure 6.4. We note that in all these cases the form of the lower convex envelope of $\tilde{c}_s(s, N_c)$ is the same as that illustrated in Figure 6.2.

Conclusions and Scope for Future Work

We first summarize the main results which are obtained in this thesis. Motivated by the already available results on the monotonicity property of any stationary deterministic optimal policy for the unconstrained tradeoff problem, in this thesis we consider the constrained tradeoff problem for the set of admissible policies, which are defined to be monotone.

In Chapters 2 and 3, for the state dependent M/M/1 model, using geometric bounds on the stationary probability distribution of the queue length for admissible policies, we obtain the asymptotic behaviour of the solution to the constrained tradeoff problem for admissible policies in the regime \mathfrak{R} , as well as asymptotic bounds on any sequence of order optimal admissible policies. We identify a case in which the average queue length grows only to a finite value in the regime \mathfrak{R} . For the cases for which the average queue length grows without bound, we show that the asymptotic behaviour of the average queue length is either $\Theta(\log(\frac{1}{V}))$, $\Theta(\frac{1}{V})$, or $\Theta(\frac{1}{\sqrt{V}})$. The asymptotic behaviour of the average queue length is determined by: (i) the nature of the service cost function, i.e., whether it is piecewise linear, a corner point, or strictly convex, at the value of λ or $u^{-1}(u_c)$, and (ii) the extent of freedom that we have in the control of the arrival rates ($\lambda(q)$) and service rates ($\mu(q)$).

Guided by the analysis in Chapters 2 and 3, we obtain asymptotic bounds for the solution of the constrained tradeoff problem for a discrete time model with a fixed environment state in Chapter 4. We again consider admissible policies, for which geometric bounds on the stationary probability distribution for the queue length are obtained. The forms of these geometric bounds are motivated by the geometric bounds which were derived for state dependent M/M/1 model in Chapters 2 and 3. Using these geometric bounds, asymptotic lower bounds are derived for the solution to the constrained tradeoff problem in the regime \mathfrak{R} . Asymptotic upper bounds are also derived, using

which a complete asymptotic characterization of the tradeoff in the regime \mathfrak{R} is obtained in two cases. Asymptotic bounds on any sequence of order-optimal admissible policies are also derived.

In Chapter 5, we first consider the tradeoff of average power and average delay for a noisy point-to-point link with fast fading. Asymptotic bounds for the solution to the constrained tradeoff problem, for admissible policies, in the regime \mathfrak{R} are then derived by extending the results of Chapter 4. We also compare the asymptotic lower bounds which are obtained from a real valued approximate queueing model to those from the original integer valued queueing model. We find that the real valued approximate queueing model with a strictly convex cost function underestimates the average service cost and average delay for the original model. We show that a more appropriate approximate real valued queueing model is one in which the service cost function is the piecewise linear lower convex envelope of the service cost function for the original model. We also obtain asymptotic lower bounds for the constrained tradeoff problem for: (i) a model with admission control, and (ii) a single hop network model. Asymptotic lower bounds are also obtained for the case when the arrival process and fading process are ergodic.

In Chapter 6, we consider the tradeoff of average error rate and average delay for a noisy point-to-point link. We obtain the exponential decay rate of average error rate with respect to average delay, in the regime of large average delay, for fixed length random block coding schemes, with control on the codeword length parameter N_c . Using results from Chapter 4, we then obtain an asymptotic characterization of the tradeoff of average error rate and average delay, for the set of admissible policies, for a fixed N_c . One of the main contributions in Chapter 6 is the analysis of the constrained tradeoff problem for a non-convex service cost function, unlike the convex service cost functions in the earlier chapters. We show that the asymptotic behaviour of the tradeoff in the regime \mathfrak{R} is determined by the lower convex envelope for any non-convex service cost function.

7.1 Scope for future work

We note that all the asymptotic bounds obtained in this thesis are order bounds. From numerical computation, we have observed that the constants which are involved in the upper and lower bounds are weak. Tight non-asymptotic bounds for the constrained tradeoff problem are still not available.

We now discuss some specific problems that arise from the analysis in the previous chapters, for which some initial results have been obtained.

7.1.1 Order optimality of policies obtained from a fluid model

A simpler deterministic model for the evolution of the queue length, for the state dependent M/M/1 model in Chapter 2, is the trajectory $q(t)$ obtained from the fluid model:

$$\frac{dq(t)}{dt} = \lambda - \mu(t), q(0) = q_0, \quad (7.1)$$

where $\mu(t)$ is the service rate at time t . The fluid model can be interpreted as a *limiting* form of the evolution of the integer valued queue length, when both time as well as queue length are scaled. Such fluid models can also be obtained for the discrete time model in Chapter 4, as well as for general networks, e.g. [36, Chapter 10]. In [36, Chapter 10] and [17], it is shown that the policy $\mu(t)$ which minimizes the total cost $J(q_0)$ for the deterministic fluid model, can be used to obtain *good* policies for the unconstrained MDP for the original stochastic state dependent M/M/1 model. Our objective is to investigate the order optimality of policies which are obtained from fluid models in the asymptotic regime \mathfrak{R} .

The state transition diagram for the unconstrained MDP is the same as that in Figure 2.1, with the action at each state q being the service rate $\mu(q) \in \{0, \mu_1, \dots, \mu_K\}$. We also note that the optimal policy for the MDP with single stage cost $q + \beta c(\mu)$ is the same as that for the MDP with single stage cost $q + \beta(c(\mu) - c(\lambda))$. We construct a heuristic policy for the model, when $c(\mu) = \mu^2$, for $\mu \in \mathcal{S} = \{0, \mu_1, \dots, \mu_K\}$, based upon the analysis in Chen et al. [17]. As in Chen et al. [17], we consider the fluid model (7.1), with a modified single stage cost $q + \beta [c(\mu) - c(\lambda)]^+$ and the service rate at time t , $\mu(t) \in \mathbb{R}_+$. Then

$$J(q_0) = \inf_{\mu(t)} \int_0^\infty \{q(t) + \beta [c(\mu(t)) - c(\lambda)]^+\} dt.$$

If $c(\mu) = \mu^2$, then from [17] we have that $J(q) = \frac{2}{3} q^{\frac{3}{2}} \left(\frac{1}{\beta}\right)^{\frac{1}{2}}$. Furthermore the optimal service rate function $\mu^*(t)$ for which

$$J(q_0) = \int_0^\infty \{q(t) + \beta [c(\mu^*(t)) - c(\lambda)]^+\} dt,$$

is given by the state dependent policy $\mu^*(q) = (q/\beta)^{\frac{1}{2}} + c(\lambda)$. Motivated by this development, we define the fluid policy $\gamma_{F,\beta}$ for our M/M/1 model as choosing the service rate

$$\mu_{F,\beta}(q) = \left[(q/\beta)^{\frac{1}{2}} \right]_{\mathcal{S}},$$

where $[x]_{\mathcal{S}}$ means that we pick a service rate in \mathcal{S} which is closest to the argument x . Interestingly, the policy $\gamma_{F,\beta}$ does not depend on λ but varies only as a function of β . We note that $\gamma_{F,\beta}$ is admissible for any $\lambda < \mu_K$, since there exists a queue length q such that $\mu_{F,\beta}(q) = \mu_K$ and from

the definition, $\mu_{F,\beta}(q)$ is non-decreasing in q . We illustrate the order optimality of $\gamma_{F,\beta}$ only for case 2. We first obtain an upper bound on $\bar{Q}(\gamma_{F,\beta})$ using Proposition 2.A.1. It can be shown that for $\epsilon > 0$, if there exists a q_ϵ such that $\mu_{F,\beta}(q_\epsilon) - \lambda \geq \epsilon$, then $\bar{Q}(\gamma_{F,\beta}) \leq \frac{q_\epsilon(\epsilon+\lambda)}{\epsilon} + \frac{\lambda+\mu_K}{2\epsilon}$.

We note that if $\lambda < \mu_K$, for case 2, we can pick $\epsilon = \mu_u - \lambda$. Then $q_\epsilon \leq \beta \left(\frac{\mu_{k_u-1} + \mu_u}{2} \right)^2 + 1$. Thus we obtain that $\bar{Q}(\gamma_{F,\beta}) = \mathcal{O}(\beta)$.

We have that $\bar{C}(\gamma_{F,\beta})$

$$\begin{aligned} &= \sum_{k < k_l} \pi_\mu(k) c(\mu_k) + \sum_{k > k_u} \pi_\mu(k) c(\mu_k) + \sum_{k=k_l}^{k_u} \pi_\mu(k) c(\mu_k), \\ &\leq c(\mu_{k_l-1}) Pr \{ \mu_{F,\beta}(Q) < \mu_l \} + c(\mu_K) Pr \{ \mu_{F,\beta}(Q) > \mu_u \} \\ &\quad + c(\lambda) + \sum_{k=k_l}^{k_u} \pi_\mu(k) m(\lambda - \mu_k), \end{aligned} \tag{7.2}$$

where m is the slope of the line joining $(\mu_l, c(\mu_l))$ and $(\mu_u, c(\mu_u))$. We note that since $\gamma_{F,\beta}$ is admissible, $\sum_{k=0}^K \pi_\mu(k) \mu_k = \lambda$. Hence, we have that $\sum_{k < k_l} \pi_\mu(k) (\mu_k - \lambda) + \sum_{k > k_u} \pi_\mu(k) (\mu_k - \lambda) = \sum_{k=k_l}^{k_u} \pi_\mu(k) (\lambda - \mu_k)$. Since for $k < k_l$, $\mu_k < \lambda$, we have that $\sum_{k=k_l}^{k_u} \pi_\mu(k) (\lambda - \mu_k) \leq (\mu_K - \lambda) Pr \{ \mu_{F,\beta}(Q) > \mu_u \}$. We note that $Pr \{ \mu_{F,\beta}(Q) > \mu_u \} = \sum_{q > q_{k_u}} \pi(q)$ and $Pr \{ \mu_{F,\beta}(Q) < \mu_l \} = \sum_{q \leq q_{k_l-1}} \pi(q)$. From the birth-death structure of $Q(t)$ under $\gamma_{F,\beta}$, we can show that

$$\begin{aligned} \sum_{q > q_{k_u}} \pi(q) &\leq \pi(q_{k_u} + 1) \left(\frac{\mu_{k_u+1}}{\mu_{k_u+1} - \lambda} \right), \\ \sum_{q \leq q_{k_l-1}} \pi(q) &\leq \pi(q_{k_l-1}) \left(\frac{1 - \left(\frac{\mu_{k_l-1}}{\lambda} \right)^{q_{k_l-1}+1}}{1 - \frac{\mu_{k_l-1}}{\lambda}} \right). \end{aligned}$$

We note that $\pi(q_{k_u} + 1) = \pi(q_{k_u-1}) \left(\frac{\lambda}{\mu_u} \right)^{q_{k_u} - q_{k_u-1} + 1}$ and $\pi(q_{k_l-1}) = \pi(q_{k_l}) \left(\frac{\mu_l}{\lambda} \right)^{q_{k_l} - q_{k_l-1}}$. From the definition of $\mu_{F,\beta}$ we have that $q_{k_u} - q_{k_u-1}$ and $q_{k_l} - q_{k_l-1}$ are both $\Theta(\beta)$, as $\beta \uparrow \infty$. Since $\pi(q_{k_u-1})$ and $\pi(q_{k_l-1})$ are both bounded above by one, we have that $Pr \{ \mu_{F,\beta}(Q) > \mu_u \}$ as well as $Pr \{ \mu_{F,\beta}(Q) < \mu_l \}$ are both $\rho^{\Theta(\beta)}$, where $0 < \rho < 1$. Then we have that $\bar{C}(\gamma_\beta) - c(\lambda) = \rho^{\Theta(\beta)}$ and $\bar{Q}(\gamma_\beta) = \mathcal{O}(\beta)$. In summary, we have shown a new optimality property of heuristic policies, obtained from a fluid model analysis. This raises the question whether such fluid policies are order optimal even for general network scenarios.

7.1.2 Extensions to systems with service time control

We note that the service rate control variable has been the service batch size for all the discrete time queueing models considered in this thesis. However, there are scenarios where both batch size and batch service time can be controlled, e.g., for a noisy point to point link which uses block

coding, both the number of encoded message symbols (the batch size) and the codeword length can be dynamically controlled (the service time for the batch) to tradeoff the average error rate with the average delay of the message symbols. Such models also arise as a special case, when *single decision policies with observed initial information* are used, for the general class of renewal models considered by Neely [46]. We note that for such models the service cost is then modelled as a function of both batch size and batch service time.

We now comment on how asymptotic lower bounds can be derived for such models, using the methods in Chapter 4. We note that for stationary policies, that decide on the batch size and batch service time as a function of the current queue length, the queue length evolution sampled at the decision epochs evolves as a semi-Markov process. By uniformization [75] the average cost and average queue length for a semi-Markov process can be obtained via an equivalent Markov process. The stagewise drift of this equivalent Markov process depends on both the batch size as well as the batch service time. We expect that the analysis of this equivalent Markov process is similar to the analysis carried out in Chapter 4, but with the above drift function, under the restriction to admissible policies.

We consider a simplified model in [52], wherein asymptotic lower bounds are derived for a continuous time queueing model with Poisson arrivals and service time control with service batch size fixed to be 1. For this model, as the service batch size is fixed to be 1, whenever there is service, the service cost is a function only of the batch service time, which is real valued. For the above continuous time model, where the queue evolution is on the integers, if the service cost per unit batch service time is a strictly convex function of the batch service time, then we obtain that the minimum average queue length grows as $\Omega\left(\frac{1}{\sqrt{V}}\right)$ when the average service cost constraint is V more than the minimum average service cost required for stability. Obtaining asymptotic bounds for queueing models where both batch size and batch service time can be controlled, is a problem which has scope for future work.

Bibliography

- [1] M. Agarwal, V.S. Borkar, and A. Karandikar. Structural properties of optimal transmission policies over a randomly varying channel. *IEEE Transactions on Automatic Control*, 53(6), 2008.
- [2] E. Altman. *Constrained Markov decision processes*. Chapman and Hall, 1999.
- [3] B. Ata. Dynamic power control in a wireless static channel subject to a quality-of-service constraint. *Operations Research*, 53(5), 2005.
- [4] B. Ata, J. M. Harrison, and L. A. Shepp. Drift rate control of a Brownian processing system. *The Annals of Applied Probability*, 15(2), 2005.
- [5] B. Ata and S. Shneorson. Dynamic control of a M/M/1 service system with adjustable arrival and service rates. *Management Science*, 52(11), 2006.
- [6] R. A. Berry. *Power and delay tradeoffs in fading channels*. PhD thesis, LIDS, Massachusetts Institute of Technology, 2000.
- [7] R.A. Berry and R.G. Gallager. Communication over fading channels with delay constraints. *IEEE Transactions on Information Theory*, 48(5), May 2002.
- [8] D. Bertsimas, D. Gamarnik, and J.N. Tsitsiklis. Geometric bounds for stationary distributions of infinite Markov chains via Lyapunov functions. *Available online at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.31.9210>*, 1998.
- [9] D. Bertsimas, D. Gamarnik, and J.N. Tsitsiklis. Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *The Annals of Applied Probability*, 11(4), 2001.

- [10] I. Bettesh and S. Shamai. Optimal power and rate control for minimal average delay: the single-user case. *IEEE Transactions on Information Theory*, 52(9), sep. 2006.
- [11] T. Bonald. Insensitive queueing models for communication networks. In *Proceedings of the 1st International Conference on Performance Evaluation Methodologies and Tools*. ACM, 2006.
- [12] A.A. Borovkov. *Stochastic processes in queueing theory*. Springer-Verlag, 1976.
- [13] S. Borst. User-level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Networking*, 13(3), 2005.
- [14] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [15] P. Chaporkar, K. Kar, Xiang Luo, and S. Sarkar. Throughput and fairness guarantees through maximal scheduling in wireless networks. *IEEE Transactions on Information Theory*, 54(2), 2008.
- [16] P. Chaporkar and A. Proutiere. Adaptive network coding and scheduling for maximizing throughput in wireless networks. In *Mobicom*, 2007.
- [17] Wei Chen, Dayu Huang, A.A. Kulkarni, J. Unnikrishnan, Quanyan Zhu, P. Mehta, S. P. Meyn, and A. Wierman. Approximate dynamic programming using fluid and diffusion approximations with applications to power management. In *Proceedings of the 48th IEEE Conference on Decision and Control*, 2009.
- [18] B. Collins and R. Cruz. Transmission policies for time varying channels with average delay constraints. In *Proceedings of 1999 Allerton Conf. Communication, Control, and Computing*.
- [19] D. Denteneer, A.J.E.M. Janssen, and J.S.H. van Leeuwen. Moment inequalities for the discrete-time bulk service queue. *Mathematical Methods of Operations Research*, 61, 2005.
- [20] W. F. Donoghue. *Distributions and Fourier Transforms*. Academic Press, 1969.
- [21] N. Ehsan and T. Javidi. Delay optimal transmission policy in a wireless multiaccess channel. *IEEE Transactions on Information Theory*, 54(8), Aug 2008.
- [22] R. G. Gallager. *Information Theory and Reliable Communication*. Wiley, 1968.
- [23] R. G. Gallager. The random coding bound is tight for the average code (corresp.). *IEEE Transactions on Information Theory*, 19(2), March 1973.
- [24] A. E. Gamal and Y. H. Kim. *Network Information Theory*. Cambridge University Press, 2011.
- [25] J. M. George and J. M. Harrison. Dynamic control of a queue with adjustable service rate. *Operations Research*, 49(5), 2001.

- [26] L. Georgiadis, M.J. Neely, and L. Tassiulas. *Resource allocation and cross layer control in wireless networks*. Foundations and trends in networking, NOW publishers, 2006.
- [27] J. Gonzalez-Hernandez and C. E. Villarreal. Optimal policies for constrained average-cost Markov decision processes in Borel spaces. *SIAM Journal on Control and Optimization*, 42(2), May 2003.
- [28] J. Gonzalez-Hernandez and C. E. Villarreal. Optimal policies for constrained average-cost Markov decision processes. *TOP Journal of Spanish Society of Statistics and Operations Research*, 19(1), July 2011.
- [29] M. Goyal, A. Kumar, and V. Sharma. Optimal cross-layer scheduling of transmissions over a fading multiaccess channel. *IEEE Transactions on Information Theory*, 54(8), 2008.
- [30] L. Huang and M.J. Neely. Max-Weight Achieves the Exact $[O(1/V), O(V)]$ Utility-Delay Tradeoff Under Markov Dynamics. *ArXiv e-prints, available online at <http://arxiv.org/abs/1008.0200/>*, August 2010.
- [31] T. Javidi and R.N. Swamy. Optimal code length for bursty sources with deadlines. In *IEEE International Symposium on Information Theory*, 2009.
- [32] G. Koole. *Monotonicity in Markov reward and decision Chains : theory and applications*. Foundations and Trends in Stochastic Systems : NOW Publishers, 2007.
- [33] A Kumar, D. Manjunath, and J Kuri. *Communication networking - An analytical approach*. Morgan Kaufmann Publishers, 2004.
- [34] R.M. Loynes. The stability of a queue with non-independent interarrival and service times. *Mathematical proceedings of the Cambridge Philosophical Society*, 58(3), July 1962.
- [35] D.J. Ma, A. M. Makowski, and A. Shwartz. Estimation and optimal control for constrained Markov chains. *IEEE Conference on Decision and Control*, 25, 1986.
- [36] S. P. Meyn. *Control techniques for complex networks*. Cambridge University Press, 2007.
- [37] S.P. Meyn. Stability, performance evaluation, and optimization. In *Handbook of Markov Decision Processes : Methods and Applications*. Editors : A. Schwartz and E. Feinberg, Kluwer, 2001.
- [38] U. Mukherji, S.V. Ramdurg, K.C.V. Sayee, V. Dua, and T.N. Krishnan. Multi-access Poisson traffic communication with random coding, independent decoding and unequal powers. *Proceedings of the IEEE Information Theory Workshop*, Oct. 2002.
- [39] S. Musy. *Delay and coding in multiuser communication*. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2007.

- [40] S. Musy and E. Telatar. On the transmission of bursty sources. In *Proceedings on the IEEE International Symposium on Information Theory*, 2006.
- [41] M.J. Neely. *Dynamic power allocation and routing for satellite and wireless networks with time varying channels*. PhD thesis, LIDS, Massachusetts Institute of Technology, 2003.
- [42] M.J. Neely. Super-fast delay tradeoffs for utility optimal fair scheduling in wireless networks. *IEEE Journal on Selected Areas in Communications*, 24(8), 2006.
- [43] M.J. Neely. Optimal energy and delay tradeoffs for multiuser wireless downlinks. *IEEE Transactions on Information Theory*, 53(9), Sept. 2007.
- [44] M.J. Neely. Intelligent packet dropping for optimal energy-delay tradeoffs in wireless downlinks. *IEEE Transactions on Automatic Control*, 54(3), March 2009.
- [45] M.J. Neely. *Stochastic network optimization with application to communication and queueing systems*. Morgan and Claypool, 2010.
- [46] M.J. Neely. Dynamic optimization and learning for renewal systems. *IEEE Transactions on Automatic Control*, 58(1), Jan 2013.
- [47] R. Negi and S. Goel. An information-theoretic approach to queuing in wireless channels with large delay bounds. In *Proceedings of the IEEE Global Telecommunications Conference, 2004*, volume 1, 2004.
- [48] J.R. Perkins and R. Srikant. The role of queue length information in congestion control and resource pricing. In *Proceedings of the 38th IEEE Conference on Decision and Control*. IEEE, 1999.
- [49] Y. Polyanskiy. *Channel coding: non-asymptotic fundamental limits*. PhD thesis, Dept. of Electrical Engineering, Princeton University, 2010.
- [50] V. Ramaiyan. *Topics in modelling, analysis and optimization of wireless networks*. PhD thesis, Dept. of Electrical Communication Engineering, Indian Institute of Science, 2009.
- [51] V. Ramaiyan, E. Altman, and A. Kumar. Delay optimal scheduling in a two-hop vehicular relay network. *Mobile Networks and Applications*, 15(1), 2010.
- [52] Vineeth B. S. Optimal tradeoff of average service cost and average delay for single server queues with service time control. *Unpublished report; available at <http://ece.iisc.ernet.in/~vineeth/>*.
- [53] Vineeth B. S. and U. Mukherji. Average-delay optimal policies for the point-to-point channel. In *Proceedings of the 7th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2009.

- [54] Vineeth B. S. and U. Mukherji. Average-delay optimal policies for the point-to-point channel. Technical report, DRDO-IISc Programme on Mathematical Engineering, TR-PME-2010-13, available at <http://www.pal.ece.iisc.ernet.in/PAM/pubs-techreports.php>, 2010.
- [55] Vineeth B. S. and U. Mukherji. Bounds on the optimal queueing delay for variable rate point-to-point channels with fixed block error probability. Technical report, DRDO-IISc Programme on Mathematical Engineering, TR-PME-2011-11, available at <http://www.pal.ece.iisc.ernet.in/PAM/pubs-techreports.php>, 2011.
- [56] Vineeth B. S. and U. Mukherji. Queueing delay - error probability tradeoff for point-to-point channels using block codes. Technical report, DRDO-IISc Programme on Mathematical Engineering, TR-PME-2011-12, available at <http://www.pal.ece.iisc.ernet.in/PAM/pubs-techreports.php>, 2011.
- [57] Vineeth B. S. and U. Mukherji. Queueing delay - error probability tradeoff for point-to-point channels using streaming codes. Technical report, DRDO-IISc Programme on Mathematical Engineering, TR-PME-2011-13, available at <http://www.pal.ece.iisc.ernet.in/PAM/pubs-techreports.php>, 2011.
- [58] Vineeth B. S. and U. Mukherji. Queueing delay - error probability tradeoff for point-to-point channels with fixed length block codes. In *Proceedings of the IEEE International Symposium on Information theory*, 2011.
- [59] Vineeth B. S. and U. Mukherji. Tradeoff of average power and average delay for slow fading point-to-point channels under reliability constraints. Technical report, DRDO-IISc Programme on Mathematical Engineering, TR-PME-2011-14, available at <http://www.pal.ece.iisc.ernet.in/PAM/pubs-techreports.php>, 2011.
- [60] Vineeth B. S. and U. Mukherji. Delay optimal scheduling of a discrete time batch service queue for point to point channel code rate selection. In *Proceedings of the 2012 National Conference on Communications*, 2012.
- [61] Vineeth B. S. and U. Mukherji. Tradeoff of average power and average delay for a point-to-point link with fading. In *Proceedings of the National Conference on Communications (NCC)*, 2013.
- [62] Vineeth B. S. and U. Mukherji. Tradeoff of average service cost and average delay for the state dependent M/M/1 queue. In *Proceedings of the National Conference on Communications (NCC)*, 2013.
- [63] KCV Kalyanarama Sesha Sayee. *Scheduling for stable and reliable communication over multiaccess channels and degraded broadcast channels*. PhD thesis, Department of Electrical Communication Engineering, Indian Institute of Science, 2006.

- [64] KCV Kalyanarama Sesha Sayee and U. Mukherji. Stability of scheduled multi-access communication over quasi-static flat fading channels with random coding and independent decoding. *International Symposium on Information Theory*, Sept. 2005.
- [65] KCV Kalyanarama Sesha Sayee and U. Mukherji. A multiclass discrete-time processor-sharing queueing model for scheduled message communication over multiaccess channels with joint maximum-likelihood decoding. *Proceedings of the 44th Annual Allerton Conference*, pages 615–622, Sept 2006.
- [66] KCV Kalyanarama Sesha Sayee and U. Mukherji. Stability of scheduled message communication over degraded broadcast channels. *IEEE International Symposium on Information Theory*, July 2006.
- [67] L. I. Sennott. Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Operations Research*, 37(4), 1989.
- [68] L. I. Sennott. Constrained average cost Markov decision chains. *Probability in the Engineering and Informational Sciences*, 7, 1993.
- [69] L. I. Sennott. *Stochastic dynamic programming and the control of queues*. Wiley- IEEE, 1999.
- [70] S. Stidham and R.R. Weber. Monotonic and insensitive optimal policies for the control of queues with undiscounted costs. *Operations Research*, 37, 1989.
- [71] R. N. Swamy and T. Javidi. Delay analysis of block coding over a noisy channel with limited feedback. In *Asilomar Conference on Signals, Systems and Computers*.
- [72] Bo Tan and R. Srikant. Online advertisement, optimization and stochastic networks. *CoRR*, abs/1009.0870, 2010.
- [73] L. Tassiulas and A Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12), 1992.
- [74] L. Tassiulas and A Ephremides. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory*, 39(2), 1993.
- [75] C. H. Tijms. *A first course in stochastic models*. Wiley, 2003.
- [76] D. Tse. *Variable-rate lossy compression and its effects on communication networks*. PhD thesis, LIDS, Massachusetts Institute of Technology, 1994.
- [77] E. Uysal-Biyikoglu, B. Prabhakar, and A. El Gamal. Energy-efficient packet transmission over a wireless link. *IEEE/ACM Transactions on Networking*, 10(4), 2002.

- [78] W. Whitt. A review of $L = \lambda W$ and extensions. *Queueing Systems*, (9), 1991.
- [79] E. M. Yeh. Delay-optimal rate allocation in multiaccess communications: A cross-layer approach. *Proceedings of the International Workshop on Multimedia Signal Processing*, 2002.
- [80] W.S. Yoon. Streaming transmission of poisson traffic. In *Proceedings of IEEE International Symposium on Information Theory*, 2008.