FUNDAMENTAL LIMITS OF RANDOM ACCESS
IN WIRELESS NETWORKS

BY

JAVAD GHADERI DEHKORDI

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2013

Urbana, Illinois

Doctoral Committee:

Professor Rayadurgam Srikant, Chair
Professor Bruce Hajek
Assistant Professor Angelia Nedich
Professor Ness Shroff
Professor Pramod Viswanath

Abstract

Random access schemes are simple and inherently distributed, yet could provide the striking capability to match the optimal throughput performance (maximum stability region) of centralized scheduling mechanisms. The throughput optimality however has been established for activation rules that are relatively sluggish, and may yield excessive queues and delays. More aggressive/persistent access schemes have the potential to improve the delay performance, but it is not clear if they can offer any universal throughput optimality guarantees. In this thesis, we identify a fundamental limit on the aggressiveness of nodes, beyond which instability is bound to occur in a broad class of networks.

We will mainly consider adapting transmission lengths by considering a weight for each node as a function of its queue size. The larger the weight, the longer the node will hold on to the channel once it starts a transmission. We first show that it is sufficient for weights to behave as logarithmic functions of the queue sizes, divided by an arbitrarily slowly increasing function. This result indicates that the maximum-stability guarantees are preserved for weights that are essentially logarithmic for all practical queue sizes, although asymptotically the weight must grow slower than any logarithmic function of the queue size. We then demonstrate instability for weights that grow faster than logarithmic functions of queue sizes in networks with sufficiently many nodes. Our stability and instability results hence imply that the "near-logarithmic growth condition" on the weights is a fundamental limit on the aggressiveness of nodes to ensure maximum stability in any general topology. We will conduct simulation experiments to illustrate and validate the analytical results. Finally, we will combine the random access scheme with window-based flow control mechanisms to provide maximum throughput and Quality-of-Service in multihop wireless networks with dynamic flows.

Acknowledgments

First and foremost, I would like to thank my advisor Professor R. Srikant for his support in the course of my PhD studies. Without his guidance, carrying out this research would have been impossible. I will miss our numerous meetings and discussions that shaped my ideas. I have been also very fortunate to collaborate with Dr. Sem Borst and Dr. Phil Whiting at Alcatel-Lucent Bell Labs. I am grateful to them for their insightful discussions and contributions to my research. Part of this thesis was done during my internship at Alcatel-Lucent Bell Labs where I was supervised by Sem and Phil. I would like to thank the members of my doctoral committee, Professor Bruce Hajek, Professor Angelia Nedich, Professor Ness Shroff, and Professor Pramod Viswanath for their time and thoughtful comments. I express my deep appreciation to the great teachers at the University of Illinois with impressive teaching styles. My career path has been influenced by such teachers and I can still feel their impact on my thinking and problem solving methodology.

My sincere thanks go to my parents for their support and dedication through my life. There is no way that I can express my gratitude to them. I am also really grateful to Majid, my brother, for his inspiration and guidance.

I would also like to thank my colleagues at the Coordinated Science Lab (CSL) for their support and friendship. They provided an excellent environment for me during my studies at the University of Illinois.

I would like to thank the University of Illinois, Alcatel-Lucent, the US National Science Foundation, Army Research Office, and the Air Force Office of Scientific Research for their financial support.

Table of Contents

Chapter 1

Introduction

Scheduling in wireless networks is of fundamental importance due to the inherent broadcast nature of the wireless medium. Two radios that are close to each other might not be able to transmit simultaneously, as they can interfere with each other. In other words, the simultaneous transmissions of such radios may cause the *Signal-to-Noise-plus-Interference-Ratio* (SINR) at their corresponding receivers to go below the required threshold for successful decoding of data packets. Therefore, in order to operate wireless systems efficiently, scheduling algorithms are needed to facilitate simultaneous transmissions of different radios.

The metrics used to evaluate the performance of a scheduling algorithm are *throughput* and *delay*. Throughput is characterized by the largest set of service rates that can be provided to the radio nodes. Delay is characterized by the average time that packets spend in the network's buffers, once they enter the network until they reach their destinations. It is also essential for the scheduling algorithm to be *distributed* and to have low *complexity/overhead*. This is because in many wireless networks there is no centralized entity and the resources at the nodes are very limited.

Scheduling algorithms for wireless networks have been widely studied since Tassiulas and Ephremides [1] proposed the *Max Weight Scheduling* (MWS) algorithm. MWS algorithm assigns a weight to each link as a function of the number of packets queued at the link, and then, at each instant of time, selects the schedule with the maximum weight, where the weight of a schedule is computed by summing the weights of the links that the schedule will serve. Tassiulas and Ephremides establish that the MWS algorithm is throughput optimal in the sense that it can stabilize the queues of the network for the largest set of arrival rates possible, without actually knowing the arrival rates. However, finding the maximum weight schedule is a complex combinatorial problem, and hence, the MWS algorithm is not typically implementable in practice. This has led to a rich amount of literature on the design of approximate algorithms to alleviate the computational complexity of the MWS algorithm.

A popular mechanism for distributed scheduling is provided by the so-called *Carrier Sense Multiple Access* (CSMA) protocol. In the CSMA protocol, each node attempts to access the medium after a certain back-off time, but nodes that sense activity of interfering nodes freeze their back-off timer until the medium is sensed idle. Due to their simplicity of implementation, CSMA schemes have been widely used in practice, e.g., in WLANs (IEEE 802.11 Wi-Fi) or emerging wireless mesh networks. From a local perspective, the CSMA algorithm might seem easy to understand, but, at a global perspective, interactions among different nodes can lead to a very complicated behavior that makes the performance characterization difficult.

In recent years, fairly simple models have been proposed that are useful in predicting the throughput of the CSMA algorithm [2, 3, 4, 5, 6]. Although the representation of the IEEE 802.11 back-off mechanism in these models is less detailed than in [7], they accommodate a general interference graph and thus cover a broad range of topologies. Experimental results in [8] demonstrate that these models, while idealized, provide throughput estimates that match remarkably well with measurements in actual systems.

Despite their asynchronous and distributed nature, CSMA-like algorithms have been shown to offer the remarkable capability of achieving the full capacity region and thus match the optimal throughput performance of centralized scheduling mechanisms operating in slotted time [9, 10, 11]. More specifically, any throughput vector in the interior of the convex hull associated with the independent sets in the underlying interference graph can be achieved through suitable back-off rates and/or transmission lengths. Based on this observation, various ingenious algorithms have been developed for finding the back-off rates that yield a particular target throughput vector or that optimize a certain concave throughput utility function in scenarios with saturated buffers [9, 11, 12]. In particular, Jiang and Walrand [11] develop an algorithm that adaptively chooses the back-off rates under a time-scale separation assumption, i.e., the CSMA dynamics converges to its equilibrium instantaneously compared to the time-scale of adaptation of the back-off rates. This time-scale separation assumption was later verified by a stochastic approximation type argument [9, 10]. In the same spirit, several effective approaches have been devised for adapting the transmission lengths based on queue-length information, and been shown to guarantee maximum stability [13, 14, 15, 16].

Roughly speaking, the maximum-stability guarantees were established under the condition that the various nodes are relatively sluggish in attempting to access the channel or in continuing transmission once they obtain the channel. In particular, a very exciting work by Rajagopalan, Shah, and Shin [13] demonstrates the maximum stability under the condition that the mean transmission length is chosen to be a logarithmic function of the queue length. In the language of [13], each link has a weight of the form $\log\log(Q)$ (Q is the queue length) that the link uses to determine its transmission length.[1]

Unfortunately, such activation rules can induce excessive queue lengths and delays, as the resulting scheduling algorithm reacts very slowly to changes in queue lengths. This issue has triggered a strong interest in developing approaches for improving the delay performance [17, 18, 19, 20, 21, 22] and has also provided the motivation of this thesis. More aggressive/persistent access schemes have the potential to improve the delay performance, but it is not clear if they can offer any universal maximum-stability guarantees.

In this thesis, we explore the scope for using more aggressive activation rules (more aggressive weights) in order to improve the delay performance while preserving the maximum-stability guarantees. To this end, we will analyze stability and instability of the network under the random access mechanism with different activation rules.

In Chapter 3, we first tighten the condition required for CSMA-like algorithms to achieve maximum throughput (maximum stability region). We show that it is in fact sufficient for weights to be logarithmic functions of the queue lengths, divided by an arbitrarily slowly increasing, unbounded function. For example, weights of the form $\log^{1-\epsilon} Q$, with $\epsilon > 0$ arbitrary small, are sufficient to ensure maximum throughput in any general network topology. This result indicates that the maximum-throughput guarantees are preserved for weight functions that are essentially logarithmic for all practical queue lengths, although asymptotically weights must grow slower than any logarithmic function of the queue length.

Since the "near-logarithmic growth condition" is only a sufficient condition,

[1] In this thesis, all logarithms are in base e.

it is not clear to what extent it is actually a strict requirement for maximum stability to be maintained. In Chapter 4, we will consider the fluid limits of the system where dynamics are scaled in both space and time. For weights which grow slower than any logarithmic function of queue lengths, "fast mixing" is guaranteed in any general topology, where the activity process evolves on a much faster time scale than the scaled queue lengths. Qualitatively similar fluid limits can arise for more aggressive weight functions as well, provided the topology is benign. However, aggressive weight functions can cause "sluggish mixing," where the activity process evolves on a much slower time scale than the scaled queue lengths, yielding random oscillatory fluid limits. Such fluid limits can force the system into inefficient states for extended periods of time and produce instability. We will demonstrate instability for weights that grow faster than $\gamma \log Q$ (Q is the queue length), for any $\gamma > 1$, but our proof arguments suggest that it can occur for any $\gamma > 0$, in networks with sufficiently many nodes. In other words, "the near-logarithmic growth condition" on the weights is a fundamental limit on the aggressiveness of nodes to ensure maximum stability (throughput optimality) in any general topology.

In Chapter 5, we will investigate the stability of the system when flows/users randomly arrive and depart. To achieve flow-level stability, prior works [23, 24, 25] require that a specific form of congestion control based on α-fair utility functions has to be used, namely the rate at which a flow generates packets into its ingress queue must maximize an α-fair utility subject to a linear penalty (price). We will show that α-fair congestion control *is not* necessary for flow-level stability, and, in fact, very general congestion control mechanisms are sufficient to ensure flow-level stability. In establishing this result, we will use link weights which are log-differentials of queue lengths, i.e., the weight of link (i, j) is roughly in the form of $\log(1+Q_i) - \log(1+Q_j)$, with Q_i and Q_j the queue lengths of nodes i and j. The use of such weights naturally suggests the use of a random access mechanism, as in Chapter 3, to implement the algorithm in a distributed fashion. We will indeed show that the maximum-stability result of the random access mechanism can be easily extended to multihop flows with log-differential weights and very general congestion control mechanisms.

Our stochastic model in this thesis considers the dynamics of flows, packets, and random access simultaneously, with no time-scale separation assumptions among the dynamics.

The remainder of this thesis is organized as follows. In Chapter 2, we introduce our model for wireless networks and describe the random access mechanism in its most general form. Chapter 3 is devoted to the stability result for near-logarithmic weight functions. In Chapter 4, we will demonstrate instability of the system for weight functions which grow faster than the near-logarithmic functions. In Chapter 5, we will investigate stability of the system with flow arrivals and departures. In each chapter, we will conduct simulation experiments to verify our analytical results.

Chapter 2

System Model and Description of Random Access Mechanism

2.1 System model

For now, we consider a singlehop communication model, i.e., each source is directly connected to its destination via a wireless link. In Chapter 5, we will show how to extend the results to the multihop scenario.

Suppose the total number of source-destination pairs (wireless links) is N. We use the notion of the *conflict graph* $G(V, E)$ to represent the interference/operating constraints. Each node of the conflict graph is a communication link in the wireless network and there is an edge $(l, k) \in E$ between nodes l and k if simultaneous transmissions over communication links l and k cannot be successful. Therefore, at each time instant, the active links should form an independent set of G, i.e., no two scheduled nodes can share an edge in G. Formally, a schedule can be represented by a vector $X = [x_s : s = 1, ..., N]$ such that $x_s \in \{0, 1\}$ and $x_i + x_j \leq 1$ for all $(i, j) \in E$. We use $\mathcal{M}$ to denote the set of all feasible schedules.

Packets arrive at link l according to some stochastic process with rate $\lambda = [\lambda_l; l = 1, \ldots, N]$. To be specific, assume packets arrive according to the Poisson process for the continuous-time system and the Bernoulli process for the discrete-time system. Each link l is associated with a queue $Q_l(t)$, representing the number of packets of link l waiting for transmission at time t. The vector of queue lengths at each time t is denoted by $Q(t) = [Q_l(t) : l = 1, ..., N]$.

A scheduling algorithm is a policy to determine which schedule to be used in each time instant; correspondingly, which links can transmit a packet in each time instant. The capacity region of the network is defined to be the largest set of arrival rates that can be supported by the network, i.e., for which there exists a scheduling algorithm that can stabilize the queues. It is known, e.g. [1], that the capacity region is given by

$$\Lambda = \{\lambda \geq 0 : \exists \mu \in \mathrm{Co}(\mathcal{M}),\ \lambda < \mu\}, \tag{2.1}$$

where $\mathrm{Co}(\cdot)$ is the convex hull operator. When dealing with vectors, inequalities are interpreted component-wise.

A scheduling algorithm is throughput-optimal if it can stabilize the network for any arrival rate in Λ. An important class of the throughput-optimal algorithms are the Max Weight Scheduling (MWS) algorithms where at any time t, the scheduling decision $X(t)$ satisfies

$$X(t) = \arg\max_{X \in \mathcal{M}} \sum_{l=1}^{N} x_l w_l(t),$$

where $w_l(t)$ is the weight of link l at time t. In [1], it was proved that the MWS algorithm is throughput-optimal for $w_l(t) = Q_l(t)$. A natural generalization of the MWS algorithm in [26] uses a weight $f(Q_l)$ instead of Q_l with the following properties:

1. $f : [0, \infty] \to [0, \infty]$ is a nondecreasing continuous function with

$$\lim_{x \to \infty} f(x) = \infty.$$

2. Given any $M_1, M_2 > 0$, and $0 < \epsilon < 1$, there must exist a $Q < \infty$ such that for $x > Q$:

$$(1-\epsilon) f(x) \leq f(x - M_1) \leq f(x + M_2) \leq (1+\epsilon) f(x).$$

The following lemma is fairly easy to prove [27] and thus its proof is omitted.

Lemma 2.1. *Suppose f is a strictly concave and increasing function, with $f(0) = 0$, then it satisfies the conditions (1) and (2) above.*

2.2 Description of random access mechanism

The various nodes in the conflict graph share the medium in accordance with a random access mechanism. The random access mechanism in its most general form can be described as follows.

Consider the conflict graph $G(V, E)$ of the network. Denote the neighbors of i by a set $C(i) = \{k \in V : (i, k) \in E\}$. When a node ends an activity period (consisting of possibly several back-to-back packet transmissions), it starts a back-off period. At the end of the back-off period at time t, the node can start a new packet transmission only if none of its neighbors are active, with probability $p_i(Q_i(t))$, $p_i(0) = 0$, and begins a next back-off period otherwise. When a transmission of node i ends at time t, it releases the medium and begins a back-off period with probability $\psi_i(Q_i(t^-))$, or starts the next transmission otherwise, with $\psi_i(1) = 1$. For conciseness, the probabilities $p_i(\cdot)$ and $\psi_i(\cdot)$ will be referred to as *activation* and *de-activation probabilities*, respectively.

The mechanism can be implemented in continuous-time or discrete-time, as we explain in Sections 2.2.1 and 2.2.2.

2.2.1 Continuous-time system

For simplicity, we assume packet arrivals are independent Poisson processes with rate $\lambda = [\lambda_i; i = 1, \ldots, N]$. The packet sizes of node i are exponentially distributed with mean $1/\mu_i$. Let $\rho_i := \lambda_i/\mu_i$.

In continuous-time mechanism, the back-off times of node i are independent and exponentially distributed with mean $1/\nu_i$. Equivalently, node i may be thought of as activating at an exponential rate $r_i(Q_i(t))$, with $r_i(\cdot) = \nu_i p_i(\cdot)$,[1] whenever it senses the channel idle at time t, and de-activating at rate $\hat{r}_i(Q_i(t))$, with $\hat{r}_i(\cdot) = \mu_i \psi_i(\cdot)$, whenever it is active at time t.

There are two special cases worth mentioning that correspond to continuous-time random access schemes considered in the literature before. First, in case $p_i(Q_i) = 1$ and $\psi_i(Q_i) = 0$ for all $Q_i \geq 1$, node i starts a transmission each time a back-off period ends and its neighbors are silent, and does not release the medium until its entire queue has been cleared. This corresponds to the random-capture scheme considered in [28]. In case $\mu_i = 1$, $\nu_i = 1$, $p_i(Q_i) = 1-\psi_i(Q_i)$, and $\psi_i(Q_i) = 1/(1+\exp(\tilde{w}_i(Q_i)))$, node i may be thought

[1] It is in general possible to consider ν_i to be also queue dependent, however, in reality the back-off period cannot be made arbitrarily small so the mentioned version seems more practical.

of as becoming (or continuing to be) active with probability

$$p_i(Q_i(t)) = \frac{\exp(\widetilde{w}_i(Q_i(t)))}{1 + \exp(\widetilde{w}_i(Q_i(t)))}, \tag{2.2}$$

each time a unit-rate Poisson clock ticks and its neighbors are silent. This roughly corresponds to the scheme considered in [27, 16, 13, 14, 15] based on Glauber dynamics with a "weight" function $\tilde{w}_i(Q_i)$. The weight of node i is chosen to be

$$\widetilde{w}_i(Q_i(t)) = \max\left(f(Q_i(t)), \frac{\epsilon}{2N} f(Q_{max}(t)) \right), \tag{2.3}$$

where $Q_{max}(t)$ is the length of the largest queue in the network at time t which is assumed to be known. The function f is a strictly concave and monotonically increasing function, with $f(0) = 0$, as in Lemma 2.1.

2.2.2 Discrete-time system

Time is slotted and packets arrive at each node according to a discrete-time process. Let $A_i(t)$ be the number of packets arriving at node i in time slot t. For simplicity, assume that $\{A_i(t)\}_{t=0}^{\infty}$, for $i = 1, \dots, N$, are independent Bernoulli processes with parameter $\lambda = [\lambda_l; l = 1, \dots, N]$. In each time slot, one packet could be successfully transmitted over a link.

2.2.3 Discrete-time mechanism with one-node update

Consider an activation probability $p_i(t) \equiv p_i(Q_i(t))$ for node i at time slot t as in (2.2) with weight $\tilde{w}$ as in (2.3).

At each time slot t, a node i is chosen uniformly at random, with probability $\frac{1}{N}$, then

(i) If all the neighbors of i are silent, i.e., $x_j(t-1) = 0$ for all $j \in C(i)$, then $x_i(t) = 1$ with probability $p_i(t)$, and $x_i(t) = 0$ with probability $\bar{p}_i(t) = 1 - p_i(t)$. Otherwise, $x_i(t) = 0$.

(ii) $x_j(t) = x_j(t-1)$ for all $j \neq i$.

2.2.4 Discrete-time mechanism with multi-node update

The previous mechanism is based on Glauber-dynamics with one node update at each time. For distributed implementation, we need a randomized mechanism to select a node uniformly at each time slot. We use the Q-CSMA idea [20] to perform the link selection as follows. Each time slot is divided into a control slot and a data slot. In the control slot, each node i that wishes to transmit data transmits a short message called INTENT message with some probability a_i. Those nodes that transmit INTENT messages and do not hear any INTENT messages from the neighboring nodes constitute a decision schedule. In the data slot, each node i that is included in the decision schedule can transmit a data packet with probability $p_i(t)$, as in (2.2), only if none of its neighbors has been transmitting in the previous data slot (see the description of the algorithm below).

(i) In the control slot, randomly select a decision schedule $m(t) \subseteq \mathcal{M}$ by using access probabilities $\{a_i\}_{i=1}^N$.

(ii) $\forall$ i in $m(t)$:

If no links in $C(i)$ were active in the previous data slot, i.e., $\sum_{j \in C(i)} x_j(t-1) = 0$:

- $x_i(t) = 1$ with probability $p_i(t)$.
- $x_i(t) = 0$ with probability $\bar{p}_i(t) = 1 - p_i(t)$.

Else $x_i(t) = 0$.

$\forall i \notin m(t)$: $x_i(t) = x_i(t-1)$.

(iii) In the data slot, use $X(t)$ as the transmission schedule.

Chapter 3

Stability of Random Access for Tame Weight Functions

Consider the continuous-time/discrete-time random access mechanism, described in Chapter 2, based on activation probability (2.2) and weight (2.3). We are interested to determine under what conditions the system is stable, i.e., the process $\{(X(t), Q(t))\}_{t\geq 0}$ is positive-recurrent, for all arrival rates λ in the capacity region Λ.

3.1 Statement of stability result

The following theorem states the main result regarding the throughput optimality of such random access mechanisms.

Theorem 3.1. *Consider any $\epsilon > 0$. The random access mechanism can stabilize the network for any $\lambda \in (1-\epsilon)\Lambda$, if the weight function is chosen to be in the form of $f(x) = \frac{\log(1+x)}{g(x)}$. The function $g(x)$ is a strictly increasing function chosen such that f is a strictly concave increasing function. In particular, the algorithm with the following weight functions is throughput-optimal: $f(x) = \frac{\log(1+x)}{\log(e+\log(1+x))}$ or $f(x) = \log^{1-\theta}(1+x)$ for any arbitrary small $\theta > 0$.*

To determine the weight at each node i, $Q_{max}(t)$ is needed. Instead, each node i can maintain an estimate of it. We can use a gossip procedure, as suggested in [13], to estimate $Q_{max}(t)$, and use Lemma 2 of [13] to complete the stability proof. So we do not pursue this issue here. In practical networks $\frac{\epsilon}{2N}\log(1+Q_{max})$ is small and we can use the weight function f directly, and thus, there may not be any need to know $Q_{max}(t)$.

Corollary 3.1. *Under the weight function f specified in Theorem 5.2, the discrete-time mechanism with multi-node update can stabilize the network for any $\lambda \in (1-\epsilon)\Lambda$.*

3.2 Proof of stability result

We present the proof for the discrete-time model. The proof of the continuous-time model follows similarly. The queue dynamics for each link l is given by

$$Q_l(t) = (Q_l(t-1) - x_l(t))^+ + A_l(t),$$

for $t \geq 0$ and $l = 1, ..., N$ where $(\cdot)^+ = \max\{\cdot, 0\}$. For notational convenience, we define

$$\begin{aligned} w_i(t) &= f(Q_i(t)), \\ w_{min}(t) &= \frac{\epsilon}{2N} f(Q_{max}(t)). \end{aligned}$$

Before we start the proof, some preliminaries, regarding stationary distribution and mixing time of Glauber dynamics, are needed.

3.2.1 Preliminaries

Consider a time-homogeneous discrete-time Markov chain over the finite state space $\mathcal{M}$. For simplicity, we index the elements of $\mathcal{M}$ by $1, 2, ..., r$, where $r = |\mathcal{M}|$. Assume the Markov chain is irreducible and aperiodic, so that a unique stationary distribution $\pi = [\pi(1), ..., \pi(r)]$ always exists, with $\pi(i) > 0$ for all $1 \leq i \leq r$.

Distance between probability distributions

First, we introduce two convenient norms on $\mathbb{R}^r$ that are linked to the stationary distribution. Let $\ell^2(\pi)$ be the real vector space $\mathbb{R}^r$ endowed with the scalar product

$$\langle z, y \rangle_\pi = \sum_{i=1}^{r} z(i) y(i) \pi(i).$$

Then, the norm of z with respect to π is defined as

$$\|z\|_\pi = \left(\sum_{i=1}^{r} z(i)^2 \pi(i) \right)^{1/2}.$$

We shall also use $\ell^2(\frac{1}{\pi})$, the real vector space $\mathbb{R}^r$ endowed with the scalar product

$$\langle z, y\rangle_{\frac{1}{\pi}} = \sum_{i=1}^{r} z(i)y(i)\frac{1}{\pi(i)},$$

and its corresponding norm. For any two strictly positive probability vectors μ and π, the following relationship holds

$$\|\mu - \pi\|_{\frac{1}{\pi}} = \|\frac{\mu}{\pi} - 1\|_{\pi} \geq 2\|\mu - \pi\|_{TV},$$

where $\|\pi - \mu\|_{TV}$ is the total variation distance

$$\|\pi - \mu\|_{TV} = \frac{1}{2}\sum_{i=1}^{r} |\pi(i) - \mu(i)|.$$

Glauber dynamics

Consider a graph $G(V, E)$. Glauber dynamics is a Markov chain to generate the independent sets of G. So, the state space $\mathcal{M}$ consists of all independent sets of G. Let $|V| = N$. Given a weight vector $\tilde{W} = [\tilde{w}_1, \tilde{w}_2, ..., \tilde{w}_N]$, at each time t, a node i is chosen uniformly at random, with probability $\frac{1}{N}$, then

(i) If $x_j(t-1) = 0$ for all nodes $j \in \mathcal{N}(i)$, then $x_i(t) = 1$ with probability $\frac{\exp(\tilde{w}_i)}{1+\exp(\tilde{w}_i)}$, or $x_i(t) = 0$ with probability $\frac{1}{1+\exp(\tilde{w}_i)}$.
Otherwise, $x_i(t) = 0$.

(ii) $x_j(t) = x_j(t-1)$ for all $j \neq i$.

The corresponding Markov chain is irreducible, aperiodic, and reversible over $\mathcal{M}$, and its stationary distribution is given by

$$\pi(X) = \frac{1}{Z}\exp\left(\sum_{i\in X} \tilde{w}_i\right); \quad X \in \mathcal{M}, \tag{3.1}$$

where Z is the normalizing constant (It is easy to check that (3.1) indeed satisfies the detailed balance equations and thus is the stationary distribution).

The random access mechanism with one-node update uses a time-varying version of the above Glauber dynamics, where the weights change with time.

This yields a time-inhomogeneous Markov chain but we will see that, for the proper choice of weights, it behaves similarly to the Glauber dynamics.

Mixing time of Glauber dynamics

The convergence to steady state distribution is geometric with a rate equal to the *second largest eigenvalue modulus* (SLEM) of the transition probability matrix as it is described next (see e.g. Chapter 6 of [29]).

Lemma 3.1. *Let P be an irreducible, aperiodic, and reversible transition matrix on the finite state space $\mathcal{M}$ with the stationary distribution π. Then, the eigenvalues of P are ordered in such a way that*

$$\lambda_1 = 1 > \lambda_2 \geq ... \geq \lambda_r > -1,$$

and for any initial probability distribution μ_0 on $\mathcal{M}$, and for all $n \geq 1$

$$\|\mu_0 \mathbf{P}^n - \pi\|_{\frac{1}{\pi}} \leq \sigma^n \|\mu_0 - \pi\|_{\frac{1}{\pi}}, \tag{3.2}$$

where $\sigma = \max\{\lambda_2, |\lambda_r|\}$ is the SLEM of $\mathbf{P}$.

The following Lemma gives an upper bound on the SLEM $\sigma(P)$ of Glauber dynamics.

Lemma 3.2. *For the Glauber dynamics with the weight vector $\tilde{W}$ on a graph $G(V, E)$ with $|V| = N$,*

$$\sigma \leq 1 - \frac{1}{16^N \exp(4N\tilde{w}_{max})},$$

where $\tilde{w}_{max} = \max_{i \in V} \tilde{w}_i$.

The proof is provided at the end of the chapter. We define *the mixing time* as $T = \frac{1}{1-\sigma}$, so

$$T \leq 16^N \exp(4N\tilde{w}_{max}). \tag{3.3}$$

Simple calculation, based on Lemma 3.1, reveals that the amount of time needed to get close to the stationary distribution is proportional to T.

3.2.2 A key lemma

At any time slot t, given the weight vector $\tilde{W}(t) = [\tilde{w}_1(t), ..., \tilde{w}_N(t)]$, the MWS algorithm should solve

$$\max_{X \in \mathcal{M}} \sum_{i \in X} \tilde{w}_i(t),$$

instead, our algorithm tries to simulate a distribution

$$\pi_t(X) = \frac{1}{Z} \exp(\sum_{i \in X} \tilde{w}_i(t)); \quad X \in \mathcal{M}, \tag{3.4}$$

i.e., the stationary distribution of Glauber dynamics with the weight vector $\tilde{W}(t)$ at time t.

Let P_t denote the transition probability matrix of Glauber dynamics with the weight vector $\tilde{W}(t)$. Also let μ_t be the true probability distribution of the inhomogeneous-time chain, over the set of schedules $\mathcal{M}$, at time t. Therefore, we have $\mu_t = \mu_{t-1} P_t$. Let π_t denote the stationary distribution of the time-homogeneous Markov chain with $P = P_t$ as in (3.1). By choosing proper w_{min} and $f(.)$, we aim to ensure that μ_t and π_t are close enough, i.e.,

$$\|\pi_t - \mu_t\|_{TV} \leq \delta/4$$

for some δ arbitrary small.

Let $w_{max}(t) = f(Q_{max}(t))$. The following lemma gives a sufficient condition under which the probability distribution of the inhomogeneous Markov chain is close to the stationary distribution of the homogeneous chain.

Lemma 3.3. *Given any $\delta > 0$, $\|\pi_t - \mu_t\|_{TV} \leq \frac{\delta}{4}$ holds for all $t \geq t^*$, if*

$$\alpha_t T_{t+1} \leq \delta/16 \textit{ for all } t > 0, \tag{3.5}$$

where

(i) $\alpha_t = 2N f'(f^{-1}(w_{min}(t+1)) - 1)$,

(ii) t^* *is the smallest* t *such that*

$$\sum_{k=1}^{t} \frac{1}{T_k^2} \geq \ln(4/\delta) + N(w_{max}(0) + \log 2)/2, \tag{3.6}$$

and T_k is the mixing time of the Glauber dynamics with the weight vector $\tilde{W}(k)$.

The proof of Lemma 3.3 is provided at the end of the chapter. Lemma 3.3 states a condition under which $\|\pi_t - \mu_t\|_{TV} \leq \frac{\delta}{4}$ for all $t \geq t^*$. The key idea in the proof is that, for α_t small, the weights change at the rate α_t while the system responds to these changes at the rate $1/T_{t+1}$. Condition (3.5) is to ensure that the weight dynamics are slow enough compared to response time of the chain such that the chain remains close to its equilibrium (stationary distribution). Now assume that (3.5) holds only when $\|Q(t)\| \geq Q_{th}$[1] for a constant $Q_{th} > 0$. Let t_1 be the first time that $\|Q(t)\|$ hits Q_{th}. Then, after that, it takes t^* time slots for the chain to get close to π_t if $\|Q(t)\|$ remains above Q_{th} for $t_1 \leq t \leq t_1 + t^*$. Alternatively, we can say that $\|\pi_t - \mu_t\|_{TV} \leq \frac{\delta}{4}$ if $\|Q(t)\| \geq Q_{th} + t^*$ since at each time slot at most one departure can happen and this guarantees that $\|Q(t)\| \geq Q_{th}$ for, at least, the past t^* time slots. This immediately implies the following Lemma that we will use in the proof of the main result.

Lemma 3.4. *Given any $\delta > 0$, $\|\pi_t - \mu_t\|_{TV} \leq \frac{\delta}{4}$ holds when $\|Q(t)\| \geq Q_{th} + t^*$, if there exists a Q_{th} such that*

$$\alpha_t T_{t+1} \leq \delta/16 \text{ whenever } \|Q(t)\| > Q_{th}, \tag{3.7}$$

where

(i) $\alpha_t = 2N f'(f^{-1}(w_{min}(t+1)) - 1)$

(ii) $T_t \leq 16^N \exp(4N w_{max}(t))$

(iii) *t^* is the smallest t such that*

$$\sum_{k=t_1:\|Q(t_1)\|=Q_{th}}^{t_1+t^*} \frac{1}{T_k^2} \geq \ln(4/\delta) + N(f(Q_{th}) + \log 2)/2. \tag{3.8}$$

In Lemma 3.4, condition (ii) is based on the upper bound of (3.3) and the fact that $\tilde{w}_{max}(t) = w_{max}(t)$.

[1] In this section, $\|y\| = \|y\|_\infty = \max_i y_i(t) = y_{max}$.

In other words, Lemma 3.4 states that when queue lengths are large, the observed distribution of the schedules is close to the desired stationary distribution.

Remark 3.1. *We will later see that, to satisfy condition (3.7) and to find a finite t^* satisfying (3.8) in Lemma 3.4, the function $f(.)$ cannot be faster than* $\log(.)$. *In fact, the function f must be slightly slower than* $\log(.)$ *to ensure a finite t^* always exists.*

Remark 3.2. *The above Lemma is a generalization of Lemma* 12 *(Network Adiabatic Theorem) of [13]. Here we consider general functions $f(.)$, whereas [13] considers a particular function* $\log\log(.)$. *The generalization allows us to use functions which are close to* $\log(.)$ *and perform much better than* $\log\log(.)$ *in simulations. The proof of Lemma 3.3 is presented in Section 3.5.*

3.2.3 Throughput optimality

We will use the following Lemma [26] to prove the throughput-optimality of the algorithm.

Lemma 3.5. *For a scheduling algorithm, if given any $0 < \epsilon < 1$ and $0 < \delta < 1$, there exists a $B(\delta, \epsilon) > 0$ such that: in any time slot t, with probability larger than $1-\delta$, the scheduling algorithm chooses a schedule $X(t) \in \mathcal{M}$ that satisfies*

$$\sum_{i \in X(t)} w_i(t) \geq (1-\epsilon) \max_{Y \in \mathcal{M}} \sum_{i \in Y} w_i(t)$$

whenever $\|\mathbf{Q}(t)\| > B(\delta, \epsilon)$, then the scheduling algorithm is throughput-optimal.

Remark 3.3. *Throughput optimality in Lemma 3.5 means that, for all the rates inside the capacity region, system will be stable in the mean (see [26] for more details), i.e.,*

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} E\left[\left(\sum_{i=1}^{N} f^2(Q_i(t))\right)^{\frac{1}{2}}\right] < \infty. \tag{3.9}$$

In our setting, the queuing system is an irreducible and aperiodic Markov chain, and therefore stability-in-the mean property (3.9) implies that the Markov chain is also positive recurrent [30] .

Let $w^*(t) = \max_{X \in \mathcal{M}} \sum_{i \in X} w_i(t)$. Let us define the following set:

$$\chi_t = \{Y \in \mathcal{M} : \sum_{i \in Y} w_i(t) < (1-\epsilon) w^*(t)\}.$$

Therefore, we need to show that

$$\mu_t(\chi_t) = \sum_{Y \in \chi_t} \mu_t(Y) \leq \delta,$$

for $\|Q(t)\|$ large enough. Suppose $f(.)$ and w_{min} are chosen such that $\alpha_t T_{t+1} \leq \delta/16$ whenever $\|Q(t)\| > Q_{th}$ for some constant $Q_{th} > 0$ to be determined later. Then, it follows from Lemma 3.4 that whenever $\|Q(t)\| > Q_{th} + t^*$,

$$2\|\mu_t - \pi_t\|_{TV} \leq \delta/2,$$

and consequently,

$$\sum_{Y \in \mathcal{M}} |\mu_t(Y) - \pi_t(Y)| \leq \delta/2.$$

Thus,

$$\begin{aligned} |\sum_{Y \in \chi_t} (\mu_t(Y) - \pi_t(Y))| &\leq \sum_{Y \in \chi_t} |\mu_t(Y) - \pi_t(Y)| \\ &\leq \delta/2, \end{aligned}$$

which yields

$$\sum_{Y \in \chi_t} \mu_t(Y) \leq \sum_{Y \in \chi_t} \pi_t(Y) + \delta/2.$$

Therefore, to ensure that $\sum_{Y \in \chi_t} \mu_t(Y) \leq \delta$, it suffices to have

$$\sum_{Y \in \chi_t} \pi_t(Y) \leq \delta/2.$$

But

$$\sum_{Y \in \chi_t} \pi_t(Y) = \sum_{Y \in \chi_t} \frac{1}{Z_t} \exp(\sum_{i \in Y} \widetilde{w}_i(t)),$$

where

$$\widetilde{w}_i(t) = \max\{w_i(t),\ w_{min}(t)\} \leq w_i(t) + w_{min}(t).$$

So

$$\begin{aligned}
\sum_{Y\in\chi_t}\pi_t(Y) &\leq \sum_{Y\in\chi_t}\frac{1}{Z_t}\exp(\sum_{i\in Y}(w_i(t)+w_{min}(t))) \\
&= \sum_{Y\in\chi_t}\frac{1}{Z_t}\exp(\sum_{i\in Y}w_i(t))\exp(|Y|w_{min}(t)) \\
&\leq \sum_{Y\in\chi_t}\frac{1}{Z_t}\exp((1-\epsilon)w^*(t))\exp(Nw_{min}(t)),
\end{aligned}$$

and

$$Z_t=\sum_{Y\in\mathcal{M}}\exp(\sum_{i\in Y}\widetilde{w}_i(t))>\sum_{Y\in\mathcal{M}}\exp(\sum_{i\in Y}w_i(t))>e^{w^*(t)}.$$

Therefore,

$$\sum_{Y\in\chi_t}\pi_t(Y) \leq 2^N\exp(Nw_{min}(t)-\epsilon w^*(t)),$$

and $w^*(t)\geq w_{max}(t)$. So, it suffices to have

$$2^N\exp(Nw_{min}(t)-\epsilon w_{max}(t))\leq \delta/2,$$

when $\|Q(t)\|>Q_{th}+t^*$. The choice of $w_{min}(t)=\frac{\epsilon}{2N}w_{max}(t)$, satisfies the above condition for $\|Q(t)\|>B$, where

$$B=\max\left\{Q_{th}+t^*, f^{-1}\left(\frac{N\log 2+\log\frac{2}{\delta}}{\epsilon/2}\right)\right\}. \tag{3.10}$$

3.2.4 A class of weight functions with the maximum throughput property

In this section, we describe a family of weight functions f that yield maximum throughput.

The function f needs to satisfy Lemma 3.4. Roughly speaking, since the mixing time T is exponential in w_{max}, $f'(f^{-1}(w_{min}))$ must be in the form of $e^{-w_{min}}$; otherwise it will be impossible to satisfy $\alpha_t T_{t+1}<\delta/16$ for any arbitrarily small δ as $\|Q(t)\|\to\infty$. The only function with such a property is the log(.) function. In fact, it turns out that f must grow slightly slower than log(.) as we show next to satisfy (3.7), and to ensure the existence of a finite t^* in Lemma 3.4.

Consider weight functions of the form $f(x) = \frac{\log(1+x)}{g(x)}$ where $g(x)$ is a strictly increasing function, chosen such that f satisfies the conditions of Lemma 2.1. For example, by choosing functions that grow much slower than $\log(1+x)$, like $g(x) = \log(e + \log(1+x))$, we can make $f(x)$ behave approximately like $\log(1+x)$ for large ranges of x.

Assume $g(0) \geq 1$, then

$$f'(x) \leq \frac{1}{1+x}. \tag{3.11}$$

The inverse of f cannot be expressed explicitly, however, it can be written as

$$f^{-1}(x) = \exp(xg(f^{-1}(x))) - 1. \tag{3.12}$$

Therefore,

$$\begin{aligned} f'(f^{-1}(w_{min}) - 1) &\leq \frac{1}{f^{-1}(w_{min})} && (3.13)\\ &= \frac{1}{\exp(w_{min}g(f^{-1}(w_{min}))) - 1}. && (3.14)\end{aligned}$$

Using (3.14), the conditions of Lemma 3.4 are satisfied if there exists a Q_{th} large enough such that

$$2N16^N \exp(4Nw_{max})\frac{1}{\exp(w_{min}g(f^{-1}(w_{min}))) - 1} \leq \delta/16, \tag{3.15}$$

for $\|Q(t)\| \geq Q_{th}$.

Using (3.12) and noting that $w_{min} = \frac{\epsilon}{2N}w_{max}$, (3.15) can be written as

$$2N16^N \exp\left(w_{min}\left[\frac{8N^2}{\epsilon} - g(f^{-1}(w_{min}))\right]\right)\left(1 + \frac{1}{f^{-1}(w_{min})}\right) \leq \delta/16. \tag{3.16}$$

Consider fixed, but arbitrary, N and ϵ. As $Q_{max} \to \infty$, $w_{max} \to \infty$, and consequently $w_{min} \to \infty$ and $f^{-1}(w_{min}) \to \infty$. Therefore, the exponent $\frac{8N^2}{\epsilon} - g(f^{-1}(w_{min}))$ is negative for Q_{max} large enough, and thus, there is a threshold Q_{th} such that for all $Q_{max} > Q_{th}$, the condition (3.16) is satisfied. To be more accurate, it suffices to choose

$$Q_{th} = f^{-1}\left(\frac{2N}{\epsilon} \times \max\left\{\log(\frac{64N16^N}{\delta}), f(g^{-1}(\frac{16N^2}{\epsilon}))\right\}\right). \tag{3.17}$$

Then, it follows from Lemma 3.4 that $\|\pi_t - \mu_t\|_{TV} \leq \frac{\delta}{4}$, whenever $\|Q(t)\| > Q_{th} + t^*$.

Remark 3.4. *The assumption $g(0) \geq 1$ is not required, since, as we saw in the above analysis, only the asymptotic behavior of g is important. If we choose Q_{th} large enough such that*

$$g(f^{-1}(w_{min}(t)) - 1) \geq 1, \tag{3.18}$$

when $\|Q(t)\| \geq Q_{th}$, then (3.13) holds and the rest of the analysis follows exactly. In particular, in order to get an explicit formula for f^{-1}, we can choose $g(x) = \log(1+x)^\theta$ for some $0 < \theta < 1$. The weight function for such a g is $f(x) = (\log(1+x))^{1-\theta}$, and f^{-1} has the closed form

$$f^{-1}(x) = \exp(x^{\frac{1}{1-\theta}}) - 1.$$

Then (3.17) yields

$$Q_{th} = \exp\left(\max\left\{\frac{2N}{\epsilon}\log(\frac{64N16^N}{\delta}), \frac{2N}{\epsilon}(\frac{16N^2}{\epsilon})^{\frac{1}{\theta}}\right\}^{\frac{1}{1-\theta}}\right). \tag{3.19}$$

It is easy to check that for $Q(t) \geq \exp\left((\frac{2N}{\epsilon})^{\frac{1}{1-\theta}}\log(1+e)\right)$, $w_{min}(t) \geq f(e)$ which satisfies (3.18). Therefore, obviously, (3.18) also holds for Q_{th} of (3.19).

The last step of the proof is to determine the constant B in (3.10), so we need to find t^*. Let t_1 be the first time that $Q_{max}(t)$ hits Q_{th}, then

$$\begin{aligned}
\sum_{k=t_1}^{t_1+t}\frac{1}{T_k^2} &\geq 16^{-2N}\sum_{k=t_1}^{t_1+t} e^{-8Nf(Q_{max}(k))} \\
&= 16^{-2N}\sum_{k=t_1}^{t_1+t} e^{-8N\frac{\log(1+Q_{max}(k))}{g(Q_{max}(k))}} \\
&= 16^{-2N}\sum_{k=t_1}^{t_1+t}(1+Q_{max}(k))^{-\frac{8N}{g(Q_{max}(k))}} \\
&\geq 16^{-2N}\sum_{k=1}^{t}(1+Q_{th}+k)^{-\frac{8N}{g(Q_{th})}} \\
&\geq 16^{-2N}t(1+Q_{th}+t)^{-\frac{8N}{g(Q_{th})}}.
\end{aligned}$$

Therefore, by Lemma 3.4, it suffices to find the smallest t that satisfies

$$16^{-2N} t(1+Q_{th}+t)^{-\frac{8N}{g(Q_{th})}} \geq \log(4/\delta) + \frac{N}{2}\log(2(1+Q_{th})),$$

for a threshold Q_{th} large enough satisfying (3.17). Recall that $g(.)$ is an increasing function, therefore, by choosing Q_{th} large enough, $\frac{8N}{g(Q_{th})}$ can be made arbitrary small. Then a finite t^* always exists since

$$\lim_{t^* \to \infty} t^*(1+Q_{th}+t^*)^{-\frac{8N}{g(Q_{th})}} = \infty.$$

In particular, for the function $f(Q) = (\log(1+Q))^{1-\theta}$, $0 < \theta < 1$, and the choice of Q_{th} in (3.19), we have

$$\frac{8N}{g(Q_{th})} = \frac{8N}{\log(1+Q_{th})^{\theta}} < \frac{\epsilon}{2N}.$$

Note that

$$\frac{t}{(t+1+Q_{th})^{\epsilon/2N}} \geq \frac{t^{1-\epsilon/2N}}{(2+Q_{th})^{\epsilon/2N}},$$

and therefore, it is sufficient to choose t^* to be

$$t^* = \left[(2+Q_{th})^{\frac{\epsilon}{2N}} 16^N \log\left(\frac{4}{\delta}(2(1+Q_{th}))^{N/2}\right)\right]^{\frac{1}{1-\frac{\epsilon}{2N}}}. \tag{3.20}$$

3.2.5 Extension of the proof to the random access with multi-node update

The distributed algorithm is based on multiple node-update (or parallel operating) Glauber dynamics as defined next. Consider the graph $G(V, E)$ as before and a weight vector $\tilde{W} = [\tilde{w}_1, \tilde{w}_2, ..., \tilde{w}_N]$. At each time t, a decision schedule $m(t) \subseteq \mathcal{M}$ is selected at random with positive probability $\alpha(m(t))$. Then, for all $i \in m(t)$,

(i) If $x_j(t-1) = 0$ for all nodes $j \in \mathcal{N}(i)$, then $x_i(t) = 1$ with probability $\frac{\exp(\tilde{w}_i)}{1+\exp(\tilde{w}_i)}$, or $x_i(t) = 0$ with probability $\frac{1}{1+\exp(\tilde{w}_i)}$.
Otherwise, $x_i(t) = 0$.

(ii) $x_j(t) = x_j(t-1)$ for all $j \notin m(t)$.

The Markov chain $X(t)$ is aperiodic and irreducible if $\cup_{m \in \mathcal{M}_0} = V$ [20]. Also, it can be shown that $X(t)$ is reversible, and has the same stationary distribution as the one-node Glauber dynamics in (3.4). Here, we will assume that $\alpha_{min} := \min_m \alpha(m) \geq (1/2)^N$. Then, the mixing time of the chain is charachterized by the following Lemma.

Lemma 3.6. *For the multiple site-update Glauber dynamics with the weight vector $\tilde{W}$ on a graph $G(V, E)$ with $|V| = N$,*

$$T \leq \frac{64^N}{2} \exp(4N\tilde{w}_{max}), \tag{3.21}$$

where $\tilde{w}_{max} = \max_{i \in V} \tilde{w}_i$.

See Section 3.5 for the proof. The distributed algorithm uses a time-varying version of the multiple-node update Glauber dynamics, where the weights change with time. Although the upperbound of Lemma 3.6 is loose, it is sufficient to prove the optimality of the algorithm. The analysis is the same as the argument for the random access with one-node update. Let D and W denote the lengths of the data slot and the control slot. Thus, the distributed algorithm can achieve a fraction $\frac{D}{D+W}$ of the capacity region. In particular, recall the simple randomized machanism, in Section 2.2.4, where each node joins the decison schedule by sending an INTENT message with probability 1/2. Note that in this case $\alpha_{min} \geq (1/2)^N$, and also it sufficies to allocate a short mini-slot at the begining of the slot for the purpose of control. By choosing the data slot to be much larger than the control slot, the algorithm can approach the full capacity.

3.3 Simulation experiments

In this section, we evaluate the performance of different weight functions via simulations. For this purpose, we have considered the grid network of Figure 3.1, which has 16 nodes and 24 links, under one hop interference constraint. Consider the following maximal schedules:

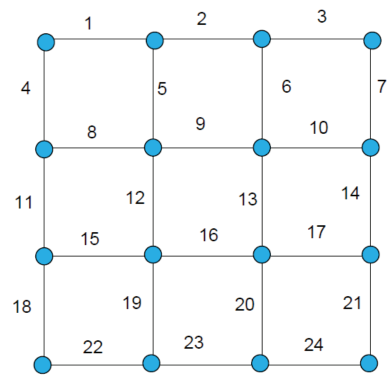

Figure 3.1: A grid network with 24 links.

$$M_1 = \{1, 3, 8, 10, 15, 17, 22, 24\},$$
$$M_2 = \{4, 5, 6, 7, 18, 19, 20, 21\},$$
$$M_3 = \{1, 3, 9, 11, 14, 16, 22, 24\},$$
$$M_4 = \{2, 4, 7, 12, 13, 18, 21, 23\}.$$

With minor abuse of notation, let M_i also be a vector that its i-th element is 1 if $i \in M_i$ and 0 otherwise. We consider arrival rates that are a convex combination of the above maximal schedules scaled by $0 \leq \rho < 1$, e.g.,

$$\lambda = \rho \sum_{i=1}^{4} c_i M_i, \;\; c = [0.2, 0.3, 0.2, 0.3].$$

Note that, as $\rho \to 1$, λ approaches a point on the boundary of the capacity region. We simulate the distributed algorithm, and use the following randomized mechanism, as in [20], similar to IEEE 802.11 DCF standard, to generate the decision schedules in the control slots. At time slot t:

1. Link i selects a random back-off time T_i uniformly in $[0, W-1]$ and waits for T_i control mini-slots.

2. If link i hears an INTENT message from a link in $\mathcal{N}(i)$ before the $(T_i + 1)$-th control mini-slot, i will not be included in $m(t)$ and will not transmit an INTENT message anymore.

3. If link i does not hear an INTENT message from any link in $\mathcal{N}(i)$ before the $(T_i + 1)$-th control mini-slot, it will broadcast an INTENT message at the beginning of the $(T_i + 1)$-th control mini-slot. Then, if there are

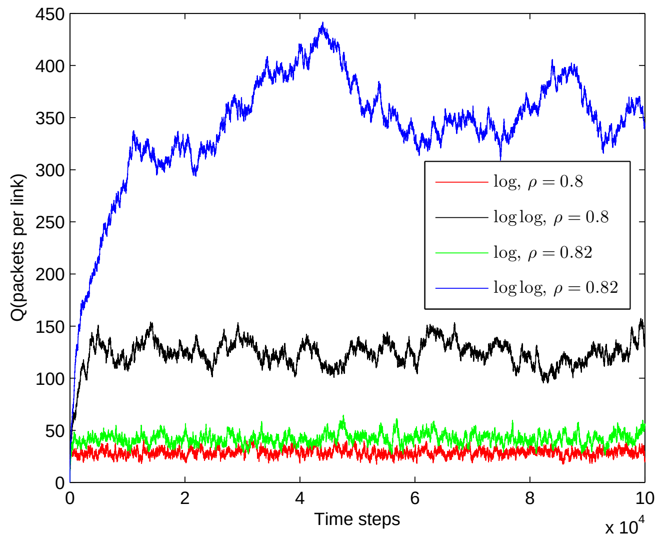

Figure 3.2: The evolution of average queue-length for $\log\log$ and $\frac{\log}{\log\log}$(called log in the plots).

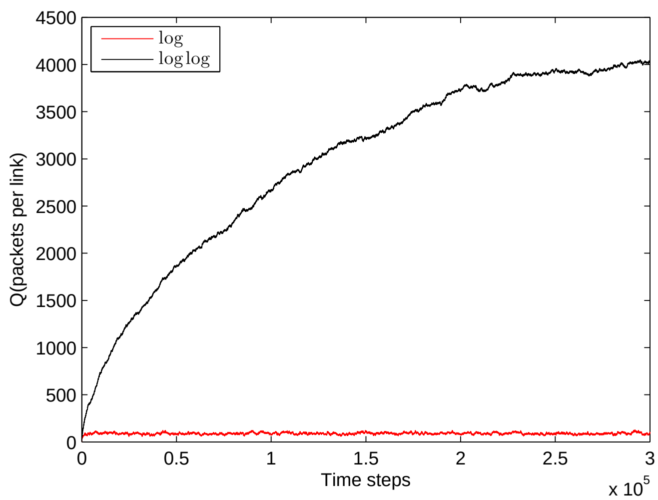

Figure 3.3: The evolution of average queue-lengths for $\rho = 0.85$.

no collisions (i.e., no other link in $\mathcal{N}(i)$ transmits an INTENT message in the same mini-slot), link i will be included in $m(t)$.

Once $m(t)$ is found, the access probabilities are determined as described in the distributed algorithm in Section 2.2.4. Here, we choose $W = 32$ (which is compatible with the back-off window size specified in IEEE 802.11 DCF).

In our simulations, the performance of $\log(1 + x)$ and $\frac{\log(1+x)}{\log(e+\log(1+x))}$ are very close to each other, so in the plots, for brevity, we use the name log

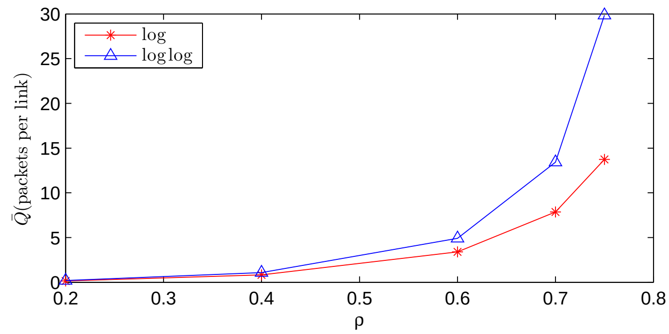

Figure 3.4: Time-average of queue-length per link for low and moderate values of ρ.

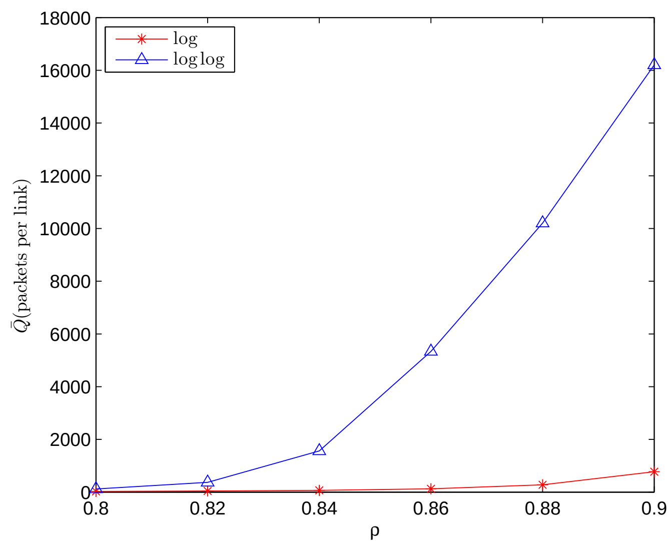

Figure 3.5: Time-average of queue-length per link for high values of ρ.

while the results actually belong to the function $\frac{\log(1+x)}{\log(e+\log(1+x))}$. Figure 3.2 shows the average queue-length evolution (total queue-length divided by the number of links), for the weight functions $f(x) = \frac{\log(1+x)}{\log(e+\log(1+x))}$ and $f(x) = \log\log(e+x)$ and for loadings $\rho = 0.8$ and 0.82. While both functions keep the queues stable, however as it is expected, the average-queue lengths for the weight function $\frac{\log}{\log\log}$ are much smaller than those for $\log\log$. Moreover, $\frac{\log}{\log\log}$ yields a faster convergence to the steady state. The performance gap of two functions, in terms of the average queue-length and the convergence speed, increases significantly for larger loadings; for example see Figure 3.3 for $\rho = 0.85$. Figures 3.4 and 3.5 show the delay performance (time-average queue-length per link) of the two weight functions under different loadings. As it is evident from the figures, log has a significantly smaller delay than what

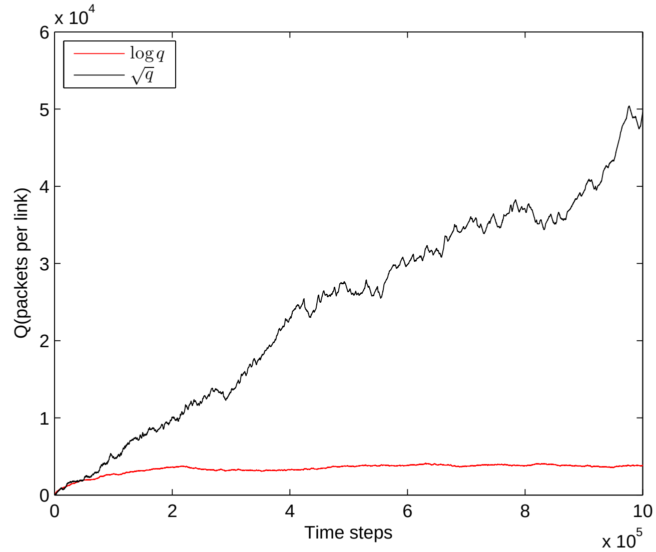

Figure 3.6: The weight function $\sqrt{Q}$ makes the system unstable ($\rho = 0.92$).

is incurred by using the weight $\log\log$. A natural question is whether there exists a function growing faster than log-type functions that still stabilizes any general network. If such a function exists, then one will expect to get a better delay performance. Our conjecture is that, since the mixing time is, in general, exponential in w_{max}, log is the fastest weight function that can make the network change in an adiabatic manner, and hence keep the system close to its equilibrium (stationary distribution). We tried faster weight functions, such as Q and $\sqrt{Q}$, but they resulted in unstable systems (for example see Figure 3.6). In Chapter 4, we will investigate this conjecture rigorously.

3.4 Conclusions

In this chapter, we considered the design of efficient random access algorithms that are throughput optimal and have a good delay performance. Activation probabilities depend on link weights, where the weight of each link is chosen to be an appropriate function $f(\cdot)$ of its queue length. We showed that weight functions of the form $f(Q) = \log(Q)/g(Q)$ (and thus $f(Q) = \log^{1-\epsilon}(Q)$) yield throughput-optimality and low delay performance. The function $g(Q)$ can grow arbitrarily slowly (and ϵ can be arbitrarily small) such that $f(Q) \approx \log(Q)$ for the range of practical queue lengths.

3.5 Additional proofs

Proof of Lemma 3.2. The upper-bound in Lemma 3.2 is based on the conductance bound [31, 29]. First, for a nonempty set $B \subset E$, define the following:

$$\pi(B) = \sum_{i \in B} \pi(i),$$

$$F(B) = \sum_{i \in B, j \in B^c} \pi(i) p_{ij}.$$

Then, conductance is defined as

$$\phi(P) = \inf_{B:\pi_t(B) \leq 1/2} \frac{F(B)}{\pi(B)}.$$

Lemma 3.7. *(Conductance Bounds)*

$$1 - 2\phi(P) \leq \lambda_2 \leq 1 - \frac{\phi^2(P)}{2}.$$

The conductance can be further lower bounded as follows:

$$\begin{aligned} \phi(P) &= \inf_{B:\pi(B) \leq 1/2} \frac{\sum_{X \in B, Y \in B^c} \pi(X) P(X,Y)}{\pi(B)} \\ &\geq 2 \inf_{B \subseteq \mathcal{M}} \sum_{X \in B, Y \in B^c} \pi(X) P(X,Y) \\ &\geq 2 \min_x \pi(X) \min_{X \neq Y} P(X,Y). \end{aligned}$$

For our Glauber dynamics, the stationary distribution is lower bounded by

$$\pi(Y) \geq \frac{1}{\sum_Y \exp(\sum_{i \in Y} \tilde{w}_i)} \geq \frac{1}{|\mathcal{M}| \exp(N \tilde{w}_{max})}.$$

In addition, X and Y can differ in at only one site, and it is easy to see that

$$P(X,Y) \geq \frac{1}{N} \frac{1}{1 + \exp(\tilde{w}_{max})}.$$

So

$$\begin{aligned}\phi(P) &\geq \frac{1}{N2^{N-1}(1+\exp(w_{max}))\exp(Nw_{max})}\\ &\geq \frac{1}{N2^N\exp((N+1)w_{max})}.\end{aligned}$$

Therefore,

$$\begin{aligned}\lambda_2(P) &\leq 1-\frac{1}{2N^2 4^N\exp(2(N+1)w_{max})}\\ &\leq 1-\frac{1}{16^N\exp(4Nw_{max})}.\end{aligned}$$

By Gershgorin's theorem (e.g. see the appendix of [29]), for a stochastic matrix $[P_{ij}]$,

$$\lambda_r \geq -1+2\min_i P_{ii}.$$

For our Glauber dynamics,

$$\begin{aligned}P_{YY} &\geq 1-\frac{1}{N}\sum_{i\in Y}\frac{1}{1+\exp(\widetilde{w}_i)}-\frac{1}{N}\sum_{i\in V\setminus Y}\frac{\exp(\widetilde{w}_i)}{1+\exp(\widetilde{w}_i)}\\ &\geq 1-\frac{1}{N}\sum_{i=1}^{N}\frac{\exp(w_{max})}{1+\exp(w_{max})}\\ &= \frac{1}{1+\exp(w_{max})}.\end{aligned}$$

So,

$$\lambda_r \geq -1+\frac{2}{1+\exp(w_{max})}=\frac{1-\exp(w_{max})}{1+\exp(w_{max})}.$$

Therefore,

$$\max\{\lambda_2,|\lambda_r|\}=\lambda_2,$$

and the SLEM of P is upperbounded by

$$\sigma \leq 1-\frac{1}{16^N\exp(4Nw_{max})}. \tag{3.22}$$

Consequently

$$T \leq 16^N\exp(4Nw_{max}). \tag{3.23}$$

□

Proof of Lemma 3.3. The corresponding stationary distributions at times t and $t+1$ are respectively given by

$$\pi_t(Y) = \frac{1}{Z_t}\exp(\sum_{i\in Y}\widetilde{w}_i(t)),$$

and

$$\pi_{t+1}(Y) = \frac{1}{Z_{t+1}}\exp(\sum_{i\in Y}\widetilde{w}_i(t+1)).$$

So

$$\frac{\pi_{t+1}(Y)}{\pi_t(Y)} = \frac{Z_t}{Z_{t+1}}\exp(\sum_{i\in Y}\widetilde{w}_i(t+1)-\widetilde{w}_i(t)),$$

where

$$\begin{aligned}\frac{Z_t}{Z_{t+1}} &= \frac{\sum_{Y\in\mathcal{M}}\exp(\sum_{i\in Y}\widetilde{w}_i(t))}{\sum_{Y\in\mathcal{M}}\exp(\sum_{i\in Y}\widetilde{w}_i(t+1))}\\ &\le \max_Y \exp(\sum_{i\in Y}\widetilde{w}_i(t)-\widetilde{w}_i(t+1))\\ &\le \exp(\sum_{i=1}^{N}(\widetilde{w}_i(t)-\widetilde{w}_i(t+1))).\end{aligned}$$

Let Q_t^* denote $f^{-1}(w_{min}(t))$, and $\widetilde{Q}(t) = \max\{Q_t^*, Q(t)\}$, where $Q(t)$ is the vector of queue lengths at time t. Recall that f is a concave and increasing function. Hence,

$$\widetilde{w}_i(t+1)-\widetilde{w}_i(t) = f(\widetilde{Q}_i(t+1))-f(\widetilde{Q}_i(t)) \le f'(\widetilde{Q}_i(t))(\widetilde{Q}_i(t+1)-\widetilde{Q}_i(t)) \le f'(\widetilde{Q}_i(t)).$$

(Note that $Q_i(t+1)$ and $Q_i(t)$ at most differ by one since there can at most one packet arrival or departure in a time slot.) Similarly,

$$\widetilde{w}_i(t)-\widetilde{w}_i(t+1) \le f'(\widetilde{Q}_i(t+1)),$$

and thus,

$$\frac{\pi_{t+1}(Y)}{\pi_t(Y)} \le \exp\left(\sum_{i=1}^{N} f'(\widetilde{Q}_i(t)) + f'(\widetilde{Q}_i(t+1))\right).$$

Similarly, we have

$$\frac{\pi_t(Y)}{\pi_{t+1}(Y)} \leq \exp\left(\sum_{i=1}^{N} f'(\widetilde{Q}_i(t)) + f'(\widetilde{Q}_i(t+1))\right).$$

Note that

$$f'(\widetilde{Q}_i(t)) + f'(\widetilde{Q}_i(t+1)) \leq 2f'(Q^*(t+1) - 1).$$

Therefore, if we define

$$\alpha_t = 2Nf'(Q^*(t+1) - 1), \tag{3.24}$$

then

$$e^{-\alpha_t} \leq \frac{\pi_{t+1}(Y)}{\pi_t(Y)} \leq e^{\alpha_t}. \tag{3.25}$$

The drift in π_t is given by

$$\begin{aligned}
\|\pi_{t+1} - \pi_t\|^2_{1/\pi_{t+1}} &= \|\frac{\pi_t}{\pi_{t+1}} - 1\|^2_{\pi_{t+1}} \\
&= \sum_Y \pi_{t+1}(Y)(\frac{\pi_t(Y)}{\pi_{t+1}(Y)} - 1)^2 \\
&\leq \sum_Y \pi_{t+1}(Y) \max\{(e^{\alpha_t} - 1)^2, (1 - e^{-\alpha_t})^2\} \\
&\leq \max\{(e^{\alpha_t} - 1)^2, (1 - e^{-\alpha_t})^2\} \\
&= (e^{\alpha_t} - 1)^2,
\end{aligned}$$

for $\alpha_t < 1$. Thus,

$$\|\pi_{t+1} - \pi_t\|_{1/\pi_{t+1}} \leq 2\alpha_t, \tag{3.26}$$

for $\alpha_t < 1$, where

$$\alpha_t = 2Nf'(f^{-1}(w_{min}(t+1)) - 1). \tag{3.27}$$

The distance between the true distribution and the stationary distribution at time t can be bounded as follows. First, by triangle inequality,

$$\begin{aligned}
\|\mu_t - \pi_t\|_{1/\pi_t} &\leq \|\mu_t - \pi_{t-1}\|_{1/\pi_t} + \|\pi_{t-1} - \pi_t\|_{1/\pi_t} \\
&\leq \|\mu_t - \pi_{t-1}\|_{1/\pi_t} + 2\alpha_{t-1}.
\end{aligned}$$

On the other hand,

$$\begin{aligned}
\|\mu_t - \pi_{t-1}\|^2_{1/\pi_t} &= \sum_Y \frac{1}{\pi_t(Y)}(\mu_t(Y) - \pi_{t-1}(Y))^2 \\
&= \sum_Y \frac{\pi_{t-1}(Y)}{\pi_t(Y)} \frac{1}{\pi_{t-1}(Y)}(\mu_t(Y) - \pi_{t-1}(Y))^2 \\
&\leq e^{\alpha_{t-1}} \|\mu_t - \pi_{t-1}\|^2_{1/\pi_{t-1}}.
\end{aligned}$$

Therefore, for $\alpha_t < 1$,

$$\|\frac{\mu_t}{\pi_t} - 1\|_{\pi_t} \leq (1+\alpha_{t-1})\|\mu_t - \pi_{t-1}\|_{1/\pi_{t-1}} + 2\alpha_{t-1}.$$

Suppose $\alpha_t \leq \delta/16$, then $\|\frac{\mu_t}{\pi_t} - 1\|_{\pi_t} \leq \delta/2$ holds for $t > t^*$, if

$$\|\mu_t - \pi_{t-1}\|_{1/\pi_{t-1}} \leq \delta/4,$$

for all $t > t^*$.

Define $a_t = \|\mu_{t+1} - \pi_t\|_{1/\pi_t}$. Then

$$\begin{aligned}
a_{t+1} &= \|\mu_{t+2} - \pi_{t+1}\|_{1/\pi_{t+1}} \\
&= \|\mu_{t+1}P_{t+1} - \pi_{t+1}\|_{1/\pi_{t+1}} \\
&\leq \sigma_{t+1}\|\mu_{t+1} - \pi_{t+1}\|_{1/\pi_{t+1}},
\end{aligned}$$

where σ_{t+1} is the SLEM of P_{t+1}, since (P_{t+1}, π_{t+1}) is reversible. Therefore,

$$a_{t+1} \leq \sigma_{t+1}[(1+\alpha_t)a_t + 2\alpha_t].$$

Suppose $a_t \leq \delta/4$. Defining $T_t = \frac{1}{1-\sigma_t}$, we have

$$a_{t+1} \leq (1 - \frac{1}{T_{t+1}})[\delta/4 + (2+\delta/4)\alpha_t].$$

Thus, $a_{t+1} \leq \delta/4$, if

$$(2+\delta/4)\alpha_t < \frac{1}{T_{t+1}}(\delta/4 + (2+\delta/4)\alpha_t),$$

or equivalently if

$$\alpha_t < \frac{\frac{\delta/4}{T_{t+1}}}{(2+\delta/4)(1-1/T_{t+1})}.$$

But

$$\frac{\frac{\delta/4}{T_{t+1}}}{(2+\delta/4)(1-1/T_{t+1})} > \frac{\frac{\delta/4}{T_{t+1}}}{4(1-1/T_{t+1})} > \frac{\delta}{16}\frac{1}{T_{t+1}},$$

so, it is sufficient to have

$$\alpha_t T_{t+1} \leq \delta/16.$$

Therefore, if there exists a time t^* such that $a_{t^*} \leq \delta/4$, then $a_t \leq \delta/4$ for all $t \geq t^*$. To find t^*, note that $a_t > \delta/4$ for all $t < t^*$. So, for $t < t^*$, we have

$$\begin{aligned}
a_t &\leq (1-\frac{1}{T_t})[(1+\alpha_{t-1})a_{t-1}+2\alpha_{t-1}] \\
&\leq (1-\frac{1}{T_t})[(1+\alpha_{t-1})a_{t-1}+2\alpha_{t-1}4\frac{a_{t-1}}{\delta}] \\
&\leq (1-\frac{1}{T_t})(1+\alpha_{t-1}+\frac{8}{\delta}\alpha_{t-1})a_{t-1} \\
&\leq (1-\frac{1}{T_t})(1+\frac{\delta/16}{T_t}(1+\frac{8}{\delta}))a_{t-1} \\
&\leq (1-\frac{1}{T_t})(1+\frac{1}{T_t})a_{t-1} \\
&= (1-\frac{1}{T_t^2})a_{t-1} \\
&\leq e^{-\frac{1}{T_t^2}}a_{t-1}.
\end{aligned}$$

Thus,

$$a_t \leq a_0 e^{-\sum_{k=1}^{t^*}\frac{1}{T_k^2}},$$

where

$$\begin{aligned}
a_0 &= \|\frac{\mu_1}{\pi_0}-1\|_{\pi_0} \\
&= \|\mu_0 P_0 - \pi_0\|_{1/\pi_0} \\
&\leq \sigma(P_0)\|\mu_0-\pi_0\|_{1/\pi_0} \\
&\leq \sqrt{\frac{1}{\pi_0^{min}}}
\end{aligned}$$

and

$$\begin{aligned}
\pi_0^{min} &= \min_Y \pi_0(Y) \\
&\geq \frac{1}{\sum_Y \exp(\sum_{i\in Y} \widetilde{w}_i(0))} \\
&\geq \frac{1}{|\mathcal{M}| \exp(N w_{max}(0))},
\end{aligned}$$

which yields

$$a_0 \leq (2e^{w_{max}(0)})^{N/2}.$$

Putting everything together, t^* must satisfy

$$(2e^{w_{max}(0)})^{N/2} e^{-\sum_{k=1}^{t^*} \frac{1}{T_k^2}} \leq \delta/4,$$

or as a sufficient condition,

$$\sum_{k=1}^{t^*} \frac{1}{T_k^2} \geq \log(4/\delta) + N(w_{max}(0) + \log 2)/2.$$

$\square$

Proof of Lemma 3.6. Let $\mathcal{M}_0 \subseteq \mathcal{M}$ be the set of all possible decision schedules. Given $X(t) = X$, for some $X \in \mathcal{M}$, the next state/schedule could be $X(t+1) = Y$ with the following transition probability

$$P(X,Y) = \sum_{m\in\mathcal{M}_0 : X\Delta Y \subseteq m} \alpha(m) \prod_{i\in m\backslash(Y\cup\mathcal{N}(X\cup Y))} \frac{1}{1+\exp(\tilde{w}_i)} \prod_{j\in m\cap Y} \frac{\exp(\tilde{w}_j)}{1+\exp(\tilde{w}_j)},$$

where $X\Delta Y = (X\backslash Y) \cup (Y \backslash X)$.

The upper-bound in Lemma 3.6 is based on the conductance bound as in the proof of Lemma 3.2. Recall that the conductance can be lower bounded as follows:

$$\phi(P) \geq \quad 2 \min_X \pi(X) \min_{X\neq Y} P(X,Y).$$

As in the regular Glauber dynamics,

$$\pi(X) \geq \frac{1}{2^N \exp(N w_{max})},$$

and

$$P(X,Y) \geq \alpha_{min} \left(\frac{1}{1+\exp(w_{max})} \right)^N,$$

where $\alpha_{min} = \min_{m \in M_0} \alpha(m) \geq \frac{1}{2^N}$. Hence,

$$\begin{aligned}
\phi(P) &\geq \frac{2}{4^N(1+\exp(w_{max}))^N \exp(N w_{max})} \\
&\geq \frac{2}{8^N \exp(2N w_{max})}.
\end{aligned}$$

Therefore, based on the conductance upperbound,

$$\lambda_2(P) \leq 1 - \frac{2}{64^N \exp(4N w_{max})},$$

and by Gershgorin's theorem,

$$\lambda_r \geq -1 + \frac{2}{2^N(1+\exp(w_{max}))^N}.$$

Therefore,

$$\max\{\lambda_2, |\lambda_r|\} = \lambda_2,$$

and the SLEM of P is upperbounded by

$$\sigma_t \leq 1 - \frac{2}{64^N \exp(4N w_{max})}.$$

Consequently

$$T \leq \frac{64^N}{2} \exp(4N w_{max}). \tag{3.28}$$

□

Chapter 4

Instability of Random Access for Aggressive Weight Functions

In Chapter 3, we showed that to ensure maximum stability, it is sufficient for weights to behave as logarithmic functions of the queue lengths, divided by an arbitrarily slowly increasing, unbounded function. The result indicated that the maximum-stability guarantees are preserved for weight functions that are essentially linear for all practical values of the queue lengths, although asymptotically the growth rate must be slower than any logarithmic function of the queue length. A careful inspection reveals that the proof arguments leave little room to weaken the stated growth condition. Since the growth condition is only a sufficient one, however, it is not clear to what extent it is actually a strict requirement for maximum stability to be maintained.

In this chapter, we explore the scope for using more aggressive weight functions in order to improve the delay performance while preserving the maximum-stability guarantees. Since the earlier proof methods do not easily extend to more aggressive weight functions, we will instead adopt fluid limits where the dynamics of the system are scaled in both space and time. Fluid limits may be interpreted as first-order approximations of the original stochastic process, and provide valuable qualitative insight and a powerful approach for establishing (in)stability properties [32, 33, 34, 35].

As observed in [36], qualitatively different types of fluid limits can arise, depending on the structure of the interference graph, in conjunction with the functional shape of the weight function. For sufficiently *tame* weight functions as in [13, 14, 15, 27], "fast mixing" is guaranteed, where the activity process evolves on a much faster time scale than the scaled queue lengths. Qualitatively similar fluid limits can arise for more *aggressive* weight functions as well, provided the topology is benign in a certain sense, which implies that the maximum-stability guarantees are preserved in those cases. In different regimes, however, aggressive weight functions can cause "sluggish mixing", where the activity process evolves on a much slower time scale than the scaled queue lengths, yielding oscillatory fluid limits that follow random trajectories. It is highly unusual for such random dynamics to occur, as in queueing networks typically the random characteristics vanish and deterministic limits emerge on the fluid scale. A few exceptions are known for various polling-type models as considered in [37, 38, 39].

The random nature of the fluid limits gives rise to several complications in the convergence proofs that are not commonly encountered. Since the random access networks that we consider are fundamentally different from the polling type-models in the above-mentioned references, the fluid limits are qualitatively different as well, and require a substantially different approach to establish convergence. Specifically, we develop an approach based on stopping time sequences to deal with the switching probabilities governing the sample paths of the fluid limit process. While these proof arguments are developed in the context of random access networks, several key components extend far beyond the scope of the present problem. Hence, we believe that the proof constructs are of broader methodological value in handling random fluid limits and of potential use in establishing both stability and instability results for a wider range of models. For example, the methodology that we develop could be easily applied to prove the stability results for the random capture scheme as conjectured in [28].

The possible oscillatory behavior of the fluid limit itself does not necessarily imply that the system is unstable, and in some situations maximum stability is in fact maintained. In other scenarios, however, the fluid limit reflects that more aggressive weight functions may force the system into inefficient states for extended periods of time and produce instability. We will demonstrate instability for weight functions of the form $\gamma \log(\cdot)$, for $\gamma > 1$, but our proof arguments suggest that it can potentially occur for any $\gamma > 0$, in networks with sufficiently many nodes. In other words, the growth conditions for maximum stability depend on the number of nodes, which seems loosely related to results in [40, 41, 42] characterizing how (upper bounds for) the mean queue length and delay scale as a function of the size of the network.

The remainder of the chapter is organized as follows. We introduce fluid limits and discuss the various qualitative regimes in Section 4.1. We then use the fluid limits to demonstrate the potential instability of aggressive activity functions in Sections 4.2 and 4.3. Simulation experiments are conducted in Section 4.4 to support the analytical results. We will focus on the continuous-time model but the results naturally hold for the discrete-time model as well.

Under the continuous-time random access scheme, the process $(X(t), Q(t))$ evolves as a continuous-time Markov process with state space $\mathcal{M} \times \mathbb{N}_0^N$. Transitions (due to arrivals) from a state (X, Q) to $(X, Q + e_i)$ occur at rate λ_i, transitions (due to activations) from a state (X, Q) with $Q_i \geq 1$, $X_i = 0$, and $X_j = 0$ for all neighbors of node i, to $(X + e_i, Q)$ occur at rate $r_i(Q_i) := \nu_i p_i(Q_i)$, transitions (due to transmission completions followed back-to-back by a subsequent transmission) from a state (X, Q) with $X_i = 1$ (and thus $Q_i \geq 1$) to $(X, Q - e_i)$ occur at rate $\mu_i(1 - \psi_i(Q_i))$, transitions (due to transmission completions followed by a back-off period) from a state (X, Q) with $X_i = 1$ (and thus $Q_i \geq 1$) to $(X - e_i, Q - e_i)$ occur at rate $\hat{r}_i(Q_i) := \mu_i \psi_i(Q_i)$.

We are interested to determine under what conditions the system is stable, i.e., the process $\{(X(t), Q(t))\}_{t \geq 0}$ is positive-recurrent. It is easily seen that $(\rho_1, \ldots, \rho_N) \in \Lambda$ is a necessary condition for that to be the case. In Chapter 3, we showed that this condition is in fact also sufficient for weight functions of the form $w_i(Q_i) = \log(1 + Q_i)/g_i(Q_i)$, where $g_i(Q_i)$ is allowed to increase to infinity at an arbitrarily slow rate. Results in [36] suggest that more aggressive choices of the functions $p_i(\cdot)$ and $\psi_i(\cdot)$, which translate into functions $w_i(\cdot)$ that grow faster to infinity, can improve the delay performance. In view of these results, we will be particularly interested in such weight functions $w_i(\cdot)$, where the stability results of Chapter 3 do not apply. In order to examine under what conditions the system will remain stable, we will examine fluid limits for the process $\{(X(t), Q(t))\}_{t \geq 0}$ as introduced in the next section.

4.1 Qualitative discussion of fluid limits

Fluid limits may be interpreted as first-order approximations of the original stochastic process, and provide valuable qualitative insight and a powerful approach for establishing (in)stability properties [32, 33, 34, 35]. In this section we discuss fluid limits for the process $\{(X(t), Q(t))\}_{t \geq 0}$ from a broad perspective, with the aim to informally exhibit their qualitative features in various regimes, and we deliberately eschew rigorous claims or proofs.

4.1.1 Fluid-scaled process

In order to obtain fluid limits, the original stochastic process is scaled in both space and time. More specifically, we consider a sequence of processes $\{(X^{(R)}(t), Q^{(R)}(t))\}_{t\geq 0}$ indexed by a sequence of positive integers R, each governed by similar statistical laws as the original process, where the initial states satisfy $\sum_{i=1}^{N} Q_i^{(R)}(0) = R$ and $Q_i^{(R)}(0)/R \to Q_i$ as $R \to \infty$. The process $\{(X^{(R)}(Rt), \frac{1}{R}Q^{(R)}(Rt))\}_{t\geq 0}$ is referred to as the fluid-scaled version of the process $\{(X^{(R)}(t), Q^{(R)}(t)\}_{t\geq 0}$. Note that the activity process is scaled in time as well but not in space. For compactness, denote $Q^R(t) = \frac{1}{R}Q^{(R)}(Rt)$. Any (possibly random) weak limit $\{q(t)\}_{t\geq 0}$ of the sequence $\{Q^R(t)\}_{t\geq 0}$, as $R \to \infty$, is called a fluid limit.

It is worth mentioning that the above notion of fluid limit based on the continuous-time Markov process is only introduced for the convenience of the qualitative discussion that follows. For all the proofs of fluid limit properties and instability results we will rely on a rescaled linear interpolation of the uniformized jump chain (as will be defined in Section 4.8.3), with a time-integral version of the $X(\cdot)$ component. This construction yields convenient properties of the fluid limit paths and allows us to extend the framework of Meyn [35] for establishing instability results for discrete-time Markov chains. (The original continuous-time Markov process has in fact the same fluid limit properties, but this is not directly relevant in any of the proofs.)

The process $\{(X^{(R)}(Rt), \frac{1}{R}Q^{(R)}(Rt))\}_{t\geq 0}$ comprises two interacting components. On the one hand, the evolution of the (scaled) queue length process $\frac{1}{R}Q^{(R)}(Rt)$ depends on the activity process $X^{(R)}(Rt)$. On the other hand, the evolution of the activity process $X^{(R)}(Rt)$ depends on the queue length process $Q^{(R)}(Rt)$ through the activation and de-activation functions $f_i(\cdot)$ and $g_i(\cdot)$. In many cases, a separation of time scales arises as $R \to \infty$, where the transitions in $X^{(R)}(Rt)$ occur on a much faster time scale than the variations in $Q^R(t) = \frac{1}{R}Q^{(R)}(t)$. Loosely phrased, the evolution of $Q^R(t)$ is then governed by the time-average characteristics of $X^{(R)}(\cdot)$ in a scenario where $Q^R(t)$ is fixed at its instantaneous value.

In other cases, however, the transitions in $X^{(R)}(Rt)$ may in fact occur on a much slower time scale than the variations in $Q^R(t)$, or there may not be a separation of time scales at all. As a result, qualitatively different types of fluid limits can arise, as observed in [36], depending on the mixing properties of the activity process. These mixing properties, in turn, depend on the functional shape of the activation and de-activation probabilities $p_i(\cdot)$ and $\psi_i(\cdot)$, in conjunction with the structure of the interference graph G.

4.1.2 Fast mixing: Stability result revisited

We first consider the case of fast mixing. In this case, the transitions in $X^{(R)}(Rt)$ occur on a much faster time scale than the variations in $Q^R(t)$, and completely average out on the fluid scale as $R \to \infty$. Informally speaking, this entails that the mixing time of the activity process in a scenario with fixed activation rates $r_i(Rq_i)$ and de-activation rates $\hat{r}_i(Rq_i)$ grows slower than R as $R \to \infty$. In order to obtain a rough bound for the mixing time, assume that $r_i(\cdot) \equiv r_i(\cdot)$, $\hat{r}_i(\cdot) \equiv \hat{r}(\cdot)$, and denote $h(x) = r(x)/\hat{r}(x)$. Further suppose that $h(R) \to \infty$ as $R \to \infty$, and $h(aR)/h(R) \to \hat{h}(a)$ as $R \to \infty$, with $\hat{h}(a) > 0$ for any $a > 0$. The latter assumptions are satisfied, for example, when $h(x) = x^\gamma$, $\gamma > 0$, with $\hat{h}(a) = a^\gamma$, or when $h(x) = \log(x)$ with $\hat{h}(a) \equiv 1$. Without proof, we claim that the mixing time then grows at most at rate $r(R)^{m^*-1}\hat{r}(R)^{-m^*}$ as $R \to \infty$, with m^* the cardinality of a maximum-size independent set. Thus, fast mixing behavior is guaranteed when $r(\cdot)$ does not grow too fast, $\hat{r}(\cdot)$ does not decay too fast, or m^* is sufficiently small, e.g.,

(i) $\hat{r}(x) = r$ and $m^* = 1$;

(ii) $r(x) = x^{1/(m^*-1)-\delta}$, $\hat{r}(x) = \hat{r}$, and $m^* \geq 2$;

(iii) $r(x) = r$ and $\hat{r}(x) \geq x^{-1/m^*+\delta}$;

(iv) $r(x) = r$, $\hat{r}(x) = 1/\log(1 + x)$;

(v) $r(x) = \log(1 + x)$ and $\hat{r}(x) = \hat{r}$.

The fluid limit then follows an entirely deterministic trajectory, which is described by a differential equation of the form

$$\frac{\mathrm{d}}{\mathrm{d}t}q_i(t) = \lambda_i - \mu_i u_i(q(t)),$$

as long as $q(t) > 0$ (component-wise), with the function $u_i(\cdot)$ representing the fraction of time that node i is active. We may write

$$u_i(q) = \sum_{s \in S} s_i \pi(s; q),$$

with $\pi(s; q)$ denoting the fraction of time that the activity process resides in state $s \in S$ in a scenario with fixed activation rates $r_j(Rq_j)$ and de-activation rates $\hat{r}_j(Rq_j)$ as $R \to \infty$. Let $S^* = \{s \in S : \sum_{i=1}^{N} s_i = m^*\}$ correspond to the collection of all maximum-size independent sets. Under the above-mentioned assumptions,

$$\begin{aligned}
\pi(s; q) &= \lim_{R \to \infty} \frac{\prod_{i=1}^{N} h(Rq_i)^{s_i}}{\sum_{u \in S^*} \prod_{i=1}^{N} h(Rq_i)^{u_i}} \\
&= \frac{\prod_{i=1}^{N} \hat{h}(q_i)^{s_i}}{\sum_{u \in S^*} \prod_{i=1}^{N} \hat{h}(q_i)^{u_i}} \\
&= \frac{\exp(\sum_{i=1}^{N} s_i \log(\hat{h}(q_i)))}{\sum_{u \in S^*} \exp(\sum_{i=1}^{N} u_i \log(\hat{h}(q_i)))},
\end{aligned}$$

for $s \in S^*$, while $\pi(s; q) = 0$ for $s \notin S^*$. In particular, if $h(x) = x^\gamma$, $\gamma > 0$, then

$$\pi(s; q) = \frac{\prod_{i=1}^{N} q_i^{\gamma s_i}}{\sum_{u \in S^*} \prod_{i=1}^{N} q_i^{\gamma u_i}} = \frac{\exp(\gamma \sum_{i=1}^{N} s_i \log(q_i))}{\sum_{u \in S^*} \exp(\gamma \sum_{i=1}^{N} u_i \log(q_i))},$$

for $s \in S^*$. Also, if $h(x) = \log(1 + x)$, then $\pi(s; q) = 1/|S^*|$ for $s \in S^*$.

When some of the components of q are zero, i.e., some of the queue lengths are zero at the fluid scale, it is considerably harder to characterize $u_i(q)$, since the competition for medium access from the queues that are zero at the fluid scale still has an impact. It may be shown though that

$$\sum_{i=1}^{N} \rho_i \mathrm{I}_{\{q_i > 0\}} \leq (1 - \epsilon) \sum_{i=1}^{N} u_i(q) \mathrm{I}_{\{q_i > 0\}}$$

for some $\epsilon > 0$, assuming that $(\rho_1, \ldots, \rho_N) < \sigma \in \mathcal{C}$. The latter inequality also holds when $q > 0$, noting that then $\sum_{i=1}^N u_i(q) = m^*$, while $\sum_{i=1}^N \rho_i \leq (1-\epsilon)m^*$ for some $\epsilon > 0$.

We conclude that almost everywhere

$$\begin{aligned}
\sum_{i=1}^N \frac{1}{\mu_i} \frac{\mathrm{d}q_i(t)}{\mathrm{d}t} &\leq \sum_{i=1}^N (\rho_i - u_i(q(t)))\mathrm{I}_{\{q_i(t) > 0\}} \\
&\leq -\epsilon \sum_{i=1}^N \rho_i \mathrm{I}_{\{q_i(t) > 0\}},
\end{aligned}$$

as long as $q(t) \neq 0$. This means that $q(t) = 0$ for all $t \geq T$ for some finite $T < \infty$, which implies that the original Markov process is positive-recurrent [32, 34].

For the discrete-time/continuous-time random access mechanism based on Glauber dynamics, $p(x) = 1 - \psi(x)$, and $\psi(x) = 1/(1+\exp(w(x)))$. Thus the mixing time scales as $e^{m^* w(R)}$ as $R \to \infty$, which is consistent with our mixing time calculations in Chapter 3. Thus, fast mixing behavior is guaranteed when

(i) $w(x) = \gamma \log(1 + x)$, and $\gamma \leq \frac{1}{m^*} - \epsilon$;

(ii) $w(x) = \log(1 + x)/g(x)$, when $g(\cdot)$ can grow at an arbitrary slow rate;

(iii) $w(x) = \log^{1-\epsilon}(x)$;

(iv) $w(x) = \log\log(e + x)$.

This agrees with the stability results in Chapter 3 and suggests that these results in fact hold without the need to know the maximum queue size $Q_{\max}$.

Of course, in order to convert the above arguments into an actual stability proof, the informal characterization of the fluid limit needs to be rigorously justified. This is a major challenge, and not the real goal of this chapter, since we aim to demonstrate the opposite, namely that more aggressive activity or de-activation functions can cause instability. Strong evidence of the technical complications in establishing the fluid limits is provided by recent work of Robert and Véber [43]. Their work focuses on the simpler case of a single work-conserving resource (which corresponds to a full interference graph in the present setting) without any back-off mechanism, where the service rates

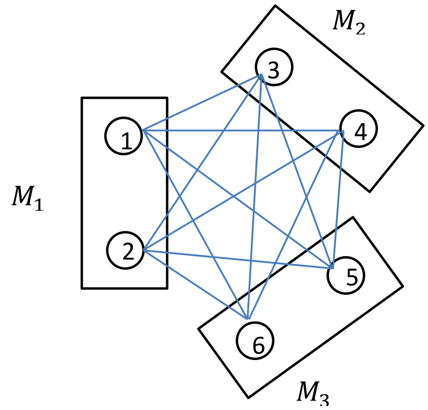

Figure 4.1: The *diamond network*: A complete partite graph with $K = 3$ components, each containing two nodes.

of the various nodes are determined by a logarithmic function of their queue lengths.

4.1.3 Sluggish mixing: Random oscillatory fluid limits

With the aim of demonstrating instability for more aggressive schemes, we now turn to the case of sluggish mixing. In this case, the transitions in $X^{(R)}(Rt)$ occur on a much slower time scale than the variations in $Q^R(t)$, and vanish on the fluid scale as $R \to \infty$, except at time points where some of the queues hit zero. The detailed behavior of the fluid limit in this case depends delicately on the specific structure of the interference graph G and the shape of the functions $r_i(\cdot)$ and $\hat{r}_i(\cdot)$. This prevents a characterization in any degree of generality, and hence we focus attention on some particular scenarios.

In order to show that sluggish mixing behavior itself need not imply instability, we first examine a complete K-partite graph as considered in [28], where the nodes can be partitioned into $K \geq 2$ components. All nodes are connected except those belonging to the same component. Figure 4.1 depicts an example of a complete partite graph with $K = 3$ components, each containing two nodes. We will refer to this network as the *diamond network*, since the edges correspond to those of an eight-faced diamond structure, with the node pairs constituting the three components positioned at the opposite ends of three orthogonal axes.

Denote by $M_k \subseteq \{1, \ldots, N\}$ the subset of nodes belonging the k-th component. Once one of the nodes in component M_k is active, other nodes within M_k can become active as well, but none of the nodes in the other components M_l, $l \neq k$, can be active. The necessary stability condition then takes the form $\rho = \sum_{k=1}^{K} \hat{\rho}_k < 1$, with $\hat{\rho}_k = \max_{i \in M_k} \rho_i$ denoting the maximum traffic intensity of any of the nodes in the k-th component.

Now consider the case that each node operates with an activation function $r(x)$ with $\lim_{x \to \infty} r(x) > 0$ and a de-activation function $\hat{r}(x) = \mathrm{o}(x^{-\gamma})$, with $\gamma > 1$. This subsumes the Glauber dynamics with weight functions of the form $w_i(x) = \gamma \log(1 + x)$, with $\gamma > 1$. The random-capture scheme of [28] is a special case of our activation/decativation functions with $\hat{r}(x) \equiv 0$ for all $x \geq 1$. Since the de-activation rate decays so sharply, the probability of a node releasing the medium once it has started transmitting with an initial queue length of order R, is vanishingly small, until the queue length falls below order R or the total number of transmissions exceeds order R (but the latter implies the former). Hence, in the fluid limit, a node must completely empty its queue almost surely before it releases the medium. Because of the interference constraints, it further follows that once the activity process enters one of the components, it remains there until all the queues in that component have entirely drained (on the fluid scale), and then randomly switches to one of the other components. For conciseness, the fluid limit process is said to be in an M_k-period during time intervals when at least one of the nodes in component M_k is served at full rate (on the fluid scale).

Based on the aforementioned informal observations, we now proceed with a more detailed description of the dynamics of the fluid limit process. We do not aim to provide a proof of the stated properties, since the main goal of this chapter is to demonstrate the potential for instability rather than establish stability. However, the proof arguments that we will develop for a similar but more complicated interference graph in the remainder of this chapter, could easily be applied to provide a rigorous justification of the fluid limit and establish the claimed stability results.

Assume that the system enters an M_k-period at time t, then

(a) It spends a time period $T_k(t) = \max_{i \in M_k} \frac{q_i(t)}{\mu_i - \lambda_i}$ in M_k.

(b) During this period, the queues of the nodes in M_k drain at a linear rate

(or remain zero)

$$q_i(t+u) = \max\{q_i(t) + (\lambda_i - \mu_i)u, 0\}, \ \forall i \in M_k,$$

while the queues of the other nodes fill at a linear rate

$$q_i(t+u) = q_i(t) + \lambda_i u, \ \forall i \notin M_k,$$

for all $u \in [0, T_k(t)]$.

(c) At time $t + T_k(t)$, the system switches to an M_l-period, $l \neq k$, with probability

$$p_{kl}(t + T_k(t)) = \lim_{R\to\infty} \frac{\sum_{i\in M_l} r(Rq_i(t+T_k(t)))}{\sum_{l'\neq k,l} \sum_{i\in M_{l'}} r(Rq_i(t+T_k(t)))}.$$

Thus the fluid limit follows a piece-wise linear sample path, with switches between different periods governed by the transition probabilities specified above. Figure 4.2 depicts an example of the fluid limit sample path for the network of Figure 4.1 with $r(x) = 1$, for all $x \geq 1$.

Now define the Lyapunov function $L(t) := \sum_{k=1}^{K} \hat{q}_k(t)$, with

$$\hat{q}_k(t) = \max_{i\in M_k} q_i(t)/\mu_i.$$

Then,

$$\frac{\mathrm{d}}{\mathrm{d}t} L(t) \leq \sum_{k=1}^{K} \hat{\rho}_k - 1 = \rho - 1 < 0$$

almost everywhere when $\rho < 1$, as long as $L(t) > 0$. Therefore, $L(t) = 0$, and hence $q(t) = 0$, for all $t \geq T$, with $T = \frac{L(0)}{1-\rho} < \infty$, implying stability [32, 34], even though the fluid limit behavior is not smooth at all.

4.2 Fluid limits for broken-diamond network

In Section 4.1 we discussed qualitative features of fluid limits in various scenarios, and in particular for so-called complete partite graphs. We now proceed to consider a "nearly" complete partite graph, and will demonstrate that if some of the edges between two components M_k and M_l are removed (thus

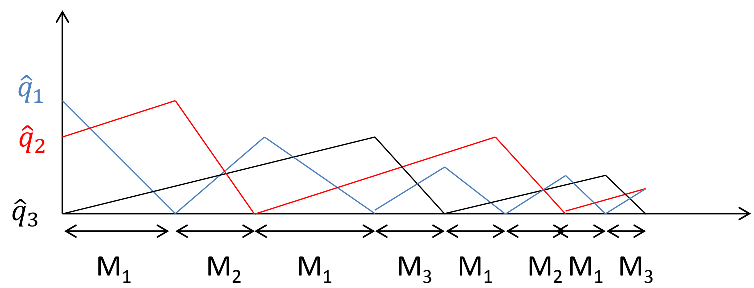

Figure 4.2: A fluid limit sample path for the *diamond network* of Figure 4.1.

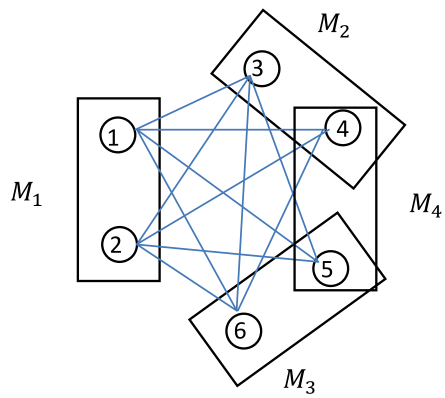

Figure 4.3: The *broken-diamond network*, obtained by removing one edge from the *diamond network* of Figure 4.1, yielding an additional schedule M_4.

reducing interference), the network might become unstable for "aggressive" activation and/or de-activation functions! Specifically, we will consider the *diamond network* of Figure 4.1, and remove the edge between nodes 4 and 5 to obtain a *broken-diamond network* with an additional component/maximal schedule M_4, as depicted in Figure 4.3.

The intuitive explanation for the potential instability may be described as follows. Denote $\rho_0 = \max\{\rho_1, \rho_2\}$, and assume $\rho_3 \geq \rho_4$ and $\rho_6 \geq \rho_5$. It is easily seen that the fraction of time that at least one of the nodes 1, 2, 3 and 6 is served, must be no less than $\rho = \rho_0 + \rho_3 + \rho_6$ in order for these nodes to be stable. During some of these periods nodes 4 or 5 may also be served, but not simultaneously, i.e., schedule M_4 cannot be used. In other words, the system cannot be stable if schedule M_4 is used for a fraction of the time larger than $1 - \rho$. As it turns out, however, when the de-activation function is sufficiently aggressive, e.g., $\hat{r}(x) = \mathrm{o}(x^{-\gamma})$, with $\gamma > 1$ (correspondingly weight functions of the form $w_i(x) = \gamma \log(1 + x)$, $\gamma > 1$), schedule M_4 is in fact persistently used for a fraction of the time that does not tend to 0 as ρ approaches 1, which forces the system to be unstable.

Although the above arguments indicate that invoking schedule M_4 is a recipe for trouble, the reason may not be directly evident from the system dynamics, since no obvious inefficiency occurs as long as the queues of nodes 4 and 5 are nonempty. However, the fact that the Lyapunov function

$$L(t) = \sum_{k=1}^{3} \max_{i \in M_k} q_i(t)$$

may increase while serving nodes 4 and 5, when $q_3(t) \geq q_4(t)$ and $q_6(t) \geq q_5(t)$, is already highly suggestive. (Such an increase is depicted in Figure 4.4 during the M_4 period of the switching sequence $M_1 \to M_2 \to M_1 \to M_4 \to M_3 \to M_1$.) Indeed, serving nodes 4 and 5 may make their queues smaller than those of nodes 3 and 6, leaving these queues to be served by themselves at a later stage, at which point inefficiency inevitably occurs.

In the sequel, the fluid limit process is said to be in a natural state when $q_3(t) \geq q_4(t)$ and $q_6(t) \geq q_5(t)$, with equality only when both sides are zero. We will assume $\lambda_3 > \lambda_4$ and $\lambda_6 > \lambda_5$, and will show that the process must always reside in a natural state after some finite amount of time. As described above, instability is bound to occur when schedule M_4 is used repeatedly for substantial periods of time while the fluid limit process is in a natural state. Since the process is always in a natural state after some finite amount of time, it is intuitively plausible that such events occur repeatedly with positive probability, but a rigorous proof that this leads to instability is far from simple. Such a proof requires detailed analysis of the underlying stochastic process (in our case via fluid limits), and its conclusion crucially depends on the de-activation function. Indeed, the stability results in [27, 16, 13, 14, 15] indirectly indicate that the *broken-diamond network* is *not* rendered unstable for sufficiently cautious de-activation functions.

Just like for the complete partite graphs, the fluid limit process is said to be in an M_1-period when node 1 or node 2 (or both) is served at full rate. The process is in an M_2- or M_3-period when node 3 or 6 is served at full rate, respectively. The process is in an M_4-period when nodes 4 and 5 are both served at full rate simultaneously.

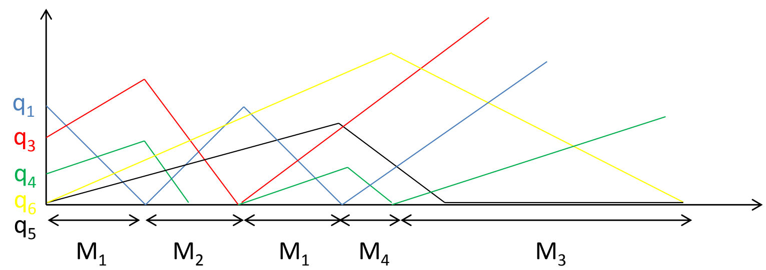

Figure 4.4: A fluid limit sample path for the *broken-diamond network* of Figure 4.3, corresponding to the switching sequence $M_1 \to M_2 \to M_1 \to M_4 \to M_3$.

In Section 4.2.1 we will provide a detailed description of the dynamics of the fluid limit process once it has reached a natural state and entered an M_1-, M_2-, M_3- or M_4-period. The justification for the description follows from a collection of lemmas and propositions which are stated and proved in Sections 4.8.3–4.8.6, with a high-level outline provided in Section 4.2.2. In Section 4.3 we will exploit the properties of the fluid limit process in order to prove that the harmful behavior described above indeed occurs for sufficiently aggressive de-activation functions, implying instability of the fluid limit process as well as the original stochastic process.

4.2.1 Description of the fluid limit process

We now provide a detailed description of the dynamics of the fluid limit process once it has reached a natural state and entered an M_1-, M_2-, M_3- or M_4-period. For sufficiently high load, i.e., ρ sufficiently close to 1, a natural state and such a period occur in uniformly bounded time almost surely for any initial state. As will be seen, for de-activation functions $\hat{r}_i(x) = \mathrm{o}(x^{-\gamma})$ (correspondingly weight functions $w_i(x) = \gamma \log(1+x)$), with $\gamma > 1$, the fluid limit process then follows similar piece-wise linear trajectories, with random switches, as described in the previous section for complete partite graphs and further illustrated in Figure 4.4. For notational convenience, we henceforth assume $\mu_i \equiv 1$, so that $\rho_i \equiv \lambda_i$, for all $i = 1, \dots, N$, and additionally assume activation functions $r_i(x) \equiv 1$, for $x \geq 1$, for all $i = 1, \dots, N$.

M_1-period

Assume the system enters an M_1-period at time t, then

(a) It spends a time period $T_1(t) = \max\left\{\frac{q_1(t)}{1-\rho_1}, \frac{q_2(t)}{1-\rho_2}\right\}$ in M_1.

(b) During this period, the queues of nodes 1 and 2 drain at a linear rate (or remain zero)

$$q_i(t+u) = \max\{q_i(t) - (1-\rho_i)u, 0\}, \text{ for } i = 1, 2,$$

while the queues of nodes 3, 4, 5, 6 fill at a linear rate

$$q_i(t+u) = q_i(t) + \rho_i u, \text{ for } i = 3, 4, 5, 6,$$

for all $u \in [0, T_1(t)]$. In particular, $q_1(t + T_1(t)) = q_2(t + T_1(t)) = 0$.

(c) At time $t + T_1(t)$, the system switches to an M_2-, M_3- or M_4-period with transition probabilities $p_{12} = \frac{3}{8}$, $p_{13} = \frac{3}{8}$, and $p_{14} = \frac{1}{4}$, respectively.

M_2-period

Assume that the system enters an M_2-period at time t, then

(a) The system spends a time period $T_2(t) = \frac{q_3(t)}{1-\rho_3}$ in M_2.

(b) During this period, the queues of nodes 3 and 4 drain (or remain zero)

$$q_i(t+u) = \max\{q_i(t) - (1-\rho_i)u, 0\}, \text{ for } i = 3, 4,$$

while the queues of nodes 1, 2, 5, 6 fill at a linear rate

$$q_i(t+u) = q_i(t) + \rho_i u, \text{ for } i = 1, 2, 5, 6,$$

for all $u \in [0, T_2(t)]$. In particular, $q_3(t + T_2(t)) = 0$.

(c) At time $t + T_2(t)$, the system switches to an M_1- or M_3-period. Note that $\frac{q_3(t)}{1-\rho_3} > \frac{q_4(t)}{1-\rho_4}$ by the assumption that $\lambda_3 > \lambda_4$ and that the process has reached a natural state, so that $q_3(t) > q_4(t)$ (since $q_3(t) = q_4(t) = 0$ cannot occur at the start of an M_2-period). Thus node 4 has emptied before time $t+T_2(t)$, and remained empty (on the fluid scale) since then, precluding a switch to an M_4-period except for a negligible duration on the fluid scale), only allowing the system to switch to either an M_1- or M_3-period. The corresponding transition probabilities can be formally expressed in terms of certain stationary distributions, but are difficult to obtain in explicit form. Note that in order for any of the nodes 1, 2, 5 or 6 to activate, node 3 must be inactive. In order for nodes 1, 2 or 6 to activate, node 4 must be inactive as well, but the latter is not necessary in order for node 5 to activate. Since node 4 may be active even when it is empty on the fluid scale, it follows that node 5 enjoys an advantage in competing for access to the medium over nodes 1, 2 and 6. While it may be argued that node 4 is active with probability ρ_4 by the time node 3 becomes inactive for the first time, the resulting probabilities for the various nodes to gain access to the medium first do not seem to allow a simple expression.

Remark 4.1. *If the process had not yet reached a natural state, the case $\frac{q_3(t)}{1-\rho_3} \leq \frac{q_4(t)}{1-\rho_4}$ could also arise. In case that inequality is strict, i.e., $\frac{q_3(t)}{1-\rho_3} < \frac{q_4(t)}{1-\rho_4}$, the queue of node 4 is still nonempty by time $t + T_2(t)$, simply forcing a switch to an M_4-period with probability 1.*

In case of equality, i.e., $\frac{q_3(t)}{1-\rho_3} = \frac{q_4(t)}{1-\rho_4}$, however, the situation would be much more complicated, which serves as the illustration for the significance of the notion of a natural state. In order to describe these difficulties, note that the queues of nodes 3 and 4 both empty at time $t+T_2(t)$, barring a switch to an M_4-period, and permitting only a switch to either an M_1- or M_3-period. Just like before, node 5 is the only one able to activate during periods where node 3 is inactive while node 4 is active, and hence enjoys an advantage in competing for access to the medium. In fact, node 5 will gain access to the medium first almost surely if node 3 is the first one to become inactive (in the pre-limit). The probability of that event, and hence the transition probabilities to an M_1- or M_3-period, depends on queue length differences between nodes 3 and 4 at time t that can be affected by the history of the process and are not visible on the fluid scale.

M_3-period

The dynamics for an M_3-period are entirely symmetric to those for an M_2-period, but will be replicated below for completeness.

Assume that the system enters an M_3-period at time t, then

(a) The system spends a time period $T_3(t) = \frac{q_6(t)}{1-\rho_6}$ in M_3.

(b) During this period, the queues of nodes 5 and 6 drain (or remain zero)

$$q_i(t+u) = \max\{q_i(t) - (1-\rho_i)u, 0\}, \text{ for } i = 5, 6,$$

while the queues of nodes 1, 2, 3, 4 fill at a linear rate

$$q_i(t+u) = q_i(t) + \rho_i u, \text{ for } i = 1, 2, 3, 4,$$

for all $u \in [0, T_3(t)]$. In particular, $q_6(t + T_3(t)) = 0$.

(c) At time $t + T_3(t)$, the system switches to an M_1- or M_2-period. Note that $\frac{q_5(t)}{1-\rho_5} < \frac{q_6(t)}{1-\rho_6}$ by the assumption that $\lambda_5 > \lambda_6$ and that the process has reached a natural state, so that $q_5(t) < q_6(t)$ (since $q_5(t) = q_6(t) = 0$ cannot occur at the start of an M_3-period).

Thus node 5 has emptied before time $t+T_3(t)$, and remained empty (on the fluid scale) since then, precluding a switch to an M_4-period (except for a negligible period on the fluid scale), only allowing the system to switch to either an M_1- or M_2-period. The corresponding transition probabilities are difficult to obtain in explicit form for similar reasons as mentioned in case 2(c).

Remark 4.2. *If the process had not yet reached a natural state, the case* $\frac{q_5(t)}{1-\rho_5} \geq \frac{q_6(t)}{1-\rho_6}$ *could also arise. In case that inequality is strict, i.e.,* $\frac{q_5(t)}{1-\rho_5} < \frac{q_6(t)}{1-\rho_6}$*, the queue of node 5 is still nonempty by time* $t + T_3(t)$*, forcing a switch to an* M_4*-period with probability 1.*

In case of equality, i.e., $\frac{q_5(t)}{1-\rho_5} = \frac{q_6(t)}{1-\rho_6}$*, the queues of nodes 5 and 6 both empty at time* $T_3(t)$*, barring a switch to an* M_4*-period, and permitting only a switch to either an* M_1*- or* M_2*-period. For similar reasons as mentioned in case 2(c), the corresponding transition probabilities depend on queue length differences that are affected by the history of the process and are not visible on the fluid scale.*

M_4-period

Assume that the system enters an M_4-period at time t, then

(a) It spends a time period $T_4(t) = \min\left\{\frac{q_4(t)}{1-\rho_4}, \frac{q_5(t)}{1-\rho_5}\right\}$ in M_4.

(b) During this period, the queues of nodes 4 and 5 drain at a linear rate

$$q_i(t+u) = q_i(t) - (1-\rho_i)u, \text{ for } i = 4, 5,$$

while the queues of nodes 1, 2, 3, 6 fill at a linear rate

$$q_i(t+u) = q_i(t) + \rho_i u, \text{ for } i = 1, 2, 3, 6,$$

$u \in [0, T_4(t)]$. In particular, $\min\{q_4(t+T_4(t)), q_5(t+T_4(t))\} = 0$.

(c) At time $t + T_4(t)$, the system switches to either an M_2- or M_3-period. In order to determine which of these events can occur, we need to distinguish between three cases, depending on whether $\frac{q_4(t)}{1-\rho_4}$ is (i) larger than, (ii) equal to, or (iii) smaller than $\frac{q_5(t)}{1-\rho_5}$.

In case (i), i.e., $\frac{q_4(t)}{1-\rho_4} > \frac{q_5(t)}{1-\rho_5}$, we have $q_4(t+T_4(t)) > 0$, i.e., the queue of node 4 is still nonempty by time $t + T_4(t)$, causing a switch to an M_2-period with probability 1.

In case (ii), i.e., $\frac{q_4(t)}{1-\rho_4} = \frac{q_5(t)}{1-\rho_5}$, we have $q_4(t+T_4(t)) = q_5(t+T_4(t)) = 0$, i.e., the queues of nodes 4 and 5 both empty at time $t + T_4(t)$. Even though both queues empty at the same time on the fluid scale, there will with overwhelming probability be a long period in the pre-limit where one of the nodes has become inactive for the first time while the other one has yet to do so. Since both nodes 4 and 5 must be inactive in order for nodes 1 and 2 to activate, these nodes have no chance to activate during that period, but either node 3 or node 6 does, depending on whether node 5 or node 4 is the first one to become inactive. As a result, the system cannot switch to an M_1-period, but only to an M_2- or M_3-period. In fact, a switch to M_2 will occur almost surely if node 5 is the first one to become inactive, while a switch to M_3 will occur almost surely if node 4 is the first one to become inactive. The probabilities of these two scenarios, and hence the transition probabilities to M_2 and M_3, depend on queue length differences between nodes 4 and 5 at time t that are affected by the history of the process and are not visible on the fluid scale.

In case (iii), i.e., $\frac{q_4(t)}{1-\rho_4} < \frac{q_5(t)}{1-\rho_5}$, we have $q_5(t + T_4(t)) > 0$, i.e., the queue of node 5 is still nonempty by time $t + T_4(t)$, forcing a switch to an M_3-period with probability 1.

Remark 4.3. *As noted in the above description of the fluid limit process, in cases 2(c), 3(c), and 4(c)(ii) the transition probabilities from an M_2-period to an M_1- or M_3-period, from an M_3-period to an M_1- or M_2-period, and from an M_4- to an M_2- or M_3-period, depend on queue length differences that are affected by the history of the process and are not visible on the fluid scale. Depending on whether or not the initial state and parameter values allow for these cases to arise, it may thus be impossible to provide a probabilistic description of the evolution of the resulting fluid limit process, even in terms of its entire own history.*

4.2.2 Overview of fluid limit proofs

In the previous subsection we provided a description of the dynamics of the fluid limit process once it has reached a natural state and entered an M_1, M_2-, M_3- or M_4-period. As was further stated, for ρ sufficiently close to 1, a natural state and such a period occurs in uniformly bounded time almost surely for any initial state. The justification for all these properties follows from a series of lemmas and propositions stated and proved in Sections 4.8.3–4.8.6. In this subsection we present a high-level outline of the fluid limit statements and proofs.

First of all, recall that the description of the fluid limit process referred to the continuous-time Markov process representing the system dynamics as introduced at the beginning of this chapter. For all the proofs of fluid limit properties and instability results however we consider a rescaled linear interpolation of the uniformized jump chain (as defined in Section 4.8.3). This construction yields convenient properties of the fluid limit paths and allows us to extend the framework of Meyn [35] for establishing instability results for discrete-time Markov chains. (The original continuous-time Markov process has in fact the same fluid limit properties, but this is not directly relevant in any of the proofs.)

The proofs of the fluid limit properties consist of four main parts. Part A identifies several basic properties of the fluid limit paths, and in particular establishes that the queue length trajectory of each of the individual nodes exhibits "sawtooth" behavior. This fundamental property in fact holds in arbitrary interference graphs, and only requires an exponent $\gamma > 1$ in the backoff probability. Part B of the proof shows a certain dominance property, saying that if all the interferers of a particular node also interfere with some other node that is currently being served at full rate, then the former node must be empty or served at full rate (on the fluid scale) as well. Under the assumption $\lambda_3 > \lambda_4$, $\lambda_5 < \lambda_6$, the dominance property implies that after a finite amount of time the fluid limit process for the broken-diamond network must always reside in a natural state as defined in the previous subsection. Part C of the proof centers on the M_1-, M_2-, M_3- and M_4-periods, and establishes that at the end of any such period, the process immediately switches to one of the other types of periods with the probabilities indicated in the previous subsection. In particular, it is deduced that an M_4-period cannot be entered from an M_2- or M_3-period, and must always be preceded by an M_1-period once the process has reached a natural state. The combination of the sawtooth queue length trajectories and the switching probabilities provides a probabilistic description of the dynamics of the fluid limit once the process has reached a natural state and entered an M_1-, M_2-, M_3- or M_4-period. Part B already established that the process must always reside in a natural state after a finite amount of time, but it remains to be shown that the process will inevitably enter an M_1-, M_2-, M_3- or M_4- period, which constitutes the final Part D of the proof. The core argument is that interfering empty and nonempty queues can not coexist, since the empty nodes will frequently enter back-off periods, offering the nonempty nodes abundant opportunities to gain access, drain their queues, and cause the empty nodes to build queues in turn.

Part A of the proof starts with the simple observation that, by the "skip-free" property of the original pre-limit process, the sample paths of the interpolated version of the uniformized jump chain are Lipschitz continuous, and hence so are the sample paths of the fluid-scaled process. The fluid limit paths inherit the Lipschitz continuity, and are thus differentiable almost everywhere with probability one.

Then fluid limit paths are determined by a countable set of "entrance" times and “exit” times of $(0, \infty)$ with probability one. The proof then proceeds to show that if a nonempty node (on the fluid scale) receives any amount of service during some time interval, then it must in fact be served at the full rate until it has completely emptied (on the fluid scale), assuming $\gamma > 1$. This implies that when node i is nonempty (on the fluid scale), its queue must either increase at rate λ_i or decrease at rate $1 - \lambda_i$ until it has entirely drained. In other words, the queue length trajectory of each of the individual nodes exhibits sawtooth behavior (Theorem 4.5).

Part B of the proof pertains to the joint behavior of the fluid limit trajectories of the various queue lengths. First of all, the natural property is proved that whenever a particular node is served, none of its interferers can receive any service (Lemma 4.3). Second, it is established that whenever a particular node is served, any node whose interferers are a subset of those of the node served, must either be empty or be served at full rate as well (on the fluid scale) (Corollary 4.3). For example, in the broken-diamond network, whenever node 3 is served, node 4 must either be empty or be served at full rate as well, and similarly for nodes 5 and 6. These two properties combined yield a dominance property, saying that if all the interferers of a particular node also interfere with some other node that is currently being served at full rate, then the former node must be empty or served at full rate (on the fluid scale) as well. In the case of the broken-diamond network, under the assumption $\lambda_3 > \lambda_4$, the queue of node 3 will therefore never be smaller than that of node 4 after some finite amount of time, and similarly for nodes 4 and 5. Thus the fluid limit process will always reside in a natural state after some finite amount of time.

Part C of the proof focuses on the M_1-, M_2-, M_3- and M_4-periods as described above. Because of the sawtooth behavior, an M_1-period can only end when both nodes 1 and 2 are empty (on the fluid scale). Likewise, an M_2- or M_3-period can only end when node 3 or node 6 is empty, respectively. An M_4-period can only end when node 4 or node 5 (or both) is empty. It is then proven that at the end of an M_1-period, the fluid limit process immediately switches to an M_2-, M_3- or M_4-period with the probabilities specified in the previous subsection (Theorem 4.7). When the process resides in a natural state, an M_2-period is always instantaneously followed by an M_1- or M_3-period, while an M_3-period is always instantaneously followed by an M_1- or M_2-period. In particular, it is concluded that an M_4-period cannot be entered from an M_2- or M_3-period, and must always be preceded by an M_1-period once the process has reached a natural state. After an M_4-period, the process always immediately switches to an M_2- or M_3-period.

There is no reason a priori, however, that the process is guaranteed to actually ever enter an M_1-, M_2-, M_3- or M_4- period. In fact, the process may very well spend time in different kinds of states, but the final Part D of the proof establishes that these kinds of states are transient, and cannot occur once a natural state has been reached, which is forced to happen in a finite amount of time for particular arrival rates as was already shown in Part B. Note that an M_1-, M_2-, M_3- or M_4- period occurs as soon as node 1, node 2, node 3, node 6 or nodes 4 and 5 simultaneously are served at full rate. In other words, the only ways for the process to avoid an M_1-, M_2-, M_3- or M_4-period, are: (i) for node 4 to be served at full rate, but not nodes 3 and 5; (ii) for node 5 to be served at full rate, but not nodes 4 and 6; (iii) for none of the nodes to be served at full rate. Scenario (i) requires node 3 to be empty (on the fluid scale) and node 4 to be nonempty, which cannot occur in a natural state. Likewise, scenario (ii) cannot arise in a natural state either. Scenario (iii) requires that every empty node i is served at rate ρ_i (on the fluid scale), while all nonempty nodes are served at rate 0. Such a scenario is not particularly plausible, but a rigorous proof turns out to be quite involved. The insights rely strongly on the specific properties of the broken-diamond network, and an extension to arbitrary graphs does not seem straightforward. The core argument is that interfering empty and nonempty queues can not coexist, since the empty nodes will frequently enter back-off periods, offering the nonempty nodes abundant opportunities to gain access, drain their queues, and cause the empty nodes to build queues in turn.

4.3 Instability results for broken-diamond network

In Section 4.2, we provided a detailed description of the dynamics of the fluid limit process, once it has reached a natural state and entered an M_1-, M_2-, M_3-, or M_4-period. In this section we exploit the properties of the fluid limit process in order to prove that it is unstable for ρ sufficiently close to 1, and then show how the instability of the original stochastic process can be deduced from the instability of the fluid limit process.

4.3.1 Instability of the fluid limit process

In order to prove instability of the fluid limit process, we first revisit the intuitive explanation discussed earlier in Section 4.2, see Figure 4.4 for an illustration. Denote $\rho_0 = \max\{\rho_1, \rho_2\}$, and recall that $\rho_3 \geq \rho_4$ and $\rho_5 \leq \rho_6$ by assumption. Since nodes 1, 2, 3 and 6 are only served during M_1-, M_2- and M_3-periods, and not during M_4-periods, it is easily seen that the fraction of time that the system spends in M_1-, M_2- and M_3-periods must be no less than $\rho = \rho_0 + \rho_3 + \rho_6$ in order for these nodes to be stable. Thus, the system cannot be stable if it spends a fraction of the time larger than $1 - \rho$ in M_4-periods. As it turns out, however, when the de-activation function is sufficiently aggressive, e.g., $\hat{r}(x) = \mathrm{o}(x^{-\gamma})$ (correspondingly weight functions $w_i(x) = \gamma \log(1 + x)$), with $\gamma > 1$, M_4-periods in fact persistently occur for a fraction of time that does not tend to 0 as ρ approaches 1, which forces the system to be unstable.

Figure 4.4 shows a fluid-limit sample path corresponding to the switching sequence $M_1 \to M_2 \to M_1 \to M_4 \to M_3 \to M_1$. The aggregate queue size starts building up in the M_3-period that follows the M_4-period.

In order to prove instability of the fluid limit process, we adopt the Lyapunov function $L(t) = \sum_{k=1}^{3} \max_{i \in M_k} q_i(t)$, and will show that the load $L(t)$ grows without bound almost surely. Note that the load $L(t)$ increases during M_4-periods while the process is in a natural state.

In preparation for the instability proof, we first state two auxiliary lemmas. It will be convenient to view the evolution of the fluid limit process, and in particular the Lyapunov function $L(t)$, over the course of *cycles*. The i-th *cycle* is the period from the start of the $(i-1)$-th M_1-period to the start of the i-th M_1-period *once the fluid limit process has reached a natural state*. Denote by t_i the start time of the i-th cycle, $i = 1, 2, \ldots$. Each t_i is finite almost surely for ρ sufficiently close to 1, and in particular an infinite number of cycles must occur almost surely. In order to see that, recall that the fluid limit process will reach a natural state and enter an M_1-, M_2-, M_3- or M_4-period in finite time almost surely for any initial state as stated in Section 4.2.1. The description of the dynamics of the fluid limit process provided in that section then implies that M_1-periods and hence cycles must occur infinitely often (and if only finitely many M_1-periods occurred, then at least one of the nodes would in fact never be served again after some finite time, implying that the fluid limit process is unstable regardless).

The next lemma shows that the duration of a cycle and the possible increase in the load over the course of a cycle are linearly bounded in the load at the start of the cycle.

Lemma 4.1. *The duration of the i-th cycle, $\Delta t_i = t_{i+1} - t_i$, and the increase in the load over the course of the i-th cycle, $L(t_{i+1}) - L(t_i) = L(t_i + \Delta t_i) - L(t_i)$, are bounded from above by*

$$\Delta t_i \leq C_T L(t_i) \text{ and } L(t_{i+1}) - L(t_i) \leq C_L L(t_i),$$

for all $\rho \leq 1$, where $C_T = \frac{1}{1-\rho_3-\rho_6}\left(\frac{1}{1-\rho_0} + \frac{1}{1-\max\{\rho_4,\rho_5\}}\right)$ and $C_L = \frac{\rho}{1-\max\{\rho_4,\rho_5\}}$.

The proof of Lemma 4.1 is presented in Section 4.8.1.

In order to establish that the durations of M_4-periods are non-negligible, it will be useful to introduce the notion of "weakly-balanced" queues, ensuring that the queues of nodes 4 and 5 are not too small compared to the queues of nodes 3 and 6.

Definition 4.1. *Let $\beta^{\min}$ and $\beta^{\max}$ be fixed positive constants. The queues are said to be weakly balanced in a given cycle (with respect to $\beta^{\min}$ and $\beta^{\max}$) if $\beta^{\min} \leq \frac{q_3(t)}{q_5(t)}, \frac{q_6(t)}{q_4(t)} \leq \beta^{\max}$, with t denoting the time when the M_1-period ends that initiated the cycle.*

The next lemma shows that over two consecutive cycles, the queues will be weakly balanced with probability at least 1/3.

Lemma 4.2. *Let*

$$\epsilon = \frac{\rho_2}{2\left(\rho_2 + (\rho_3+\rho_6)\frac{1-\min\{\rho_4,\rho_5\}}{1-\max\{\rho_4,\rho_5\}}\right)} \geq \frac{\rho_2}{\rho}\frac{1-\max\{\rho_4,\rho_5\}}{1-\min\{\rho_4,\rho_5\}}.$$

Then over two consecutive cycles, with probability at least 1/3, the queues will be weakly balanced in at least one of these cycles with

$$\beta^{\max} = \frac{\max\{\rho_3,\rho_6\} + (1-\rho_2)(1-\epsilon)/\epsilon}{\min\{\rho_4,\rho_5\}},$$

and $\beta^{\min} = \frac{1}{\beta^{\max}}$.

The proof of Lemma 4.2 is presented in Section 4.8.2.

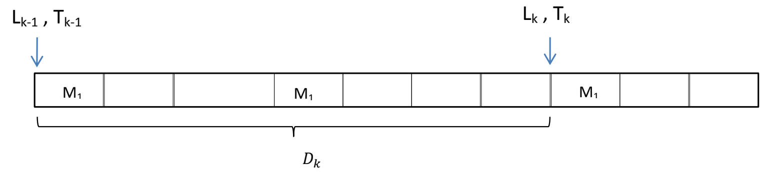

Figure 4.5: A cycle D_k consisting of a pair of consecutive cycles.

As suggested by Lemma 4.2, it will be convenient to consider pairs of two consecutive cycles in order to prove instability of the fluid limit process.

Let D_k be the pair of cycles consisting of cycles $2k-1$ and $2k$ as in Figure 4.5, $k = 1, 2, \ldots$. With minor abuse of notation, denote by $T_k = t_{2k-1}$ the start time of D_k and $L_k = L(T_k)$. Denote by $\Delta T_k = T_{k+1} - T_k$ the duration of D_k and by $\Delta L_k = L_{k+1} - L_k$ the increase in $L(t)$ over the course of D_k.

The next proposition shows that for ρ sufficiently close to 1 the load cannot significantly decrease over a pair of cycles and will increase by a substantial amount with non-zero-probability. We henceforth assume $(\rho_1, \rho_2, \rho_3, \rho_4, \rho_5, \rho_6) = \rho(\kappa_1, \kappa_2, \kappa_3, \kappa_3 - \alpha, \kappa_6 - \alpha, \kappa_6)$, with $\max\{\kappa_1, \kappa_2\} + \kappa_3 + \kappa_6 = 1$ and $0 < \alpha < \min\{\kappa_3, \kappa_6\}$, so that $\rho = \rho_0 + \rho_3 + \rho_6$.

Proposition 4.1. *Let $C_{LT} = C_T(2 + C_L)$, with C_T and C_L as specified in Lemma 4.1, $\theta = 1 - (1-\rho)C_{LT}$, $p = 1/12$. Over cycle pairs D_k, $k = 1, 2, \ldots$,*

(i) $\Delta T_k \leq C_{LT} L_k$;

(ii) $L(t) \geq \theta L_k$ *for all* $t \in [T_k, T_{k+1}]$;

(iii) $\mathbb{P}\big(L_{k+1} - \theta L_k \geq \delta(\rho)\theta L_k | L_k\big) \geq p$,

with $\delta(\rho)$ a constant, depending on ρ, and $\delta(\rho) \uparrow \delta = \frac{1}{\beta^{\max}(1+\beta^{\max})(1+\alpha-\min\{\kappa_3,\kappa_6\})}$, as $\rho \uparrow 1$.

Proof. We first show part (i). Using Lemma 4.1, we find

$$\Delta T_k = \Delta t_{2k-1} + \Delta t_{2k} \leq C_T(L(t_{2k-1}) + L(t_{2k})) \leq C_T(2 + C_L)L_k.$$

In order to prove part (ii), note that $L(t)$ cannot decrease at a larger rate than $1 - \rho$, so that in view of part (i),

$$L(t) \geq L_k - (1-\rho)(t - T_k) \geq L_k - (1-\rho)\Delta T_k \geq (1 - (1-\rho)C_{LT})L_k = \theta L_k,$$

for all $t \in [T_k, T_{k+1}]$.

We now turn to part (iii). Suppose that the following event occurs: the queues are weakly balanced at the end of an M_1-period, say time τ, during D_k (which according to Lemma 4.2) happens with at least probability 1/3) and the system then enters an M_4-period (which happens with probability 1/4). Recalling that $\rho_3 > \rho_4$, $q_3(t) \geq q_4(t)$, $\rho_5 < \rho_6$ and $q_5(t) \leq q_6(t)$, we find that during the M_4-period $L(t)$ increases by

$$\rho \min\left\{\frac{q_4(\tau)}{1-\rho_4}, \frac{q_5(\tau)}{1-\rho_5}\right\} \geq \rho \frac{\min\{q_4(\tau), q_5(\tau)\}}{1-\rho\min\{\kappa_3,\kappa_6\}+\rho\alpha}.$$

Since the queues are weakly balanced, we deduce $q_3(\tau) \leq \beta^{\max} q_5(\tau) \leq \beta^{\max} q_6(\tau) \leq (\beta^{\max})^2 q_4(\tau)$ and $q_6(\tau) \leq \beta^{\max} q_4(\tau) \leq \beta^{\max} q_3(\tau) \leq (\beta^{\max})^2 q_5(\tau)$. Noting that $q_1(\tau) = q_2(\tau) = 0$, we obtain

$$L(\tau) = q_3(\tau) + q_6(\tau) \leq (1+\beta^{\max}) q_6(\tau) \leq \beta^{\max}(1+\beta^{\max}) q_4(\tau),$$

and also

$$L(\tau) = q_3(\tau) + q_6(\tau) \leq (1+\beta^{\max}) q_3(\tau) \leq \beta^{\max}(1+\beta^{\max}) q_5(\tau).$$

So

$$L(\tau) \leq \beta^{\max}(1+\beta^{\max}) \min\{q_4(\tau), q_5(\tau)\},$$

and thus the increase in $L(t)$ during the M_4-period is no less than $\delta(\rho)L(\tau)$, with

$$\delta(\rho) = \frac{\rho}{\beta^{\max}(1+\beta^{\max})(1-\rho\min\{\kappa_3,\kappa_6\}+\rho\alpha)}.$$

Using part (i) once again, we conclude that with at least probability 1/12,

$$\begin{aligned}
L_{k+1} &\geq L_k + \delta(\rho)L(\tau) - (1-\rho)\Delta T_k \geq L_k + \delta(\rho)(L_k - (1-\rho)\Delta T_k) - (1-\rho)\Delta T_k \\
&= (1+\delta(\rho))(L_k - (1-\rho)\Delta T_k) \geq (1+\delta(\rho))(L_k - (1-\rho)C_{LT}L_k) \\
&= (1+\delta(\rho))\theta L_k.
\end{aligned}$$

□

Armed with the above proposition, we now proceed to prove that the fluid limit process is unstable, in the sense that $L(T) \to \infty$ as $T \to \infty$. In fact, $L(T)$ grows faster than any sub-linear function $T^{\frac{1}{m}}$, $m > 1$, as stated in the next theorem.

Theorem 4.1. *For any $m > 1$, there exists a constant $\rho^* = \rho^*(\kappa, m) < 1$, such that for all $\rho \in (\rho^*, 1]$,*

$$\limsup_{T\to\infty} \mathbb{E}\left[\frac{T}{L^m(T)}\right] = 0,$$

for any initial state $\mathbf{q}(0)$ with $||\mathbf{q}(0)|| = 1$, and $||\cdot||$ denoting the L_1-norm.

Proof. Consider the cycle pairs D_k, $k = 1, 2, \ldots$, as defined right before Proposition 4.1. Assume $\rho \in (1 - \frac{1}{C_{LT}}, 1]$, so that $\theta \in (0, 1]$ in Proposition 4.1. For any time $t > T_1$, we can define a stopping time N_t such that $T_{N_t} < t \leq T_{N_t+1}$, i.e., t is within the N_t-th cycle pair. (This is possible almost surely, since $T_k \to \infty$ as $k \to \infty$ almost surely, as will be proven below.) Recall that $T_{N_t+1} \leq T_{N_t} + C_{LT} L_{N_t}$ and $L(t) \geq \theta L_{N_t}$ by parts (i) and (ii) of Proposition 4.1, respectively, and trivially $L_{N_t} \leq L(0) + \rho T_{N_t} \leq 2T_{N_t}$ for t sufficiently large. Thus,

$$\begin{aligned}
\limsup_{t\to\infty} \mathbb{E}\left[tL^{-m}(t)\right] &\leq \limsup_{t\to\infty} \mathbb{E}\left[T_{N_t+1}\theta^{-m} L_{N_t}^{-m}\right] \\
&\leq \theta^{-m} \limsup_{t\to\infty} \mathbb{E}\left[T_{N_t} L_{N_t}^{-m}\right] + \theta^{-m} C_{LT} \limsup_{t\to\infty} \mathbb{E}\left[L_{N_t}^{-m+1}\right] \\
&\leq \theta(1 + 2C_{LT}) \limsup_{t\to\infty} \mathbb{E}\left[T_{N_t} L_{N_t}^{-m}\right]. \qquad (4.1)
\end{aligned}$$

So it suffices to prove that there exists $\rho^* = \rho^*(\kappa, m) < 1$ such that (4.1) is zero for $\rho > \rho^*$, which we now proceed to show.

First of all, by Proposition 4.1, for any $m > 0$,

$$\begin{aligned}
\mathbb{E}\left[L_{k+1}^{-m} | \mathcal{F}_k\right] &\leq (1-p)(\theta L_k)^{-m} + p((\theta + \delta)L_k)^{-m} \\
&= \alpha_m L_k^{-m}, \qquad (4.2)
\end{aligned}$$

where $\mathcal{F}_k$ is a suitable filtration and $\alpha_m := (1-p)\theta^{-m} + p(\theta + \delta)^{-m}$.

Since $\theta(\rho) \to \theta(1) = 1$ and $\delta(\rho) \to \delta(1) = \delta > 0$ as $\rho \uparrow 1$, $\alpha_m(\rho)$ is a continuous function of ρ in the vicinity of 1. Because $\alpha_m(1) < 1$, there must exist a $\rho_m^* = \rho^*(\kappa, m) < 1$ such that $\alpha_m < 1$ for all $\rho > \rho^*$. This shows that, for $\rho > \rho_m^*$, L_k^{-m} is a positive (geometric) super-martingale with parameter $\alpha_m < 1$. Taking expectations on both sides of (4.2) yields

$$\mathbb{E}\left[L_k^{-m}\right] \leq \alpha_m^k L_0^{-m}. \qquad (4.3)$$

with $L_0 = L(t_{i_0}) > 0$ as noted earlier. In particular, $\lim_{k\to\infty} \mathbb{E}\left[L_k^{-m}\right] = 0$, and $1/L_k \to 0$ almost surely as $k \to \infty$ by the Doob's super-martingale-convergence theorem (page 147 of [44]). This implies that $T_k \to \infty$ almost surely because $L_k \le \rho T_k + 1 \le T_k + 1$. Therefore, the stopping time T_{N_t} is well-defined.

Next, consider the sequence of random variables $T_k L_k^{-m}$. Using Proposition 4.1,

$$\begin{aligned} \mathbb{E}\left[T_k L_k^{-m} | \mathcal{F}_{k-1}\right] &\le (T_{k-1} + C_{LT} L_{k-1}) \mathbb{E}\left[L_k^{-m} | \mathcal{F}_{k-1}\right] \\ &\le (T_{k-1} + C_{LT} L_{k-1}) \alpha_m L_{k-1}^{-m} \\ &= \alpha_m T_{k-1} L_{k-1}^{-m} + \alpha_m C_{LT} L_{k-1}^{-m+1}. \end{aligned} \tag{4.4}$$

Define $\epsilon_k := C_{LT} \alpha_m L_k^{-m+1}$, then, by (4.2) and (4.3), ϵ_k is a positive (geometric) super-martingale with parameter $\alpha_{m-1} < 1$ for $\rho > \rho^*_{m-1} = \rho^*(\kappa, m-1)$. Then, $\sum_{k=1}^{\infty} \mathbb{E}\left[\epsilon_k\right] \le C_{LT} \alpha_m \sum_{k=1}^{\infty} \alpha_{m-1}^k < \infty$, which shows that

$$\lim_{k\to\infty} T_k L_k^{-m} = 0,$$

almost surely. In particular, define $\alpha := \max(\alpha_m, \alpha_{m-1})$ and $\rho^* = \max\{\rho^*_m, \rho^*_{m-1}\}$, then taking expectations on both sides of (4.4) yields

$$\mathbb{E}\left[T_k L_k^{-m}\right] \le \alpha \mathbb{E}\left[T_{k-1} L_{k-1}^{-m}\right] + \alpha C_{LT} \alpha^{k-1}, \tag{4.5}$$

which, by induction, shows that

$$\mathbb{E}\left[T_k L_k^{-m}\right] \le \alpha^{k-1} (\mathbb{E}\left[T_1 L_1^{-m}\right] + C_{LT}(k-1)\alpha), \tag{4.6}$$

for $\rho \in (\rho^*, 1]$. Now observe that T_1 is strictly bounded and L_1 is bounded away from zero, since a natural state is reached in finite time, before the system can empty, almost surely. It then follows that $\lim_{k\to\infty} \mathbb{E}\left[T_k L_k^{-m}\right] = 0$.

The fact that $T_k L_k^{-m}$ converges in $\mathcal{L}_1$ implies that the sequence of random variables $T_k L_k^{-m}$ is *Uniformly Integrable* (UI) (page 147, Theorem 50.1 of [44]). It therefore follows, by adapting the arguments of Doob's optional sampling theorem (page 159 of [44]), that the family of random variables $\{T_{N_t} L_{T_{N_t}}^{-m}\}$ is also UI. Thus by definition, given $\varepsilon > 0$, there exists K_ε such that

$$\mathbb{E}\left[T_{N_t}L_{N_t}^{-m}\mathrm{I}_{\left\{T_{N_t}L_{N_t}^{-m}\geq K_\varepsilon\right\}}\right]\leq\varepsilon,\quad\forall t>0.$$

We deduce

$$\begin{aligned}\mathbb{E}\left[T_{N_t}L_{N_t}^{-m}\right]&\leq\sum_{k=1}^{\infty}\mathbb{E}\left[T_kL_k^{-m}\mathrm{I}_{\left\{N_t=k\right\}}\mathrm{I}_{\left\{T_{N_t}L_{N_t}^{-m}\leq K_\varepsilon\right\}}\right]+\varepsilon\\&\leq K_\varepsilon\mathbb{P}\left\{N_t\leq D\right\}+\sum_{k=D+1}^{\infty}C_{LT}k\alpha^{k-1}+\varepsilon.\end{aligned}$$

Fixing ε and D, we find that

$$\limsup_{t\to\infty}\mathbb{E}\left[T_{N_t}L_{N_t}^{-m}\right]\leq(D+1)\frac{\alpha^D}{1-\alpha}+\varepsilon$$

by the Monotone Convergence Theorem [45], and thus, letting $D\to\infty$ and $\varepsilon\to 0$, we have $\limsup_{t\to\infty}\mathbb{E}\left[T_{N_t}L_{N_t}^{-m}\right]=0$ for $\rho>\rho^*$. □

Corollary 4.1. *For any $m>1$, there exists a constant $\rho^*=\rho^*(\kappa,m)<1$, such that for all $\rho\in(\rho^*,1]$,*

$$\liminf_{T\to\infty}\frac{L(T)}{T^{1/m}}=\infty,$$

almost surely for any initial state $\mathbf{q}(0)$ with $||\mathbf{q}(0)||=1$.

Proof. Note that for any initial state $\mathbf{q}(0)$ with $||\mathbf{q}(0)||=1$,

$$\liminf_{T\to\infty}\frac{L(T)}{T^{1/m}}\geq\liminf_{k\to\infty}\frac{\theta L_k}{T_{k+1}^{1/m}},$$

as can be seen from Proposition 4.1, and so it suffices to show that

$$\limsup_{k}T_{k+1}L_k^{-m}=0.$$

But $T_{k+1}\leq T_k+C_{LT}L_k$, thus,

$$\limsup_{k\to\infty}T_{k+1}L_k^{-m}\leq\limsup_{k\to\infty}T_kL_k^{-m}+C_{LT}\limsup_{k\to\infty}L_k^{-m+1}.\tag{4.7}$$

The right-hand side is zero because, as we saw in the proof of Theorem 4.1, both $T_kL_k^{-m}$ and L_k^{-m+1} converge to zero almost surely for $\rho\in(\rho^*(\kappa,m),1]$. □

4.3.2 Instability of the original stochastic process

In Theorem 4.1 we established that the fluid limit process in unstable, in the sense that $L(T) \to \infty$ as $T \to \infty$. We now proceed to show how the instability of the original stochastic process can be deduced from the instability of the fluid limit process. The original stochastic process is said to be unstable when $\{(X(t), Q(t))\}_{t \geq 0}$ is transient, and $\|Q(t)\| \to \infty$ almost surely for any initial state $Q(0)$.

We will exploit similar arguments as developed in Meyn [35]. A notable distinction is that the result in [35] requires that a suitable Lyapunov function exhibits strict growth over time. In our setting the fluid limit is random, and the growth behavior as stated in Theorem 4.1 is not strict, but only in expectation and in an asymptotic sense, which necessitates a somewhat delicate extension of the arguments in [35].

The next theorem states the main result of the present section, indicating that aggressive de-activation functions cause the network of Figure 4.3 to be unstable for load values ρ sufficiently close to 1.

Theorem 4.2. *Consider the network of Figure 4.3, and suppose that* $r_i(x) \equiv 1$, $x \geq 1$, *and* $\hat{r}_i(x) = \mathrm{o}(x^{-\gamma})$, *with* $\gamma > 1$. *Let* $(\rho_1, \rho_2, \rho_3, \rho_4, \rho_5, \rho_6) = \rho(\kappa_1, \kappa_2, \kappa_3, \kappa_3 - \alpha, \kappa_6 - \alpha, \kappa_6)$, *with* $\max\{\kappa_1, \kappa_2\} + \kappa_3 + \kappa_6 = 1$, *and* $0 < \alpha < \min\{\kappa_3, \kappa_6\}$. *Then there exists a constant* $\rho^\star(\kappa, \alpha) < 1$, *such that for all* $\rho \in (\rho^\star(\kappa, \alpha), 1]$*:*

$$\lim_{\|Q(0)\| \to \infty} \mathbb{P}_{Q(0)}\{\liminf_{t \to \infty} \|Q(t))\| = \infty\} = 1.$$

Since our Markov chain is irreducible, the theorem immediately implies that it is transient.

Remark 4.4. *Recall that the class of de-activation functions* $\hat{r}_i(x) = \mathrm{o}(x^{-\gamma})$ *includes the random-capture scheme with* $\hat{r}(x) \equiv 0$, $x \geq 1$, *as considered in [28]. The result in Theorem 4.2 thus disproves the conjecture that the random-capture scheme is throughput-optimal in arbitrary topologies.*

The proof of Theorem 4.2 relies on similar arguments as developed in the proof of Theorem 3.2 in [35]. A crucial role is played by Theorem 3.1 of [35], which is reproduced for completeness in the following.

Theorem 4.3. *Suppose that for a Markov chain $\{Y(n); n = 0, 1, 2, \dots\}$ with discrete state space S, there exist positive functions $W(\cdot)$ and $\Delta(\cdot)$ on S, and a positive constant c_0, such that*

$$\mathbb{E}\left[W(Y(n+1))|\mathcal{F}_n\right] \leq W(Y(n)) - \Delta(Y(n)), \tag{4.8}$$

whenever $Y(n) \in S_{c_0} = \{y \in S : W(y) \leq c_0\}$, with $\mathcal{F}_n := \sigma(Y(0), Y(1), \dots, Y(n))$. Then for all $x \in S$,

$$\mathbb{P}_y \left\{ \sum_{n=0}^{\infty} \Delta(Y(n)) < \infty \right\} \geq 1 - W(y)/c_0.$$

In order to apply Theorem 4.3, we need to construct suitable functions $W(\cdot)$ and $\Delta(\cdot)$. The proof details follow.

Proof of Theorem 4.2. Let $(X(n), Q(n))$ denote the jump chain obtained from the continuous-time Markov process by uniformization according to a Poisson clock of rate β as described in Section 4.8.3. In order to prove Theorem 4.2 for the original stochastic process, it suffices to establish a similar result for the jump chain:

$$\lim_{\|Q(0)\| \to \infty} \mathbb{P}_{Q(0)}\{\liminf_n \|Q(n)\| = \infty\} = 1. \tag{4.9}$$

In order to apply Theorem 4.3, consider the function $W(y) = \mathbb{E}\left[\mathcal{W}|Q(0) = y\right]$, where the random variable $\mathcal{W}$ is defined as

$$\mathcal{W} := \sum_{n=0}^{\|Q(0)\|T} [1 + \|Q(0)\| + a\|Q(n)\|]^{-m}$$

for some positive constants a and T to be determined later and $m > 1$. Note that, with minor abuse of notation, $W(Q(0) = y, X(0) = x) = W(y)$, i.e., W only depends on the queue and not on the activity vector. The function $W(y)$ may be interpreted as the following approximation to a Lyapunov function for the fluid limit process

$$\|y\|^{m-1} W(y) \approx \mathbb{E}_{\hat{y}} \left[\int_0^T (1 + a\|q_{\hat{y}}(t/\beta)\|)^{-m} \mathrm{d}t \right] = V(q_{\hat{y}}(t)), \tag{4.10}$$

with equality when $\|y\| \to \infty$, and $\hat{y} = \frac{y}{\|y\|}$ is the initial state of the fluid limit process. Then it follows from the instability of the fluid limit process that we can choose a and T large enough such that $V(q_{\hat{y}}(t+r)) < V(q_{\hat{y}}(t))$ for any $r > 0$ and any initial state $\hat{y}$. This implies that

$$\|y\|^m \mathbb{E}\left[W(Q(n+1)) - W(Q(n))|\mathcal{F}_n\right] \leq -\text{constant}$$

when $y = Q(n)$ and $\|y\|$ is sufficiently large. Thus, we can apply Theorem 4.3.

The detailed arguments may be described as follows. First of all, note that

$$\mathbb{E}\left[W(Q(1)) - W(Q(0))|Q(0) = y, X(0) = x\right] = \mathbb{E}\left[\theta^1 \mathcal{W} - \mathcal{W}|Q(0) = y\right],$$

where θ^1 is the usual backward shift operator on the sample path space [35]. We write $\theta^1\mathcal{W} - \mathcal{W} = A + B + C$, where

$$A = -[1 + \|Q(0)\| + a\|Q(0)\|]^{-m},$$

$$\begin{aligned} B = \sum_{n=1}^{\|Q(0)\|T} & \left\{[1 + \|Q(1)\| + a\|Q(n)\|]^{-m}\right. \\ & \left. - [1 + \|Q(0)\| + a\|Q(n)\|]^{-m}\right\}, \end{aligned}$$

and

$$C = \sum_{n=\|Q(0)\|T+1}^{\|Q(1)\|T} [1 + \|Q(1)\| + a\|Q(n)\|]^{-m}.$$

The term "A" provides the negative drift and the other terms can be bounded as follows. Using the fact that $\|Q(1)\| \geq \|Q(0)\| - 1$, and noting that $[\cdot]^{-m}$ is a convex decreasing function, we have

$$B \leq \sum_{n=1}^{\|Q(0)\|T} m[\|Q(0)\| + a\|Q(n)\|]^{-m-1}. \tag{4.11}$$

Multiplying both sides by $\|Q(0)\|^m$, we see that

$$\|Q(0)\|^m B \leq \frac{m}{\|Q(0)\|} \sum_{n=1}^{\|Q(0)\|T} \left(1 + a\frac{\|Q(n)\|}{\|Q(0)\|}\right)^{-m-1}. \tag{4.12}$$

Let $Q(0) = y$ and $\hat{y} := y/\|y\|$. For any y, the random variable in the right-hand side (RHS) of (4.12) is bounded by mT, and hence

$$\limsup_{\|y\|\to\infty} \mathbb{E}_{\hat{y}} \left[\|y\|^m B\right] \leq \mathbb{E}_{\hat{y}} \left[m \int_0^T [1 + a\|q(s/\beta)\|]^{-m-1} \mathrm{d}s\right],$$

because of the weak limit convergence of $\frac{1}{\|y\|} Q^{(\|y\|)}(\|y\|t) \Rightarrow q(t/\beta)$ over $[0, T]$ and uniform integrability of the random variables of the form RHS of (4.12).

Next, for "C", it is sufficient to consider the case that $\|Q(1)\| = \|Q(0)\| + 1$, where

$$C \leq T[1 + \|y\| + 1 + a(\|Q(\|y\|T)\| - T)]^{-m}.$$

Similarly to "B", multiplying both sides with $\|y\|^m$ and taking the limit gives

$$\limsup_{\|y\|\to\infty} \mathbb{E}_{\hat{y}} \left[\|y\|^m C\right] \leq \mathbb{E}_{\hat{y}} \left[T[1 + a\|q(T/\beta)\|]^{-m}\right], \tag{4.13}$$

again, because $\|y\|^m C < T$ (thus, uniform integrability holds) and by the weak limit convergence. Putting the bounds together, we obtain

$$\begin{aligned}\limsup_{\|y\|\to\infty} \|y\|^m \mathbb{E}_{\hat{y}} \left[\theta^1 \mathcal{W} - \mathcal{W}\right] \quad \leq \quad & -(1 + a)^{-m} + m\mathbb{E}_{\hat{y}} \left[\int_0^\infty (1 + aL(s/\beta))^{-m-1} \mathrm{d}s\right] \\ & + \mathbb{E}_{\hat{y}} \left[T(1 + aL(T/\beta))^{-m}\right],\end{aligned}$$

because $\|q(s)\| \geq L(s)$ based on our notation with some initial state $q(0) = \hat{y}$ such that $\|\hat{y}\| = 1$. Consider the cycle pairs D_k, $k = 1, 2, \ldots$, as defined for Theorem 4.1. Then,

$$\begin{aligned}\mathbb{E}_{\hat{y}} \left[\int_0^\infty (1 + aL(s/\beta))^{-m-1} \mathrm{d}s\right] \quad & \leq \quad \beta\mathbb{E}_{\hat{y}} \left[\sum_{k=0}^{\infty} \int_{T_k}^{T_{k+1}} (1 + aL(s))^{-m-1} \mathrm{d}s\right] \\ & \leq \quad \beta\mathbb{E}_{\hat{y}} \left[\sum_{k=0}^{\infty} \int_{T_k}^{T_{k+1}} (1 + a\theta L_k)^{-m-1} \mathrm{d}s\right] \\ & \leq \quad \beta\mathbb{E}_{\hat{y}} \left[\sum_{k=0}^{\infty} \Delta T_k (a\theta L_k)^{-m-1}\right] \\ & \leq \quad \beta C_{LT} (a\theta)^{-m-1} \sum_{k=0}^{\infty} \mathbb{E}_{\hat{y}} \left[L_k^{-m}\right].\end{aligned}$$

Note that the times T_k are random variables in general and we have used the fact that $L_{k+1} \geq \theta L_k$ with $0 < \theta \leq 1$. As we saw in the proof of Theorem 4.1, for $\rho \in (\rho^*, 1]$, $\mathbb{E}\left[L_k^{-m}\right] \leq \alpha^k$. Therefore,

$$m\mathbb{E}_{\hat{y}}\left[\int_0^\infty (1+aL(s/\beta))^{-m-1}\mathrm{d}s\right] \leq m\beta C_{LT}(a\theta)^{-m-1}\frac{1}{1-\alpha}. \tag{4.14}$$

So, we can choose a large enough to ensure that the RHS of (4.14) is less than $\frac{1}{3}(1+a)^{-m}$. Next we show that we can choose T large enough such that

$$\mathbb{E}_{\hat{y}}\left[T[1+aL(T/\beta)]^{-m}\right] \leq \frac{1}{3}(1+a)^{-m}. \tag{4.15}$$

Note that

$$\mathbb{E}_{\hat{y}}\left[T[1+aL(T/\beta)]^{-m}\right] \leq a^{-m}\mathbb{E}_{\hat{y}}\left[TL^{-m}(T/\beta)\right], \tag{4.16}$$

and by Theorem 4.1, $\limsup_{T\to\infty}\mathbb{E}_{\hat{y}}\left[TL^{-m}(T)\right] = 0$, for $\rho \in (\rho^*, 1]$. Hence, we can choose T large enough such that (4.15) holds.

Therefore,

$$\limsup_{\|y\|\to\infty} \|y\|^m \mathbb{E}\left[W(Q(1)) - W(Q(0))|(Q(0), X(0)) = (y,x)\right] \leq -\frac{1}{3}(1+a)^{-m}.$$

This means that there exists a a positive constant $\|y_0\|$ such that,

$$\mathbb{E}\left[W(Y(1)) - W(Y(0))|Y(0) = (y,x)\right] \leq -\frac{1}{6}(1+a)^{-m}\|y\|^{-m},$$

whenever $\|y\| > \|y_0\|$. Let $c_0 = W(y_0) = W(\|y_0\|)$. On the other hand, it follows from (4.10) that $\limsup_{\|y\|\to\infty} W(y) = 0$, which means that A_{c_0} is well-defined and also c_0 can be made arbitrary small by letting $\|y_0\| \to \infty$. Therefore, the conditions of Theorem 4.3 are satisfied with $\Delta(y) = \frac{1}{6}(1+a)^{-m}\|y\|^{-m}$. This shows that

$$\mathbb{P}_{Q(0)}\left\{\sum_{n=0}^{\infty}\frac{\text{constant}}{\|Q(n)\|^m} < \infty\right\} \to 1, \tag{4.17}$$

as $\|Q(0)\| \to \infty$, which implies (4.9). $\square$

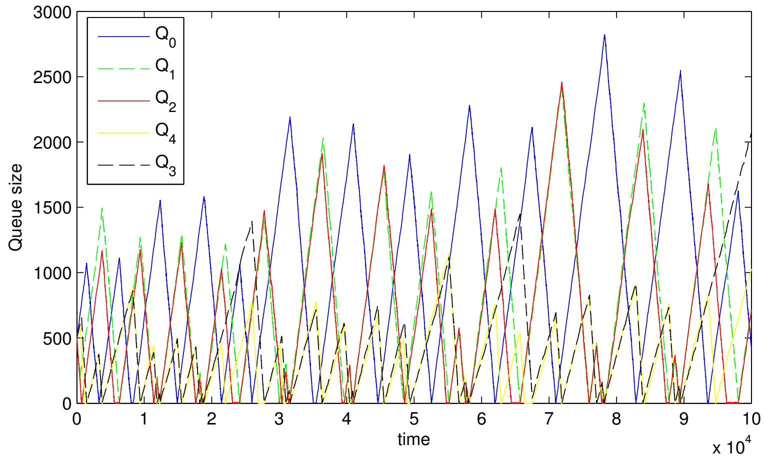

Figure 4.6: Queue sizes at the various nodes as a function of time for the network of Figure 4.3.

4.4 Simulation experiments

We now discuss the simulation experiments that we have conducted to support and illustrate the analytical results. Consider the broken-diamond network as depicted Figure 4.3 and considered in the previous sections. In the simulation experiments, the relative traffic intensities are assumed to be $\kappa_1 = \kappa_2 = 0.4$, $\kappa_3 = 0.4$, and $\kappa_6 = 0.2$ with $\alpha = 0$, for the components M_1, M_2, and M_3, respectively, with a normalized load of $\rho = 0.97$. At each node i, the initial queue size is $Q_i(0) = 500$, the activation function is $r_i(x) \equiv 1$, $x \geq 1$, and the de-activation function is $\hat{r}_i(x) = (1 + x)^{-\gamma}$ (or the weight function is $w_i(x) = \gamma \log(1 + x)$), where we set $\gamma = 2$.

Figure 4.6 plots the evolution of the queue sizes at the various nodes over time, and shows that once a node starts transmitting, it will continue to do so until the queue lengths of all nodes in its component have largely been cleared. This characteristic, and the associated oscillations in the queues, strongly mirror the qualitative behavior displayed by the fluid limit.

Although Figure 4.6 suggests an upward trend in the overall queue lengths, the fluctuations make it hard to discern a clear picture. Figure 4.7 therefore plots the evolution of the node-average queue size over time, and reveals a distinct growth pattern. Evidently, it is difficult to make any conclusive statements concerning stability/instability based on simulation results alone. However, the saw-tooth type growth pattern in Figure 4.7 demonstrates strong signs of instability, and corroborates the qualitative growth

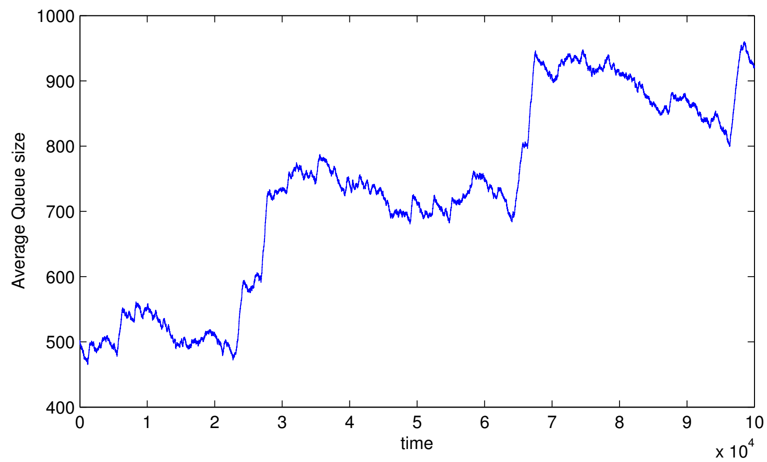

Figure 4.7: Node-average queue size as function of time for the network of Figure 4.3.

behavior exhibited by the fluid limit. Indeed, careful inspection of the two figures confirms that the large increments in the node-average queue size occur immediately after M_4-periods, exactly as predicted by the fluid limit. We further observe that in between these periods, the node-average queue size tends to follow a slightly downward trend, consistent with the negative drift of rate $(\rho - 1)/3$ in the fluid limit.

4.5 Instability in general interference graphs

We have used fluid limits to demonstrate the potential instability of queue-based random access algorithms. For the sake of transparency, we focused on a specific six-node network. Similar instability issues, however, can arise in a far broader class of interference graphs, as we will discuss in the following.

Consider a general interference graph $G = (V, E)$. Without loss of generality, we can assume G is connected, because otherwise we can consider each connected subgraph separately. For $\gamma > 1$, the fluid limit sample paths still exhibit the *sawtooth behavior*, i.e., when a node starts transmitting, it does not release the channel until its entire queue is cleared (on the fluid scale). Let $\bar{\mathcal{M}} = \{M_1, \ldots, M_K\}$ denote the set of maximal independent sets (maximal schedules) of G. We say the network operates in M_i if a subset $W \subseteq M_i$ of nodes are served at full rate (on the fluid scale), and W does not belong to any other maximal schedules M_j, $j \neq i$. Under the random access algorithm, at any point in time the network operates in one of the maximal schedules and switches to another maximal schedule when one or several of the queues in the current maximal schedule drain (on the fluid scale). More specifically, assume the network operates in a maximal schedule M_i. If M_i interferes with all other maximal schedules, i.e., $M_i \cap M_j = \emptyset$ for all $1 \leq j \leq K$, $j \neq i$, then a transition from M_i to any maximal schedule M_j, $j \neq i$, is possible when all the queues in M_i hit zero (on the fluid scale). On the other hand, if M_i overlaps with a subset of maximal schedules $\bar{\mathcal{M}}_i' := \{M_j \in \bar{\mathcal{M}} : M_i \cap M_j \neq \emptyset\}$, then the activity process can make a transition to $M_j \in \bar{\mathcal{M}}_i'$ when all the queues in $M_i \backslash M_j$ drain (on the fluid scale).

The capacity region of the network Λ is full-dimensional because all the basis vectors of $\mathbb{R}^N$ belong to that set. The incidence vectors of the sets $\bar{\mathcal{M}}$ correspond to the extreme points of $\mathcal{C}$ as they cannot be expressed as convex combinations of other points. Consider a covering of $V = \{1, 2, \ldots, N\}$ using the maximal schedules. Formally, a set cover C of V is a collection of maximal schedules such that $V \subseteq \cup_{M_i \in C} M_i$. A set cover C is minimal if removal of any of the elements $M_i \in C$ leaves some nodes of V uncovered. Consider the class of graphs in which $|C| \leq K - 1$ for some minimal set cover C, i.e., we do not need all M_i's for covering V. Without loss of generality, let $\bar{\mathcal{M}}^* = \{M_1, M_2, \ldots, M_{K^*}\}$ denote such a minimal cover with $K^* \leq K - 1$. Consider a (strictly positive) vector of arrival rates $\lambda = \rho \sum_{i=1}^{K^*} \sigma_i 1_{M_i}$ where $\sigma_i > 0$, $1 \leq i \leq K^*$, such that $\sum_{i=1}^{K^*} \sigma_i = 1$, and $0 < \rho < 1$ is the load factor. Hence, a centralized algorithm can stabilize the network by scheduling each $M_i \in \bar{\mathcal{M}}^*$ for at least a fraction $\rho\sigma_i$ of the time. However, under the random access algorithm, the network might spend a non-vanishing fraction of time in the schedules $\bar{\mathcal{M}} \backslash \bar{\mathcal{M}}^*$, which can cause instability as ρ approaches 1. This phenomenon is easier to observe in graphs with a *unique* minimal set cover $\bar{\mathcal{M}}^*$ and with a maximal schedule M_1 interfering with all the other maximal schedules, hence $M_1 \in \bar{\mathcal{M}}^*$.

This means any valid covering of V must contain $\bar{\mathcal{M}}^*$. Therefore, considering arrival rate vectors of the form $\lambda = \rho \sum_{i=1}^{K^*} \sigma_i 1_{M_i}$, $\sigma_i > 0$, $\sum_{i=1}^{K^*} \sigma_i = 1$, *the only way* to stabilize the network is to use M_i for a time fraction greater than $\rho\sigma_i$. Visits to M_1 have to occur infinitely often, otherwise the network is trivially unstable, and at the end of such visits, a transition to any other maximal schedule is possible, including the schedules in $\bar{\mathcal{M}} \backslash \bar{\mathcal{M}}^*$ with positive probability. Then, upon entrance to schedules in $\bar{\mathcal{M}} \backslash \bar{\mathcal{M}}^*$, the network

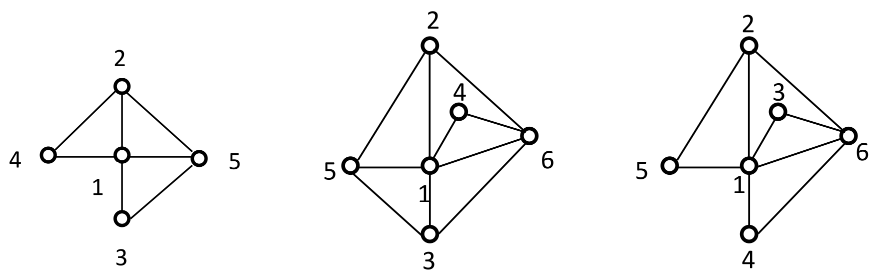

(a) $\{1\}, \{2,3\}, \{4,5\}$ (b) $\{1\}, \{2,3,4\}, \{5,6\}$ (c) $\{1\}, \{2,3,4\}, \{5,6\}$

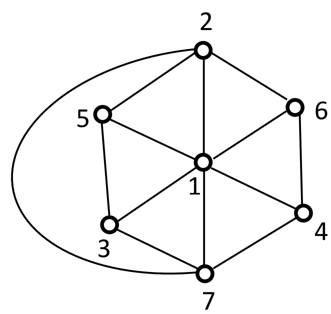

(d) $\{1\}, \{2,3,4\}, \{5,6,7\}$

Figure 4.8: A few unstable networks with their unique minimal cover $\bar{\mathcal{M}}^*$ using the maximal schedules.

spends a positive time in such schedules because the queues in $\bar{\mathcal{M}} \backslash \{M_1\}$ build up during visits to M_1. Hence, the arguments in the instability proof of the broken-diamond network can be extended to such networks, although a rigorous proof of the fluid limits in such general cases remains a formidable task. Figure 4.8 shows a few examples of such unstable networks with unique minimal set covers.

4.6 Instability for de-activation functions with polynomial decay

Our instability arguments suggest that instability can in fact occur for any activity factor that grows as a positive power $1/k$ of the queue length for network sizes of order k. In terms of weight functions of Chapter 3, the results imply that weight functions $\gamma \log(\cdot)$, with $\gamma > 1/k$, can cause instability for network sizes of order k, as will be described below.

Consider any unstable network $G = (V, E)$, for example the broken-diamond network or a graph as described in Section 4.5. Let $\mathcal{I}(i)$ denote the set of neighbors of node i in G. We construct a *k-duplicate graph* $G^{(k)}$, $k \in \mathbb{N}$, of G as follows. For each node $i \in V$, add k duplicate nodes $d_1^{(i)}, \ldots, d_k^{(i)}$ to the graph, with the same arrival rate λ_i and the same initial queue length $Q_i(0)$, such that each node $d_j^{(i)}$ is connected to all the neighbors of node i and their duplicates, i.e., $\mathcal{I}(d_j^{(i)}) = \mathcal{I}(i) \cup_{l \in \mathcal{I}(i)} \{d_1^{(l)}, \cdots, d_k^{(l)}\}$, for all $1 \leq j \leq k$. For notational convenience, we define $D_i^{(k)} := \{i, d_1^{(i)}, \ldots, d_k^{(i)}\}$ and call it the duplicate collection of node i. Note that the duplicate graph has the same number of maximal schedules as the original graph. In fact, each maximal schedule $M_i^{(k)}$ of $G^{(k)}$ consists of nodes in the maximal schedule M_i of G and their duplicates, i.e., $M_i^{(k)} = \cup_{l \in M_i} D_l^{(k)}$. Next, we show that the duplicate graph is unstable for de-activation functions that decay as $o(x^{-\gamma})$, for $\gamma > 1/(k+1)$. Essentially, for such a range of γ, each duplicate collection acts as a super node with $\gamma > 1$, i.e., (i) if one of the queues in a duplicate collection $D_i^{(k)}$ starts growing, all the queues in $D_i^{(k)}$ grow linearly at the same rate λ_i (on the fluid scale), (ii) if a nonempty queue in $D_i^{(k)}$ starts draining, then all the queues in $D_i^{(k)}$ drain at full rate until they all hit zero (on the fluid scale). Then the instability follows from that of the original network, as we can simply regard the duplicate collections as super nodes. An informal proof of claims (i) and (ii) is presented next.

Claim (i) is easy to prove as all the queues in a duplicate collection share the same set of conflicting neighbors and the fact that one of the queues grows, over a small time interval, implies that some conflicting neighbors are transmitting over such interval. To show (ii), note that if one of the queues in the duplicate collection drains over a nonzero time interval, no matter how small the interval is, all the conflicting neighbors must be in backoff for $O(R)$ units of time in the pre-limit process. This guarantees that all the queues in the duplicate collection will start a packet transmission during such interval almost surely. As long as the duplicate collection does not lose the channel, each queue of the collection follows the fluid limit trajectory of an M/M/1 queue. Suppose all the queues of the duplicate collection are above a level ϵ on the fluid scale for some fixed small $\epsilon > 0$. Thus, in the pre-limit process, the amount of time required for the queues to fall below a threshold ϵR is $O(R)$ with high probability as $R \to \infty$. The duplicate collection loses the channel if and only if all $k + 1$ nodes in the collection are in backoff and a conflicting node acquires the channel by winning the competition between the backoff timers. The probability that a node goes into backoff at the end of a packet transmission is $O((\epsilon R)^{-\gamma})$, or approximately the fraction of time that a node spends in backoff is $O((\epsilon R)^{-\gamma})$. Therefore, the fraction of time that all $k+1$ nodes of the duplicate collection are simultaneously in backoff is $O((\epsilon R)^{-k\gamma})$ because the nodes in the duplicate collection act independently from each other. Therefore, over an interval of duration $O(R)$, the amount of time that all $k+1$ nodes are in backoff is $O(R^{1-(k+1)\gamma})$, which goes to zero as $R \to \infty$ if $\gamma > 1/(k+1)$. Thus, the nodes in the duplicate collection follow the fluid limits of an M/M/1 queue until their backlog is below ϵ on the fluid scale. Since ϵ could be made arbitrarily small, we can view the duplicate collection as a super node that does not release the channel until its backlog hits zero. This demonstrates the instability of fluid limits for the initial queue lengths described above for the duplicate network.

To rigorously prove instability of the original process using the framework of Meyn [35], we need to show instability of the fluid limit for any initial state. Handling arbitrary initial states for general activity functions and interference graphs is more involved than in the specific broken-diamond network considered here. An alternative option would be to extend the methodology and develop a proof apparatus where it suffices to show instability of the fluid limit for one particular initial state. The framework of Dai [33] offers the advantage that instability of the fluid limit only needs to be shown for an all-empty initial state. However the characterization of the fluid limit for an all-empty initial state appears to involve additional complications.

The above proof arguments suggest that instability can in fact occur for any $\gamma > 0$ as k can be chosen arbitrarily large. This indicates that the growth conditions in Ghaderi and Srikant [27] (Chapter 3) are sharp in the sense that the weight functions of the form $w(x) = \log(1+x)/g(x)$, where $g(x)$ is an arbitrarily slowly increasing function, are essentially the most aggressive weight functions that guarantee maximum stability in any general topology.

4.7 Conclusions

We have used fluid limits to demonstrate the potential instability of queue-based random access mechanisms. For the sake of transparency, we focused on instability for weights that grow faster than $\gamma \log(\cdot)$, for any $\gamma > 1$, but proof arguments suggest that instability can occur for any $\gamma > 0$, in networks with sufficiently many nodes. In other words, the "near-logarithmic growth condition" on the weights is a fundamental limit on the aggressiveness of nodes to ensure maximum stability in any general topology.

4.8 Additional proofs

4.8.1 Proof of Lemma 4.1

The proof relies on basic sample path properties of the fluid limit process $\{q(t)\}$ as described in Section 4.2.1. First of all, the M_1-period that initiates the i-th cycle ends at time $t_i + T_{i1}$, with

$$T_{i1} = \max\left\{\frac{q_1(t_i)}{1-\rho_1}, \frac{q_2(t_i)}{1-\rho_2}\right\} \leq \frac{\max\{q_1(t_i), q_2(t_i)\}}{1-\rho_0} \leq \frac{L(t_i)}{1-\rho_0}.$$

Define $K(t) = \max\{q_3(t), q_4(t)\} + \max\{q_5(t), q_6(t)\}$ and recall that $\rho = \rho_0 + \rho_3 + \rho_6$. Then

$$\begin{aligned} K(t_i + T_{i1}) &\leq K(t_i) + (\rho_3 + \rho_6) T_{i1} \\ &\leq L(t_i) - \max\{q_1(t_i), q_2(t_i)\} + (\rho_3 + \rho_6)\frac{\max\{q_1(t_i), q_2(t_i)\}}{1-\rho_0} \\ &= L(t_i) - \frac{(1-\rho)\max\{q_1(t_i), q_2(t_i)\}}{1-\rho_0} \\ &= L(t_i) - (1-\rho) T_{i1}, \end{aligned}$$

which may also be seen from the fact that $L(t)$ decreases at a rate $1-\rho$ or larger during the time interval $[t_i, t_i + T_{i1}]$ and $K(t_i + T_{i1}) = L(t_i + T_{i1})$ since $q_1(t_i + T_{i1}) = q_2(t_i + T_{i1}) = 0$.

Define $T_0 = \frac{K(t_i + T_{i1})}{1-\rho_3-\rho_6}$. We distinguish between two cases, depending on whether an M_4-period starts before time $t_i + T_{i1} + T_0$ or not.

If no M_4-period occurs before time $t_i + T_{i1} + T_0$, then $K(t)$ decreases at a rate $1 - \rho_3 - \rho_6$ or larger for all $t \in [t_i + T_{i1}, t_i + T_{i1} + T_0]$ and reaches zero no later than time $t_i + T_{i1} + T_0$, unless an M_1-period intervenes. This implies that the next M_1-period must start no later than time $t_i + T_{i1} + T_0$.

Using the above results, a simple calculation shows that

$$\Delta t_i \le T_{i1} + T_0 \le T_{i1} + \frac{L(t_i) - (1-\rho)T_{i1}}{1-\rho_3-\rho_6} \le \frac{L(t_i)}{(1-\rho_0)(1-\rho_3-\rho_6)} \le C_T L(t_i).$$

Also, $L(t)$ has continuously decreased during the cycle, so $L(t_{i+1}) - L(t_i) \le 0$.

Now suppose that an M_4-period does start at some time $t_0 \in [t_i + T_{i1}, t_i + T_{i1} + T_0]$, and ends at time u_0.

Since $K(t)$ decreases at a rate $1-\rho_3-\rho_6$ or larger during the time interval $[t_i + T_{i1}, t_0]$, it follows that

$$K(t_0) \le K(t_i + T_{i1}) - (1-\rho_3-\rho_6)(t_0 - t_i - T_{i1}).$$

Noting that $q_4(t_0), q_5(t_0) \le K(t_0)$, we conclude that the duration of the M_4-period is no longer than

$$u_0 - t_0 \le \min\left\{\frac{q_4(t_0)}{1-\rho_4}, \frac{q_5(t_0)}{1-\rho_5}\right\} \le \frac{K(t_0)}{1-\max\{\rho_4,\rho_5\}}.$$

Since $K(t)$ increases at a rate no larger than $\rho_3+\rho_6$ during the time interval $[t_0, u_0]$, it follows that

$$K(u_0) \le K(t_0) + (\rho_3+\rho_6)(u_0 - t_0).$$

The M_4-period will cause the queue of node 4 to empty at some point and become smaller than the queue of node 3, and likewise the queue of node 5 must empty at some point and become smaller than the queue of node 6. Because M_4-periods can no longer be initiated from M_2 and M_3, $K(t)$ decreases at a rate $1-\rho_3-\rho_6$ or larger from time u_0 onward, and reaches zero no later than time $u_0 + \frac{K(u_0)}{1-\rho_3-\rho_6}$, unless an M_1-period intervenes. This implies that the next M_1-period must start no later than time $u_0 + \frac{K(u_0)}{1-\rho_3-\rho_6}$.

Combining the above results, we obtain

$$\begin{aligned}
\Delta t_i &\leq u_0 + \frac{K(u_0)}{1-\rho_3-\rho_6} - t_i = T_{i1} + (t_0 - t_i - T_{i1}) + (u_0 - t_0) + \frac{K(u_0)}{1-\rho_3-\rho_6} \\
&\leq T_{i1} + (t_0 - t_i - T_{i1}) + \frac{K(t_0) + (u_0 - t_0)}{1-\rho_3-\rho_6} \\
&\leq T_{i1} + (t_0 - t_i - T_{i1}) + \left(1 + \frac{1}{1-\max\{\rho_4,\rho_5\}}\right)\frac{K(t_0)}{1-\rho_3-\rho_6} \\
&\leq T_{i1} - \frac{t_0 - t_i - T_{i1}}{1-\max\{\rho_4,\rho_5\}} + \frac{(2-\max\{\rho_4,\rho_5\})K(t_i+T_{i1})}{(1-\max\{\rho_4,\rho_5\})(1-\rho_3-\rho_6)} \\
&\leq T_{i1} + \frac{(2-\max\{\rho_4,\rho_5\})(L(t_i) - (1-\rho)T_{i1})}{(1-\max\{\rho_4,\rho_5\})(1-\rho_3-\rho_6)} \\
&\leq \frac{(2-\max\{\rho_4,\rho_5\})L(t_i) + \rho_0(1-\max\{\rho_4,\rho_5\})T_{i1})}{(1-\max\{\rho_4,\rho_5\})(1-\rho_3-\rho_6)} \\
&= \frac{L(t_i)}{(1-\max\{\rho_4,\rho_5\})(1-\rho_3-\rho_6)} + \frac{L(t_i) + \rho_0 T_{i1}}{1-\rho_3-\rho_6} \\
&\leq \frac{L(t_i)}{(1-\max\{\rho_4,\rho_5\})(1-\rho_3-\rho_6)} + \frac{L(t_i)}{(1-\rho_0)(1-\rho_3-\rho_6)} \\
&= \frac{L(t_i)}{1-\rho_3-\rho_6}\left(\frac{1}{1-\rho_0} + \frac{1}{1-\max\{\rho_4,\rho_5\}}\right) = C_T L(t_i).
\end{aligned}$$

Also, $L(t)$ has only increased during the M_4-period at a rate no larger than $\rho = \rho_0 + \rho_3 + \rho_6$, so

$$L(t_{i+1}) - L(t_i) \leq \rho(u_0 - t_0) \leq \frac{\rho K(t_0)}{1-\max\{\rho_4,\rho_5\}} \leq \frac{\rho L(t_i)}{1-\max\{\rho_4,\rho_5\}} = C_L L(t_i).$$

4.8.2 Proof of Lemma 4.2

Denote by t_1 and t_2 the times that the cycles start and by u_1 and u_2 the times that the M_1-periods end. First assume $\max\{q_1(t_1), q_2(t_1)\} \leq \epsilon L(t_1)$. Then, $\max\{q_3(t_1), q_4(t_1)\} + \max\{q_5(t_1), q_6(t_1)\} \geq (1-2\epsilon)L(t_1)$, so we must have $\max\{q_3(t_1), q_4(t_1)\} \geq (1-2\epsilon)L(t_1)/2$ or $\max\{q_5(t_1), q_6(t_1)\} \geq (1-2\epsilon)L(t_1)/2$. In the former scenario, with probability 3/8 the M_1-period is followed by an M_2-period, which will last for an amount of time no less than $\max\{\frac{q_3(t_1)}{1-\rho_3}, \frac{q_4(t_1)}{1-\rho_4}\} \geq \frac{\max\{q_3(t_1), q_4(t_1)\}}{1-\rho_4} \geq \frac{(1-2\epsilon)L(t_1)}{2(1-\rho_4)}$. Likewise, in the latter scenario, with probability 3/8 the M_1-period is followed by an M_3-period, which will last for an amount of time no less than $\max\{\frac{q_5(t_1)}{1-\rho_5}, \frac{q_6(t_1)}{1-\rho_6}\} \geq \frac{\max\{q_5(t_1), q_6(t_1)\}}{1-\rho_5} \geq \frac{(1-2\epsilon)L(t_1)}{2(1-\rho_5)}$. Thus, in either scenario, with probability at least 3/8, the time until the start of the next cycle is at least $\frac{(1-2\epsilon)L(t_1)}{2(1-\min\{\rho_4,\rho_5\})}$, so that

$$\max\{q_1(t_2), q_2(t_2)\} \geq q_2(t_2) \geq \frac{\rho_2(1-2\epsilon)L(t_1)}{2(1-\min\{\rho_4,\rho_5\})}.$$

Invoking the fact that $L(t_2) \leq C_L L(t_1)$, with C_L as defined in Lemma 4.1, we find that

$$\max\{q_1(t_2), q_2(t_2)\} \geq \epsilon L(t_2),$$

with ϵ as specified in the statement of the lemma.

Now consider a cycle with $\max\{q_1(t_k), q_2(t_k)\} \geq \epsilon L(t_k)$, $k = 1, 2$. Then

$$q_i(u_k) = q_i(t_k) + \rho_i \frac{\max\{q_1(t_k), q_2(t_k)\}}{1-\rho_2}, \text{ for } i = 3, 4, 5, 6.$$

Note that $0 \leq q_i(t_k) \leq (1-\epsilon)L(t_k)$, $i = 3, 4, 5, 6$, and $\epsilon L(t_k) \leq \max\{q_1(t_k), q_2(t_k)\} \leq L(t_k)$. Then it is easily verified that the queues are weakly balanced at time u_k with $\beta^{\min}$ and $\beta^{\max}$ as given in the statement of the lemma.

4.8.3 Fluid limit proofs: Part A

Prelimit model

We start with the time-homogeneous Markov process $(\mathbf{X}(t), \mathbf{Q}(t)), t \geq 0$ with state space $\mathcal{M} \times \mathbb{N}_0^N$ where $N = 6$ and $\mathcal{M} \subseteq \{0,1\}^N$ is the set of feasible schedules. We recap to state that service times are unit exponential as are backoff periods. In addition the Poisson arrival processes are determined by the vector of arrival rates $\boldsymbol{\lambda}$ and the probability of backoff is determined as a function of queue length $1/(1+Q)^\gamma$ with $\gamma \in (1, \infty)$.

The fluid limit will not be obtained directly from the above process but rather via the *jump chain* of a *uniformized version* with "clock ticks" from a Poisson clock with constant rate,

$$\beta \doteq \sum_{\ell=1}^{N} \lambda_\ell + N, \tag{4.18}$$

independent of state, with null (dummy) events introduced as needed.

With minor abuse of notation, denote by $(\mathbf{X}(n), \mathbf{Q}(n)) \in \mathcal{S}$ to be the state of the jump chain at nth clock tick. For our subsequent construction, it will be convenient to replace $\mathbf{X}(n)$ with the cumulative state $\mathbf{I}(n) = \sum_{k=0}^{n} \mathbf{X}(n) \in \mathbb{N}_0^N$, which is by definition increasing. It determines and is determined by the sequence $\mathbf{X}(n)$ and the associated jump chain is Markov if the state is altered to be $(\mathbf{I}(n), \mathbf{I}(n-1), \mathbf{Q}(n))$ with $\mathbf{I}(-1) = 0$. From the jump chain, we obtain a continuous stochastic process in $C[0, \infty)$ by linear interpolation and by accelerating time by a factor β. To be specific, at an arbitrary intermediate time $t > 0$ between two clock ticks $t_l = (k-1)/\beta \leq t \leq t_u = k/\beta,\ k \in \mathbb{N}$, the interpolated process takes the values

$$\begin{aligned} \overline{\mathbf{Q}}(t) &\doteq \beta(t_u - t)\mathbf{Q}(k-1) + \beta(t - t_l)\mathbf{Q}(k), \\ \bar{\mathbf{I}}(t) &\doteq \beta(t_u - t)\mathbf{I}(k-1) + \beta(t - t_l)\mathbf{I}(k). \end{aligned}$$

From this construction we can obtain a sequence of such processes, indexed by $R \in \mathbb{N}$, with the usual fluid limit scaling

$$\left(\mathbf{Q}^R(t), \mathbf{I}^R(t)\right) \doteq \left(\frac{1}{R}\overline{\mathbf{Q}}^{(R)}(Rt), \frac{1}{R}\bar{\mathbf{I}}^{(R)}(Rt)\right). \tag{4.19}$$

This is obtained together with a corresponding sequence of initial queue lengths

$$\mathbf{Q}^R(0) = \frac{1}{R}\overline{\mathbf{Q}}^{(R)}(0) \to \mathbf{q}(\mathbf{0}). \tag{4.20}$$

Recall that the underlying jump chain $(\mathbf{Q}(n), \mathbf{X}(n))_{n \geq 0}$ is affected only through the initial state. Its transition probabilities are unaffected. The convergence in (4.20) is with respect to the Euclidean norm and without loss of generality we may take $||\mathbf{q}(\mathbf{0})|| = 1$.

For every R and time $t \geq 0$, $\left(\mathbf{Q}^R(t), \mathbf{I}^R(t)\right)$ take values in $E \doteq \mathbb{R}_+^N \times \mathbb{R}_+^N$, which is therefore the *state space* of the process. E has the usual Euclidean metric and associated topology and we will denote the Borel sets by $\mathcal{B}_E$. Furthermore the underlying jump chain $(\mathbf{Q}(n), \mathbf{X}(n))_{n \geq 0}$ of the uniformized Markov process satisfies the "skip-free property" [35] which ensures that the jumps between states are bounded in $\mathcal{L}_1$. It follows that the interpolated paths are Lipschitz continuous with Lipschitz constant $3\beta < \infty$. This property is conferred on the sample paths ω themselves as stated below

$$||\mathbf{Q}^R(t, \omega) - \mathbf{Q}^R(s, \omega)|| + ||\mathbf{I}^R(t, \omega) - \mathbf{I}^R(s, \omega)|| \leq 3\beta\,(t - s)\,, \tag{4.21}$$

which holds $\forall \omega, 0 \leq s < t, R \in \mathbb{N}$. The factor 3 appears since at most two queues can be active at the same time and at each clock tick at most one queue can be in(de)cremented.

To summarize, the scaled sequence of processes as defined in (4.19) take values in the space $C[0,\infty)$ of continuous paths taking values in E, endowed with the supnorm topology, and σ-algebra $\mathcal{C}$ generated by the open sets. This is obtained through the usual metric ρ_C as defined in [46], page 6. This space is both separable and complete, see [46] Theorem 2.1. The probability measure induced on $\mathcal{C}$ by the Rth interpolated process (4.19) is denoted μ_R so that $\mu_R(A)$ is the probability of an event $A \in \mathcal{C}$. Finally, it is of course the case that the jump chain sequence determines and is determined by the corresponding interpolated path. Hence μ_R and the jump chain probabilities are equivalent, given the initial conditions.

Fluid limit

If there is an infinite subsequence, $R_{k_1}, R_{k_2}, \ldots$ such that $\mu_{R_{k_n}} \Rightarrow \mu$ where $\Rightarrow$ denotes weak convergence, then μ is said to be a fluid limit measure. If such a fluid limit exists then the corresponding process can be defined as follows. Its *state space* is again E with underlying sample space $C[0,\infty)$ and corresponding σ-algebra $\mathcal{C}$ generated by the open sets under the metric, ρ_C, as mentioned earlier. This is the same space as for the sequence of prelimit processes. With the fluid limit measure μ (including the deterministic initial conditions) we have an underlying probability space $(C[0,\infty), \mathcal{C}, \mu)$. The stochastic process, $(\mathbf{q}, \mathbf{I})$ is the mapping $[0,\infty) \times C[0,\infty) \to E$ with values $(\mathbf{q}(t,\omega), \mathbf{I}(t,\omega)) \in E$. The curves $(\mathbf{q}(.,\omega), \mathbf{I}(.,\omega))$ and ω itself are the same. While these definitions are somewhat redundant, nevertheless in what follows, it will be convenient to think of a sample path as either a point ω or as a random function. Finally, on some occasions, we will use the notation $X \in m\mathcal{C}$ to indicate that $X : C[0,\infty) \to \mathbb{R}$ is measurable.

The proof of the next theorem is standard and follows from Lipschitz continuity, Theorem 8.3 of [1], and Lemma 3.1 of [40]. The details are omitted for brevity.

Theorem 4.4. *The sequence of measures μ_R defined on $(C[0,\infty), \mathcal{C})$ is tight.*

Thus, it follows from Prohorov's Theorem (Theorem 6.1 of [45]) that the sequence μ_R is relatively compact and fluid limit measure μ must exist. We suppose without loss of generality that $\mu_R \Rightarrow \mu$. The sample paths under μ have the same Lipschitz constant 3β. It follows that the sample paths of μ are differentiable a.e., almost surely [47].

Lipschitz continuity also implies that there are only a countable number of closed intervals $[a,b]$, $0 \leq a < b$, such that $q_\ell(a,\omega) = q_\ell(b,\omega) = 0$, $\ell = 1,\cdots,N$, and $q_\ell(x,\omega) > 0, \forall x \in (a,b)$, $\ell = 1,\cdots,N$, holding almost surely.

We denote by $\{\mathcal{F}_t\}_{t\in[0,\infty)}$, $\mathcal{F}_t \subset \mathcal{C}$, the filtration of sub σ-algebras generated by the open sets restricted to the interval $[0,t]$. The process $(\mathbf{q},\mathbf{I})$ is adapted to $\{\mathcal{F}_t\}_{t\in[0,\infty)}$.

By consideration of the weak law of large numbers and the existence of the fluid limit measure μ, it holds that

$$\mathbf{q}(t) = \mathbf{q}(0) + \boldsymbol{\lambda} t - \frac{1}{\beta}\mathbf{I}(t),\ t \geq 0. \tag{4.22}$$

This equation can be thought of as an accounting identity. If queue ℓ is active for a unit interval then I_ℓ increases by β, which corresponds (almost surely) to departures at unit rate. During the same period the arrival rate is λ_ℓ of course.

Since $I_\ell(t+h) - I_\ell(t) \leq \beta h$ for any node ℓ, and any times $t \geq 0$ and $h > 0$, it follows from (4.22) that

$$q_\ell(t+h) \geq q_\ell(t) + \lambda_\ell h - h,\ \mu\ a.s. \tag{4.23}$$

We now derive an elementary property of the fluid limit process. Given $t \geq 0, h > 0$, define

$$Y^\ell_{t,h} \doteq \{\omega : I_\ell(t+h,\omega) - I_\ell(t,\omega) = \beta h\} \tag{4.24}$$

to be the event that queue ℓ is being served (at maximum rate) during the interval $[t,t+h]$, i.e., the node is fully active during the given interval. Since many of the events that we consider later are in terms of activity, we adapt the following notation throughout the paper. In the case of (4.24),

$$Y^\ell_{t,h} = J^{(\ell)}_{=}(t,h,\beta h), \tag{4.25}$$

where the superscript "ℓ" denotes the node, "t" time and "h" duration. "βh" is the amount of activity which must be met with equality here, as indicated by the subscript "=". The subscript "=" may be replaced by $>$, $\geq$, $<$, or $\leq$, depending on the event.

Lemma 4.3 (No Conflict Lemma)**.** *Let $\ell_1 \neq \ell_2 \in \{1, \ldots, N\}$ be two neighbors in the interference graph G, and $h > 0, t \geq 0$, then*

$$\mu \left\{ Y_{t,h}^{\ell_1} \cap Y_{t,h}^{\ell_2} \right\} = 0. \tag{4.26}$$

Proof. This follows by definition, and the existence of the fluid limit. The event $Y_{t,h}^{\ell_1} \cap Y_{t,h}^{\ell_2}$ contradicts the inequality that for all $t \geq 0, h > 0$,

$$[I_{\ell_1}(t+h,\omega) - I_{\ell_1}(t,\omega)] + [I_{\ell_2}(t+h,\omega) - I_{\ell_2}(t,\omega)] \leq \beta h,$$

which holds μ almost surely. □

To obtain more detailed information with respect to the sample paths of μ, we proceed to the construction of sequences of stopping times.

Sequences of stopping times

The following definition is in connection with the amount of time a sample path for q_ℓ is positive, immediately prior to a time $z > 0$.

Definition 4.2. *Given a time $z > 0$, and $v, 0 < v \leq z$, and an $\ell = 1, \ldots, N$, define*

$$K_{z,v}^{(\ell)} \doteq \{\omega : q_\ell(z-s,\omega) > 0,\ \forall s \in (0,v)\}.$$

In words, $K_{z,v}^{(\ell)}$ is the set of sample paths for q_ℓ which are strictly positive in the interval $(z-v, z)$; if $z = 0$, $K_{z,v}^{(\ell)}$ is taken to be $\emptyset$.

Observe that it could be the case that either $q_\ell(z,\omega) = 0$ or $q_\ell(z-v,\omega) = 0$ (or both) and still $\omega \in K_{z,v}^{(\ell)}$. Finally note that it is possible for a given ω that no such v can be found, which requires that $q_\ell(z,\omega) = 0$ on account of continuity.

Definition 4.3. *Given a time $z \geq 0$ and a path ω, we define the mapping $A^{(\ell)}(z,\omega) : C[0,\infty) \to [0,z]$ to be*

$$A^{(\ell)}(z,\omega) \doteq \sup \left[\left\{ v : \omega \in K_{z,v}^{(\ell)} \right\} \cup \{0\} \right],$$

which is the time for which q_ℓ was positive immediately prior to z.

By definition, if $z \geq u > 0$ then $\{\omega : A^{(\ell)}(z,\omega) \geq u\} = K^{(\ell)}_{z,u}$, from which it follows that $A^{(\ell)}(z,\omega) \in m\mathcal{F}_z$. So far z has been fixed. However $A^{(\ell)}$: $[0,\infty) \times C[0,\infty) \to \mathbb{R}_+$ is a stochastic process carried by the underlying probability space $(C[0,\infty), \mathcal{C}, \mu)$ and $\mathcal{F}_t$-adapted. This process is piecewise linear and left-continuous (it falls to 0 immediately after q_ℓ returns to 0 from being positive). It follows that $A^{(\ell)}$ is $\mathcal{F}_t$-progressive, see for example [48]. We are now in a position to make the following definition.

Definition 4.4. *Given an $\mathcal{F}_t$ stopping time σ, a queue $\ell \in 1, \cdots, N$ and $m \in \mathbb{Z}_0 \doteq \mathbb{Z} - \{0\}$, define $T_{\ell,m}(\omega,\sigma) : C[0,\infty) \times [0,\infty] \to [0,\infty]$ as follows*

$$T_{\ell,m}(\omega,\sigma) \doteq \inf\left\{z \geq \sigma(\omega) : q_\ell(z) = 0, A^{(\ell)}(z,\omega) \in (e_m, f_m]\right\} \leq \infty,$$

where

$$\begin{aligned} f_m &= \frac{1}{m},\ e_m = \frac{1}{m+1};\ \text{for } m \in \mathbb{Z}_0,\ m > 0, \\ f_m &= |m-1|,\ e_m = |m|;\ \text{for } m \in \mathbb{Z}_0,\ m < 0, \end{aligned}$$

where again empty sets have an infinite infimum.

In words, $T_{\ell,m}$ is the earliest right-hand end of an open interval, with value z, such that q_ℓ is positive for a period $A^{(\ell)}(z,\omega) \in (e_m, f_m]$, immediately prior to $T_{\ell,m}$. If $z - f_m$ is the first time prior to z that $q_\ell = 0$, then z is in the set on the RHS. However, if this occurs at $z - e_m$, this is not the case.

It is plausible that $T_{\ell,m}$ is also an $\mathcal{F}_t$ stopping time, and we will subsequently prove this with particular choices for σ. We now state a construction lemma using a sequence of stopping times. These are returns to 0 following a fixed positive interval, in which we wait for a particular event A_k to occur.

Lemma 4.4 (Stop and Look Back)**.** *Let $\sigma \geq 0$ be an $\mathcal{F}_t$ stopping time and $a > 0$ a constant. Proposition 1.5 in Ethier and Kurtz [48] ensures that the following inductively defined sequence is a sequence of $\mathcal{F}_t$ stopping times: $s_0, s_1, s_2, \ldots$,*

$$\begin{aligned} s_0 &\doteq \sigma \\ s_k &\doteq \tau_c(\{0\}, s_{k-1} + a),\ k = 1, 2, \ldots \end{aligned} \tag{4.27}$$

Here, given an $\mathcal{F}_t$ stopping time $\sigma_1 > 0$, $\tau_c(\{0\}, \sigma_1) = \inf\{t \geq \sigma_1, q(t,\omega) = 0\}$. Now let $A_k \in \mathcal{F}_{s_k}, k = 1, 2, \ldots$ be a sequence of events in the pre-T σ-algebras *of the above stopping time sequence. Finally, define $\tau \doteq s_k$ if A_k occurs for the first time at step k and $\tau = \infty$ otherwise. Then τ is an $\mathcal{F}_t$ stopping time.*

Note that we do not check to see if A_k has occurred if $s_k = \infty$ at any stage, as τ is assigned this value regardless.

We now proceed to show the following.

Lemma 4.5. *Let $\sigma_0 \geq 0$ be an $\mathcal{F}_t$ stopping time such that $q_\ell(\sigma_0(\omega), \omega) = 0$ or $\sigma_0 = \infty$ and suppose ℓ, m are given. Let $a = e_m$ and $\sigma \doteq a + \sigma_0$ which is therefore an $\mathcal{F}_t$ stopping time, and $T_{\ell,m}$ be the mapping given in Definition 4.4. Then $T_{\ell,m}(\omega, \sigma)$ is an $\mathcal{F}_t$ stopping time.*

Proof. Given σ we will obtain a sequence of stopping times as in the first part of Lemma 4.4. However as we have already discussed, $A^{(\ell)}$ is $\mathcal{F}_t$-progressive, from which it follows by Proposition 1.4 of [48] that

$$A_k := A^{(\ell)}(s_k(\omega), \omega) \in m\mathcal{F}_{s_k}, \forall\ k = 1, 2, 3, \dots$$

so that $A^{(\ell)}(s_k(\omega), \omega) \in (e_m, f_m] \in \mathcal{F}_{s_k}$. Hence τ as defined in Lemma 4.4 is an $\mathcal{F}_t$ stopping time.

It remains to show that τ coincides with $T_{\ell,m}$ as defined. First suppose $\tau < \infty$, and immediately, $\tau \geq \sigma, q_\ell(\tau, \omega) = 0, A^{(\ell)}(\tau, \omega) \in (e_m, f_m]$ by definition. The fact that there is no earlier time satisfying these conditions follows since each s_k is a zero of q_ℓ and the construction rules out that the event could have taken place at any earlier time. The case $\tau = \infty$ coincides with there being no zero satisfying the required conditions. □

We now make the following recursive definitions.

Definition 4.5. *Given $m \in \mathbb{Z}_0$ and queue $\ell \in \{1, \dots, N\}$, let τ_0 be the first entry of $q_\ell(t, \omega)$ into 0 (τ_0 is an $\mathcal{F}_t$ stopping time). Then $Z^\ell_{m,0}$ is defined as*

$$\begin{aligned}
Z^\ell_{m,0} &\doteq T_{\ell,m}(\omega, 0);\ \textit{if}\ q_\ell(0,\omega) = 0 \\
Z^\ell_{m,0} &\doteq \tau_0;\ \textit{if}\ q_\ell(0,\omega) > 0,\ \tau_0 \in (e_m, f_m] \\
Z^\ell_{m,0} &\doteq T_{\ell,m}(\omega, \tau_0);\ \textit{if}\ q_\ell(0,\omega) > 0,\ \tau_0 \notin (e_m, f_m],
\end{aligned}$$

and subsequent stopping times are defined as

$$Z^\ell_{m,n} \doteq T_{\ell,m}(\omega, Z^\ell_{m,n-1}),\ n = 1, 2, 3, \dots.$$

The value of the stopping time is taken to be ∞ if the events do not occur. With obvious notation, we also define

$$A^{\ell}_{m,n}(\omega) \doteq A^{(\ell)}(Z^{\ell}_{m,n}, \omega)$$

to be the actual amount of time that the queue ℓ is positive prior to $Z^{\ell}_{m,n}$, and $A^{\ell}_{m,n} = \infty$ in case $Z^{\ell}_{m,n} = \infty$. Finally define the time at which q_ℓ last enters $(0, \infty)$ prior to $Z^{\ell}_{m,n}$ to be

$$V^{\ell}_{m,n} \doteq Z^{\ell}_{m,n} - A^{\ell}_{m,n},$$

when $Z^{\ell}_{m,n} < \infty$ and $V^{\ell}_{m,n} = \infty$ otherwise. Note that $V^{\ell}_{m,n} \in m\mathcal{F}_{Z^{\ell}_{m,n}}$ and is thus a non-negative random variable but not *a stopping time.*

The following corollary follows immediately from Lemma 4.5 and Definition 4.5.

Corollary 4.2. *$Z^{\ell}_{m,n}, n \in \mathbb{N}_0$ is a strictly increasing sequence of $\mathcal{F}_t$ stopping times, $\forall \ell \in N, m \in \mathbb{Z}_0$.*

This completes our goal of constructing sequences of stopping times for the queue processes. For any queue ℓ, by construction and by Lipschitz continuity, it follows that the set of stopping times, $Z^{\ell}_{m,n}$ determine all intervals where q_ℓ is positive for any sample path almost surely.

For each $m \in \mathbb{Z}_0$, and $\ell \in \{1, \cdots, N\}$, define

$$B^{\ell}_{m} \doteq \sup_{n} \left\{ Z^{\ell}_{m,n} : Z^{\ell}_{m,n} < \infty \right\}$$

to be the supremum of the finite stopping times for positive intervals with duration in $(e_m, f_m]$. If there is an m such that $B^{\ell}_{m} = \infty$, then q_ℓ returns to 0 infinitely often. Otherwise there is a $B > 0, B > B^{\ell}_{m}, \forall m \in \mathbb{Z}_0$. In this case, either q_ℓ remains at 0 as $t \to \infty$, or queue ℓ never returns to 0.

Piecewise linearity and no backoff until empty

So far the backoff exponent $\gamma > 1$ has not been taken into consideration, but from now on it will be. The following lemma bounds the probability, for the jump chain, that node ℓ has a backoff before its queue gets "small" provided that it was active earlier. Given some number $Q_T \in \mathbb{N}$ define,

$$K_{Q_T} \doteq \{\exists n : Q_\ell(k) \geq Q_T, 0 \leq k \leq n, X_\ell(n) = 0, X_\ell(n-1) = 1\}$$

to be the event that queue ℓ has remained above Q_T and has had a backoff at step n. We then have the following lemma.

Lemma 4.6 (No Early Backoff)**.** *Given any $Q_0 > Q_T$,*

$$\mathbb{P}\left(K_{Q_T} | Q_\ell(0) \geq Q_0\right) \leq \frac{1}{1-\lambda_\ell} \times \sum_{r=Q_T}^{\infty} \frac{1}{(1+r)^\gamma} = \varepsilon_{Q_T}. \tag{4.28}$$

Notice that since the sum is convergent, $\varepsilon_{Q_T} \downarrow 0$ as $Q_T \uparrow \infty$.

Proof. It is convenient to consider the packets being served in generations. That is given a target packet, suppose that we serve the packets which arrive during its service with preemptive priority up to and including the target packet. This makes no difference to queue behavior as the service times are exponential and we are only interested in the *first occasion* when node ℓ goes into backoff.

Suppose there are $Q_T > 0$ packets in the queue at the time the service of a given target packet starts. Consider the busy period of this packet, i.e. the time to serve the target packet and the subsequent high-priority packets (without backoff). It is easy to see that the mean number of packet arrivals during this busy period is $\frac{1}{1-\lambda_\ell}$, including the target packet itself.

The service of each packet ends with a random decision to backoff with probability less than $\frac{1}{(1+Q_T)^\gamma}$, since the queue length is never shorter than Q_T until the target packet has departed. Thus, by the union bound, the probability of a backoff occurring before or immediately after the target packet departs, is less than $\frac{1}{(1+Q_T)^\gamma} \times \frac{1}{1-\lambda_\ell}$. The probability of a backoff, starting with $Q_0 > Q_T$ packets, and before the queue drops below the level Q_T, is therefore smaller than $\frac{1}{1-\lambda_\ell} \sum_{r=Q_T}^{r=Q_0} \frac{1}{(1+r)^\gamma}$ which implies the statement of the lemma.

$\square$

The following lemma will also be useful. First given times $t_2 > t_1 \geq 0$ on the fluid scale, let $B_\ell([t_1, t_2])$ be the event that node ℓ *starts a backoff* in the interval $[t_1, t_2]$. This event occurs in the prelimit process $\left(Q_\ell^R, I_\ell^R\right)$ if for some jump chain index $n, X_\ell(n) = 1, X_\ell(n+1) = 0$ with $\lfloor Rt_1\beta \rfloor \le n \le \lceil Rt_2\beta \rceil$. Let $D_{\ell,\varsigma}([t_1, t_2])$ be the event that $Q_\ell^R(u) > \varsigma$ (or equivalently $\bar{Q}_\ell(Ru) \ge R\varsigma$) for all u in the interval $[t_1, t_2]$.

Lemma 4.7. *Given the above definitions,*

$$\lim_{R\to\infty} \mu_R \left\{ B_\ell[t_1, t_2] \cap D_{\ell,\varsigma}([t_1, t_2]) \right\} = 0. \tag{4.29}$$

Proof. First we may suppose node ℓ becomes active at some stage or there is nothing to prove. The lemma then follows from the union bound. Since there are at most $R_{t_1,t_2} \doteq \lceil R(t_2 - t_1)\beta \rceil + 2$ departures in the entire interval, the union bound implies that the probability of a backoff is smaller than,

$$\mu_R \left\{ B_\ell[t_1, t_2] \cap D_{\ell,\varsigma}([t_1, t_2]) \right\} \le \frac{R_{t_1,t_2}}{(1 + R\varsigma)^\gamma} \to 0.$$

This completes the proof. $\square$

Definition 4.6. *Given a queue ℓ, a time $t \in [0, \infty)$ on the fluid scale, and a queue length $q_\ell(t) = Q > 0$, we say that t is a point of increase for the activity process of queue ℓ if the event*

$$P_{t,Q}^{(\ell)} \doteq \cap_{M=1}^{\infty} \left\{ J_{>}^{(\ell)}(t, \frac{1}{M}, 0) \right\} \cap Q_{t,Q}^{(\ell)}$$

occurs, with $Q_{t,Q}^{(\ell)} \doteq \{\omega : q_\ell(t, \omega) > Q\}$. In words, queue ℓ is active in any arbitrarily small interval $(t, t + 1/M)$ and q_ℓ is greater than Q at time t.

Furthermore, given a time $s \in [0, \infty)$ and $h > 0$, we say that queue ℓ is under active, with duration $h > 0$ if the following event occurs

$$G_{s,h}^{(\ell)} \doteq J_{<}^{(\ell)}(s, h, \beta h). \tag{4.30}$$

Points of increase rule out that there is a sequence $t_n \downarrow t$ such that $I_\ell(t_n, \omega) = I_\ell(t, \omega)$, as there is activity no matter how small the interval. Under activity means that there was some idling during the interval. Given our choice of γ, it will be shown that a point of increase cannot be followed by a period of under activity until queue q_ℓ has drained. This is because the probability of even a single backoff once service has begun, is effectively 0 until the queue has drained on the fluid scale.

Lemma 4.8. *Suppose $s \in (t, t+Q/(1-\lambda_\ell))$. Then $\forall h, 0 < h < t+Q/(1-\lambda_\ell) - s$, and for all sufficiently large M,*

$$\mu\left\{J_>^{(\ell)}(t, 1/M, 0) \cap Q_{t,Q}^{(\ell)} \cap G_{s,h}^{(\ell)}\right\} = 0. \tag{4.31}$$

Proof. Consider the sequence of prelimit processes. We will choose M large enough and ς small enough so that $[s, s+h] \subset (t+1/M, t+(Q-\varsigma)/(1-\lambda_\ell)]$ (for some small constant $\varsigma > 0$). Then for R, M sufficiently large, and then by definition, occurrence of $G_{s,h}^{(\ell)}$, for the R-th prelimit process, implies occurrence of $B_\ell([t+1/M, s+h])$. Hence we obtain,

$$\begin{aligned} \mu_R\left\{J_>^{(\ell)}(t, 1/M, 0) \cap Q_{t,Q}^{(\ell)} \cap G_{s,h}^{(\ell)}\right\} &\leq \mu_R\Big\{J_>^{(\ell)}(t, 1/M, 0) \cap Q_{t,Q}^{(\ell)} \\ &\qquad \cap B_\ell([t+1/M, s+h])\Big\} \\ &\leq \mu_{B,R} + \mu_{F,D,R}, \end{aligned}$$

where

$$\mu_{B,R} \doteq \mu_R\left\{J_>^{(\ell)}(t, 1/M, 0) \cap B^{(\ell)}(t+1/M, s+h) \cap D_{\ell,\varsigma}([t, s+h])\right\}, \tag{4.32}$$

and

$$\mu_{F,D,R} \doteq \mu_R\left\{Q_{t,Q}^{(\ell)} \cap (D_{\ell,\varsigma}([t, s+h]))^c\right\}. \tag{4.33}$$

Thus, in order to prove the lemma, it is sufficient to show that both $\mu_{B,R} \to 0$ and $\mu_{F,D,R} \to 0$, as $R \to \infty$, because then we may conclude (4.31) by applying Theorem 2.1 in [45] and since the sets $J_>^{(\ell)}(t, 1/M, 0), Q_{t,Q}^{(\ell)}, G_{s,h}^{(\ell)}$ are all open.

The fact that $\mu_{F,D,R} \to 0$ follows from (4.22) and then by definition of $Q_{t,Q}^{(\ell)}$ and additionally by the choice of s, h, ς. As far as $\mu_{B,R}$ is concerned, the event $J_>^{(\ell)}(t, 1/M, 0)$ implies that service has started during the interval $[t, t+1/M]$. On the other hand, the event $B_\ell([t+1/M, s+h])$ implies that at some time in $[t, s+h]$ node ℓ starts to backoff. Setting $t_1 = t$ and $t_2 = s+h$, we may invoke Lemma 4.7 as by definition the event $D_{\ell,\varsigma}([t, s+h])$ implies q_ℓ did not go below ς in the interval $[t_1, t_2]$. It follows that $\mu_{B,R} \to 0$ as required. □

The implication of Lemma 4.8 is that any positive period of transmission, no matter how short, must be followed by full activity until the queue has drained on the fluid scale. This implies that there is no period of under activity, until the queue has drained, with probability 1.

Piecewise linear paths with probability 1

The aim of this section is to show that the queue sample paths follow a certain bilinear path during the interval prior to the queue becoming zero again. The bilinear path depends on the duration of the interval and on the arrival rate for the given queue.

To make the above statements precise, given $\ell \in \{1, \cdots, N\}$ define the bilinear path $\Phi^{\ell}_{t_0,t_1}$ for the interval $[t_0, t_1]$ to be,

$$\Phi^{\ell}_{t_0,t_1}(s) = \begin{cases} \lambda_\ell\,(s - t_0)\,; & t_0 \leq s \leq s_0, \\ \lambda_\ell\,(s - t_0) - (1 - \lambda_\ell)(s - s_0); & s_0 \leq s \leq t_1, \end{cases} \tag{4.34}$$

where $s_0 \doteq t_1 - \lambda_\ell(t_1 - t_0)$. In words, q_ℓ builds up linearly in the interval $[t_0, s_0]$ at rate λ_ℓ and drains at rate $1 - \lambda_\ell$ in the interval $[s_0, t_1]$.

Given $\eta > 0$, and $\ell \in \{1, \cdots, N\}$, define $\mathbb{1}^{(\eta,\ell)}_{t_0,t_1}$ to be the indicator for the event

$$\left\{ \omega : \sup_{s\in[t_0,t_1]} |q_\ell(s,\omega) - \Phi^{\ell}_{t_0,t_1}(s)| < \eta \right\} \in \mathcal{F}_{t_1}.$$

In words, $\mathbb{1}^{(\eta,\ell)}_{t_0,t_1}(\omega) = 1$ iff the absolute difference between $\Phi^{\ell}_{t_0,t_1}$ and the sample path for q_ℓ is smaller than η in supnorm over the interval $[t_0, t_1]$.

We now examine the conditional probability that $\mathbb{1}^{(\eta,\ell)}_{V^\ell_{m,n},Z^\ell_{m,n}}(\omega) = 1$, given $Z^\ell_{m,n} < \infty$ and $A^\ell_{m,n}$ (the case $Z^\ell_{m,n} = \infty$ is irrelevant). Define, $\mathcal{Z}^\ell_{m,n} \doteq \sigma\left(Z^\ell_{m,n}, A^\ell_{m,n}\right) \subset \mathcal{C}$ and also $\mathcal{Z}^{\ell,\infty}_{m,n} = \mathcal{Z}^\ell_{m,n} \cap \left\{\omega : Z^\ell_{m,n}(\omega) < \infty\right\}$.

It will be enough to show that the sample paths lie in an arbitrarily small tube around $\Phi^{\ell}_{V^\ell_{m,n},Z^\ell_{m,n}}$ conditional on $A^\ell_{m,n}, Z^\ell_{m,n}$ lying in some small rectangle $Z^{(a,b)}_{(s,t)} \doteq \left\{\omega : A^\ell_{m,n}(\omega) \in (a, b], Z^\ell_{m,n}(\omega) \in (s, t]\right\} \in \mathcal{Z}^{\ell,\infty}_{m,n}$.

Theorem 4.5. *Given $n \geq 1, m \in \mathbb{Z}_0$, then $\forall \eta > 0$,*

$$\mu\left\{ \mathbb{1}^{(\eta,\ell)}_{V^\ell_{m,n},Z^\ell_{m,n}} = 1 | \mathcal{Z}^{\ell,\infty}_{m,n} \right\} = 1 \ a.s. \tag{4.35}$$

In words, given the stopping time $Z^\ell_{m,n}$ and the time prior to this that q_ℓ was positive, $A^\ell_{m,n}$, the probability that $\Phi^{\ell}_{V^\ell_{m,n},Z^\ell_{m,n}}$ is followed, starting at $V^\ell_{m,n}$ and ending at $Z^\ell_{m,n}$, is 1 under the fluid limit measure μ.

Proof. For any given $\varepsilon > 0$, the sets $Z^{(a,b)}_{(s,t)}$ $0 < s < t$, $0 < a < b$, $0 < t - s, b - a < \varepsilon$, are a π-system [49] (i.e. closed under finite intersections) and which generate $\mathcal{Z}^{\ell,\infty}_{m,n}$. Hence we only need to show that

$$\mu\left\{\mathbb{1}^{(\eta,\ell)}_{V^{\ell}_{m,n},Z^{\ell}_{m,n}}(\omega) = 1; Z^{(a,b)}_{(s,t)}\right\} = \mu\left\{Z^{(a,b)}_{(s,t)}\right\},$$

for suitably chosen ε given $\eta > 0$. Let $B^{\ell}_{m,n}(\omega) \leq A^{\ell}_{m,n}$ be the additional time, following strict entry of q_ℓ into $(0,\infty)$ at $V^{\ell}_{m,n}$, until the first point of increase of I_ℓ is reached. $B^{\ell}_{m,n} \in m\mathcal{F}_{Z^{\ell}_{m,n}}$ as can be seen on consideration of its definition,

$$B^{\ell}_{m,n}(\omega) \doteq \inf\left\{u \in (0, A^{\ell}_{m,n}(\omega)) \cap \mathbb{Q} \ : I_\ell(V^{\ell}_{m,n} + u, \omega) - I_\ell(V^{\ell}_{m,n}, \omega) > 0\right\}, \tag{4.36}$$

when $Z^{\ell}_{m,n} < \infty$.

By definition of $B^{\ell}_{m,n}$, Lemma 4.8 and then (4.22), we may deduce that for $\omega \in Z^{(a,b)}_{(s,t)}$

$$\begin{aligned} q_\ell(V^{\ell}_{m,n}(\omega) + u, \omega) &= \lambda_\ell u,\ u \in [0, B^{\ell}_{m,n}(\omega)], \\ q_\ell(V^{\ell}_{m,n}(\omega) + u, \omega) &= \lambda_\ell B^{\ell}_{m,n}(\omega) - (1-\lambda_\ell)(u - B^{\ell}_{m,n}),\ u \in [B^{\ell}_{m,n}(\omega), \frac{B^{\ell}_{m,n}(\omega)}{1-\lambda_\ell}], \end{aligned}$$

μ almost surely. Moreover $B^{\ell}_{m,n}(\omega)$ must satisfy

$$t - s + b \geq \frac{B^{\ell}_{m,n}(\omega)}{1-\lambda_\ell} \geq s - t + a,\ \mu\ a.s.$$

in order to reach 0 in $[s,t]$.

Therefore, given any $\eta > 0$, we may choose $\varepsilon_\eta > 0$ such that for all $v \in [s-b, t-a], z \in [s,t]$ with $b - a, t - s < \varepsilon_\eta$

$$\sup_{u \in [v,z]} |q_\ell(u,\omega) - \Phi^{\ell}_{v,z}(u)| < \eta,$$

μ almost surely, using Lipschitz continuity. Since $\omega \in Z^{(a,b)}_{(s,t)}$ implies $V^{\ell}_{m,n} \in [s-b, t-a], Z^{\ell}_{m,n} \in [s,t]$, we obtain that

$$\mu\left\{\mathbb{1}^{(\eta,\ell)}_{V^{\ell}_{m,n},Z^{\ell}_{m,n}}(\omega) = 1; Z^{(a,b)}_{(s,t)}\right\} = \mu\left\{Z^{(a,b)}_{(s,t)}\right\},$$

for all such a, b, s, t as required. $\square$

A similar result can be obtained when $n = 0$, where the possibility occurs that $q_\ell(V^\ell_{m,0}) > 0$.

Theorem 4.5 applies to general networks and relies only on the assumption that $\gamma > 1$. The theorem implies that the sample paths are more or less determined given the sequences of stopping times $Z^\ell_{m,n}$. Since there are only countably many stopping times, and since for each finite $Z^\ell_{m,n} < \infty$ the queue sample paths follow $\Phi^\ell_{\cdot,\cdot}$ for some finite interval with probability 1, we may confine sample path realizations to countable successions of such intervals. These either determine the entire sample path; or the queue remains at 0 following the final return; or as the final alternative, the queue remains zero for some interval and then increases linearly at rate λ_ℓ thereafter. We define the set of such sample paths by $P \subset C[0, \infty)$. The probability of any event $F \in \mathcal{C}$ can as well be taken as

$$\mu\{F\} = \mu\{F \cap P\},$$

and, therefore, we suppose that the probability space is defined on $(P, \mathcal{C}_P)$ with topology relativized in the usual way to P which is a subset of $\mathcal{C}[0, \infty)$. This establishes that the queue-length trajectory of each of the individual nodes exhibits *sawtooth* behavior in the fluid limit. This concludes Part A.

In Part B, we will show that we can in fact confine ourselves to a smaller set of paths which reflect the constraints resulting from the underlying interference graph.

4.8.4 Fluid limit proofs: Part B

No idling property and zero delay capture

Given a node ℓ, let $\mathcal{I}_\ell$ be the set of its interfering nodes, i.e. the set of its neighbors in the interference graph G. The following lemma shows that if $q_\ell(s) > Q$, and all its interferers are idle in some interval $[s, t]$ then node ℓ is fully active until its queue drains.

Lemma 4.9 (No Idling Property). *Given a node ℓ with interference set $\mathcal{I}_\ell$*

and an interval $[s,t]$, define

$$D^{(\ell)}_{s,t} \doteq \cap_{j\in\mathcal{I}_\ell} J^{(j)}_{=}(s, t-s, 0),$$

that is, no activity for any node in $\mathcal{I}_\ell$ during $[s,t]$. Further, given $Q>0$, define $h^{(\ell)}_{s,t,Q} \doteq Q/(1-\lambda_\ell)\wedge(t-s)$ so that the queue at most empties over this period, and let

$$S^{(\ell)}_{s,t} \doteq J^{(\ell)}_{<}(s, h^{(\ell)}_{s,t,Q}, \beta h^{(\ell)}_{s,t,Q}),$$

which implies that node ℓ is under active. Then

$$\mu\left\{D^{(\ell)}_{s,t}\cap S^{(\ell)}_{s,t}\cap Q^{(\ell)}_{s,Q}\right\} = 0,$$

where $Q^{(\ell)}_{s,Q}$ is the event $\{\omega : q_\ell(s,\omega) > Q\}$, as defined earlier.

Proof. Given $n\in\mathbb{N}$ such that $n > 1/(t-s)$, fix an arbitrary ζ, $0<\zeta<\frac{1}{2N}$ (recall that N is the number of nodes in the network). Clearly,

$$D^{(\ell)}_{s,t} \subset \tilde{D}^{(\ell)}_{\zeta,n} := \cap_{j\in\mathcal{I}_\ell} J^{(j)}_{<}(s, 1/n, \zeta/n).$$

Hence, for arbitrary $\epsilon_n > 0$ depending on n, to be fixed later,

$$D^{(\ell)}_{s,t} \subseteq J^{(\ell)}_{>}(s,1/n,0) \cup \left(J^{(\ell)}_{<}(s,1/n,\epsilon_n)\cap \tilde{D}^{(\ell)}_{\zeta,n}\right).$$

Next, observe that for all $n_S\in\mathbb{N}$ sufficiently large,

$$S^{(\ell)}_{s,t} = \cup_{n>n_S} G_n,$$

with $G_n \doteq G^{(\ell)}_{s+2/n, h^{(\ell)}_{s,t,Q}-2/n}$ and $G^{(\ell)}_{s,h}$ as defined in (4.30). The union bound thus implies that

$$\begin{aligned}\mu\left\{D^{(\ell)}_{s,t}\cap S^{(\ell)}_{s,t}\cap Q^{(\ell)}_{s,Q}\right\} \;&\le\; \sum_{n>n_S}\mu\left\{G_n\cap Q^{(\ell)}_{s,Q}\cap J^{(\ell)}_{>}(s,1/n,0)\right\} \qquad (4.37)\\ &+\; \sum_{n>n_S}\mu\left\{\tilde{D}^{(\ell)}_{\zeta,n}\cap J^{(\ell)}_{<}(s,1/n,\epsilon_n)\cap Q^{(\ell)}_{s,Q}\right\}.\end{aligned}$$

Provided n_S is sufficiently large, each term in the first sum must be 0, else Lemma 4.8 is contradicted. To complete the proof, it is therefore sufficient to show that each of the terms in the second sum is 0 as well by suitable choice of ϵ_n.

Given n, it is sufficient to find $\epsilon_n > 0$ so that

$$\lim_{R\to\infty} \mu_R \left\{ \tilde{D}^{(\ell)}_{\zeta,n} \cap J^{(\ell)}_{<}(s, 1/n, \epsilon_n) \cap Q^{(\ell)}_{s,Q} \right\} = 0,$$

because $\tilde{D}^{(\ell)}_{\zeta,n}$, $J^{(\ell)}_{<}(s, 1/n, \epsilon)$, and $Q^{(\ell)}_{s,Q}$ are all open, so that Theorem 2.1 [45] implies that

$$\mu \left\{ \tilde{D}^{(\ell)}_{\zeta,n} \cap J^{(\ell)}_{<}(s, 1/n, \epsilon_n) \cap Q^{(\ell)}_{s,Q} \right\} = 0.$$

The event $\tilde{D}^{(\ell)}_{\zeta,n}$ implies that there must have been at least

$$\frac{R\beta}{n}(1 - N\zeta) > \frac{R\beta}{2n} \tag{4.38}$$

steps in the jump chain (if we allow for no overlap between active periods and since $|\mathcal{I}_\ell| < N$) at which all queues in $\mathcal{I}_\ell$ are in backoff for the interval $[s, s + 1/n]$. Also,

$$Q^R_\ell > Q - \frac{\beta}{n} > \varsigma > 0, \tag{4.39}$$

throughout $[s, s + 1/n]$ since there can be at most $R\beta/n$ departures.

But if (4.38) occurs, we may suppose that node ℓ becomes active within $R\beta/(4n)$ such steps, as the probability converges to 1 as $R \to \infty$ that it does so. But if we take $0 < \epsilon_n < \beta/(4n)$ the implication is that there is a subsequent backoff. Since (4.39) also occurs, Lemma 4.7 with $t_1 = s, t_2 = s + 1/n$ and ς above shows that the probability of a subsequent backoff goes to 0 which establishes the result. □

Since s, t, Q are arbitrary in Lemma 4.9, it follows from continuity that node ℓ begins service the instant its interferers become idle, if it has a positive queue-length.

Lemmas 4.3 and 4.9 carry an implication for the node pairs $(1, 2), (3, 4), (5, 6)$ in our network. We say that node ℓ_1 dominates node ℓ_2, $\ell_1 \neq \ell_2$ if $\mathcal{I}_{\ell_2} \subseteq \mathcal{I}_{\ell_1}$. Hence, if (say) queue 3 (the dominant queue) is draining, then no other queue than 4 may be active as a consequence of Lemma 4.3. But this implies all interferers of queue 4 are inactive. Hence, if $q_4 > 0$, it will therefore begin to drain immediately, i.e. if queue 3 is draining so is queue 4. Also if q_4 becomes 0 before q_3, then it must remain at 0, until queue 3 drains.

This result is formally stated in the following corollary, the proof of which is omitted for brevity.

Given any node $k \in \{1, \cdots, N\}$, $q_k \geq 0$, and time t define,

$$\Psi^k_{t,q_k}(u) \doteq [q_k - (u-t)(1-\lambda_k)]_+ \,, \; u \geq t, \tag{4.40}$$

and given $v > t$, let $F^k_{t,v,\eta}$ be the event that $|\Psi^k_{t,q_k(t,\omega)}(u) - q_k(u,\omega)| < \eta$ for $u \in [t, v]$.

Corollary 4.3. *Given a queue ℓ, let k be any other queue with $\mathcal{I}_k \subseteq \mathcal{I}_\ell$. $\forall t \geq 0, Q > 0, \eta > 0$, define $v = t + Q/(1-\lambda_\ell)$, then with $P^{(\ell)}_{t,Q}$ as in Definition 4.6, it holds that,*

$$\mu\left\{P^{(\ell)}_{t,Q} \cap \left(F^{(k)}_{t,v,\eta}\right)^c\right\} = 0.$$

Corollary 4.3 implies that μ almost surely the dominated node k follows Ψ^k the moment that dominating node ℓ becomes active.

In case the arrival rates satisfy, $\lambda_1 = \lambda_2 = \lambda > 0$, $\lambda_4 = \lambda_5$, $\lambda_6 > \lambda_5$, and $\lambda_3 > \lambda_4$, we show that the network enters a *natural state* (as defined in Section 4.3) μ a.s. This result is proved in the following theorem.

Theorem 4.6 (Almost Sure Natural State)**.** *Given the initial condition $\mathbf{q}(0)$ with $||\mathbf{q}(0)|| = 1$, there exits a $T_N < \infty$ such that μ a.s. for all $t \geq T_N$,*

$$\begin{aligned} q_3(t) &\geq q_4(t), \\ q_6(t) &\geq q_5(t). \end{aligned}$$

Moreover (recalling the definition of ρ given in Section 4.2) $\exists \rho^* < 1$ such that for all $\rho \in [\rho^*, 1)$, $\vee_\ell q_\ell(T_N) > 0$ i.e. the network is nonempty at time T_N.

Proof. This result follows from Lipschitz continuity and more particularly from the fact that the sample paths are piecewise linear. Hence, apart from a set of measure 0, the derivatives of all queue lengths exist.

Consider now nodes 3 and 4. Where the derivatives exist and $q_4 > 0$, it holds that

$$\frac{dq_3}{dt} > \frac{dq_4}{dt},$$

since $\lambda_3 > \lambda_4$ and since q_4 is decreasing at linear rate whenever $q_4 > 0$ and q_3 is decreasing at a linear rate, as shown in Lemma 4.9. We may therefore deduce μ a.s. and where differentiability holds that,

$$\frac{d\left[q_4(t)-q_3(t)\right]_+}{dt} \le \lambda_4 - \lambda_3 < 0,$$

until some time T_3, such that $\left[q_4(t)-q_3(t)\right]_+ = 0,\ t \ge T_3$. The same holds for queues 5 and 6, with corresponding time T_6 and the following inequalities are satisfied,

$$T_3 \le \frac{\left[q_4(0)-q_3(0)\right]_+}{\lambda_3-\lambda_4},\ T_3 \le \frac{\left[q_5(0)-q_6(0)\right]_+}{\lambda_6-\lambda_5}.$$

We may therefore take

$$T_N = T_3 \vee T_6,$$

and by taking worst case values in the above inequalities, we obtain a uniform bound on T_N. This concludes the first part of the lemma.

We now show that T_E, the time to empty, can be taken arbitrarily large. Define $L_P(t) \doteq \max\left(q_1(t), q_2(t)\right) + q_3(t) + q_6(t)$. Then L_P can be reduced at most at rate 1, since service of nodes (1, 2), 3 and 6 is mutually exclusive, and grows at rate $\rho = \rho_0 + \rho_3 + \rho_6$, which can be made arbitrarily close to 1. Hence $T_E \to \infty$ as $\rho \uparrow 1$ if $L_P(0) > 0$. It can be the case that $L_P(0) = 0$ but then $q_4(0) + q_5(0) = 1$, so that $L_P(1/2) = \rho_0/2$ and $T_E \ge \frac{1}{2}\left(1 + \frac{\rho_0}{1-\rho}\right)$ and again $T_E \to \infty$ as $\rho \uparrow 1$. $\square$

This shows that a nonempty natural state can be reached in finite time. Given Theorem 4.6, we can and will suppose that the state is natural at time 0, without loss of generality.

We define the set of paths which additionally satisfy the constraints of Lemmas 4.3 and 4.9 to be $P_L \subset P \subset C[0,\infty)$. We now restrict the set of sample paths to P_L, so that the probability of an event $F \in \mathcal{C}$ can be determined as $\mu\left\{F\right\} = \mu\left\{F \cap P_L\right\}$.

We now give a largely informal description of the paths in P_L. Section 4.2.1 gives a detailed description of the periods $M_k, k = 1, 2, 3, 4$. The ends of periods M_1, M_2, M_3 are marked by the corresponding stopping times $Z^{(1,2)}_{m,n}, Z^{(3)}_{m,n}, Z^{(6)}_{m,n}$. For M_4 periods, the following construction is needed. (It is needed because $Z^{(4)}_{m,n}$ stopping times may be part of an M_2 period and hence do not mark the end of a M_4 period.)

We first define $P^{(\ell)}_{m,n} = V^{(\ell)}_{m,n} + B^{(\ell)}_{m,n} \in m\mathcal{F}_{Z^{(\ell)}_{m,n}}$ to be the time prior to $Z^{(\ell)}_{m,n}$

when service begins (recall the definition of $B^{(\ell)}_{m,n}$ in (4.36)).

Definition 4.7. *A stopping time $Z^{(4)}_{m,n}$ is a $M^{(5)}_4$ stopping time, denoted $Z^{4,M^{(5)}_4}_{m,n}$ if the following holds,*

$$\begin{aligned} q_5(Z^4_{m,n} - P^{(4)}_{m,n}) &\geq q_4(Z^4_{m,n} - P^{(4)}_{m,n}), && (4.41) \\ I_\ell(Z^4_{m,n} - P^{(4)}_{m,n}) &= I_\ell(Z^4_{m,n}),\ \ell = 3, 6. && (4.42) \end{aligned}$$

That $Z^{4,M^{(5)}_4}_{m,n}$ is an $\mathcal{F}_t$ stopping time follows as both the above events lie in $\mathcal{F}_{Z^{(4)}_{m,n}}$.

This is consistent with an M_4-period taking place in which queue 4 emptied first (or at the same time as queue 5) by (4.41). If this is a strict inequality then we say this is a strict $M^{(5)}_4$ stopping time. Equation (4.42) ensures that queue 5 is being served throughout $[P^{(4)}_{m,n}, Z^{(4)}_{m,n}]$ as a consequence of Lemma 4.9. Similary we may define $Z^{5,M^{(4)}_4}_{m,n}$.

This concludes Part B. In Part C, we will derive the probabilities according to which one period is followed by another with no delay (on the fluid scale) in switching from one period to the next.

4.8.5 Fluid limit proofs: Part C

We begin with some preliminary results. The first is for measures constructed from closed continuity sets. Given a set of sample paths G, define the measures,

$$\mu_G\{F\} \doteq \mu\{F \cap G\},\quad \mu^{(R)}_G(F) = \mu_R\{F \cap G\}.$$

The following lemma shows that weak convergence is conferred on $\mu^{(R)}_G$ provided G is closed and a μ-continuity set.

Lemma 4.10. *Suppose $\mu^{(R)}$ is a sequence of probability measures on a metric space, $(\Omega, \mathcal{F})$, such that*

$$\mu^{(R)} \Rightarrow \mu,$$

where μ is also a probability measure on the same space. Let $G \in \mathcal{F}$ be closed and a μ-continuity set. Then it holds that

$$\mu^{(R)}_G \Rightarrow \mu_G.$$

In particular, the weak convergence definitions (iii), (iv), and (v), in Theorem 2.1 of [45], all equivalently hold.

Suppose a pair of non-interfering queues in the network are operating in isolation e.g. queues (1, 2). Then each queue will be empty and in fact will then subsequently be empty infinitely often, almost surely. Given that the evolutions of the two queues are independent, we prove next that the total number of steps in the jump chain for which both queues are backed off together increases to infinity in a period which is negligible on the fluid scale.

Given a start time taken to be 0, define $W^R(u)$ to be the total number of steps that queues 1 and 2 are both in backoff, starting at time 0 and ending at time $u > 0$ on the fluid scale, in $\left(\mathbf{Q}^R(t), \mathbf{I}^R(t)\right)$.

Lemma 4.11 (Total Backoff). *Given $Q > 0$, define $t \doteq Q/(1-\lambda)$, and suppose that $Q^R_\ell(0) \leq Q$, $\ell = 1, 2$, and both queues are active at time 0. Then for any $Q, \xi > 0$,*

$$\lim_{R\to\infty} \mu_R \left\{ W^R(t + 2\xi) \geq 2\sqrt{R} \right\} = 1.$$

Proof. Let $\tau_{0,0}$ be the stopping index in the jump chain for the first occurrence of

$$Q_1(\tau_{0,0}) = Q_2(\tau_{0,0}) = 0. \tag{4.43}$$

Given any $\xi > 0$, define $p^R_{\xi,Q} \doteq \mathbb{P}\left(\tau_{0,0} \leq \lfloor \beta(t+\xi)R \rfloor | Q_\ell(0) \leq RQ\right)$. It will be enough to show that $p^R_{\xi,Q} \to 1$ as $R \to \infty$. To see this, note that any queue in isolation is positive recurrent, as a consequence of Lemma 4.6. Thus, the jump chain restricted to nodes 1 and 2 in isolation (i.e., with remaining queues barred from gaining the medium) is also positive recurrent. Let m_0 be the mean number of steps between indices k such that (4.43) is again satisfied. Also let K^R_ξ be the random number of such steps in the next interval of $\lfloor \beta\xi R \rfloor$ steps. It is easily seen from the weak law of large numbers that

$$\lim_{R\to\infty} \mu_R \left\{ K^R_\xi > \frac{\lfloor \beta\xi R \rfloor}{2m_0} \right\} = 1,$$

which implies the statement of the lemma.

Thus, to complete the proof, we just need to show that $p^R_{\xi,Q} \to 1$. Fix $\varepsilon_{Q_T} > 0$ and choose $Q_T := Q_T(\lambda, \gamma) < \infty$ as in (4.28) so that the probability of even a single backoff before either queue reaches Q_T is no more than ε_{Q_T}.

Moreover let $\tau_{T,\ell}, \ell = 1, 2$ be the stopping indices for $Q_\ell(\tau_{T,\ell}) = Q_T$. Then, given any $\eta > 0$, and $\varepsilon_{R,\eta} > 0$, it can be seen that $\tau_{T,1} \vee \tau_{T,2} \leq \beta R(t+\eta)$ occurs with probability larger than $1 - 2\varepsilon_{R,\eta} - 2\varepsilon_{Q_T}$, with $\varepsilon_{R,\eta} \to 0$ as $R \to \infty$ by the weak law of large numbers.

Next, given any $\varepsilon_L > 0$, there exists a Q_L large enough such that

$$\mathbb{P}\left(Q_\ell(\tau_{T,\ell} + k) \leq Q_L\right) > 1 - \varepsilon_L$$

for all $k \in \mathbb{N}$. This follows from the fact that the jump chain in isolation is positive recurrent, and thus the corresponding sequence of infinite probability vectors is tight as they are converging to the steady-state distribution. Hence, with probability larger than $1 - 2\varepsilon_{R,\eta} - 2\varepsilon_{Q_T} - 2\varepsilon_L$, $Q_\ell(\lfloor \beta R(t+\eta) \rfloor \leq Q_L, \ell = 1, 2$.

Moreover, again by the positive recurrence of the isolated jump chain, the mean number of steps for queues 1 and 2 both to become 0, starting from any state with $Q_\ell \leq Q_L$, $\ell = 1, 2$, is bounded by some constant $m_L := m_L(Q_L) < \infty$. Thus, by Markov's inequality, with a probability less than $m_L/(\eta R)$, in a further ηR steps both queues will become 0 (and thus inactive).

Finally, given any $\epsilon > 0$, choose Q_T and Q_L large enough so that $\varepsilon_{Q_T} < \epsilon/8$ and $\varepsilon_L < \epsilon/8$ and then R sufficiently large so that $\varepsilon_{R,\eta} < \epsilon/8$ and $m_L/(\eta R) < \epsilon/8$. Hence, with probability larger than $1 - \epsilon$, $\tau_{0,0} < (t + 2\eta)R$ for all R sufficiently large. Since ϵ and η are arbitrary, the proof is complete. □

Transition from an M_1-period

In what follows we will further suppose that the lengths of queues 1 and 2 and their activity are both equal, as the following arguments are readily modified where this is not the case. We therefore denote their common queue length as $q(u) = q_1(u) = q_2(u)$ in what follows and similarly for the activity $I(u) = I_1(u) = I_2(u)$. Finally, in the following t, c and hence s are fixed,

$$\begin{aligned} s &\doteq t - \frac{c}{1-\lambda}, \\ \delta_k &\doteq \alpha_k c,\ 0 < \alpha_k < 1,\ k = 0, 1, \\ h &\doteq \nu c, \\ \zeta &\doteq \chi c,\ \nu > \chi > 0, \end{aligned}$$

for some small positive constants α_k, ν, and χ to be determined later. We define the following closed set of paths that correspond to an M_1-period

$$\begin{aligned} G_{c,t} &\doteq \{\omega : 0 < c - \delta_0 \leq q(s,\omega) \leq c + \delta_1\} \\ &\cap \{\omega : I(s+h,\omega) - I(s,\omega) \geq \beta(h-\zeta)\}. \end{aligned} \tag{4.44}$$

Now given $0 < s_1 < s_2$, and $\bar{t}$ (which will be specified later), define

$$I_{c,t}^{(3,4)} \doteq J_{=}^{(3)}(\bar{t}+s_1, s_2-s_1, \beta(s_2-s_1)) \cap J_{=}^{(4)}(\bar{t}+s_1, s_2-s_1, \beta(s_2-s_1)). \tag{4.45}$$

$I_{c,t}^{(3,4)}$ is a (closed) set of paths for which queue 3 (and also queue 4) are fully active during the interval $[\bar{t}+s_1, \bar{t}+s_2]$. Similar definitions, using the same s_1, s_2, and $\bar{t}$, can be made for $I_{c,t}^{(4,5)}, I_{c,t}^{(5,6)}$.

Note all sample paths must pass through the interval $[c-\delta_0, c+\delta_1]$ at time s, but may continue to increase for a brief period at the beginning. After $s+h$ the two queues must be draining at rate $1-\lambda$ almost surely as shown in Lemma 4.8. One of the other three queue pairs are expected to have the medium during the interval $[\bar{t}+s_1, \bar{t}+s_2]$. Only one such pair will be active during this period as a result of the forthcoming construction.

The following is the earliest time that queues 1 and 2 can drain if the sample paths are constrained to lie in $G_{c,t}$,

$$\underline{t} \doteq t - \frac{\alpha_0}{1-\lambda} c. \tag{4.46}$$

As far as additional queue build up is concerned, under the fluid limit,

$$q(s+h,\omega) \leq q(s,\omega) + \lambda h = q(s,\omega) + \lambda\nu c$$

holds for sample paths in $G_{c,t}$ (see (4.22)). It then follows that queues 1 and 2 will reach 0 under the fluid limit no later than

$$\bar{t} \doteq t + \frac{\alpha_1 + \lambda\nu}{1-\lambda} c, \tag{4.47}$$

which is the definition for $\bar{t}$. We thus conclude that, under the fluid limit, queues 1 and 2 will reach 0 in the interval $(\underline{t}, \bar{t})$ (for the first time after $s+h$ on occurrence of the event $G_{c,t}$). We formalize the above in the following lemma.

Lemma 4.12 (Queue Bounds). *Let* $\tau^0_{c,s} \doteq \tau_c(s, \{0\}) = \inf\{t \geq s : q(t) = 0\}$ *be the first contact time with 0 for* $q = q_1 = q_2$. *Then,*

$$\mu\left\{G_{c,t} \cap \left\{\omega : \tau^0_{c,s}(\omega) \notin [\underline{t}, \overline{t}]\right\}\right\} = 0.$$

Additionally, $\forall \ell = 3, 4, 5, 6$,

$$\mu\left\{G_{c,t} \cap \{\omega : q_\ell(\underline{t}, \omega) < \Delta t \lambda_\ell\}\right\} = 0, \tag{4.48}$$

where,

$$\Delta t \doteq \underline{t} - (s + h) = c(1 - \alpha_0 - \nu(1 - \lambda))/(1 - \lambda).$$

Proof. By definition of $G_{c,t}$, $q(s, \omega) \geq c - \delta_0, \forall \omega \in G_{c,t}$. It follows that q cannot reach 0 before $\underline{t}$, as sample paths by definition lie in P_L (see Section 4.8.4, following Theorem 4.5). A similar argument applies to $\overline{t}$.

For the last part, Lemma 4.3, shows that nodes $3, 4, 5, 6$ must be idle in the period $[s + h, \underline{t}]$. Since the sample paths are restricted to lie in P_L, it follows that their queues must satisfy the stated inequality at time $\underline{t}$. The proof is complete. □

The time for queue ℓ to reach 0 following $\underline{t}$ is therefore at least,

$$f_\ell \doteq \Delta t \times \frac{\lambda_\ell}{1 - \lambda_\ell}, \ \ \ell = 3, 4, 5, 6.$$

Clearly $f_\ell \to (c \times \lambda_\ell)/\left((1 - \lambda) \times (1 - \lambda_\ell)\right)$ as $\alpha_0, \nu \downarrow 0$, and so this expression is bounded from below as $\alpha_0, \alpha_1, \nu > \chi$ are made arbitrarily small. For future use, we define

$$\underline{f} \doteq \Delta t \wedge_{\ell=3}^{6} \lambda_\ell/(1 - \lambda_\ell),$$

as a lower bound on the time needed to drain any queue $\ell = 3, 4, 5, 6$.

Our results thus far do not rule out the possibility that there is an idle period during which queues 3, 4, 5, 6 fail to obtain the medium. In order to make allowance for this, we introduce a period ξc, $\xi > 0$, which comes following queues 1 and 2 draining, and to be definite, we set $\xi c = \underline{f}/8$. Hence, if it is the case that

$$\overline{t} - \underline{t} < \underline{f}/4 \tag{4.49}$$

and that service of queue ℓ cannot start before $\underline{t} - \xi c$ and must have started no later than $\overline{t} + \xi c$, then it follows that service will continue throughout the interval $[\bar{t}+\xi c, \bar{t}+\xi c+\underline{f}/2]$. In this case, we may take $s_1 = \underline{f}/4, s_2 = \underline{f}/2$ again to be definite. Further, set $s_3 = \xi c + \underline{f}/2$. To summarize, if (4.49) holds, on occurrence of $G_{c,t}$ and that service of queues 3 and 4 commences in the interval $[\underline{t}-\xi c, \bar{t}+\xi c]$, then the event $I^{(3,4)}_{c,t}$ must take place. The same is true in case service commences for either queue pair $(4,5)$ or $(5,6)$ in $[\underline{t}-\xi c, \bar{t}+\xi c]$.

Let $\hat{C}_k, k = 3,4,5,6$, be the residual time to backoff for queues 3, 4, 5, 6, at time $s+h$, with $\hat{C}_1 = \hat{C}_2 = 0$ as these queues will be almost surely active. Define S_M to be the number of steps in the jump chain before one of these nodes gains the medium and also define

$$\begin{aligned} W^{(3,4)} &\doteq \left\{\hat{C}_3 < \wedge_{k=4}^{6}\hat{C}_k\right\} \cup \left(\left\{\hat{C}_4 < \hat{C}_3 \wedge \hat{C}_5 \wedge \hat{C}_6\right\} \cap \left\{\hat{C}_3 < \hat{C}_5\right\}\right), \\ C^{(3,4)}_{c,t,R} &\doteq W^{(3,4)} \cap \left\{S_M \leq \sqrt{R}\right\}. \end{aligned}$$

$W^{(3,4)}$ is the event that queues 3 and 4 win the backoff competition to take the medium first from queues 1 and 2. Similar definitions can be made for queues $(4,5)$ and for queues $(5,6)$ in addition. The probabilities of these events are

$$\mathbb{P}\left(W^{(3,4)}\right) = \frac{3}{8} = \mathbb{P}\left(W^{(5,6)}\right),\ \mathbb{P}\left(W^{(4,5)}\right) = \frac{1}{4}, \tag{4.50}$$

as the backoff periods are unit mean i.i.d. exponential random variables. $C^{(3,4)}_{c,t,R}$ is the event that queues $(3,4)$ win the backoff competition and that it does so in no more than $\sqrt{R}$ of the jump chain steps when queues 1 and 2 are in backoff together.

Next let

$$B^{(1,2)}_R \doteq N^{(1,2)}_R(s+h, \underline{t}-\xi c) \cap \{W^R(s+h, \bar{t}+\xi c) \geq 2\sqrt{R}\}$$

be the intersection of the event $N^{(1,2)}_R(s+h, \underline{t}-\xi c)$ that neither queue 1 nor queue 2 starts to backoff during the time interval $[s+h, \underline{t}-\xi c]$ and the event $\{W^R(s+h, \bar{t}+\xi c) \geq 2\sqrt{R}\}$ that queues 1 and 2 operating in isolation would be simultaneously in backoff for a cumulative period of time of at least $2\sqrt{R}$ during the interval $[s+h, \bar{t}+\xi c]$. Informally speaking, the event $B^{(1,2)}_R$ ensures that there is sufficient backoff by queues 1 and 2 and that they do not begin to backoff while there are a significant number of queue 1 or queue 2 packets remaining.

Next define c_Q to be,

$$c_Q \doteq \frac{s_3 - s_2}{2} \times (1 - \lambda_3) > 0,$$

which is at least half the content of queues 3 and 4 on the fluid scale at time $\bar{t} + s_2$, given our construction. Further, define the following event

$$Q_R^{(3,4)}(\underline{t}, \bar{t} + s_2) \doteq \left\{\omega : \inf\left\{Q_m^R(u,\omega), u \in [\underline{t}, \bar{t} + s_2]\right\} > c_Q, m = 3, 4\right\} \in \mathcal{F}_{\bar{t}+s_2},$$

for which we obtain the following corollary.

Corollary 4.4.

$$\lim_{R\to\infty} \mu_R \left\{G_{c,t} \cap \left(Q_R^{(3,4)}\right)^c\right\} = 0$$

Proof. Lemma 4.12 implies that for all n sufficiently large,

$$\limsup_{R\to\infty} \mu_R \left\{G_{c,t} \cap \left\{\omega : Q_\ell^R(\underline{t},\omega) \le \Delta t \lambda_\ell - 1/n\right\}\right\} = 0, \;\; \ell = 3, 4,$$

on using Theorem 2.1 in [45] and that both the above sets are closed. Hence we need only show that,

$$\lim_{R\to\infty} \mu_R \left\{\left\{\omega : Q_\ell^R(\underline{t},\omega) > \Delta t \lambda_\ell - 1/n, \ell = 3, 4\right\} \cap \left(Q_R^{(3,4)}\right)^c\right\} = 0, \quad (4.51)$$

for sufficiently large n. However (4.51) follows from the weak law of large numbers, from the definition of Δt, c_Q, and the event $Q_R^{(3,4)}$. ☐

Finally, define $N_R^{(3,4)}(\underline{t}, \bar{t} + s_2)$ to be the event that neither queue 3 nor queue 4 has a backoff during the time interval $[\underline{t}, \bar{t} + s_2]$ (on the fluid scale). Clearly, equivalent definitions for this and the above corollary can be made for queue pairs $(4, 5)$ and $(5, 6)$.

In what follows it will be convenient to write $G := G_{c,t}$.

Our aim now is to show that no matter what trajectory the fluid limit path followed earlier, if it lies in G so that queues 1 and 2 almost surely reach 0 in the interval $[\underline{t}, \bar{t}]$, marking the end of an M_1 period, then the probability of the next period depends only on the residual backoff times, which is a Markov property.

Lemma 4.13. *Suppose that G is a set of paths as defined in (4.44), with parameter values so that (4.49) holds, and is also a μ-continuity set. In* *addition, let $F \in \mathcal{F}_s$ be an arbitrary closed, finite-dimensional set of paths defined by times s and earlier. It then holds that*

$$
\begin{aligned}
\mu_G\left\{F \cap I_{c,t}^{(3,4)}\right\} &\geq \frac{3}{8}\mu_G\left\{F^o\right\} \\
\mu_G\left\{F \cap I_{c,t}^{(5,6)}\right\} &\geq \frac{3}{8}\mu_G\left\{F^o\right\} \\
\mu_G\left\{F \cap I_{c,t}^{(4,5)}\right\} &\geq \frac{1}{4}\mu_G\left\{F^o\right\}.
\end{aligned}
$$

In case F is a μ-continuity set, the interior can be dropped and $\geq$ replaced with equality.

Proof. We first show the last part of the lemma, assuming the first part to be true. If F is a μ-continuity set, then by definition, $0 = \mu\{\partial F\} \geq \mu\{G \cap \partial F\}$ and it follows that F is a μ_G-continuity set as well. Since the factors sum to 1 and the events on the left are almost surely exclusive as a consequence of Lemma 4.3, we can now replace the inequality sign with equality.

We move to the first part of the lemma, which we will prove for queues 3 and 4. The proof for the other queue pairs is similar.

First observe that

$$
C_{c,t,R}^{(3,4)} \cap B_R^{(1,2)} \cap N_R^{(3,4)}(\underline{t}, \bar{t} + s_2) \subseteq I_{c,t}^{(3,4)},
$$

since $C_{c,t,R}^{(3,4)} \cap B_R^{(1,2)}$ implies that queues 3 and 4 activate before time $\bar{t} + s_1$, while $N_R^{(3,4)}(\underline{t}, \bar{t} + s_2)$ ensures that neither queue 3 nor queue 4 has a backoff during the time interval $[\underline{t}, \bar{t} + s_2]$. We thus obtain the following chain of inequalities

$$
\begin{aligned}
\mu_G^{(R)}\left\{F \cap I_{c,t}^{(3,4)}\right\} &\geq \mu_G^{(R)}\left\{F \cap C_{c,t,R}^{(3,4)} \cap B_R^{(1,2)} \cap N_R^{(3,4)}\right\} \\
&\geq \mu_G^{(R)}\left\{F \cap W^{(3,4)}\right\} \\
&\quad -\mu_G^{(R)}\left\{\left(\left\{S_M \leq \sqrt{R}\right\} \cap B_R^{(1,2)} \cap N_R^{(3,4)}\right)^c\right\} \\
&\geq \frac{3}{8}\mu_G^{(R)}\left\{F\right\} - \mu_G^{(R)}\left\{S_M > \sqrt{R}\right\} \\
&\quad -\mu_G^{(R)}\left\{\left(B_R^{(1,2)}\right)^c\right\} - \mu_G^{(R)}\left\{\left(N_R^{(3,4)}\right)^c\right\},
\end{aligned}
\tag{4.52}
$$

with $N_R^{(3,4)} \equiv N_R^{(3,4)}(\underline{t}, \bar{t} + s_2)$ for compactness. The first inequality follows by inclusion, the second using $\mu_G\{A \cap B\} \geq \mu_G\{A\} - \mu_G\{B^c\}$, and the third from (4.50) by independence of the back-off clocks and by using the union bound in conjunction with de Morgan's laws. We now proceed to show that

$$
\begin{aligned}
\mu_{G,1}^{(R)} &\doteq \mu_R \left\{ S_M > \sqrt{R} \right\} \to 0 \\
\mu_{G,2}^{(R)} &\doteq \mu_R \left\{ \left(B_R^{(1,2)} \right)^c \cap G_{c,t} \right\} \to 0 \\
\mu_{G,3}^{(R)} &\doteq \mu_R \left\{ \left(N_R^{(3,4)} \right)^c \cap G_{c,t} \right\} \to 0.
\end{aligned}
$$

The first limit is immediate. In order to deal with the second limit, define the event

$$
Q_R^{(1,2)}(s+h, \underline{t} - \xi c) \doteq \left\{ \omega : \inf \left\{ Q_m^R(u, \omega), u \in [s+h, \underline{t} - \xi c] \right\} > \varsigma, m = 1, 2 \right\}
$$

for some small constant $\varsigma > 0$, and use the upper bound

$$
\begin{aligned}
\mu_{G,2}^{(R)} \leq\ & \mu_R \left\{ \left(B_R^{(1,2)} \right)^c \cap Q_R^{(1,2)}(s+h, \underline{t} - \xi c) \cap G_{c,t} \right\} \\
& + \mu_R \left\{ \left(Q_R^{(1,2)}(s+h, \underline{t} - \xi c) \right)^c \cap G_{c,t} \right\}.
\end{aligned}
$$

The limit of the second term is 0 by definition of $\underline{t}$ as the earliest time that queues 1 and 2 can drain under the event $G_{c,t}$ and on making a suitable choice for ς. It suffices then to show that the limit of the first term is 0. In order to prove this, we invoke the definition of the event $B_R^{(1,2)}$ to obtain that the first term is bounded from above by

$$
\begin{aligned}
& \mu_R \left\{ \left(N_R^{(1,2)}(s+h, \underline{t} - \xi c) \right)^c \cap Q_R^{(1,2)}(s+h, \underline{t} - \xi c) \right\} \\
& \quad + \mu_R \left\{ \{ W^R(s+h, \bar{t} + \xi c) \leq 2\sqrt{R} \} \cap G_{c,t} \right\}.
\end{aligned}
$$

That the first term converges to 0 follows by definition of the events and Lemma 4.7. Lemma 4.11 shows that the limit of the second term (i.e. the event there is insufficient backoff by queues 1 and 2 on occurrence of $G_{c,t}$) is 0.

In order to handle the third limit, we apply the upper bound

$$
\mu_{G,3}^{(R)} \leq \mu_R \left\{ \left(N_R^{(3,4)}(\underline{t}, \bar{t} + s_2) \right)^c \cap Q_R^{(3,4)} \right\} + \mu_R \left\{ \left(Q_R^{(3,4)} \right)^c \cap G_{c,t} \right\}.
$$

Lemma 4.7 shows that the limit of the first term is 0, while the statement of Corollary 4.4 is that the limit of the second term is 0.

Taking limits in (4.52) with respect to R, and using Lemma 4.10, it follows that,

$$\begin{aligned}\mu_G\left\{F \cap I_{c,t}^{(3,4)}\right\} &\geq \frac{3}{8}\limsup_R \mu_G^{(R)}\{F\} \\ &\geq \frac{3}{8}\liminf \mu_G^{(R)}\{F^o\} \\ &\geq \frac{3}{8}\mu_G\{F^o\},\end{aligned} \tag{4.53}$$

where the first inequality follows from the fact that F and $I_{c,t}^{(3,4)}$ are both closed and the third since F^o is open and again from Lemma 4.10. $\square$

Let μ be the fluid limit measure and proceed to define for any given $t \geq 0$ the following class of sets, the *finite-dimensional continuity rectangles* $\mathcal{K}_{\mu,t}$ which are a subset of the *finite-dimensional sets*, $\mathcal{H}_t$.

Definition 4.8. *Define the class of finite closed rectangles* $\mathcal{R}$ *to be the sets,*

$$\left(\prod_{j=1}^{N}[s_{j,L}, s_{j,H}]\right) \times \left(\prod_{j=1}^{N}[r_{j,L}, r_{j,H}]\right) \subset \mathbb{R}_+^N \times \mathbb{R}_+^N,$$

where $s_{j,L} \leq s_{j,H}, r_{j,L} \leq r_{j,H}$*, otherwise we obtain the empty set.*

Given times $0 \leq t_1 < t_2 < \cdots < t_J \leq t$*, define* $\pi_{J,t} : C[0,\infty) \to E^J$ *to be the (continuous) projection map taking the sample path to its position at times* $t_1, \cdots, t_J$

$$\pi_{K,t}(\omega) = \Big(\left(\mathbf{q}(t_1,\omega), \mathbf{I}(t_1,\omega)\right), \cdots, \left(\mathbf{q}(t_J,\omega), \mathbf{I}(t_J,\omega)\right) \Big).$$

Finally, take $\mathcal{R}^J$ *to be* J*-products of closed rectangles. Define* $\mathcal{K}_t$ *to be sets of the form* $\pi_{K,t}^{-1}R_J, R_J \in \mathcal{R}^J$ *and finally* $\mathcal{K}_{\mu,t} \subset \mathcal{K}_t$ *to be those* $H \in \mathcal{K}_t$ *such that* $\mu\{\partial H\} = 0$*. Clearly* $\mathcal{K}_{\mu,t} \subset \mathcal{K}_t \subset \mathcal{F}_t$*.*

Returning to Lemma 4.13, we see that it is satisfied by all sets $F \in \mathcal{K}_{\mu,s}$ with equality since they are by definition closed μ-continuity sets. Furthermore, since the terms on the left and on the right are measures and since $\mathcal{K}_{\mu,s}$ generates $\mathcal{F}_s$, the following corollary holds.

Corollary 4.5. *$\forall F \in \mathcal{F}_s$, Lemma 4.13 holds with equality, i.e.*

$$\begin{aligned}\mu_G \left\{F \cap I_{c,t}^{(3,4)}\right\} &= \frac{3}{8}\mu_G\{F\} \\ \mu_G \left\{F \cap I_{c,t}^{(5,6)}\right\} &= \frac{3}{8}\mu_G\{F\} \\ \mu_G \left\{F \cap I_{c,t}^{(4,5)}\right\} &= \frac{1}{4}\mu_G\{F\}.\end{aligned}$$

Proof. First note that the measures on the LHS and RHS are both finite and are therefore both σ-finite, with respect to the sets in $\mathcal{K}_{\mu,s}$. It is readily shown that $\mathcal{K}_{\mu,s}$ is a π-system and $\sigma(\mathcal{K}_{\mu,s}) = \mathcal{F}_s$. Theorem 10.3 in [49] thus shows that LHS and RHS agree on $\mathcal{F}_s$. □

To continue toward Theorem 4.7, we now define paths that one of which is followed immediately on completion of a (positive) M_1-period at time s, μ a.s. First define

$$\vartheta_{s,t}^{(\mathbf{q},M_k)}(u), \;\; u \in [s,t], k = 2,3,4, \tag{4.54}$$

to be the path which is at $\mathbf{q}$ at time s and then follows M_k until time t, e.g., if $k = 1$, queues 1 and 2 are decreasing linearly at rate $(1-\lambda)$ and any other queue $\ell = 3,4,5,6$ is increasing at rate λ_ℓ. Precise definitions we omit as the form of the sample paths have already been discussed. The next definition is for an indicator function that the above path is being followed in an interval $[s, s+h], h > 0$.

$$\mathbb{1}_{\mathrm{M_k},\mathbf{q}}^{(\mathrm{s,h},\eta)} \doteq \mathbb{1}\left\{\omega : ||\mathbf{q}(v,\omega) - \vartheta_{s,s+h}^{(\mathbf{q}(s,\omega),M_k)}(v)|| < \eta, v \in [s,s+h]\right\}. \tag{4.55}$$

In words M_k is "followed" for an interval of duration h starting at s to a closeness η.

Note that the result of Corollary 4.5 applies only to events in some σ-algebra $\mathcal{F}_w$ where $w \geq 0$ is fixed. However, the equivalent results can be established for all events $F \in \mathcal{F}_{Z_{m,n}^{(1,2)}}$ as stated by the following theorem. The proof can be found in [50].

Theorem 4.7. *$\forall n \in \mathbb{N}_0, m \in \mathbb{Z}_0$, $\exists \eta_m$ such that $\forall \eta$, $\eta_m > \eta > 0$,*

$$\begin{aligned}
\mu\left\{\mathbb{1}^{(\mathrm{M}_1,\eta)}_{\mathrm{M}_2,\mathrm{m},\mathrm{n}} | \mathcal{F}^{\infty}_{Z^{(1,2)}_{m,n}}\right\} &= \frac{3}{8},\ \mu \ a.s.,\\
\mu\left\{\mathbb{1}^{(\mathrm{M}_1,\eta)}_{\mathrm{M}_3,\mathrm{m},\mathrm{n}} | \mathcal{F}^{\infty}_{Z^{(1,2)}_{m,n}}\right\} &= \frac{3}{8},\\
\mu\left\{\mathbb{1}^{(\mathrm{M}_1,\eta)}_{\mathrm{M}_4,\mathrm{m},\mathrm{n}} | \mathcal{F}^{\infty}_{Z^{(1,2)}_{m,n}}\right\} &= \frac{1}{4}.
\end{aligned}$$

Since $\eta > 0$ can be taken arbitrarily small, the conclusion is that one of the periods M_2, M_3, M_4 start immediately at $Z^{(1,2)}_{m,n}$ on occurrence of $Z^{(1,2)}_{m,n} < \infty$ and with probabilities determined solely by the residual backoff times.

Switchover from M_2, M_3, M_4

Here we will only state our results, moreover M_2- and M_3-periods are analogous and so we will only deal with the former. To state our theorem for switching out of a M_2-period, define $\mathbb{1}^{(\mathrm{M}_2,\eta)}_{\mathrm{M}_\mathrm{k},\mathrm{m},\mathrm{n}}$, $k = 1, 3$ as was done for switching out of M_1. Also define $\mathcal{F}^{\infty}_{Z^3_{m,n}} \doteq \mathcal{F}_{Z^3_{m,n}} \cap \left\{Z^3_{m,n} < \infty\right\}$.

Theorem 4.8. *$\forall n \in \mathbb{N}_0, m \in \mathbb{Z} - \{0\}$, $\exists \overline{p}, \overline{q} > 0, \overline{p} + \overline{q} = 1$ and $\exists \eta_m$ such that $\forall \eta$, $\eta_m > \eta > 0$*

$$\begin{aligned}
\mu\left\{\mathbb{1}^{(\mathrm{M}_2,\eta)}_{\mathrm{M}_1,\mathrm{m},\mathrm{n}} | \mathcal{F}^{\infty}_{Z^3_{m,n}}\right\} &= \overline{p},\ \mu \ a.s.,\\
\mu\left\{\mathbb{1}^{(\mathrm{M}_2,\eta)}_{\mathrm{M}_3,\mathrm{m},\mathrm{n}} | \mathcal{F}^{\infty}_{Z^3_{m,n}}\right\} &= \overline{q}.
\end{aligned}$$

The quantities $\overline{p}$ and $\overline{q}$ are determined as follows,

$$\begin{aligned}
\overline{p} &= \sum_{Q=0}^{\infty}\sum_{Q_4=0}^{\infty}\sum_{X_4=0,1} b^{(3)}(Q)\pi_4^{\infty}(Q_4,X_4)c_1^Q(Q_4,X_4) \qquad (4.56)\\
\overline{q} &= \sum_{Q=0}^{\infty}\sum_{Q_4=0}^{\infty}\sum_{X_4=0,1} b^{(3)}(Q)\pi_4^{\infty}(Q_4,X_4)c_5^Q(Q_4,X_4),
\end{aligned}$$

where $\pi_4^{\infty}(Q_4, X_4)$ is the equilibrium jump chain probability that node 4 is in state (Q_4, X_4) when operating in isolation (i.e., when node 4 is the only node in the network). The $b^{(3)}(Q)$ is the limiting probability as $Q_0^3 \uparrow \infty$ that a first backoff of node 3 occurs when $Q_3 = Q$, service starting with Q_0^3 packets. $c_1^Q(Q_4, X_4)$ is the probability that node 1 or 2 first gain the medium when node 3 has a first backoff with $Q_3 = Q$ packets and the state of node 4 as given. The remaining definitions for $\overline{q}$ are similar. Thus in this case there is no simple formula and $\overline{p}, \overline{q}$ depend on the backoff parameter γ as well as the arrival rates to nodes 3 and 4.

For the case of switching out of M_4 we have the following result, again making the corresponding definitions as in Theorem 4.7.

Theorem 4.9. *For any $Z_{m,n}^{(4,M_4^{(5)})}$ stopping time, there is a $\eta_m > 0$ sufficiently small so that, for all $\eta_m > \eta > 0$*

$$\mu\left\{\mathbb{1}_{\mathrm{M_4,m,n}}^{(\mathrm{M_2},\eta)} \vee \mathbb{1}_{\mathrm{M_4,m,n)}}^{(\mathrm{M_3},\eta} \mid \mathcal{F}^\infty_{Z_{m,n}^{(4,M_4^{(5)})}}\right\} = 1 \ \ \mu \ a.s.,$$

and so that

$$\mu\left\{\mathbb{1}_{\mathrm{M_4,m,n}}^{(\mathrm{M_1},\eta)} | \mathcal{F}^\infty_{Z_{m,n}^{(4,M_4^{(5)})}}\right\} \ = \ 0, \ \mu \ a.s.$$

A similar result holds for $Z_{m,n}^{(5,M_4^{(4)})}$ stopping times.

This concludes Part C.

4.8.6 Fluid limit proofs: Part D

In Parts A-C we have established (a) sawtooth properties and some constraints on those sample paths, (b) what will occur at the end of a given M_k-period, $k = 1, 2, 3, 4$ and (c) that a natural state will be entered in finite time before the network can empty almost surely. What has not been shown, is whether any M_1-period would ensue at all. We can indeed show that M_1-periods will occur μ a.s. following a natural state, provided ρ is sufficiently close to 1.

In fact, establishing this result is not strictly needed to prove instability. If there is a last visit to queues 1 and 2 (which might occur when they are both empty), then these two queues must grow linearly and therefore the fluid system is unstable. Nevertheless, we have shown in [50] that an infinite sequence of M_1-periods will occur μ almost surely and in strictly bounded time, following T_N to enter a nonempty natural state. Our main result from [50] is the following.

Denote by $\tau_p^{(1,2)} : C[0,\infty] \to [0,\infty]$ as the first point of increase of either I_1, I_2, as in Definition 4.6, following T_N. It is easily shown that $\tau_p^{(1,2)}$ is a $\mathcal{F}_{t^+}$ stopping time, and corresponds to the start of a positive M_1-period.

Theorem 4.10. *There exists a* $0 < T_V < \infty$ *such that* $\tau_p^{(1,2)} < T_N + T_V$, μ *a.s.*

Thus an M_1-period occurs within bounded fluid time following a natural state. This concludes Part D.

Chapter 5

Stability for Multihop Networks with Dynamic Flows

In Chapters 1–4, we described the Max Weight Scheduling (MWS) algorithm of Tassiulas and Ephremides [1] and our random access mechanism. The stability results so far have been concerned with *packet-level dynamics* when flows (users) are not arriving/departing. In real networks however, flows arrive randomly to the network, have only a finite amount of data, and depart the network after the data transfer is completed. Moreover, there is no notion of *congestion control* in MWS and the random access mechanism while most modern communication networks use some congestion control mechanism for fairness purposes or to avoid excessive congestion inside the network [23].

In this chapter, we will examine stability of the system in presence of flow arrivals/departures. Here, by stability, we mean that the number of flows and the queue sizes in the network remain finite is some appropriate sense. We will also show how the stability results for the single-hop model in previous chapters can be extended to the multihop model.

Stability of wireless networks under flow-level dynamics has been studied in, e.g., [23, 24, 25]. Under flow-level dynamics, if the scheduler has access to the total queue-length information at nodes, then it can use max weight/back-pressure algorithm to achieve throughput optimality, but this information is not typically available to the scheduler because it is implemented as part of the MAC layer. Moreover, without congestion control, queue sizes at different nodes could be widely different. This could lead to long periods of unfairness among flows because links with long flows/files[1] (large weights) will get priority over links with short flows/files (small weights) for long periods of time. Therefore, we need to use congestion control to provide better *Quality-of-Service* (QoS). With congestion control, only a few packets from each file are released to the MAC layer at each time instant,

[1]The terms file, flow, and user can be used interchangeably here.

and scheduling is done based on these MAC layer packets. Specifically, the network control policy consists of two parts: (a) "congestion control" which determines the rate of service provided to each flow, and (b) "packet scheduling" which determines the rate of service provided to each link in the network.

However, to achieve flow-level stability, prior works [23, 24, 25] require that a specific form of congestion control has to be used, namely, *ingress queue based rate adaptation using α-fair utility functions*. More accurately, (a) the rate at which a flow/file generates packets into its ingress queue must maximize its utility subject to a linear penalty (price). The utility function of each flow is assumed to be in the form $x^{1-\alpha}/(1-\alpha)$, for some $\alpha > 0$, with x the flow rate, and the penalty (price) charged is the number of packets queued at the ingress queue associated with the flow, (b) scheduling of packets is performed using the max weight/back-pressure algorithm, where the weight of each link is the queue size (or the queue size raised to the power α).

The main contributions of this chapter can be summarized as follows:

1. We show that α-fair congestion control is not necessary for stability, and, in fact, very general ingress queue based congestion control mechanisms are sufficient to ensure stability. A key ingredient of our result is the use of the difference between the logarithms of queue lengths as the link weights.

2. The design of efficient scheduling and congestion control algorithms can be decoupled. This separation result would allow using different congestion control mechanisms at the edge of the network for providing different fairness or QoS considerations without need to change the scheduling algorithm implemented at internal routers of the network.

3. A by-product of the weight function that we use for each link is that one can use random access (CSMA-type) mechanisms, as in Chapter 2, to implement the scheduling algorithm in a distributed fashion. In particular, unlike [51] which also considers flow-level stability, we do not have to assume time-scale separation between the dynamics of flows, packets, and CSMA algorithm.

The rest of this chapter is organized as follows. In Section 5.1, we describe our models for the multihop wireless network and file arrivals/departures. We describe our network control policy (congestion control mechanism and scheduling mechanism) in Section 5.2.2. Section 5.3 is devoted to the formal statement about the throughput-optimality of our network control policy. In Section 5.4, we consider the distributed implementation of our policy using random access mechanisms.

5.1 From single-hop to multihop

A multihop wireless network consist of a set of nodes $\mathcal{N} = \{1, 2, .., N\}$ and a set of links $\mathcal{L}$ between the nodes. There is a link from i to j, i.e., $(i, j) \in \mathcal{L}$, if transmission from i to j is allowed. There is a set of users/source nodes $\mathcal{U} \subseteq \mathcal{N}$ and each user/source transfers data to a destination over a fixed route in the network.[2] For a user/source $u \in \mathcal{U}$, we use $d(u)(\neq u)$ to denote its destination. Let $\mathcal{D} := d(\mathcal{U})$ denote the set of all destinations.

We consider a time-slotted system. At each time slot t, files of different sizes arrive at the source nodes. As in the standard congestion control algorithm, TCP, files inject packets into their MAC-layer queues. The packets then travel to their respective destinations in a multihop fashion, i.e., along links in the network with queueing in buffers at intermediate nodes. Transmission of each packet along its route is subject to physical layer constraints such as interference and limited link capacity.

Recall from Chapter 2 that $\mathcal{M}$ denotes the set of feasible schedules $X = [x_{ij} : (i, j) \in \mathcal{L}]$ at each time slot. Thus, x_{ij} is the number of packets that can be transmitted from i to j during time slot t if the schedule X is selected at time slot t. Note that each transmission schedule X corresponds to a set of node power assignments chosen by the network. Note that if $\gamma = [\gamma_{ij} : (i, j) \in \mathcal{L}]$ be the average rate of service provided to the links, then, $\gamma \in \mathrm{Co}(\mathcal{M})$.

We use $a_s(t)$ to denote the number of files that arrive at source s at time t and assume that the process $\{a_s(t); s \in \mathcal{U}\}_{\{t=1,2,\cdots\}}$ is i.i.d. over time and independent across users with rate $[\kappa_s; s \in \mathcal{U}]$ and has bounded second moments. Moreover, we assume that there are K possible file types where the files of type i are geometrically distributed with mean $1/\eta_i$ packets. The

[2]The final results can be extended to case when each source has multiple destinations or to the cases of multi-path routing and adaptive routing. Here, to expose the main features, we have considered a simpler model.

file arrived at source s can belong to type i with probability p_{si}, $i = 1, 2, .., K$, $p_{si} \geq 0$, $\sum_{i=1}^{k} p_{si} = 1$. Our motivation for selecting such a model is due to the *large variance distribution* of file sizes in the Internet. It is believed, see e.g., [52], that most of bytes are generated by long files while most of the files are short files. By controlling the probabilities p_{si}, for the same average file size, we can obtain distributions with very large variance. Let $m_s := \sum_{i=1}^{K} p_{si}/\eta_i$ denote the mean file size at node s, and define the workload at source s by $\rho_s := \kappa_s m_s$. Let $\boldsymbol{\rho} := [\rho_s : s \in \mathcal{U}]$ be the vector of such workloads in the network.

5.2 Network control policy

Upon arrival of a file at a source Transport layer, a TCP-connection is established that regulates the injection of packets into the MAC layer. Once transmission of a file ends, the file departs and the corresponding TCP-connection will be closed. The MAC-layer is responsible for making the scheduling decisions to deliver the MAC-layer packets to their destinations over their corresponding routes. Each node has a fixed routing table that determines the next hop for each destination.

At each source node, we index the files according to their arriving order such that the index 1 is given to the earliest file. This means that once transmission of a file ends, the indices of the remaining files are updated such that indices again start from 1 and are consecutive. Note that the indexing rule *is not* part of the algorithm implementation and it is used here only for the purpose of analysis.

5.2.1 Description of congestion control algorithm

We use $\mathcal{W}_{sf}(t)$ to denote the TCP *congestion window size* for file f at source s at time t. Hence, $\mathcal{W}_{sf}$ is a time-varying sequence which changes as a result of TCP congestion control. If the congestion window of file f is not full, TCP will continue injecting packets from the remainder of file f to the congestion window until file f has no packets remaining at the Transport layer or the congestion window becomes full. We consider *ingress queue-based congestion control* meaning that when a packet of congestion window departs the ingress queue, it is replaced with a new packet from its corresponding file at the Transport layer. It is important to note that the MAC layer does not know the number of remaining packets at the Transport layer, so scheduling decisions have to be made based on the MAC-layers information only. It is reasonable to assume that $1 \leq \mathcal{W}_{sf}(t) \leq \mathcal{W}_{cong}$, i.e., each file has at least one packet waiting to be transferred and all congestion window sizes are bounded from above by a constant $\mathcal{W}_{cong}$. We only require that the congestion window dynamics could be described as some function of queue lengths of the network. Even in the case that the congestion window is a function of the delayed queue lengths of the network up to T time slots earlier, due to the feedback delay of at most T from destination to source, the proof could be easily modified by including the queues up to T time slots before in the network state (details in Section 5.3).

5.2.2 Description of scheduling algorithm

At the MAC layer of each node $n \in \mathcal{N}$, we consider separate queues for the packets of different destinations. Let $q_n^{(d)}$, $d \in \mathcal{D}$, denote the packets of destination d at the MAC-layer of n. Also let $\mathbf{R}^{(d)}_{N \times N}$ be the routing matrix corresponding to packets of destination d where $R^{(d)}_{ij} = 1$ if the next hop of node i for destination d is node j, for some j such that $(i, j) \in \mathcal{L}$, and 0 otherwise. Routes are *acyclic* meaning that each packet eventually reaches its destination and leaves the network. A packet of destination d that is transmitted from i to j is removed from $q_i^{(d)}$ and added to $q_j^{(d)}$. The packet that reaches its destination is removed from the network. Note that packets in $q_n^{(d)}$ could be generated at node n itself (if n is a source with destination d) or belong to other sources that use n as an intermediate relay along their routes to destination d.

The scheduling algorithm is essentially the back-pressure algorithm [1] but it only uses the MAC-layer information. The key step in establishing the optimality of such an algorithm is using an appropriate weight function of the MAC-layer queues instead of using the total queues. In particular, consider a *log-type* function as in Chapter 3, which is

$$g(x) := \frac{\log(1 + x)}{h(x)}, \tag{5.1}$$

where $h(x)$ is an arbitrary increasing function which makes $g(x)$ an increasing concave function. Assume that $h(0) > 0$ and $g(x)$ is continuously differentiable on $(0, \infty)$: for example, $h(x) = \log(e + \log(1+x))$ or $h(x) = \log^{\epsilon}(e+x)$ for some $\epsilon > 0$. For each link (i, j) with $R_{ij}^{(d)} = 1$, define

$$w_{ij}^{(d)}(t) := g\Big(q_i^{(d)}(t)\Big) - g\Big(q_j^{(d)}(t)\Big). \tag{5.2}$$

Note that if $\{d \in \mathcal{D} : R_{ij}^{(d)} = 1\} = \emptyset$, then we can remove the link (i, j) from the network without reducing the capacity region since no packets are forwarded over it. So without loss of generality, we assume that $\{d \in \mathcal{D} : R_{ij}^{(d)} = 1\} \neq \emptyset$, for every $(i, j) \in \mathcal{L}$.

Let $x_{ij}^{(d)}(t)$ denote the scheduling variable that shows the rate at which the packets of destination d can be forwarded over the link (i, j) at time slot t. The scheduling algorithm is as follows.

At each time t:

- Each node n observes the MAC-layer queue sizes of itself and its next hop, i.e., for each $d \in \mathcal{D}$, it observes $q_n^{(d)}$ and $q_j^{(d)}$ for a j such that $R_{ij}^{(d)} = 1$.

- For each link (i, j), calculate a weight

$$w_{ij}(t) := \max_{d \in \mathcal{D}: R_{ij}^{(d)} = 1} w_{ij}^{(d)}(t), \tag{5.3}$$

  and

$$\tilde{d}_{ij}^{*}(t) := \arg\max_{d \in \mathcal{D}: R_{ij}^{(d)} = 1} w_{ij}^{d}(t). \tag{5.4}$$

- Find the optimal rate vector $\tilde{x}^* \in \mathcal{M}$ that solves

$$\tilde{x}^*(t) = \arg\max_{r \in \mathcal{M}} \sum_{(i,j) \in \mathcal{L}} r_{ij} w_{ij}(t). \tag{5.5}$$

- Finally, assign $x_{ij}^{(d)}(t) = \tilde{x}_{ij}^{*}$ if $d = \tilde{d}_{ij}^{*}(t)$, and zero otherwise (break ties at random).

5.3 System stability

In this section, we analyze the system and prove its stability under the network control policy described in Section 5.2.

For the analysis, we use $Q_n^{(d)}$ (*with capital* Q) to denote the *total per-destination queues*, i.e., the total number of packets of destination d at node n, in its MAC or Transport layer. Note that, for each node n, the MAC (or total) per-destination queues $q_n^{(d)}$ (or $Q_n^{(d)}$) fall into three cases: (i) n is source and d is its destination, (ii) n is a source but d is not its destination, and (iii) n is not a source. In the case (i), it is important to distinguish between the MAC-layer queue and the total queue associated with d, i.e., $Q_n^{(d)} \neq q_n^{(d)}$, because of the existing packets of destination d at the Transport layer of n. However, $Q_n^{(d)} = q_n^{(d)}$ holds in case (ii), and for all destinations in case (iii).

Let $z_{ij}(t)$ denote the number of packets transmitted over link $(i,j) \in \mathcal{L}$ at time t. Then, the total-queue dynamics for a destination d, at each node n, is given by

$$Q_n^{(d)}(t+1) = Q_n^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} z_{nj}^{(d)}(t) + \sum_{i=1}^{N} R_{in}^{(d)} z_{in}^{(d)}(t) + A_n^{(d)}(t), \qquad (5.6)$$

where $A_n^{(d)}(t)$ is the total number of packets for destination d that new files bring to node n at time slot t. Note that $A_n^{(d)}(t) \equiv 0$ in the cases (ii) and (iii) above. With minor abuse of notation, we write $\mathbb{E}\left[A_n^{(d)}(t)\right] = \rho_n^{(d)}$ with $\rho_n^{(d)} := \rho_n$ in the case (i) and $\rho_n^{(d)} := 0$ otherwise. Also $z_{ij}^{(d)}(t) = \min\left\{x_{ij}^{(d)}(t), q_i^{(d)}(t)\right\}$ obviously, because i cannot send more than its MAC-layer queue content at each time.

Definition 5.1. *The capacity region of the network $\mathcal{C}$ is defined as the set of all load vectors $\boldsymbol{\rho}$ that under which the total queues in the network can be stabilized. Note that under our flow-level model, stability of total queues will imply that the number of files in the network is also stable. It is well-known, see e.g. [53], that a vector $\boldsymbol{\rho}$ belongs to $\mathcal{C}$ if and only if there exits an average* *service rate vector* $\gamma \in Co(\mathcal{M})$ *such that*

$$\gamma_{ij}^{(d)} \geq 0;\ \forall d \in \mathcal{D} \text{ and } \forall (i,j) \in \mathcal{L},$$
$$\rho_n^{(d)} - \sum_{j=1}^{N} R_{nj}^{(d)} \gamma_{nj}^{(d)} + \sum_{i=1}^{N} R_{in}^{(d)} \gamma_{in}^{(d)} \leq 0;\ \forall d \in \mathcal{D},\ \forall n \neq d,$$
$$\sum_{d \in \mathcal{D}} \gamma_{ij}^{(d)} \leq \gamma_{ij};\ \forall (i,j) \in \mathcal{L}.$$

Theorem 5.1. *For any* $\boldsymbol{\rho}$ *strictly inside* $\mathcal{C}$*, the scheduling algorithm can stabilize the network independent of congestion control mechanism and the (non-idling) service discipline used to transmit packets from active nodes.*

Remark 5.1. *Theorem 5.1 holds even when* $h \equiv 1$ *in (5.1), however, for the distributed implementation of the algorithm in Section 5.4, we need* g *to grow slightly slower than* log.

Theorem 5.1 shows that it is possible to design the ingress queue-based congestion controller regardless of the scheduling algorithm implemented in the core network. This will allow using different congestion control mechanisms at the edge of the network for different fairness or QoS considerations without need to change the scheduling algorithm implemented at internal routers of the network. As we will see, a key ingredient of such a decomposition result is the use of difference between the logarithms of queue lengths, as in (5.2), for the link weights in the scheduling algorithm.

The rest of this section is devoted to proof of Theorem 5.1.

5.3.1 Proof of Theorem 5.1

Order of events

Since we use a discrete-time model, we have to specify the order in which files/packets arrive and depart, which we do below:

1. At the beginning of each time slot, a scheduling decision is made by the scheduling algorithm. Packets depart from the MAC layers of scheduled links.

2. File arrivals occur next. Once a file arrives, a new TCP connection is set up for that file with an initial pre-determined congestion window size.

3. For each TCP connection, if the congestion window is not full, packets are injected into the MAC layer from the Transport layer until the window size is fully used or there are no more packets at the Transport layer.

We re-index the files at the beginning of each time slot because some files might have been departed during the last time slot.

State of the system

Define the state of node n as

$$\mathcal{S}_n(t) = \big\{(q_n^{(d)}(t), \mathcal{I}_n^{(d)}(t)) : d \in \mathcal{D}, (\xi_{nf}(t),\ \mathcal{W}_{nf}(t), \sigma_{nf}(t)) : 1 \leq f \leq N_n(t)\big\}.$$

Here $N_n(t)$ is the number of existing files at node n at the beginning of time slot t, $\sigma_{nf}(t) \in \{1/\eta_1, \cdots, 1/\eta_K\}$ is its mean size (or type), and $\mathcal{W}_{nf}(t)$ is its corresponding congestion window size. Note that $\sigma_{nf}(t)$ is a function of time only because of re-indexing since a file might change its index from slot to slot. $\xi_{nf}(t) \in \{0, 1\}$ indicates whether file f has still packets in the Transport layer. More accurately, if $U_{nf}(t)$ is the number of remaining packets of file f at node n, then $\xi_{nf}(t) = \mathbb{1}\{U_{nf}(t) > \mathcal{W}_{nf}(t)\}$. Obviously, if n is not a source node, then we can remove $(\xi_{nf}, \mathcal{W}_{nf}, \sigma_{nf})$ from the description of $\mathcal{S}_n$. $\mathcal{I}_n^{(d)}(t)$ denotes the information required about $q_n^{(d)}(t)$ to serve the MAC-layer packets which depends on the specific service discipline implemented in MAC-layer queues. In the rest of the chapter, we consider the case of FIFO (*First In-First Out*) service discipline in MAC-layer queues. In this case, $\mathcal{I}_n^{(d)}(t)$ is simply the ordering of packets in $q_n^{(d)}(t)$ according to their entrance times. As it will turn out from the proof, the system stability will hold for any none-idling service discipline. Define the state of the system to be $\mathcal{S}(t) = \{\mathcal{S}_n(t) : n \in \mathcal{N}\}$. Now, given the scheduling algorithm in Section 5.2.2, and our system model in Section 5.1, $\mathcal{S}(t)$ evolves as a discrete-time Markov chain.

Next, we analyze the *Lyapunov drift* to show that the network Markov chain is positive recurrent and, as a result, the number of files in the system and queue sizes are stable.

Lyapunov analysis

Define $\bar{Q}_n^{(d)}(t) := \mathbb{E}\left[Q_n^{(d)}(t)|\mathcal{S}_n(t)\right]$ to be the expected total queue length at node n given the state $\mathcal{S}_n(t)$. Then, if n is a source, and d is its destination,

$$\bar{Q}_n^{(d)}(t) = q_n^{(d)}(t) + \sum_{f=1}^{N_n(t)} \Big[\sigma_{nf}(t)\xi_{nf}(t)\Big]. \tag{5.7}$$

Otherwise, if $d \neq d(n)$ or n is not a source, then $\bar{Q}_n^{(d)}(t) = q_n^{(d)}(t)$. Note that given the state $\mathcal{S}(t)$, $\bar{Q}_n^{(d)}$ is known.

The dynamics of $\bar{Q}_n^{(d)}(t)$ involves the dynamics of $q_n^{(d)}(t)$, $\xi_n(t)$, and $N_n(t)$, and, thus, it consists of: (i) departure of MAC-layer packets, (ii) new file arrivals (if n is a source), (iii) arrival of packets from previous hops that use n as an intermediate relay, (iv) injection of packets into the MAC layer (if n is a source), and (v) departure of files from the Transport layer (if n is a source). Hence,

$$\begin{aligned}\bar{Q}_n^{(d)}(t+1) \quad &= \quad \bar{Q}_n^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} z_{nj}^{(d)}(t) + \bar{A}_n^{(d)}(t) + \sum_{i=1}^{N} R_{in}^{(d)} z_{in}^{(d)}(t) \\ &\quad + \hat{A}_n^{(d)}(t) - \hat{D}_n^{(d)}(t),\end{aligned} \tag{5.8}$$

where $\bar{A}_n^{(d)}(t) = \sum_{f=N_n(t)+1}^{N_n(t)+a_n(t)} \sigma_{nf}(t)$ is the expected number of packet arrivals due to new files, $\hat{A}_n^{(d)}(t)$ is the total number of packets injected into the MAC layer to fill up the congestion window after scheduling and new file arrivals, and $\hat{D}_n^{(d)}(t) = \sum_{f=1}^{N_n(t)+a_n(t)} \sigma_{nf}(t) I_{nf}(t)$ is the Transport-layer "expected packet departure" because of the MAC-layer injections. Here, $I_{nf}(t) = 1$ indicates that the last packet of file f *leaves the Transport layer during time slot* t; otherwise, $I_{nf}(t) = 0$.[3] Note that $\mathbb{E}\left[\bar{A}_n^{(d)}(t)\right] = \rho_n^{(d)}$.

Let $B_n^{(d)}(t) := \hat{A}_n^{(d)}(t) - \hat{D}_n^{(d)}(t)$ in (5.8), and $\mathbb{E}_{\mathcal{S}}[\cdot] := \mathbb{E}_{\mathcal{S}}[\cdot|\mathcal{S}(t)]$. It should

[3]To notice the difference between the indicators $I_{nf}(t)$ and $\xi_{nf}(t)$, consider a specific file and assume that its last packet enters the Transport layer at time slot t_0, departs the Transport layer during time slot t_1 and departs the MAC layer during time slot t_2, then its corresponding indicator I is 1 at time t_1 and is 0 for $t_0 \leq t < t_1$ and $t_1 < t \leq t_2$, while its indicator ξ is 0 for all time $t_1 \leq t \leq t_2$, and 1 for $t_0 \leq t < t_1$.

be clear that when n is a source but $d \neq d(n)$, or when n is not a source, $\bar{A}_n^{(d)}(t) = \hat{A}_n^{(d)}(t) = \hat{D}_n^{(d)}(t) = B_n^{(d)}(t) \equiv 0$. Let r_{max} denote the maximum link capacity over all the links in the network. Lemma 5.1 characterizes the first and second moments of $B_n^{(d)}(t)$.

Lemma 5.1. *Let $\eta_{min} := \min_{1\leq i\leq K} \eta_i$. For the process $\{B_n^{(d)}(t)\}$,*

(i) $\mathbb{E}_{\mathcal{S}(t)}\left[B_n^{(d)}(t)\right] = 0$.

(ii) $\mathbb{E}_{\mathcal{S}(t)}\left[B_n^{(d)}(t)^2\right] \leq (\kappa_n + N^2 r_{max}^2)\max\{\mathcal{W}_{cong}^2, 1/\eta_{min}^2\}$.

Therefore, we can write

$$\bar{Q}_n^{(d)}(t+1) = \bar{Q}_n^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} z_{nj}^{(d)}(t) + \tilde{A}_n^{(d)}(t) + \sum_{i=1}^{N} R_{in}^{(d)} z_{in}^{(d)}(t), \qquad (5.9)$$

where $\tilde{A}_n^{(d)}(t) := \bar{A}_n^{(d)}(t) + B_n^{(d)}(t)$. Note that $\tilde{A}_n^{(d)}(t)$ has mean $\rho_n^{(d)}$ and finite second moment.

Let $G(u) := \int_0^u g(x)dx$ for the function g defined in (5.1). Then G is a strictly convex and increasing function. Consider a Lyapunov function

$$V(\mathcal{S}(t)) = \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} G(\bar{Q}_n^{(d)}(t)).$$

Let $\Delta V(t) := V(\mathcal{S}(t+1)) - V(\mathcal{S}(t))$. Using convexity and monotonicity of G, we get

$$\Delta V(t) \leq \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t+1))\left(\bar{Q}_n^{(d)}(t+1) - \bar{Q}_n^{(d)}(t)\right).$$

Next, observe that, using (5.9),

$$|\bar{Q}_n^{(d)}(t+1) - \bar{Q}_n^{(d)}| \leq \tilde{A}_n^{(d)}(t) + N r_{max}.$$

Hence, because g is strictly increasing,

$$\begin{aligned} g(\bar{Q}_n^{(d)}(t+1)) &\leq g\Big(\bar{Q}_n^{(d)}(t) + \tilde{A}_n^{(d)}(t) + N r_{max}\Big) \\ &\leq g(\bar{Q}_n^{(d)}(t)) + (\tilde{A}_n^{(d)}(t) + N r_{max}), \end{aligned}$$

where the last inequality follows from concavity of g and the fact that $g' \leq 1$. Hence,

$$\Delta V(t) \leq \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))(\bar{Q}_n^{(d)}(t+1) - \bar{Q}_n^{(d)}(t)) + \sum_{n=1}^{N}\sum_{d\in\mathcal{D}}(\tilde{A}_n^{(d)}(t) + Nr_{max})^2.$$

Define $u_n^{(d)}(t) := \max\left\{\sum_{j=1}^{N} R_{nj}^{(d)} x_{nj}^{(d)}(t) - q_n^{(d)}(t), 0\right\}$, to be the wasted service for packets of destination d, i.e., when n is included in the schedule but it does not have enough packets of destination d to transmit. Then, we have

$$\begin{aligned}\Delta V(t) \;\; \leq \;\; & \sum_{n=1}^{N}\sum_{d\in\mathcal{D}}\Big\{g(\bar{Q}_n^{(d)}(t))\Big[\sum_{i=1}^{N} R_{in}^{(d)} x_{in}^{(d)}(t) + \tilde{A}_n^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} x_{nj}^{(d)}(t)\Big]\Big\} \\ & + \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) u_n^{(d)}(t) + \sum_{n=1}^{N}\sum_{d\in\mathcal{D}}(\tilde{A}_n^{(d)}(t) + Nr_{max})^2.\end{aligned}$$

Taking the expectation of both sides, given the state at time t is known, yields

$$\begin{aligned}&\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) u_n^{(d)}(t)\Big] \\ &+ \sum_{n=1}^{N}\sum_{d\in\mathcal{D}}\Big\{g(\bar{Q}_n^{(d)}(t))\mathbb{E}_{\mathcal{S}(t)}[\rho_n^{(d)} + \sum_{i=1}^{N} R_{in}^{(d)} x_{in}^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} x_{nj}^{(d)}(t)]\Big\} + C_1,\end{aligned}$$

where $C_1 = \mathbb{E}\left[\sum_{n=1}^{N}\sum_{d\in\mathcal{D}}(\tilde{A}_n^{(d)}(t) + Nr_{max})^2\right] < \infty$, because $\mathbb{E}\left[\tilde{A}_n^{(d)}(t)^2\right] < \infty$.

Lemma 5.2. *There exists a positive constant C_2 such that, for all $\mathcal{S}(t)$,*

$$\sum_{n=1}^{N}\sum_{d\in\mathcal{D}} \mathbb{E}_{\mathcal{S}(t)}\Big[g(\bar{Q}_n^{(d)}(t)) u_n^{(d)}(t)\Big] \leq C_2.$$

Using Lemma 5.2 and changing the order of summations, we have

$$\begin{aligned}&\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq C_1 + C_2 + \textstyle\sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))\rho_n^{(d)} \\ &-\mathbb{E}_{\mathcal{S}(t)}\Big[\textstyle\sum_{(i,j)\in\mathcal{L}}\sum_{d\in\mathcal{D}} x_{ij}^{(d)}(t)(g(\bar{Q}_i^{(d)}(t)) - g(\bar{Q}_j^{(d)}(t)))\Big]. \qquad (5.10)\end{aligned}$$

Recall that the link weight that is actually used in the algorithm is based on the MAC-layer queues as in (5.2)-(5.3). For the analysis, we also define a new link weight based on the state as

$$W_{ij}(t) = \max_{d \in \mathcal{D}: R_{ij}^{(d)}=1} W_{ij}^{(d)}(t), \tag{5.11}$$

where, for a link $(i,j) \in \mathcal{L}$ with $R_{ij}^{(d)} = 1$,

$$W_{ij}^{(d)}(t) := g(\bar{Q}_i^d(t)) - g(\bar{Q}_j^d(t)). \tag{5.12}$$

Then, the two types of link weights only differ by a constant as stated by the following lemma.

Lemma 5.3. *Let $W_{ij}(t)$ and $w_{ij}(t)$, $(i,j) \in \mathcal{L}$, be the link weights defined by (5.11)-(5.12) and (5.2)-(5.3) respectively. Then at all times*

$$|W_{ij}(t) - w_{ij}(t)| \leq \frac{\log(1 + 1/\eta_{min})}{h(0)}.$$

Proof. Recall that, at each node n, for all destinations $d \neq d(n)$, we have $\bar{Q}_n^d(t) = q_n^d(t)$. If $d = d(n)$ is the destination of n, then $\bar{Q}_n^d(t)$ consists of: (i) packets of d received from upstream flows that use n as an intermediate relay, and (ii) MAC-layer packets received from the files generated at n itself. Since $1 \leq \mathcal{W}_{nf}(t) \leq \mathcal{W}_{cong}$, the number of files with destination d that are generated at node n or have packets at node n as an intermediate relay, is at most $q_n^{(d)}(t)$. Therefore, it is clear that $q_n^d(t) \leq \bar{Q}_n^d(t) \leq q_n^d(t) + q_n^d(t)\frac{1}{\eta_{min}}$. In the rest of the proof, we drop the dependence of queues on t for compactness. For all n and d, using a log-type function, as the function g in (5.1), yields

$$\begin{aligned}
g(q_n^d) \leq g(\bar{Q}_n^d) &\leq g\left(q_n^d(1 + 1/\eta_{min})\right) \\
&\leq \frac{\log\left((1 + q_n^d)(1 + 1/\eta_{min})\right)}{h(q_n^d(1 + 1/\eta_{min}))} \\
&\leq g(q_n^d) + \frac{\log(1 + 1/\eta_{min})}{h(0)}.
\end{aligned} \tag{5.13}$$

It then follows that, $\forall d \in \mathcal{D}$, and $\forall (i,j) \in \mathcal{L}$ with $R_{ij}^{(d)} = 1$,

$$|W_{ij}^{(d)} - w_{ij}^{(d)}| \leq \log(1 + 1/\eta_{min})/h(0). \tag{5.14}$$

Let $d^*_{ij} := \arg\max_{d:R^{(d)}_{ij}=1} W^{(d)}_{ij}$ and $\tilde{d}^*_{ij}$ as in (5.4). Then, using (5.14)

$$w_{ij} \geq w^{(d^*_{ij})}_{ij} \geq W_{ij} - \log(1 + 1/\eta_{min})/h(0),$$

and, similarly,

$$W_{ij} \geq W^{(\tilde{d}^*_{ij})}_{ij} \geq w_{ij} - \log(1 + 1/\eta_{min})/h(0).$$

This concludes the proof. □

Let $x^*(t)$ be the max weight schedule based on weights $\{W_{ij}(t) : (i, j) \in \mathcal{L}\}$, i.e.,

$$x^*(t) = \arg\max_{x \in \mathcal{M}} \sum_{(i,j)\in\mathcal{L}} x_{ij} W_{ij}(t). \tag{5.15}$$

Note the distinction between x^* and $\tilde{x}^*$ as we used $\tilde{x}^*(t)$ in (5.5) to denote the max weight schedule based on MAC-layer queues. The weights of the schedules $\tilde{x}^*$ and x^* differ only by a constant for all queue values as we show next. From definition of x^*, in (5.15),

$$\sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} W_{ij}(t) \geq 0. \tag{5.16}$$

On the other hand, we can write

$$\sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} W_{ij}(t) = A + B + C$$

where

$$\begin{aligned}
A &= \sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) \\
B &= \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} w_{ij}(t) \\
C &= \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} w_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} W_{ij}(t).
\end{aligned}$$

By Lemma 5.3, “A” and “C” are less than $N^2 r_{max} \log(1+1/\eta_{min})/h(0)$ each, and "B" is negative by definition of $\tilde{x}^*$ in (5.5). Thus,

$$\sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} W_{ij}(t) \leq 2N^2 r_{max} \log(1 + 1/\eta_{min})/h(0). \quad (5.17)$$

Hence, using (5.10), (5.11), and (5.17), *under MAC scheduling* $\tilde{x}^*$, the Lyapunov drift is bounded as follows

$$\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq \sum_{n=1}^{N} \sum_{d\in\mathcal{D}} \{g(\bar{Q}_n^{(d)}(t))\rho_n^{(d)}\} - \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j)\in\mathcal{L}} x^*_{ij}(t) W_{ij}\Big] + C,$$

where $C = C_1 + C_2 + 2N^2 r_{max} \log(1 + 1/\eta_{min})/h(0)$.

Accordingly, using (5.11)-(5.12), and changing the order of summations in the right-hand side of the above inequality yields

$$\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq$$
$$\sum_{n=1}^{N} \sum_{d\in\mathcal{D}} \left\{ g(\bar{Q}_n^{(d)}(t)) \mathbb{E}_{\mathcal{S}(t)}\big[\rho_n^{(d)} + \sum_{i=1}^{N} R_{in}^{(d)} {x^*}_{in}^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} {x^*}_{nj}^{(d)}(t)\big] \right\} + C,$$

where ${x^*}_{ij}^{(d)}(t) = x^*_{ij}(t)$ for $d = d^*_{ij}(t)$ (ties are broken at random) and is zero otherwise. The rest of the proof is standard. Since load $\boldsymbol{\rho}$ is strictly inside the capacity region, there must exist a $\epsilon > 0$ and a $\gamma \in \mathrm{Co}(\mathcal{M})$ such that

$$\rho_n^{(d)} + \epsilon \leq \sum_{j=1}^{N} R_{nj}^{(d)} \gamma_{nj}^{(d)} - \sum_{i=1}^{N} R_{in}^{(d)} \gamma_{in}^{(d)} \ ; \forall n \in \mathcal{N}, \forall d \in \mathcal{D}. \quad (5.18)$$

Hence, for any $\delta > 0$,

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq\ & \sum_{n=1}^{N} \sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) \left[\sum_{i=1}^{N} R_{in}^{(d)} {x^*}_{in}^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} {x^*}_{nj}^{(d)}(t)\right] \\
& - \sum_{n=1}^{N} \sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) \left[\sum_{i=1}^{N} R_{in}^{(d)} \gamma_{in}^{(d)}(t) - \sum_{j=1}^{N} R_{in}^{(d)} \gamma_{nj}^{(d)}(t)\right] \\
& - \epsilon \sum_{n=1}^{N} \sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) + C.
\end{aligned}$$

But from the definition of $x^*(t)$ and convexity of $\mathrm{Co}(\mathcal{M})$, $\sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) \geq$

$\sum_{(i,j)\in\mathcal{L}} \gamma_{ij} W_{ij}(t)$, $\forall \gamma \in \mathrm{Co}(\mathcal{M})$, hence,

$$\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \leq -\epsilon \sum_{n=1}^{N} \sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) + C \leq -\delta,$$

whenever $\max_{n,d} \bar{Q}_n^{(d)} \geq g^{-1}\big(\frac{C_2+\delta}{\epsilon}\big)$ or, as a sufficient condition, whenever $\max_{n,d} q_n^{(d)} \geq g^{-1}\big(\frac{C_2+\delta}{\epsilon}\big)$. Therefore, it follows that the system is stable by an extension of the Foster-Lyapunov criteria [54] (Theorem 3.1 in [1]). In particular, queue sizes and the number of files in the system are stable.

Remark 5.2. *Although we have assumed that file sizes follow a mixture of geometric distributions, our results also hold for the case of bounded file sizes with general distribution. The proof argument for the latter case is obtained by minor modifications of the proof presented in this chapter (see [55]) and, hence, has been omitted for brevity.*

5.4 Distributed implementation

The optimal scheduling algorithm in Section 5.2.2 can be implemented in a distributed manner using the discrete-time random access mechanism (Chapter 2) as we show next.

For simplicity, we consider the following criterion for successful packet reception: packet transmission over link $(i, j) \in \mathcal{L}$ is successful if none of the neighbors of node j are transmitting. Furthermore, we assume that every node can transmit to at most one node at each time, receive from at most one node at each time, and cannot transmit and receive simultaneously (over the same frequency band). This especially models the packet reception in the case that the set of neighbors of node i, i.e., $C(i) = \{j : (i, j) \in \mathcal{L}\}$, is the set of nodes that are within the transmission range of i and the interference caused by i at all other nodes, except its neighbors, is negligible. Moreover, the packet transmission over (i, j) is usually followed by an ACK transmission from receiver to sender, over (j, i). Hence, for a *synchronized data/ACK system*, we can define a *Conflict Set* (CS) for link (i, j) as

$$\mathrm{CS}_{(i,j)} \quad = \quad \{(a, b) \in \mathcal{L} : a \in C(j), \text{ or } b \in C(i), \text{or } a \in \{i, j\}, \text{ or } b \in \{i, j\}\}.$$

This ensures that when the links in $\mathrm{CS}_{(i,j)}$ are inactive, the data/ACK transmission over $(i,j)/(j,i)$ is successful. Thus, we can represent the interference constraints by using a conflict graph $\mathcal{G}(\mathcal{V},\mathcal{E})$, where each vertex in $\mathcal{V}$ is a communication link, and there is an edge between vertices (i,j) and (a,b) if simultaneous transmissions over the communication links (i,j) and (a,b) are not successful.

Furthermore, for simplicity, assume that in each time slot, at most one packet could be successfully transmitted over a link (i,j), i.e., $x_{ij}(t) \in \{0,1\}$. Let $|\mathcal{L}|$ denote the number of wireless links.

We say that a node is active if it is a sender or a receiver for some active link. Inactive nodes can *sense the wireless medium* and know if there is an active node in their neighborhood. This is possible because we use a synchronized data/ACK system and detecting active nodes can be performed by sensing the data transmission of active senders and sensing the ACK transmission of active receivers. Hence, using such carrier sensing, nodes i and j know if the channel is idle, i.e., $\sum_{(a,b)\in\mathrm{CS}_{(i,j)}} x_{ab}(t) = 0$, or if the channel is busy, i.e., $\sum_{(a,b)\in\mathrm{CS}_{(i,j)}} x_{ab}(t) \geq 1$.

Remark 5.3. *For the case of single-hop networks, the link weight (5.3) is reduced to $w_{ij}(t) = g(1+q_i(t))/h(q_i(t))$ where i is the source and j is the destination of flow over (i,j). Such a weight function is exactly the one that under which throughput optimality of random access has been established in Chapter 3. Next, we will propose a slightly modified version of the random access mechanism that is suitable for the general case of multihop flows.*

5.4.1 Basic random access mechanism for multihop networks

For our algorithm, *based on the MAC layer information*, we define a modified weight for each link (i,j) as

$$\tilde{w}_{ij}(t) = \max_{d:R_{ij}^{(d)}=1} \tilde{w}_{ij}^{(d)}(t), \tag{5.19}$$

where

$$\tilde{w}_{ij}^{(d)}(t) = \tilde{g}\left(q_i^{(d)}(t)\right) - \tilde{g}\left(q_j^{(d)}(t)\right), \tag{5.20}$$

and,

$$\tilde{g}\left(q_i^{(d)}(t)\right) = \max\left\{g\left(q_i^{(d)}(t)\right), g^*(t)\right\}, \tag{5.21}$$

where the function g is the same as (5.1) defined for the centralized algorithm, and

$$g^*(t) := \frac{\epsilon}{4N^3} g(q_{max}(t)), \tag{5.22}$$

where $q_{max}(t) := \max_{i,d} q_i^{(d)}(t)$ is the maximum MAC-layer queue length in the network at time t and assumed to be known, and ϵ is an arbitrary small but fixed positive number. Note that if we remove $g^*(t)$ from the above definition, then $\tilde{w}_{ij}$ is equal to w_{ij} in (5.2)-(5.3).

Consider the conflict graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ of the network as defined earlier. At each time slot t, a link (i, j) is chosen uniformly at random, then

(i) If $\tilde{x}_{ab}(t-1) = 0$ for all links $(a, b) \in \mathrm{CS}_{(i,j)}$, then $\tilde{x}_{ij}(t) = 1$ with probability $p_{ij}(t)$, and $\tilde{x}_{ij}(t) = 0$ with probability $1 - p_{ij}(t)$. Otherwise, $\tilde{x}_{ij}(t) = 0$.

(ii) $\tilde{x}_{ab}(t) = x_{ab}(t-1)$ for all $(a, b) \neq (i, j)$.

(iii) $x_{ij}^{(d)}(t) = \tilde{x}_{ij}(t)$ if $d = \arg\max_{d:R_{ij}^{(d)}=1} \tilde{w}_{ij}^{(d)}(t)$ (break ties at random), and zero otherwise.

We choose $p_{ij}(t)$ to be

$$p_{ij}(t) = \frac{\exp(\widetilde{w}_{ij}(t))}{1 + \exp(\widetilde{w}_{ij}(t))}. \tag{5.23}$$

The following theorem states the main result regarding the throughput optimality of the above algorithm.

Theorem 5.2. *Under the function g specified in (5.1), the basic random access mechanism, with any $\epsilon > 0$, can stabilize the network for any $\boldsymbol{\rho} \in (1-3\epsilon)\mathcal{C}$, independent of Transport-layer ingress queue-based congestion control (as long as the minimum window size is one and the window sizes are bounded) and the (non-idling) service discipline used to serve packets of active queues.*

5.4.2 Distributed implementation

The basic algorithm is based on Glauber dynamics with one node update at each time. For distributed implementation, we use the random access mechanism with multiple node updates as described in Chapter 2. Next, we describe the mechanisms for generation of decision schedules and data transmission schedules in more detail.

Generation of decision schedule

We divide the control slot into two mini-slots. In the first mini-slot, each node i chooses one of its neighbors $j \in C(i)$ uniformly at random, then it transmits a RTD (*Request-To-Decide*) packet, containing the ID(index) of node j, with probability β_i. If RTD is received successfully by j (i.e., j and none of the neighbors of j transmit RTD messages), in the second mini-slot, j sends a CTD (*Clear-To-Decide*) packet back to i, containing the ID of node i. The CTD message is received successfully at i if there is no collision with other CTD messages. Given a successful RTD/CTD exchange over the link (i, j), the link (i, j) will be included in the decision schedule m and no link from $\text{CS}_{(i,j)}$ will be included in m. Hence, m is a valid schedule. So each node i needs to maintain the following memories:

- $AS_i(t)/AR_i(t)$: Node i is included in $m(t)$ as a sender/receiver for some link.
- $ID_i(t)$: The index of the node which is paired with i as its sender (when $AR_i(t) = 1$) or its receiver (when $AS_i(t) = 1$).
- $NR_i(t)/NS_i(t)$: Carrier sense by node i, i.e., node i has an active receiver/sender in it neighborhood during data slot t.

The CTD message sent back from a node j to i also contains the carrier sense information of node j, i.e., $NR_j(t-1)$ and $NS_j(t-1)$, and the vector of MAC layer queue sizes of node j at time t, i.e, $q_j^{(d)}(t)$.

Generation of data transmission schedule

After the control slot, every node i knows if it is included in the decision schedule $m(t)$, as a sender, and also knows its corresponding receiver $ID_i = j$.

Algorithm 1 Decision schedule at control slot t

1: For every node i, set $AS_i(t) = AR_i(t) = 0$.
2: In the first mini-slot:
- $AS_i(t) = 1$ with probability β_i; otherwise $AS_i(t) = 0$.
-If $AS_i(t) = 1$, choose a node $j \in C(i)$ uniformly at random and send a RTD to j and set $ID_i(t) = j$; otherwise listen for RTD messages.
3: In the second mini-slot:
-If received a RTD from j in the first mini-slot, send a CTD to j and set $AR_i(t) = 1$ and $ID_i = j$; nodes with $AS_i(t) = 1$ listen for CTD messages.
-If $AS_i(t) = 1$ and CTD received successfully from $ID_i(t)$, include $(i, ID_i(t))$ in $m(t)$, otherwise $AS_i(t) = 0$.

The data transmission schedule at time t, i.e., $x(t)$, is generated based on $x(t-1)$ and $m(t)$. Only those links that are in $m(t)$ can change their states and the state of other links remain unchanged. A link (i, j) that is included in $m(t)$, can start a packet transmission with probability $p_{ij}(t)$ only if its conflict set has been silent during the previous time slot, as in the basic CSMA algorithm.

Algorithm 2 Data transmission schedule at slot t

1: - $\forall$ i with $AS_i(t) = 1$ and receiver $j = ID_i$:
If no links in $\mathrm{CS}_{(i,j)}$ were active in the previous data slot, i.e., $x_{ij}(t-1) = 1$ or $NR_i(t-1) = NS_j(t-1) = 0$,

- $x_{ij}(t) = 1$ with probability $p_{ij}(t)$,
- $x_{ij} = 0$ with probability $\bar{p}_{ij}(t) = 1 - p_{ij}(t)$.

Else $x_{ij}(t) = 0$.
- $\forall (i, j) \notin m(t)$: $x_{ij}(t) = x_{ij}(t-1)$.
2: In the data slot, use $x(t)$ as the transmission schedule.

Data transmission and carrier sensing

In the data slot, we use $x(t)$ for the data transmission. In this slot, every node i will perform of the following.

$x_{ij}(t) = 1$: Node i will send a data packet to node j.

$x_{ji}(t) = 1$: Node i will send an ACK to node j after receiving a data packet from j.

All other nodes are inactive and perform carrier sensing. Since the data/ACK transmissions are synchronized in our system, every inactive node i will set $NS_i(t) = 0$ is it does not sense any transmission during the data transmission period and set $NS_i(t) = 1$ otherwise. Similarly, node i will set $NR_i(t) = 0$ if it senses no signal during the ACK transmission period and set $NR_i(t) = 1$ otherwise.

Remark 5.4. *In IEEE 802.11 DCF, the RTS/CTS exchange is used to reduce the Hidden Terminal Problem. However, even with RTS/CTS, the hidden terminal problem can still occur. Since, in our synchronized system, RTD and CTD messages are sent in two different mini-slots, this completely eliminates the hidden terminal problem.*

Corollary 5.1. *Under the weight function g specified in (5.1), the distributed algorithm can stabilize the network for any $\boldsymbol{\rho} \in (1 - 3\epsilon)\mathcal{C}$, independently of the congestion control mechanism.*

The rest of this section is devoted to proof of Theorem 5.2. The proof of Corollary 5.1 is almost identical.

5.4.3 Proof of Theorem 5.2

The proof uses ideas from Chapter 3, thus we focus only on the main differences here. First, recall that when weights are fixed, the basic algorithm is essentially an irreducible, aperiodic, and reversible Markov chain to generate the independent sets of $\mathcal{G}(\mathcal{V}, \mathcal{E})$, with the stationary distribution

$$\pi(s) = \frac{1}{Z} \exp\Big(\sum_{(i,j) \in s} \tilde{w}_{ij} \Big); \quad s \in \mathcal{M}, \tag{5.24}$$

where Z is the normalizing constant.

We start with the following lemma that relates the modified link weight and the original link weight.

Lemma 5.4. *For all links $(i, j) \in \mathcal{L}$, the link weights (5.19) and (5.3) differ at most by $g^*(t)$, i.e.,*

$$|\tilde{w}_{ij}(t) - w_{ij}(t)| \leq g^*(t). \tag{5.25}$$

Proof is simple and has been omitted. Next, we characterize the amount of change in the stationary distribution as a result of queue/file evolutions.

Lemma 5.5. *For any schedule* $s \in \mathcal{M}$, $e^{-\alpha_t} \leq \frac{\pi_{t+1}(s)}{\pi_t(s)} \leq e^{\alpha_t}$, *where,*

$$\alpha_t = 2(1 + \mathcal{W}_{cong})|\mathcal{L}|g'\Big(g^{-1}(g^*(t+1)) - 1 - \mathcal{W}_{cong}\Big), \tag{5.26}$$

and $\mathcal{W}_{cong}$ *is the maximum congestion window size.*

Now, equipped with Lemmas 3.2 and 5.5, we make use of the following key proposition from Chapter 3 that is reproduced below for completeness.

Proposition 5.1. *Given any* $\delta > 0$, $\|\pi_t - \mu_t\|_{TV} \leq \delta/4$ *holds when* $q_{max}(t) \geq q_{th} + t^*$, *if there exists a* q_{th} *such that*

$$\alpha_t T_{t+1} \leq \delta/16 \text{ whenever } q_{max}(t) > q_{th}, \tag{5.27}$$

where

(i) $T_t \leq 16^{|\mathcal{L}|}\exp(4|\mathcal{L}|\tilde{w}_{max}(t))$,

(ii) t^* *is the smallest* t *such that*

$$\frac{1}{\sqrt{\min_s \pi_{t_1}(s)}} \exp\Big(-\sum_{k=t_1}^{t_1+t^*} \frac{1}{T_k^2}\Big) \leq \delta/4, \tag{5.28}$$

with $q_{max}(t_1) = q_{th}$.

We will also use the following lemma that relates the maximum queue length and the maximum weight in the network. Hence, when one grows, the other one increases as well.

Lemma 5.6. *Let* $w_{max}(t) = \max_{(i,j)\in\mathcal{L}} w_{ij}(t)$. *Then*

$$\frac{1}{N} g\left(q_{max}(t)\right) \leq w_{max}(t) \leq g\left(q_{max}(t)\right).$$

Some useful properties of the basic algorithm

Lemma 5.7. *The Basic random access algorithm, with function* g *as in (5.1), satisfies the requirements of Proposition 5.1.*

The formal proof can be found in Section 5.6. Next, the following lemma states that, with high probability, the basic algorithm chooses schedules that their weights are close to the max weight schedule.

Lemma 5.8. *Given any* $0 < \varepsilon < 1$ *and* $0 < \delta < 1$*, there exists a* $B(\delta, \varepsilon) > 0$ *such that whenever* $q_{max}(t) > B(\delta, \varepsilon)$*, the basic random access algorithm chooses a schedule* $s(t) \in \mathcal{M}$ *such that*

$$\sum_{(i,j)\in s(t)} w_{ij}(t) \geq (1-\epsilon) \max_{s\in\mathcal{M}} \sum_{(i,j)\in s} w_{ij}(t),$$

with probability larger than $1 - \delta$.

Proof. Proof is in parallel with the arguments in Section 3.2.3. Let $w^*(t) = \max_{s\in\mathcal{M}} \sum_{(i,j)\in s} w_{ij}(t)$ and define $\chi_t := \{s \in \mathcal{M} : \sum_{(i,j)\in s} w_{ij}(t) < (1-\epsilon)w^*(t)\}$. Therefore, we need to show that $\mu_t(\chi_t) \leq \delta$, for $q_{max}(t)$ large enough. For our choice of $g(\cdot)$ and g^*, it follows from Proposition 5.1 that, whenever $q_{max}(t) > q_{th} + t^*$,

$$\sum_{s\in\chi_t} \mu_t(s) \leq \sum_{s\in\chi_t} \pi_t(s) + \delta/2.$$

Therefore, it suffices to have $\sum_{s\in\chi_t} \pi_t(s) \leq \delta/2$. But, by Lemma 5.4, $\widetilde{w}_{ij}(t) \leq w_{ij}(t) + g^*(t)$, so,

$$\begin{aligned}
\sum_{s\in\chi_t} \pi_t(s) &\leq \sum_{s\in\chi_t} \frac{1}{Z_t} e^{\sum_{(i,j)\in s} w_{ij}(t)} e^{|s|g^*(t)} \\
&\leq \sum_{s\in\chi_t} \frac{1}{Z_t} e^{(1-\varepsilon)w^*(t)} e^{|\mathcal{L}|g^*(t)},
\end{aligned}$$

and

$$Z_t = \sum_{s\in\mathcal{M}} e^{\sum_{(i,j)\in s} \widetilde{w}_{ij}(t)} > \sum_{s\in\mathcal{M}} e^{\sum_{(i,j)\in s}(w_{ij}(t)-g^*(t))} > e^{w^*(t)-|\mathcal{L}|g^*(t)}.$$

Therefore,

$$\sum_{s\in\chi_t} \pi_t(s) \leq 2^{|\mathcal{L}|} e^{2|\mathcal{L}|g^*(t)-\varepsilon w^*(t)},$$

when $q_{max}(t) > q_{th} + t^*$. Note that $w^*(t) \geq w_{max}(t) \geq g(q_{max}(t))/N$, and

$g^*(t) = \frac{\epsilon}{4N^3} g(q_{max}(t))$, so

$$\sum_{s \in \chi_t} \pi_t(s) \quad \leq \quad 2^{N^2} e^{-\frac{\epsilon}{2N} g(q_{max}(t))} \leq \delta/2,$$

whenever $q_{max}(t) > B(\delta, \epsilon)$ with

$$B(\delta, \epsilon) = \max \left\{ q_{th} + t^*, g^{-1}\Big(\frac{2N}{\epsilon}(N^2 \log 2 + \log \frac{2}{\delta})\Big) \right\}.$$

$\square$

Lyapunov analysis

Now we are ready to prove the stability of the network under our network policy, when we use the basic random access mechanism instead of the centralized scheduling algorithm. Let x^* and $\tilde{x}^*$ be the optimal schedules based on total queues and MAC queues respectively, given by (5.15) and (5.5), and $\tilde{x}$ be the schedule generated by the basic random access mechanism. The proof is parallel to the stability argument of the centralized algorithm. In particular, the inequality (5.10) still holds, which is

$$\begin{aligned} \mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] &\leq \sum_{n=1}^{N} \sum_{d \in \mathcal{D}} g(\bar{Q}_n^{(d)}(t)) \rho_n^{(d)} - \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j) \in \mathcal{L}} \sum_{d \in \mathcal{D}} x_{ij}^{(d)}(t) W_{ij}^{(d)}\Big] + C_1 + C_2 \\ &= \sum_{n=1}^{N} \sum_{d \in \mathcal{D}} g(\bar{Q}_n^{(d)}(t)) \rho_n^{(d)} - \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j) \in \mathcal{L}} \tilde{x}_{ij}(t) W_{ij}(t)\Big] + C_1 + C_2. \end{aligned}$$

Next, observe that

$$\sum_{(i,j) \in \mathcal{L}} x_{ij}^* W_{ij}(t) - \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j) \in \mathcal{L}} \tilde{x}_{ij} W_{ij}(t)\Big] = A + B + C,$$

where

$$\begin{aligned} A &= \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t)\Big] \\ B &= \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}_{ij} w_{ij}(t)\Big] \\ C &= \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j)\in\mathcal{L}} \tilde{x}_{ij} w_{ij}(t) - \sum_{(i,j)\in\mathcal{L}} \tilde{x}_{ij} W_{ij}(t)\Big]. \end{aligned}$$

Each of the terms "A" and "C" are less than $|\mathcal{L}|\log(1+1/\eta_{min})/h(0)$ by Lemma 5.3. The term "B" is bounded from above, by using Lemma 5.8, as follows.

$$\begin{aligned} B &\leq \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) - (1-\delta)(1-\epsilon) \sum_{(i,j)\in\mathcal{L}} \tilde{x}^*_{ij} w_{ij}(t) \\ &\leq \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) - (1-\delta)(1-\epsilon) \sum_{(i,j)\in\mathcal{L}} x^*_{ij} w_{ij}(t) \\ &\leq (1-(1-\delta)(1-\epsilon)) \sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) + |\mathcal{L}|\log(1+1/\eta_{min})/h(0), \end{aligned}$$

whenever $q_{max}(t) \geq B(\delta,\epsilon)$, for any $\delta > 0$. Thus, using the above bounds for terms A, B and C, we get

$$\begin{aligned} \mathbb{E}_{\mathcal{S}(t)}\Big[\sum_{(i,j)\in\mathcal{L}} \tilde{x}_{ij} W_{ij}(t)\Big] &\geq (1-\delta)(1-\epsilon) \sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) \\ &\quad -3|\mathcal{L}|\log(1+1/\eta_{min})/h(0). \end{aligned} \tag{5.29}$$

Using (5.29) yields

$$\begin{aligned} \mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] &\leq \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}^{(d)}_n(t))\rho^{(d)}_n - (1-\delta)(1-\epsilon) \sum_{(i,j)\in\mathcal{L}} x^*_{ij} W_{ij}(t) \\ &\quad +C_3, \end{aligned} \tag{5.30}$$

where $C_3 := C_1 + C_2 + 3|\mathcal{L}|\log(1+1/\eta_{min})/h(0)$. Using (5.11) and rewriting the right-hand side of (5.30), by changing the order of summations, yields

$$\begin{aligned}\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \;\; \leq \;\; & \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))\Big[\rho_n^{(d)} + (1-\delta)(1-\epsilon)\Big(\sum_{i=1}^{N} R_{in}^{(d)} {x^*}_{in}^{(d)}(t) \\ & -\sum_{j=1}^{N} R_{nj}^{(d)} {x^*}_{nj}^{(d)}(t)\Big)\Big] + C_3,\end{aligned}$$

whenever $q_{max}(t) \geq B(\delta, \epsilon)$. The rest of the proof is standard. For any load $\boldsymbol{\rho}$ strictly inside $(1-3\epsilon)\mathcal{C}$, there must exist a $\gamma \in \mathrm{Co}(\mathcal{M})$ such that for all $1 \leq n \leq N$, and all $d \in \mathcal{D}$,

$$\rho_n^{(d)} < (1-3\epsilon)\Big(\sum_{j=1}^{N} R_{nj}^{(d)}\gamma_{nj}^{(d)} - \sum_{i=1}^{N} R_{in}^{(d)}\gamma_{in}^{(d)}\Big). \tag{5.31}$$

Let $\frac{\rho^*}{1-3\epsilon} = \min_{n\in\mathcal{N},d\in\mathcal{D}}\left(\sum_j R_{nj}^{(d)}\gamma_{nj}^{(d)} - \sum_i R_{in}^{(d)}\gamma_{in}^{(d)}\right)$ for some positive ρ^*. Hence,

$$\begin{aligned}\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] \;\; \leq \;\; & (1-\delta)(1-\epsilon)\sum_{n=1}^{N}\sum_{d\in\mathcal{D}}\Big\{g(\bar{Q}_n^{(d)}(t))\Big[\sum_{i=1}^{N} R_{in}^{(d)} {x^*}_{in}^{(d)}(t) - \sum_{j=1}^{N} R_{nj}^{(d)} {x^*}_{nj}^{(d)}(t)\Big]\Big\} \\ + \;\; & (1-3\epsilon)\sum_{n=1}^{N}\sum_{d\in\mathcal{D}}\Big\{g(\bar{Q}_n^{(d)}(t))\Big[\sum_{j=1}^{N} R_{nj}^{(d)}\gamma_{nj}^{(d)} - \sum_{i=1}^{N} R_{in}^{(d)}\gamma_{in}^{(d)}\Big]\Big\} + C_3.\end{aligned}$$

For any fixed small $\epsilon > 0$, we can choose $\delta < \epsilon/(1-\epsilon)$ to ensure $(1-\delta)(1-\epsilon) > 1-2\epsilon$. Moreover, from definition of $x^*(t)$ and convexity of $\mathrm{Co}(\mathcal{M})$, it follows that

$$\begin{aligned}&\sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))\Big[\sum_{j=1}^{N} R_{nj}^{(d)} {x^*}_{nj}^{(d)}(t) - \sum_{i=1}^{N} R_{in}^{(d)} {x^*}_{in}^{(d)}(t)\Big] \\ &\geq \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))\Big[\sum_{j=1}^{N} R_{nj}^{(d)}\gamma_{nj}^{(d)} - \sum_{i=1}^{N} R_{in}^{(d)}\gamma_{in}^{(d)}\Big],\end{aligned} \tag{5.32}$$

for any $\gamma \in \text{Co}(\mathcal{M})$. Hence, for any fixed $\epsilon' > 0$,

$$\begin{aligned}\mathbb{E}_{\mathcal{S}(t)}\Big[\Delta V(t)\Big] &\leq -\epsilon \sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t))\Big[\sum_{j=1}^{N} R_{nj}^{(d)}\gamma_{nj}^{(d)} - \sum_{i=1}^{N} R_{in}^{(d)}\gamma_{in}^{(d)}\Big] + C_3 \\ &\leq -\rho^* \frac{\epsilon}{1-3\epsilon}\sum_{n=1}^{N}\sum_{d\in\mathcal{D}} g(\bar{Q}_n^{(d)}(t)) + C_3 \leq -\epsilon',\end{aligned}$$

whenever $\max_{n,d} \bar{Q}_n^{(d)} \geq g^{-1}\big(\frac{C_3+\epsilon'}{\rho^*}\frac{1-3\epsilon}{\epsilon}\big)$ and $q_{max}(t) \geq B(\delta,\epsilon)$ or, as a sufficient condition, whenever

$$q_{max}(t) \geq \max\left\{B(\delta,\epsilon), g^{-1}\big(\frac{C_3+\epsilon'}{\rho^*}\frac{1-3\epsilon}{\epsilon}\big)\right\}.$$

In particular, to get negative drift, $-\epsilon'$, it suffices that

$$\max_n N_n > \max\left\{g^{-1}\big(\frac{C_3+\epsilon'}{\rho^*}\frac{1-3\epsilon}{\epsilon}\big), B(\delta,\epsilon)\right\},$$

because $q_{max}(t) \geq \max_n N_n$, and g is an increasing function. This concludes the proof of Theorem 5.2.

5.5 Conclusions

In this chapter, we showed that α-fair congestion control is not necessary for flow-level stability. In fact, by using back-pressure with link weights that are log-differentials of (MAC-layer) queue lengths, the network stability is guaranteed for very general congestion control mechanisms. Hence, one can use different congestion control mechanisms for providing different QoS, without need to change the scheduling algorithm implemented at the internal routers of the network. The choice of log-differential link weights also enables us to implement our algorithm in a distributed fashion using random access schemes of Chapter 2, without loss of throughput optimality.

Our constraining assumptions regarding the congestion control mechanisms are very mild and compatible with the standard implementations like TCP. It is observed in [56] in the context of multiclass queueing systems that a fixed congestion window size implicitly solves an optimization problem in an asymptotic regime. It would be interesting to investigate how the congestion window dynamics and the links weights impact the system QoS performance for wireless networks. Our simulation results in [55] show that log-differential link weights, with a fixed congestion window size, reduce the file transfer delays. It will be certainly interesting to establish the validity of such an observation rigorously as a future research.

5.6 Additional proofs

5.6.1 Proof of Lemma 5.1

Let $\hat{A}^{(d)}_{nf}(t)$ denote the number of packets of file f injected into the MAC layer of node n, and $\hat{D}^{(d)}_{nf}(t) = \sigma_{nf}(t) I_{nf}(t)$ denote the expected "packet departure" of file f from the Transport layer. Let $B_{nf}(t) = \hat{A}^{(d)}_{nf}(t) - \hat{D}^{(d)}_{nf}(t)$ for file f. From the definition of $B_n(t)$, we have

$$B_n(t) = \sum_{f=1}^{N_n(t)} B_{nf}(t) + \sum_{f=N_n(t)+1}^{N_n(t)+a_n(t)} B_{nf}(t).$$

Part (i)

It suffices to show that for each individual file $1 \leq f \leq N_n(t)$, $\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)\Big] = 0$. We only need to focus on files f with $\xi_{nf}(t) = 1$, i.e., existing files in the Transport layer, or new files, i.e, $f \in \big(N_n(t)+1, N_n(t)+a_n(t)\big)$, because $\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)\Big] = 0$ if file f has no packets in the Transport layer.

Let $\mathcal{W}^r_{nf}(t)$ be the remaining window size of file f at node n after MAC-layer departure but before the MAC-layer injection. We want to show that, for any $w \geq 0$,

$$\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)\Big|\mathcal{W}^r_{nf}(t) = w\Big] = 0, \tag{5.33}$$

then (5.33) implies $\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)\Big] = 0$. Because the number of remaining packets at the Transport layer at each time is geometrically distributed with mean size $\sigma_{nf}(t)$, the Transport layer will continue to inject packets into the MAC layer with probability $\varsigma_{nf}(t) = 1 - 1/\sigma_{nf}(t) = 1 - \eta_{nf}(t)$ as long as all previous packets are successfully injected and the window size is not full.

Clearly, if $w = 0$, no packet can be injected into the MAC layer. Therefore,

$\hat{A}_{nf}^{(d)}(t) = 0$ and $\hat{D}_{nf}^{(d)}(t) = 0$, and (5.33) is satisfied. Next, consider the case $w > 0$. Let

$$p_w(k,j) := \mathbb{P}\left(\hat{A}_{nf}^{(d)}(t) = k, I_{nf}(t) = j | \mathcal{W}_{nf}^r(t) = w\right),$$

for $j \in \{0,1\}$ and $k \geq 1$. For $k \leq w$, $p_w(k,1)$ directly follows the geometric distribution of the remaining packets of file f, i.e., for $1 \leq k \leq w$,

$$\begin{aligned} p_w(k,1) &= \mathbb{P}\left(\hat{A}_{nf}^{(d)}(t) = k | \mathcal{W}_{nf}^r(t) = w\right) \\ &= \varsigma_{nf}^{k-1}(t)(1 - \varsigma_{nf}(t)). \end{aligned}$$

Note that from the definition of $I_{nf}(t)$, we have

$$\mathbb{P}\left(I_{nf}(t) = 0 | \mathcal{W}_{nf}^r(t) = w\right) = 1 - \sum_{k=1}^{w} p_w(k,1) = \varsigma_{nf}^w(t).$$

Then, a simple calculation shows that

$$\begin{aligned} \mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t) | \mathcal{W}_{nf}^r(t) = w\Big] &= \sum_{k=1}^{w} p_w(k,1)\Big(k - \sigma_{nf}(t)\Big) \\ &\quad + \mathbb{P}\left(I_{nf}(t) = 0 | \mathcal{W}_{nf}^r(t) = w\right) w \\ &= \sum_{k=1}^{w} k\varsigma_{nf}^{k-1}(t)(1 - \varsigma_{nf}(t)) \\ &\quad -(1 - \varsigma_{nf}^w(t))\sigma_{nf}(t) + w\varsigma_{nf}^w(t) = 0, \end{aligned}$$

because $\varsigma_{nf}(t) = 1 - 1/\sigma_{nf}(t)$ by definition.

Part (ii)

Using the fact that new arriving files are mutually independent, and are also independent of current network state, we can write $\mathbb{E}_{\mathcal{S}(t)}\Big[B_n(t)^2\Big] =$ "G" + "H" with

$$\begin{aligned} \text{“G”} &:= \mathbb{E}_{\mathcal{S}(t)}\Big[\Big(\textstyle\sum_{f=1}^{N_n(t)} B_{nf}(t)\Big)^2\Big], \\ \text{“H”} &:= \mathbb{E}_{\mathcal{S}(t)}\Big[\textstyle\sum_{f=N_n(t)+1}^{N_n(t)+a_n(t)} B_{nf}(t)^2\Big], \end{aligned}$$

where we have also used the fact that $\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)\Big] = 0$. Note that $B_{nf}(t)^2 \leq \max\{\hat{A}_{nf}^{(d)}(t)^2, \hat{D}_{nf}^{(d)}(t)^2\}$. Since the congestion window size is bounded by $\mathcal{W}_{cong}$ and the mean file size is bounded by $1/\eta_{min}$, we get

$$\mathbb{E}_{\mathcal{S}(t)}\Big[B_{nf}(t)^2\Big] \leq \max\{\mathcal{W}_{cong}^2, 1/\eta_{min}^2\}.$$

Thus

$$\text{“H”} < \kappa_n \max\{\mathcal{W}_{cong}^2, 1/\eta_{min}^2\}.$$

Next, we bound the “G” term. Let $\mathcal{F}_n(t)$ denote the set of files at node n that are served at time t. Because $B_{nf}(t) = 0$ if the existing file is not served, we have

$$\begin{aligned}
\Big|\sum_{f=1}^{N_n(t)} B_{nf}(t)\Big| &\leq \max\Big\{\sum_{f\in\mathcal{F}_n(t)} \hat{A}_{nf}^{(d)}(t), \sum_{f\in\mathcal{F}_n(t)} \sigma_{nf}(t)\Big\} \\
&\leq \Big|\mathcal{F}_n(t)\Big| \cdot \max\Big\{\mathcal{W}_{cong}, 1/\eta_{min}\Big\}.
\end{aligned}$$

Note that $|\mathcal{F}_n(t)| \leq \sum_{j:(n,j)\in\mathcal{L}} x_{nj}(t) \leq Nr_{max}$ because the number of existing files that are served cannot exceed the sum of outgoing link capacities. Thus,

$$\text{“G”} \leq N^2 r_{max}^2 \max\Big\{\mathcal{W}_{cong}^2, 1/\eta_{min}^2\Big\}.$$

This completes the proof.

5.6.2 Proof of Lemma 5.2

Note that $u_n^{(d)}(t) = 0$ if $q_n^{(d)}(t) \geq Nr_{max}$, and $u_n^{(d)}(t) \leq Nr_{max}$ if $q_n^{(d)}(t) \leq Nr_{max}$. In the latter case, since the congestion window size for every file is at least one, there are at most Nr_{max} files in the Transport layer of node n intended for destination d. Hence, using the definition of $\bar{Q}_n^{(d)}(t)$, $\bar{Q}_n^{(d)}(t) \leq Q^0 := Nr_{max} + Nr_{max}/\eta_{min}$. So,

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}(t)}\Big[g(\bar{Q}_n^{(d)}(t))u_n^{(d)}(t)\Big] &= \mathbb{E}_{\mathcal{S}(t)}\Big[g(\bar{Q}_n^{(d)}(t))u_n^{(d)}(t)\mathbb{1}\left\{q_n^{(d)}(t) \leq Nr_{max}\right\}\Big] \\
&\leq \mathbb{E}_{\mathcal{S}(t)}\Big[g(\bar{Q}_n^{(d)}(t))Nr_{max}\mathbb{1}\left\{q_n^{(d)}(t) \leq Nr_{max}\right\}\Big] \\
&\leq Nr_{max}g(Q^0).
\end{aligned}$$

Therefore $C_2 = N^3 r_{max} g(N r_{max}(1 + 1/\eta_{min}))$.

5.6.3 Proof of Lemma 5.5

Note that

$$\frac{\pi_{t+1}(s)}{\pi_t(s)} = \frac{Z_t}{Z_{t+1}} \exp\Big(\sum_{(i,j)\in s} (\widetilde{w}_{ij}(t+1) - \widetilde{w}_{ij}(t))\Big),$$

where

$$\begin{aligned}
\frac{Z_t}{Z_{t+1}} &= \frac{\sum_{s\in\mathcal{M}} \exp(\sum_{(i,j)\in s} \widetilde{w}_{ij}(t))}{\sum_{s\in\mathcal{M}} \exp(\sum_{(i,j)\in s} \widetilde{w}_{ij}(t+1))} \\
&\leq \max_s \exp\Big(\sum_{(i,j)\in s} (\widetilde{w}_{ij}(t) - \widetilde{w}_{ij}(t+1))\Big) \\
&\leq \exp\Big(\sum_{(i,j)\in\mathcal{L}} (\widetilde{w}_{ij}(t) - \widetilde{w}_{ij}(t+1))\Big).
\end{aligned}$$

Let $q^*(t)$ denote $g^{-1}(g^*(t))$, and define $\widetilde{q}_i^{(d)}(t) := \max\{q^*(t), q_i^{(d)}(t)\}$. Then,

$$\begin{aligned}
\widetilde{w}_{ij}^{(d)}(t+1) - \widetilde{w}_{ij}^{(d)}(t) &= g(\widetilde{q}_i^{(d)}(t+1)) - g(\widetilde{q}_j^{(d)}(t+1)) - g(\widetilde{q}_i^{(d)}(t)) + g(\widetilde{q}_j^{(d)}(t)) \\
&= \Big[g(\widetilde{q}_i^{(d)}(t+1)) - g(\widetilde{q}_i^{(d)}(t))\Big] + \Big[g(\widetilde{q}_j^{(d)}(t)) - g(\widetilde{q}_j^{(d)}(t+1))\Big].
\end{aligned}$$

Recall that the link service rate is at most one and the congestion window sizes are at most $\mathcal{W}_{cong}$, thus $\forall i \in \mathcal{N}$, $\forall d \in \mathcal{D}$, $|\widetilde{q}_i^{(d)}(t+1) - \widetilde{q}_i^{(d)}(t)| \leq 1 + \mathcal{W}_{cong}$. Hence,

$$\begin{aligned}
\frac{|\widetilde{w}_{ij}^{(d)}(t+1) - \widetilde{w}_{ij}^{(d)}(t)|}{1 + \mathcal{W}_{cong}} &\leq g'(\widetilde{q}_i^{(d)}(t)) + g'(\widetilde{q}_j^{(d)}(t+1)) \\
&\leq 2g'(q^*(t+1) - 1 - \mathcal{W}_{cong}),
\end{aligned}$$

where we have also used the fact that g is a concave increasing function. Therefore,

$$\frac{\pi_{t+1}(s)}{\pi_t(s)} \leq e^{2(1+\mathcal{W}_{cong})|\mathcal{L}| g'(q^*(t+1)-1-\mathcal{W}_{cong})}.$$

A similar calculation shows that also

$$\frac{\pi_t(s)}{\pi_{t+1}(s)} \leq e^{2(1+\mathcal{W}_{cong})|\mathcal{L}|g'(q^*(t+1)-1-\mathcal{W}_{cong})}.$$

This concludes the proof.

5.6.4 Proof of Lemma 5.6

The second inequality immediately follows from definition of w_{ij}. To prove the first inequality, consider a destination d, with routing matrix $\mathbf{R}^{(d)} \in \{0,1\}^{N\times N}$, and let $\mathbf{w}^{(d)} = [w_{ij}^{(d)}(t) : R_{ij}^{(d)} = 1]$, then, based on (5.2), we have

$$\mathbf{w}^{(d)} = (\mathbf{I} - \mathbf{R}^{(d)})g(\mathbf{q}^{(d)}),$$

where $g(\mathbf{q}^{(d)}) = [g(q_i^{(d)}) : i \in \mathcal{N}]$. Note that every row of $\mathbf{R}^{(d)}$ has exactly one "1" entry except the row corresponding to d which is all zero, so $(\mathbf{R}^{(d)})^N = 0$. Therefore, $(\mathbf{I} - \mathbf{R}^{(d)})^{-1} = \mathbf{I} + \mathbf{R}^{(d)} + (\mathbf{R}^{(d)})^2 + \cdots$ exists and $\mathbf{I} - \mathbf{R}^{(d)}$ is nonsingular. So $g(\mathbf{q}^{(d)}) = (\mathbf{I} - \mathbf{R}^{(d)})^{-1}\mathbf{w}^{(d)}$. Let $\|\cdot\|_\infty$ denote the ∞-norm. Then we have

$$\begin{aligned} \|(\mathbf{I} - \mathbf{R}^{(d)})^{-1}\|_\infty &= \|\sum_{k=0}^{N}(\mathbf{R}^{(d)})^k\|_\infty \leq \sum_{k=0}^{N}\|(\mathbf{R}^{(d)})^k\|_\infty \\ &\leq \sum_{k=0}^{N}\|\mathbf{R}^{(d)}\|_\infty^k \leq N, \end{aligned}$$

where we have used the basic properties of the matrix norm, and the fact that $\|\mathbf{R}^{(d)}\|_\infty = 1$. Therefore,

$$\|g(\mathbf{q}^{(d)})\|_\infty \leq \|(\mathbf{I} - \mathbf{R}^{(d)})^{-1}\|_\infty\|\mathbf{w}^{(d)}\|_\infty \leq N\|\mathbf{w}^{(d)}\|_\infty,$$

for every $d \in \mathcal{D}$. Taking the maximum over all $d \in \mathcal{D}$, and noting that g is a strictly increasing function, yields the result.

5.6.5 Proof of Lemma 5.7

The $h(\cdot)$ is strictly increasing so $h(x) \geq 1$ for all $x \geq h^{-1}(1)$. So $g'(x) \leq \frac{1}{1+x}$ for $x \geq h^{-1}(1)$. The inverse of g cannot be expressed explicitly, however, it satisfies

$$g^{-1}(x) = \exp(xh(g^{-1}(x))) - 1. \tag{5.34}$$

Therefore,

$$\begin{aligned} \alpha_t &\leq \frac{2(1+\mathcal{W}_{cong})|\mathcal{L}|}{g^{-1}(g^*) - \mathcal{W}_{cong}} && (5.35) \\ &= \frac{2(1+\mathcal{W}_{cong})|\mathcal{L}|}{\exp(g^* h(g^{-1}(g^*))) - 1 - \mathcal{W}_{cong}}, && (5.36) \end{aligned}$$

for $g^* \geq g(1 + \mathcal{W}_{cong} + h^{-1}(1))$. Next, note that

$$\begin{aligned} T_{t+1} &\leq 16^{|\mathcal{L}|} e^{4|\mathcal{L}|(w_{max}+g^*)} \\ &\leq 16^{|\mathcal{L}|} e^{4|\mathcal{L}|(g(q_{max}) + \frac{\epsilon}{4|\mathcal{L}|N} g(q_{max}))} \\ &\leq 16^{|\mathcal{L}|} e^{8|\mathcal{L}|g(q_{max})}. && (5.37) \end{aligned}$$

Consider the product of (5.36) and (5.37) and let $K := 2(\mathcal{W}_{cong} + 1)|\mathcal{L}|16^{|\mathcal{L}|}$. Using (5.34) and (5.22), the condition (5.27) is satisfied if

$$K e^{g^* [\frac{32|\mathcal{L}|N^3}{\epsilon} - h(g^{-1}(g^*))]} \left(1 + \frac{1 + \mathcal{W}_m}{g^{-1}(g^*) - \mathcal{W}_m}\right) \leq \delta/16. \tag{5.38}$$

Consider fixed, but arbitrary, $|\mathcal{L}|$, N and ϵ. As $q_{max} \to \infty$, $g(q_{max}) \to \infty$, and consequently $g^* \to \infty$ and $g^{-1}(g^*) \to \infty$. Therefore, the exponent $\frac{32|\mathcal{L}|N^3}{\epsilon} - h(g^{-1}(g^*))$ is negative for q_{max} large enough, and thus, there is a threshold q_{th} such that for all $q_{max} > q_{th}$, the condition (5.38) is satisfied.

The last step of the proof is to determine t^*. Let t_1 be the first time that $q_{max}(t)$ hits q_{th}, then

$$\begin{aligned} \sum_{k=t_1}^{t_1+t} \frac{1}{T_k^2} &\geq 16^{-2|\mathcal{L}|} \sum_{k=t_1}^{t_1+t} e^{-16|\mathcal{L}|g(q_{max}(t))} \\ &= 16^{-2|\mathcal{L}|} \sum_{k=t_1}^{t_1+t} (1 + q_{max}(t))^{-\frac{16|\mathcal{L}|}{h(q_{max}(t))}} \\ &\geq 16^{-2|\mathcal{L}|} t(1 + q_{th} + t)^{-\frac{16|\mathcal{L}|}{h(q_{th})}}, \end{aligned}$$

and

$$\begin{aligned}
\min_s \pi_{t_1}(s) &\geq \frac{1}{\sum_s \exp(\sum_{i\in s} \widetilde{w}_{ij}(t_1))} \\
&\geq \frac{1}{|\mathcal{M}| \exp(|\mathcal{L}|(w_{max}(t_1) + g^*(t_1)))} \\
&\geq \frac{1}{2^{N^2} \exp(2N^2 g(q_{th}))}.
\end{aligned}$$

Therefore, by Proposition 5.1, it suffices to find the smallest t that satisfies

$$16^{-2N^2} t(1+q_{th}+t)^{-\frac{16N^2}{g(q_{th})}} \geq \log(4/\delta) + N^2 \log(2(1+q_{th})),$$

for a threshold q_{th} large enough. Recall that $h(.)$ is an increasing function, therefore, by choosing q_{th} large enough, $\frac{16N^2}{h(q_{th})}$ can be made arbitrary small. Then a finite t^* always exists since $\lim_{t^*\to\infty} t^*(1+q_{th}+t^*)^{-\frac{16N^2}{h(q_{th})}} = \infty$.

Chapter 6

Conclusions and Open Problems

In this thesis, we first established that CSMA-like algorithms can achieve maximum throughput in any general network topology under a "near logarithmic growth condition" on the weights. In words, to achieve throughput optimality, it is sufficient for weights to be logarithmic functions of the queue lengths, divided by an arbitrarily slowly increasing, unbounded function. For example, weights of the form $\log^{1-\epsilon}(\cdot)$, with any $0 < \epsilon \ll 1$, are sufficient to ensure throughput optimality. This result indicates that the maximum-throughput guarantees are preserved for weight functions that are essentially logarithmic for all practical queue lengths, although asymptotically weights must grow slower than any logarithmic function of the queue length.

We further demonstrated that the "near-logarithmic growth condition" is indeed tight, in the sense that weights that grow faster than $\gamma \log(\cdot)$, for any $\gamma > 0$ will cause instability in some network topology. Thus our stability and instability results imply that "the near-logarithmic growth condition" on the weights is a fundamental limit on the aggressiveness of nodes to ensure maximum stability (throughput optimality) in any general topology.

Finally, we showed that the above maximum-stability results can be easily extended to multihop networks with dynamic flows by using very general congestion control mechanisms.

In this thesis, we have enforced the requirement that the CSMA algorithm must provide maximum throughput in any arbitrary topology with possibly very large number of nodes, i.e., the algorithm is robust against worst case scenarios. The "near-logarithmic growth condition" on the weights could be potentially relaxed if one imposes some restrictions on the class of topologies or could obtain some information about the topology. For example, in complete graphs, and more generally in complete partite networks with any number of components, the "near-logarithmic growth" *is not* a restriction and nodes could use weight functions that are arbitrarily more aggressive. It is also possible that one can use more aggressive weight functions in bounded degree graphs. Moreover, it is plausible that the "near-logarithmic growth condition" on weights can be mitigated if the network operator is not interested in the maximum stability region and he/she is willing to sacrifices a fraction of the capacity region to obtain better delay performance. In general, we are still far from a comprehensive characterization of the throughput-delay-complexity tradeoff and much more research is needed to design better algorithms and to reduce the gap.

For flow-level stability, our constraining assumptions regarding the congestion control mechanisms are very mild and compatible with the standard implementations like TCP. We did not discuss what type of window dynamics to use to achieve a certain QoS metric. It would be interesting to investigate how the congestion window dynamics and the links weights impact the system QoS performance for wireless networks. It is observed in [56] in the context of multiclass queueing systems that a fixed congestion window size implicitly solves an optimization problem in an asymptotic regime. Our simulation results in [55] show that log-differential link weights, with a fixed congestion window size, reduce the file transfer delays. It will be certainly interesting to establish the validity of such an observation rigorously as a future research.

References

[1] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximal throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, pp. 1936–1948, 1992.

[2] R. R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin, "Throughput analysis in multihop CSMA packet radio networks," *IEEE Transactions on Communications*, vol. 35, pp. 267–274, 1987.

[3] M. Durvy, O. Dousse, and P. Thiran, "Modeling the 802.11 protocol under different capture and sensing capabilities," in *Proc. Infocom 2007 Conf.*, 2007.

[4] M. Durvy and P. Thiran, "A packing approach to compare slotted and non-slotted medium access control," in *Proc. Infocom Conf.*, 2006.

[5] M. Garetto, T. Salonidis, and E. W. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 16, pp. 864–877, 2008.

[6] X. Wang and K. Kar, "Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks," in *Proc. IEEE Infocom Conference*, vol. 1, 2005, pp. 23–34.

[7] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, 2000.

[8] S. C. Liew, C. H. Kai, J. Leung, and B. Wong, "Back-of-the-envelope computation of throughput distributions in CSMA wireless networks," *IEEE Trans. Mob. Comp.*, vol. 9, pp. 319–1331, 2010.

[9] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *IEEE/ACM Trans. Networking*, vol. 18, pp. 960–972, 2010.

[10] J. Liu, Y. Yi, and A. Proutiere, "Towards utility-optimal random access without message passing," *Wireless Commun. Mobile Comput.*, vol. 10, pp. 115–128, 2010.

[11] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," in *Proc. Allerton Conf.*, 2008.

[12] P. Marbach and A. Eryilmaz, "A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks," in *Proc. Allerton Conf.*, 2008.

[13] S. Rajagopalan, D. Shah, and J. Shin, "Network adiabatic theorem: An efficient randomized protocol for content resolution," in *Proc. ACM SIGMETRICS/Performance Conf.*, 2009.

[14] D. Shah and J. Shin, "Randomized scheduling algorithms for queueing networks," *Ann. Appl. Prob*, vol. 22, pp. 28–171, 2012.

[15] D. Shah, J. Shin, and P. Tetali, "Medium access using queues," in *Proc. FOCS Conf.*, 2011.

[16] L. Jiang, D. Shah, J. Shin, and J. Walrand, "Distributed random access algorithm: Scheduling and congestion control," *IEEE Trans. Inf. Theory*, vol. 56, pp. 6182–6207, 2010.

[17] J. Ghaderi and R. Srikant, "The impact of access probabilities on the delay performance of Q-CSMA algorithms in wireless networks," *IEEE/ACM Trans. Networking*, 2012, to appear.

[18] C. H. Lee, D. Y. Eun, S. Y. Yun, and Y. Yi, "From Glauber dynamics to metropolis algorithm: Smaller delay in optimal CSMA," in *Proc. ISIT Conf.*, 2012.

[19] M. Lotfinezhad and P. Marbach, "Throughput-optimal random access with order-optimal delay," in *Proc. Infocom Conf.*, 2011.

[20] J. Ni, B. Tan, and R. Srikant, "Q-CSMA: Queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks," in *Proc. Infocom Mini-Conf.*, 2010.

[21] D. Shah and J. Shin, "Delay-optimal queue-based CSMA," in *Proc. ACM SIGMETRICS Conf.*, 2010.

[22] P.-K. Huang and X. Lin, "Improving the delay performance of CSMA algorithms: A virtual multi-channel approach," in *Proc. Infocom Conf.*, 2013.

[23] X. Lin, N. Shroff, and R. Srikant, "On the connection-level stability of congestion-controlled communication networks," *IEEE Transactions on Information Theory*, vol. 54, pp. 2317–2338, 2008.

[24] J. Liu, A. Proutiere, Y. Yi, M. Chiang, and V. Poor, "Flow-level stability of data networks with non-convex and time-varying rate regions," in *Proc. ACM SIGMETRICS*, 2007.

[25] C. Moallemi and D. Shah, "On the flow-level dynamics of a packet-switched network," in *Proc. ACM SIGMETRICS*, 2010.

[26] A. Eryilmaz, R. Srikant, and J. R. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Transactions on Networking*, vol. 13, pp. 411–424, 2005.

[27] J. Ghaderi and R. Srikant, "On the design of efficient CSMA algorithms for wireless networks," in *Proc. CDC Conf.*, 2010.

[28] M. Feuillet, A. Proutiere, and P. Robert, "Random capture algorithms: Fluid limits and stability," in *Proc. ITA Workshop*, 2010.

[29] P. Bremaud, *Markov chains, Gibbs fields, Monte Carlo simulation, and queues.* Springer-Verlag, 1999.

[30] S. P. Meyn and R. L. Tweedie, "Criteria for stability of Markovian processes I: Discrete time chains," *Advances in Applied Probability*, vol. 24, pp. 542–574, 1992.

[31] M. Jerrum and A. Sinclair, "Approximating the permanent," *SIAM Journal of Computing*, vol. 18, pp. 1149–1178, 1989.

[32] J. Dai, "On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models," *Ann. Appl. Prob.*, vol. 5, pp. 49–77, 1995.

[33] J. G. Dai, "A fluid limit model criterion for instability of multiclass queueing networks," *Ann. Appl. Prob.*, vol. 6, pp. 751–757, 1996.

[34] J. G. Dai and S. P. Meyn, "Stability and convergence of moments for multiclass queueing networks via fluid limit models," *IEEE Trans. Aut.Control*, vol. 40, pp. 1889–1904, 1995.

[35] S. P. Meyn, "Transience of multiclass queueing networks via fluid limit models," *Ann. Appl. Prob.*, vol. 5, pp. 946–957, 1995.

[36] N. Bouman, S. C. Borst, J. S. H. van Leeuwaarden, and A. Proutiere, "Backlog-based random access in wireless networks: Fluid limits and delay issues," in *Proc. ITC Conf.*, 2011.

[37] S. G. Foss and A. P. Kovalevskii, "A stability criterion via fluid limits and its application to a polling model," *Queueing Systems*, vol. 32, pp. 131–168, 1999.

[38] M. Frolkova, S. G. Foss, and B. Zwart, "Random fluid limit of an overloaded polling model," 2013, Submitted for publication.

[39] A. P. Kovalevskii, V. A. Topchii, and S. G. Foss, "On the stability of a queueing system with uncountable branching fluid limits," *Prob. Inf. Trans.*, vol. 41, pp. 254–279, 2005.

[40] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand, "Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling," in *Proc. Infocom Mini-Conf.*, 2010.

[41] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis, "Hardness of low delay network scheduling," *IEEE Trans. Inf. Theory*, vol. 57, pp. 7810–7817, 2011.

[42] V. Subramanian and M. Alanyali, "Delay performance of CSMA in networks with bounded degree conflict graphs." in *Proc. ISIT Conf.*, 2011.

[43] P. Robert and A. Veber, "On the fluid limits of a resource sharing algorithm with logarithmic weights," *arXiv:1211.5968v1*, 2012.

[44] L. C. G. Rogers and D. Williams, *Diffusions, Markov Processes and Martingales II.* Wiley, London, 1989.

[45] P. Billingsley, *Weak Convergence of Probability Measures.* Wiley-Interscience, 1986.

[46] W. Whitt, "Weak convergence of probability measures on the function space $C[0,\infty)$." Stanford University, Tech. Rep., 1969.

[47] H. L. Royden, *Real Analysis.* Prentice Hall, 3rd edition, 1987.

[48] N. Ethier and T. Kurtz, *Markov Processes: Characterization and Convergence.* Wiley-Interscience, 1985.

[49] P. Billingsley, *An Introduction to Probability and Measure.* Wiley-Interscience, 1995.

[50] J. Ghaderi, S. C. Borst, and P. A. Whiting, "Queue-based random-access algorithms: Fluid limits and stability issues," *Stochastic Systems*, 2013, Submitted for publication.

[51] T. Bonald and M. Feuillet, "On the stability of flow-aware CSMA," *Performance Evaluation*, vol. 67, pp. 1219–1229, 2010.

[52] M. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 835–846, 1997.

[53] M. J. Neely, E. Modiano, and C. E. Rohrs, "Dynamic power allocation and routing for time varying wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 89–103, 2005.

[54] S. Asmussen, *Applied Probability and Queues.* Springer, 2003.

[55] J. Ghaderi, T. Ji, and R. Srikant, "Connection-level scheduling in wireless networks using only MAC-layer information," in *Proc. IEEE Infocom Mini-Conf.*, 2012.

[56] N. S. Walton, "Utility optimization in congested queueing networks," *Applied Probability*, vol. 48, pp. 68–89, 2011.