Part II

Achieving Fair Rates with Ingress Policing

Chapter 4

Introduction

We consider a queuing network of d single-server stations carrying g flows. The stations are equipped with per-flow queues, and the network uses an ingress policing mechanism in which a flow's packets are discarded at the network ingress whenever any of that flow's queues exceed a threshold. Each flow f has a weight w_f, and each station serves a flow in proportion to its weight using a weighted round robin or similar queueing discipline like weighted fair queueing or generalized head of line processor sharing. Each flow's packets arrive according to an independent renewal process of some mean rate, and the service times at each station are also independent. Without discarding at the ingress, the utilization of some stations may exceed 1, and therefore the policing mechanism is necessary for the stability of the network.

Our main result is to show that there exist large enough policing thresholds so that each flow's long term average throughput can be made arbitrarily close to the weighted max-min fair allocation. Our result requires that there be unique bottlenecks for each flow, but this can always be achieved by perturbing the flow weight and/or the service rates of the stations. The other key contribution of this work is to show how the analysis of such a network's flow rates can be reduced to analyzing the rates of a corresponding fluid model.

4.1 Motivation

Achieving fairness between flows is one of the principal design challenges of any packet switched communication network. In today's Internet architecture, the Transmission Control Protocol (TCP) bears most of the burden in achieving fairness [21, 22, 23]. TCP is an end-to-end congestion protocol that infers congestion when packets are lost and responds to that congestion with a mechanism known as additive increase multiplicative decrease. That mechanism halves the rate of user's session when congestion is perceived (multiplicative decrease), and slowly increases the rate when it is not (additive increase). TCP has been tremendously successful in coping with congestion and in particular in avoiding a congestion collapse as the Internet has grown in size and speed by orders of magnitude over the years. However, as powerful as TCP is, TCP cannot achieve fairness if the underlying network is grossly unfair, or if a large proportion of the traffic does not use a transport protocol that responds to congestion as TCP does. For example, if the underlying network constrains two groups of N sessions to very different rates, where perhaps each group of sessions originate in a different geographic area, then TCP might succeed in achieving fairness between sessions of a group, but not in achieving fairness between the groups.

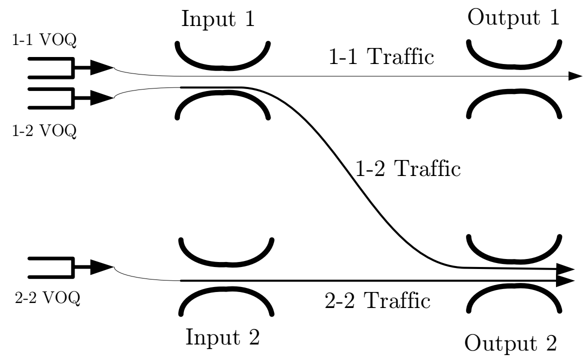

Figure 4.1: A two input, two output switch with virtual output queues.

Indeed designing a network that is not unfair to particular groups of sessions is becoming increasingly difficult. The speed of the optical links which make up much of the Internet has grown tremendously over the past 10 years, faster even then the pace of progress in the IC technology used to build the switches and routers that glue the links together [24]. This has pressured designers of switches, routers, and aggregation points to move to decentralized architectures that switch and process many packets in parallel. The architects of these new designs have been primarily concerned with maximizing throughput, but unfortunately throughput can come at the price of fairness.

We consider an example of a two input and two output switch to demonstrate the tradeoff between throughput and fairness. Suppose all input and output links of the switch have capacity 1, and label the input and output interfaces In-1, In-2, and Out-1, Out-2 respectively. Also suppose that groups of N TCP sessions are sending traffic between each of the following pairs: In-1 to Out-1 (1-1 for short); In-2 to Out-2; and In-1 to Out-2. This situation is depicted in Figure 4.1. In this situation, a switch that favors the 1-1 and 2-2 groups can achieve a higher throughput. In the extreme case the switch can offer the “unfair” allocation of 1 to groups 1-1 and 2-2, and 0 to the group 1-2 to achieve a throughput of 2. In contrast, a more fair allocation would be to offer 0.5 to all three groups, yielding a throughput of only 1.5. While the unfair allocation scenario is extreme, it is possible that a switch with a reasonable design might tend toward this extreme case.

Consider the common switch architecture in which each input interface splits its traffic into fixed length cells, and then sorts the traffic into “virtual” output queues corresponding to the traffic’s final destination [25]. Then in every time slot corresponding to the transmission time of one cell, a scheduler matches pairs of virtual queues to corresponding output interfaces. In our scenario, a scheduler optimized for throughput would favor matching 1-1 and 2-2 in each slot so as to move two cells, and not just one as would happen for a 1-2 match.

The end-to-end TCP protocol would not correct the problem. The 1-2 sessions would see congestion because their virtual queue is not being served fast and they would adapt their rates by slowing down. Consequently, the 1-1 and 2-2 sessions would find less competition on the interfaces they share with 1-2 traffic, adapt their rates higher allowing them to get even more of their traffic to their virtual queues, where that traffic will again be favored by the scheduler over the 1-2 traffic in the 1-2 virtual output queue.

Examples like this show that it is important to design switches, routers and aggregation points that treat traffic from different input output interface pairs fairly, even with the presence of rate adaptive transport protocols like TCP.

At the same time, the pressure to design faster and faster switch architectures is making it less feasible to design switches with central schedulers that coordinate packet movements at the sum of all the incoming line rates. As switch designers are pressured into eliminating the central scheduler entirely, the design of a switch is becoming more and more like the design of a network, where the “nodes” of the network are the interfaces and the interconnections between interfaces – the switch fabric – are the “links.”

This trend toward decentralized architectures, combined with the importance of fairness in switches, are transforming the switch fabric design problem into a problem of designing a queueing network that achieves fair rates with an extremely simple control policy. The control policy must be simple so that it can be implemented in silicon at the extremely fast line rates of today’s networks.

Surprisingly, while switch design is becoming more and more like network design, there are pressures to make the network look more and more like a switch. In particular, Mckeown and Zhang-Shen [27] propose using the Clos topology, popular in switch fabric design, as the basis of a next generation backbone for the proposed “100×100 network” – a network envisioned to interconnect 100 million users with an access rate of 100 Mbps [26]. The Clos topology consists of some number of input nodes which are fully interconnected with some number of middle nodes, which in turn are fully interconnected with some number of output nodes [3]. The Clos topology is popular in switch design because when combined with a simple load balancing algorithm, it can achieve full throughput between inputs and outputs for any feasible demand rate matrix – meaning any demand rates which do not exceed the capacity of any output. This feature would make it useful in the 100×100 network backbone, because it would automatically achieve full throughput without a complicated routing algorithm. However, the same issues that apply for a switch also apply for a backbone designed to look like a switch – fairness is important, and is not automatic in a throughout maximizing architecture. Therefore a methodology for designing a queuing network that achieves fairness with a simple control law would be extremely useful in the context of networks like the 100×100.

The policing scheme that we analyze in this work is extremely simple, achieves fair rates as we show, and therefore would be useful in a network like 100×100 or in a switch design. Yet the analysis of such a simple scheme could become extremely complex and involve detailed case by case analysis of sequences of discrete events in order to prove that the scheme achieves fair rates.

Therefore we feel that the principal contribution of this work is not the scheme itself, but rather the methodology we use to analyze it. We show that the analysis of the discrete stochastic system can be reduced to analysis of a fluid model obeying deterministic differential equations. Though we present the methodology in the context of our policing scheme, the same techniques could be used to analyze other simple rate control schemes for queueing networks.

While the primary application area that motivates this work is communication networks, there are other notable application areas. For example it might be necessary to control the flows of jobs in a manufacturing system, where perhaps the flows correspond to different products or different customers, and the stations correspond to the capacity constrained pieces of equipment – a lithography stepper at a semiconductor manufacturing plant for example. In such a context, it would be very desirable to have a control scheme that ensures that the different flows of jobs get fair rates at the bottleneck stations and at the same time ensure that the queues at each station remain bounded. In our scheme, the queues would be kept bounded by blocking new jobs from entering the manufacturing plant when a bottleneck station is congested, but other types of blocking schemes could be analyzed with the methodology presented in this work.

Another natural application area is for call centers. In a call center, customers are queued until a server can take the customer's call. Oftentimes the first server may not be able to fulfill all of the customer's needs and therefore the server may need to transfer the call to a second server, where the customer may have to enter yet another queue before seeing the second server. In turn, the second server might need to transfer the call again to a third server and so on. Also, the same call center maybe answering calls of different types. One example would be a contracted call center that answers calls for several different corporate clients simultaneously. Here the set of incoming calls for a particular client would be the "flow," and a call center would want to treat the different flows fairly in times of congestion. Furthermore, it might be better for the call center to block calls at the ingress, rather than having a customer put on hold, transferred, and put on hold again, only eventually to be dropped. Therefore a flow-control scheme like the one we study in this work is a logical choice for such an application.

4.2 Related Work

As we have stated, our objective is to show that our flow control scheme with large enough thresholds achieves rates close to the max-min fair allocation. A max-min fair flow allocation is a feasible allocation such that one flow's rate cannot be increased without decreasing the rate of another flow that already has a smaller allocation [28]. There have been a number of schemes proposed and analyzed that aim to achieve max-min fairness, and most such investigations have involved detailed case by case analysis of the possible state trajectories of the network. For example, Hahne shows that a network with per-flow queued stations using round-robin service, and with hop by hop window flow control achieves max-min fairness [29]. Round-robin service is the now familiar queuing discipline in which the server polls the queues in cyclic order, offering service to queues that are nonempty [30, 31]. We adopt a qualitatively different approach to our analysis. We consider a fluid limit of the network obtained by scaling time and space (including the thresholds) by a factor n that increases – an approach similar to that found in [32, 33, 34] and others. This allows us to exploit the Strong Law of Large Numbers to show that the arrival and service process converge to smooth, linear growths and also that the other processes describing the system converge to a trajectory of a fluid model described by simple differential equations.

In addition to the difference in methodology from previous investigations, our network and flow control model is also different than those studied in previous work. The model that most closely resembles ours is the network analyzed in [29]. However, that network differs from our network in that the flow control is hop by hop rather than done at the ingress, and the arrival and service times assumptions are more restrictive than the assumption that they are renewal as we assume. In particular [29] assumes arrivals are either deterministic or Bernoulli, and service times are deterministic.

Another related work is by Hayden, who proposes an iterative mechanism for finding max-min fair rates. In the mechanism, each station measures the rates of the flows passing through it, computes a fair rate allocation for these flows, and then explicitly tells sources what rates to use [35]. Gafni also proposes an improvement to Hayden's original algorithm [36]. The analysis of our fluid model resembles Hayden's proof in that, as in [35], we step through the stations of the network, at each step finding the most congested station, assign rates to the flows passing through it, subtract those flows from the network, and proceed to find the next most congested station in the reduced network. However, our problem requires extra analysis to show that within each of these steps the fluid model, described by a set of ODEs, converges. Our proof that the stochastic network converges to a fluid model is not related to [35].

Our fluid limit proof technique is closely related to known results in the literature. In particular, we borrow heavily from work by Dai [32]. Dai shows that for networks without policing, stability of the fluid model implies stability of the stochastic network – stability meaning the system is positive Harris recurrent. Our work differs in that we use the fluid model not only to analyze stability, but also to find the flow rates of the stochastic network. Also, the fact that the policing modifies the arrival processes also requires specialized treatment to show that our network converges to a fluid limit trajectory.

There is also a difference how the fluid limit is taken. Dai's proof considers a sequence of initial conditions x, with $|x| \to \infty$, and then obtains a fluid limit by scaling time and space by $|x|$. Dai uses this result to construct a Lyapunov function that shows that expected state of the system contracts, for initial states far enough from the origin. In contrast our proof requires consideration of two different types of fluid limits. In the first we consider a sequence indexed by n, and then obtain a fluid limit by scaling the policing thresholds, as well as space and time by $n \to \infty$. Analysis of this fluid limit leads us to conclude that the rates of the stochastic system are close to that of the fluid model, over a finite time interval, if two conditions are met: the stochastic system's thresholds n are large enough, and its *relative* initial condition $y = x/n - h\mathbf{e}$ is small, where $h\mathbf{e}$ is the equilibrium of the fluid model. In the second fluid limit we consider a sequence of pairs $\{(n, y)\}$, scale the policing thresholds by n, while scaling space and time by $n|y| \to \infty$. As in [32], we can construct a Lyapunov function that shows that the expected state of the system contracts for $n|y|$ large enough. Consequently the contraction property holds for smaller and smaller y as n is increased. The combination of both fluid limit results shows that for large enough n the invariant distribution of the system is concentrated in a region of the state space where the expected rates are close to those of the fluid model, and in turn the fluid model rates can be shown to be equal to a max-min allocation.

Another closely related work is by Mandelbaum, Massey, and Reiman [34]. In [34], the authors study the fluid limit of a queueing network with state dependent routing, where the function describing the arrivals to each queue can scale with n and or $\sqrt{n}$, in a manner similar to the scaling of our thresholds. The authors prove a functional strong law of large numbers and a functional central limited theorem in the context of their model. However, the authors assume that the network is driven by poisson processes, rather than just the renewal assumption that we make. An earlier work by Konstantopoulos, Papadakis, and Walrand derives a functional strong law of large numbers and a functional central limit theorem for networks with state dependent service rates [37].

There are also several works that use reflected Brownian motion models to study queueing networks with blocking [38, 39, 40]. Typically the objective of most such investigations is to approximate the distribution of the queue occupancy with a diffusion approximation, whereas in this work our objective is to show almost sure convergence using a strong law of large numbers scaling.

Chapter 5

Model

We state precisely our network model in this section, introducing the necessary notation along the way. Much of our notation is borrowed from [32].

5.1 Flows and Classes

As mentioned, we consider a queueing network of d single server stations carrying g flows. In addition to the notion of flow, each packet also has a class $k \in \{1, ..., K\}$ that is indicative of both the packet's flow and location in the network. Thus the class of a flow f packet changes as the packet progresses from station to station, but with the restriction that a flow f packet must always have a class in the set $K(f)$. Conversely, each class k is associated with one and only one flow, and we denote the class to flow mapping by $\mathrm{f}_{(k)} = \{f : k \in K(f)\}$. We also adopt the numbering convention that flow f packets enter the network as class $k = f$, and thus $f \in K(f)$.

Each class k is served by a unique station denoted $s(k)$. When a packet of class

k finishes service at station $s(k)$, it becomes a packet of class l with probability P_{kl} or exists the network with probability $1 - \sum_l P_{kl}$. To preserve a packet's flow, $P_{kl} = 0$, if $\mathrm{f}_{(k)} \neq \mathrm{f}_{(l)}$.

More precisely, the jth packet of class k to complete service has associated with it a K dimensional "Bernoulli" random column vector $\phi^k(j)$ which consists of all zeros with probability $1 - \sum_l P_{kl}$, or has a single 1 at location l with probability P_{kl}. The random vectors $\phi^k(j)$ are independent.

In Chapter 8, we will consider networks for which each flow's packets follow loop-free paths without "splitting." The matrix P is binary and nilpotent for such networks.

Because each class k is served at a unique station $s(k)$, we may define the *Class Constituency Matrix* C with $C_{ik} = 1$ if $s(k) = i$ and $C_{ik} = 0$ otherwise. At times it is more convenient to represent the same information in set rather than matrix notation, and therefore we define the *Class-Constituency* $\mathbf{C}[i] = \{k : s(k) = i\}$ for each station i. Finally, because each class belongs to a unique flow $\mathrm{f}_{(k)}$, we may define the *Flow-Constituency* $\mathbf{F}[i] = \{f : \exists k \in \mathbf{C}[i] : f = \mathrm{f}_{(k)}\}$ for each station i.

5.2 Queues

The state of the network's queues is the K dimensional column vector Q where Q_k is the number of packets of class k packets queued in the system.

The exogenous arrivals to the network for flow f are described by a renewal process $E_f(t)$ where the inter-arrival times $\{\xi_f(j), j \geq 1\}$ are i.i.d. and have mean α_f^{-1}

where α_f is the mean arrival rate. We denote the remaining waiting time for the next flow f customer arrival at time t as $U_f(t)$. Thus the process $E_f(t)$ is described by

$$E_f(t) \triangleq \max\{j : U_f(0) + \xi_f(1) + ... + \xi_f(j-1) \leq t\}.$$

The service times of each class $\{\eta_k(j), j \geq 1\}$ are also i.i.d. and have mean $m_k = \mu_k^{-1}$, where μ_k is the mean service rate. We also define the $K \times K$ diagonal matrix M where the kth entry on the diagonal is m_k. $V_k(t)$ denotes the remaining service time of the class k customer in service, if there is one at time t, otherwise $V_k(t) = 0$. We define a service process $S_k(t)$ as

$$S_k(t) \triangleq \max\{j : \tilde{V}_k(0) + \eta_k(1) + ... + \eta_k(j-1) \leq t\}$$

where $\tilde{V}_k(0) = V_k(0)$ if $V_k(0) > 0$, otherwise $\tilde{V}_k(0) = \eta_k(0)$ is a fresh service time with the same distribution as $\eta_k(1)$ and independent of all other service times.

In principle, our assumption that the service times are independent does not allow for service times that depend on a packet's size. Dependence on packet size would make the service times of stations dependent on each other. To model this explicitly would require a much more complicated model. However we feel that our results this work would still hold if this assumption were relaxed.

5.3 Policing Points

Arriving packets of each flow f first pass through a per-flow policing point at the ingress of the network. Whenever any queue in the *Control-Set*, $\mathcal{C}(f)$, exceeds a high threshold h_u, the policing point drops flow f packets as they arrive. Conversely, when all of the queues in $\mathcal{C}(f)$ are below a lower threshold h_l, flow f packets are permitted to enter the network. When the queues in $\mathcal{C}(f)$ are between the two thresholds, the policing follows a hysteresis law that we will make precise below. Also note that typically $\mathcal{C}(f) = K(f)$, the set of classes that flow f packets through, but in general, $\mathcal{C}(f) \subseteq K(f)$.

To describe the policing hysteresis, we define a binary-policing state $H_k(t)$ for each class k as

$$H_k(t) = \begin{cases} 1 & \text{if } H_k(t-) = 0 \text{ and } Q_k(t) \geq h_u \\ 0 & \text{if } H_k(t-) = 1 \text{ and } Q_k(t) \leq h_l \\ H_k(t-) & \text{otherwise.} \end{cases}$$

We allow h_l to equal h_u if we choose not to have hysteresis. We define the process $\Lambda_f(t)$, which we call the thinned arrival process, to count the flow f packets that are allowed beyond the policing point at the ingress. Therefore,

$$\Lambda_f(t) = \sum_{j=1}^{E_f(t)} \prod_{\mathcal{C}(f)} (1 - H_k(\tau_j)) \tag{5.1}$$

where $\tau_j = U_f(0) + \sum_{m=1}^{j-1} \xi_f(m)$ is the time of the jth arrival to the policing point.

5.4 A Simple Example

After having introduced the general model, it is useful to consider an example of such a network. Our example is illustrated by Figure 5.1. The example network consists of 4 stations, and carries 3 flows. Flow 1 passes only through station 1, where it is queued as class 1. Flow 2 passes through stations 1, 2, and 3, and where it is queued as class 2, 4, and 5 respectively. Flow 3 passes through stations 2 and 4, where it is queued as class 3,

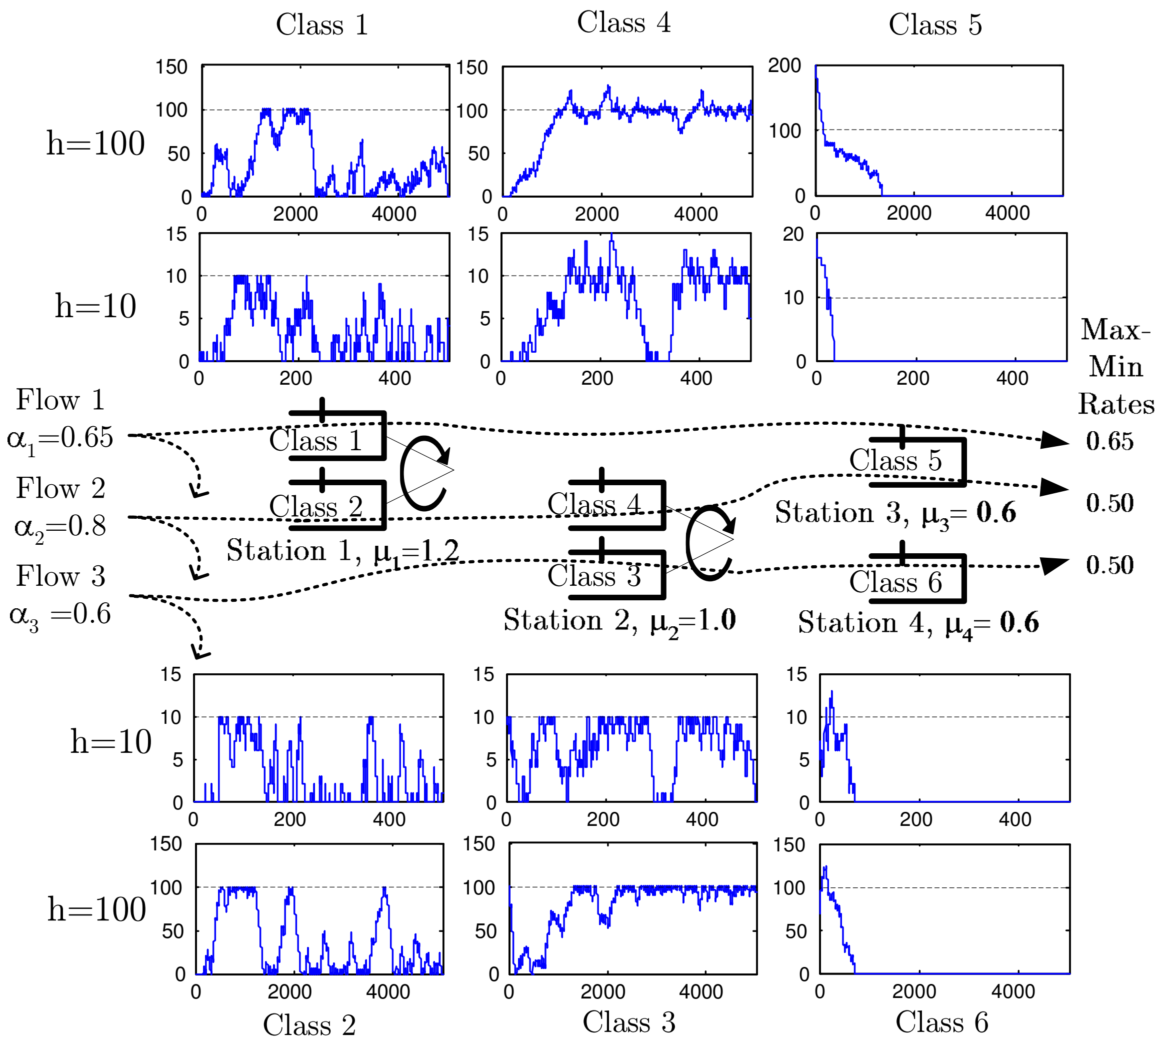

Max-Min Rates	Avg. Throughput h=10	Avg. Throughput h=100
0.65	0.6425	0.6180
0.50	0.4025	0.4740
0.50	0.5175	0.4796

Figure 5.1: Simulated trajectories of an example network, for two choices of policing threshold: $h = 10$ and $h = 100$.

and 6 respectively. All three flows have a weight $w_f = 1$. The station service rates are: $\mu_1 = 1.2$, $\mu_2 = 1.0$, $\mu_3 = \mu_4 = 0.6$. The service times are exponential, and the queueing discipline is round-robin. The exogenous flow arrival rates are $\alpha_1 = 0.65$, $\alpha_2 = 0.8$, and $\alpha_3 = 0.6$. The inter-arrival time distributions are chosen to have "heavy tails." In particular, they are Pareto distributed, with complementary cumulative distribution function given by

$$\mathrm{P}(\xi_f(j) > s) = \frac{1}{(\alpha_f s + 1)^2}$$

for each flow f and for $s \geq 0$. Thus the inter-arrival times have mean $1/\alpha_f$ and infinite variance.

Figure 5.1 shows simulated trajectories for two choices of thresholds. The inside plots, the 3 plots immediately above and below the network drawing, correspond to a threshold of $h = 10$, while the outside plots are for $h = 100$. In this example we have no hysteresis and thus $h_u = h_l = h$.

We begin by considering the $h = 10$ case. In this case, the queues start from an initial condition $\bar{Q}_2(0) = 10$, $\bar{Q}_5(0) = 20$, $\bar{Q}_6(0) = 7$, and the other queues start from 0.

The natural way to start analyzing such an example is to find the station that is the most constrictive bottleneck. In this case station 2 has a capacity of 1.0, carries two flows, and thus has a capacity of 0.5 per flow, which is less than the capacity per flow of the other stations, and less than the arrival rate of any of the flows. Thus one should expect station 2 to be the governing constraint, eventually having its queues fill up to threshold and then applying policing to throttle flows 2 and 3 to rates of around

0.5 each. Once this occurs, we can treat flows 2 and 3 as constant rate flows, "subtract" them from the network by deducting their final rates from the capacities of the other stations they cross, and then analyze the flows of the reduced network. By doing this, we find that station 1 has a remaining capacity of 0.7, but since flow 1 arrives at a slower rate, flow 1's final rate is "demand" limited to 0.65. By doing this analysis, we can conclude that the max-min fair rates for this example are $(0.65, 0.5, 0.5)$.

Indeed, Figure 5.1 shows that the queues at station 2 fill to their thresholds and tend to stay near them. However, the class 2 queue empties and remains close to empty between times 300 and 350. Because flow 2's queue at station 2 is "starved" during this epoch, flow 2 packets miss opportunities to be served at station 2, which is the most constrictive bottleneck of the network. Indeed the table at the bottom of Figure 5.1 shows the average rate, averaged over the last 80% of the simulation time to reduce some of the initial transient effect, is 0.4025. This is substantially below the max-min fair rate of 0.5. Most likely, a string of long inter-arrival times of flow f, caused the flow 2 queues at station 2 to starve.

We can make a more general observation that the fluctuations of the inter-arrival and service times can cause queues that we expect to remain near their thresholds to starve occasionally. If this starvation happens non-negligibly often over the long term, the long term rates will differ from the ideal max-min fair share rates.

To prevent or at least reduce starvation of queues at bottleneck stations, the natural thing to do is to raise the thresholds. This would ensure that bottleneck queues have enough backlog to smooth over fluctuations in the arrival process and the service processes of upstream queues. To test that intuition, we simulate the network with policing thresholds of $h = 100$. As we have scaled the thresholds by 10, we also scale the initial condition by 10, and simulate for 10 times longer– this decision will make this example useful to refer back to later when we talk about the fluid scaling. Figure 5.1 shows that the queues at station 2 do not starve after the first 20% of the simulation run, as they did in the $h = 10$ case. The table at the bottom of Figure 5.1 shows the average throughput of the flows over the last 80% of the simulation run for both the $h = 10$ and $h = 100$ cases. Flow 2's average throughput is indeed better in the $h = 100$ case than in the $h = 10$ case, suggesting that the reduction in starving has helped. However, the average throughput of the other two flows actually declined a little.

This example does offer some support to our hypothesis that increasing thresholds makes long term average rates approach max-min fair share. Yet, it shows that it is far from obvious that the long term average rates can be made close to the max-min rates with large enough thresholds. Consequently, this example has bolstered our motivation to validate the hypothesis analytically.

5.5 Threshold Scaling

Because we want to show that there exist large enough thresholds to achieve fair flow rates, we study the system behavior as we increase the policing threshold. A change in threshold changes the dynamics of the system, so we are actually considering the behavior of a series of systems indexed by $n = 1, 2, ...$, where the thresholds h_u and h_l

for system-n are given by the following assignments:

$$h_u := nh$$

$$h_l := nh - \sqrt{n}\varsigma$$

for some constants $h > \varsigma \geq 0$, where $\varsigma = 0$ if we choose not to have hysteresis. It is worth emphasizing that h is a scalar, and thus all queues use the same upper and lower thresholds.

The other aspects of the systems $1, ..., n$: routing, service disciplines, inter-arrival, and service time distributions of systems are all identical.

5.6 Candidate Equilibrium and "Relative" Initial Condition

The vector $X(t) = (Q(t); U(t); V(t); H(t))$ is the full system state of the network, where the ";" symbol denotes column concatenation. In the sections that follow, we try to show that the state of the stochastic network tends toward the equilibrium of its fluid model. We therefore define a vector e which we will later assign to be equal to the equilibrium queue depths of the fluid model with unit thresholds. However we have not yet derived the fluid model, so at this point e is left to be an arbitrary K dimensional vector.

The example in the preceding section has given us some idea of what the fluid model equilibrium will look like. Roughly speaking, we expect that in equilibrium queues that are "bottlenecks" fill up to their policing thresholds and then "chatter" there as policing turns on and off so as to thin the arrivals process to match the bottleneck queue rate. Other queues tend to remain empty. In this case, e equals a vector of 1's and 0's with 1's corresponding to queues that fill to their thresholds in the fluid model.

Because system-n has upper thresholds of nh we should expect the bottleneck queues of system-n to tend towards nhe. The full system state also contains U, V, and H so a complete candidate equilibrium should specify these states as well. We therefore define the full *candidate equilibrium* as $nh\mathbf{e} = (nhe; 0_U; 0_V; 0_H)$, where 0_U, 0_V, and 0_H are zero vectors of appropriate dimension.

We shall find that it useful to specify an initial condition relative to this candidate equilibrium, and scaled by the system size. We therefore define the *relative initial condition*, y for system-n as

$$y = \frac{1}{n}(x - nh\mathbf{e}) \tag{5.2}$$

where $x = X(0)$.

5.7 Dynamics of System-n

We now write a series of equations to describe the dynamics of system-n, starting from initial condition $y = n^{-1}(x - nh\mathbf{e})$. The equations have the same form as those originally derived by Harrison [41], [39] for Brownian models of queueing systems. Because all of the random processes involved depend both on relative initial condition y, and system scale n we define

$$\mathbf{n} \triangleq \begin{bmatrix} n \\ y \end{bmatrix} = \begin{bmatrix} \text{Threshold Scale Factor} \\ \text{Relative Initial Condition} \end{bmatrix}$$

for notational convenience. The overall state $X^{\mathbf{n}}(t)$ of the system-n, at time t, having started from relative initial condition y, is written as

$$X^{\mathbf{n}}(t) \triangleq \left[Q^{\mathbf{n}}(t); U^{\mathbf{n}}(t); V^{\mathbf{n}}(t); H^{\mathbf{n}}(t) \right].$$

The **n** superscript emphasizes a process's dependence on **n**.

5.7.1 Arrivals, Departures, and Routing

We begin by defining a process $\Phi^k(j)$ to track how often a class k packet is routed to class l for each l. Thus $\Phi^k(j)$ satisfies

$$\Phi^k(j) = \sum_{i=1}^{j} \phi^k(i).$$

We define $A_k^{\mathbf{n}}(t)$ to be the cumulative arrivals, both exogenous and internal, up to time t for class k. $D_l^{\mathbf{n}}(t)$ denotes the cumulative departures from class l from time 0 to t and therefore,

$$A_k^{\mathbf{n}}(t) = \sum_{l=1}^{K} \Phi^l(D_l^{\mathbf{n}}(t)) + \Lambda_k^{\mathbf{n}}(t). \tag{5.3}$$

For $k \in \{1, ..., g\}$, the thinned exogenous arrival process $\Lambda_k^{\mathbf{n}}(t)$ is defined by (5.1). Otherwise $\Lambda_k^{\mathbf{n}}(\cdot) \equiv 0$, reflecting that these classes do not receive exogenous arrivals.

The population of $Q_k^{\mathbf{n}}$ for $k = 1, ..., K$ evolves according to

$$Q_k^{\mathbf{n}}(t) = Q_k^{\mathbf{n}}(0) + A_k^{\mathbf{n}}(t) - D_k^{\mathbf{n}}(t), \tag{5.4}$$

$$Q_k^{\mathbf{n}}(t) \geq 0. \tag{5.5}$$

We define the K dimensional column vector $T^{\mathbf{n}}(t)$ such that the kth element $T_k^{\mathbf{n}}(t)$ is the cumulative time spent serving class k up to time t. Thus we have

$$T_k^{\mathbf{n}}(t) \text{ is nondecreasing and } T_k^{\mathbf{n}}(0) = 0. \tag{5.6}$$

Using the notation C_i to denote the ith row of the constituency matrix C, we define the cumulative idle time $I_i^{\mathbf{n}}(t)$ of station i as

$$I_i^{\mathbf{n}}(t) = t - C_i T^{\mathbf{n}}(t) \quad \text{is nondecreasing and} \quad I_i^{\mathbf{n}}(0) = 0. \tag{5.7}$$

The service disciplines that we consider are work conserving therefore,

$$\int_0^\infty C_i Q^{\mathbf{n}}(t) dI_i^{\mathbf{n}}(t) = 0. \tag{5.8}$$

5.7.2 Queueing Discipline

For a station i, the departures of class $k \in C(i)$ are determined by the composition of the service time counting process $S^{\mathbf{n}}(t)$, and the process $T^{\mathbf{n}}(t)$ as described by the equation

$$D_k^{\mathbf{n}}(t) = S_k^{\mathbf{n}}(T_k^{\mathbf{n}}(t)). \tag{5.9}$$

The queueing discipline of a station i serves each flow in proportion to the flow weights over long time intervals. More precisely for some constant c,

$$\begin{aligned} w_{\mathrm{f}(k)}^{-1} D_k^{\mathbf{n}}(t, t+\tau) &\geq w_{\mathrm{f}(l)}^{-1} D_l^{\mathbf{n}}(t, t+\tau) - c \\ &\text{whenever } Q_k(s) > 0 \ \forall s \in [t, t+\tau] \end{aligned} \tag{5.10}$$

for all $k, l \in C(i)$, where the notation $D_k^{\mathbf{n}}(t, t+\tau) \triangleq D_k^{\mathbf{n}}(t+\tau) - D_k^{\mathbf{n}}(t)$. We also assume that the instantaneous service rate of any queue is a function of the current state.

$$\dot{T}_k^{\mathbf{n}}(t) = f(X^{\mathbf{n}}(t)) \quad \text{For some function } f(\cdot).$$

If the instantaneous rates depend on other information, like the position in the polling cycle of a round robin discipline, that information may be appended to the H portion of the state vector X in some way. This also ensures that X is a true state in that it contains sufficient information to compute statistics of the future evolution of the network. We assume that $|H(t)^{\mathbf{n}}| \leq c$ for some constant c for all n, y, t. (Even this assumption can be relaxed to allow this encoding to scale with n, at the price of more complexity in the proof of Theorem 7.5.)

5.7.3 Trajectory Notation

The Markovian state of the system-n is

$$X^{\mathbf{n}}(t) \triangleq \Big[Q^{\mathbf{n}}(t); U^{\mathbf{n}}(t); V^{\mathbf{n}}(t); H^{\mathbf{n}}(t) \Big].$$

We claim that the process $X^{\mathbf{n}}$ is Markov by the following argument which follows that given by Dai [32] whose argument in turn followed from Kaspi and Mandelbaum [42]. Consider the evolution of $X^{\mathbf{n}}(t)$ starting from a particular time t^*. Each component of the residual time vector $U^{\mathbf{n}}(t^* + s)$ declines deterministically at rate 1, while the components of $V^{\mathbf{n}}(t^* + s)$ decline according to $\dot{T}^{\mathbf{n}}(\cdot)$, which is a deterministic function of $X^{\mathbf{n}}$. This continues until one of these components hits 0, say at time $t^* + s'$. Because this evolution is deterministic, s' was predictable from the value of $X^{\mathbf{n}}(t^*)$. At time $t^* + s'$, some components of Q, and perhaps H are incremented or decremented in a fashion that was also predictable with knowledge of $X^{\mathbf{n}}(t^*)$. Also at time $t^* + s'$, the component of $U^{\mathbf{n}}$ or $V^{\mathbf{n}}$ that hit zero takes a jump by a fresh inter-arrival or service time that is independent of the past. Hence, the conditional distribution of the future state after the jump, given the state $X^{\mathbf{n}}(t^*)$ is the same as given knowledge of the entire past. As the process continues to evolve in this manner, the probability distribution of the state at any future time t conditioned on the entire past before t^* is the same as it conditioned on $X^{\mathbf{n}}(t^*)$. This shows that the process is Markov. Furthermore, because the process behaves deterministically between jump times, $X^{\mathbf{n}}(t^*)$ is a type of process termed a piecewise deterministic Markov (PDM) process by Davis [43]. Davis shows that a PDM process whose expected number of jumps on $[0, t]$ is finite for each t is strong Markov [43]. As we assume that the inter-arrival and service times have a positive and finite mean, the expected number of jumps of $X^{\mathbf{n}}(t^*)$ in any closed time interval is finite. Therefore $X^{\mathbf{n}}(t^*)$ has the strong Markov property.

At times, we need to consider not only the trajectory of the state $X^{\mathbf{n}}(t)$ but also the evolution of the companion processes $T^{\mathbf{n}}(t)$, $\Lambda^{\mathbf{n}}(t)$, and the system-n threshold nh. Thus we define

$$\mathfrak{X}^{\mathbf{n}}(t) \triangleq \left[X^{\mathbf{n}}(t); T^{\mathbf{n}}(t); \Lambda^{\mathbf{n}}(t); nh \right].$$

Chapter 6

Proof Strategy

To obtain our desired result we need to use two different types of fluid limits. In the first type, we consider a sequence of systems indexed by n in which:

- System-n has policing thresholds nh.
- Time and space are scaled by the factor n.
- The initial condition x of each system-n is within some neighborhood of a candidate equilibrium ($|x - nh\mathbf{e}| \leq n\zeta$). Equivalently, $|y| \leq \zeta$.
- $n \to \infty$

We call the resulting fluid limit a TFL (Threshold Fluid Limit), while we call the sequence $\{(n, y)\}$ a TFL sequence. We show that such a limit converges along some subsequence, uniformly on compact sets, to a trajectory of a model satisfying a set of equations. We then show that if all trajectories of the fluid model converge, then the fluid limit converges (strongly) along the original sequence. Using this we conclude that there exist large enough thresholds such that whenever the initial condition is within a neighborhood of the candidate equilibrium, the expected rates of the stochastic network are close to the rates predicted by the fluid model over a compact time interval of some length t_0 seconds of n-scaled-time. This alone is not enough to obtain our desired result, because we want to show that the network achieves the desired rate over the long term, not just a compact time interval. This motivates the second type of fluid limit.

In the second type of fluid limit, we consider the limit along a sequence of system scale and relative-initial condition pairs $\{(n, y)\}$ such that:

- System-n has thresholds nh.
- Time and space are scaled by the factor $n|y|$.
- The initial condition x of each system-n is outside some neighborhood of a candidate equilibrium ($|x - h\mathbf{e}| > n\zeta$). Equivalently $|y| > \zeta$.
- $n|y| \to \infty$

We call this fluid limit a JFL (Joint threshold & initial condition Fluid Limit), and the sequence $\{(n, y)\}$ a JFL sequence. As with the other fluid limit, we will show that such a limit converges along some subsequence to a fluid model trajectory and that we can upgrade that convergence to convergence of the original sequence if the fluid model converges. In a similar fashion as [32], we exploit this convergence to construct a Lyapunov function which shows that the expected state of system-n contracts whenever $n|y| > L$ for some constant L. Thus by choosing n larger than L/ζ, we can make the Lyapunov function "active" for all y outside of the ζ-neighborhood (which we call

B) of the origin. This allows us to show that the invariant distribution of the system is concentrated in B. In particular, we show that the expected first return to B that happens at least t_0 seconds (in n-scaled-time) after starting in B, can be made arbitrarily close to t_0 for large n.

To complete the proof of the main result, we combine the results: that rates are close to desired for t_0 seconds after starting in B and that the expected first return to B after t_0 seconds is close to t_0 to conclude that the long term rates are close to the fluid model rates. A stopping time argument using the strong Markov property formalizes this reasoning.

In Chapter 8 we show that for networks in which flows follow loop-free paths without splitting the fluid model converges to an equilibrium and the rates approach the max-min fair share rates. However, this result holds only if each flow has a unique bottleneck station, a notion that we make precise later. The result allows us to invoke the fluid limit results to conclude that the corresponding stochastic network has rates close to max-min fair share.

Much of the analysis of the JFL and TFL fluid limits are identical. To avoid repetition, we define a scaling function $\langle \mathbf{n} \rangle$ that allows us to consider both fluid limits at once. The function $\langle \mathbf{n} \rangle$ is defined by

$$\langle \mathbf{n} \rangle_\zeta = \begin{cases} n & |y| \le \zeta \\ \zeta^{-1} n |y| & |y| > \zeta \end{cases}$$

where the ζ parameter is an arbitrary positive constant. To make the notation less cumbersome, we will omit the ζ subscript when either the choice of ζ does not matter, or when its value is clear from the context.

6.1 Result Summary

The upcoming Chapter 7 consists of a series of lemmas and theorems that culminate in Theorem 7.10, which shows that the long term rates of the stochastic network an be made close to the fluid model rates for large n. We briefly survey the steps that lead to that result:

Theorem 7.5: For both TFL or JFL sequences $\{\mathbf{n}_m\}$, there exists a subsequence $\{\mathbf{n}_j\}$ for which

$$\langle \mathbf{n}_j \rangle^{-1} \mathfrak{X}^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{\mathfrak{X}}(t)$$

for some trajectory $\bar{\mathfrak{X}}(t)$ that is dependent on sample path ω, and choice of subsequence, but must satisfy a set of fluid model equations (which are (7.5)-(7.19)).

Lemma 7.6: Suppose we have a functional $\mathfrak{F}$ that operates on system trajectories. (In the proof of the subsequent theorem we will pick $\mathfrak{F}(t)$ to extract the difference between the actual service rates vs. a vector of desired rates over a compact time interval, and later we will pick $\mathfrak{F}(t)$ to extract the distance between the current state and the candidate equilibrium.) Also suppose that any solution $\bar{\mathfrak{X}}(t)$ to the fluid model equations (7.5)- (7.19) (we call such a $\bar{\mathfrak{X}}(t)$ a *fluid trajectory*) is such that $\mathfrak{F} \circ [\bar{\mathfrak{X}}(\cdot)](t)$ goes to zero in a time proportional to the distance of the initial condition's distance to the fluid model equilibrium $\bar{h}\mathbf{e}$. Note that $\bar{h}$ is the threshold of the fluid model, which we will see later may not be h. Under these suppositions, $|\mathfrak{F} \circ [\langle \mathbf{n}_m \rangle^{-1} \mathfrak{X}(\langle \mathbf{n}_m \rangle t)](t)| \to 0$ a.s. for each t greater than some critical time ζt_0. (The functional $\mathfrak{F}$ must be continuous on the topology of uniform convergence on compact sets for this to hold.)

Theorem 7.7: Suppose that any fluid trajectory $\bar{\mathfrak{X}}(t)$ that solves equations (7.5)-(7.19) has service rates equal to some vector R (which we later show to be the max-min fair rates), after some time that is proportional to the starting distance from the fluid model's candidate equilibrium $\bar{h}\mathbf{e}$. Then for any positive numbers $\zeta < 1$ and any $\gamma < 1$, there exists a critical scale $L_1(\zeta, \gamma)$ such that for any system with $n > L_1$, and initial condition $|y| \leq \zeta$

$$\mathrm{E}[M^{-1}T^{\mathbf{n}}(nt_0)] \geq R(1-\zeta)(1-\gamma)(nt_0)$$

where t_0 is the time the fluid model takes to reach the ideal rates after starting one unit distance away from the equilibrium.

Theorem 7.8: Suppose that any fluid trajectory is such that the state goes to the equilibrium in a time proportional to the starting distance from the candidate equilibrium, then for any $\zeta < 1$, and $\delta < 1$ their exists a critical scale $L_2(\zeta, \gamma)$ such that for any system $n > L_2/\zeta$ and initial condition $|y| > \zeta$,

$$\mathrm{E}[Y^{n,y}(t_0|y|)] < \delta|y|.$$

where $Y^{n,y}$ is a "relative" state given by

$$Y^{n,y}(t) \triangleq \frac{1}{n}\left(X^{\mathbf{n}}(nt) - nh\mathbf{e}\right).$$

(Roughly, states outside the ζ neighborhood tend to contract.) Also shown in Theorem 7.8, is that for any $b < 1$, there exists n large enough so that

$$\mathrm{E}[Y^{n,y}(t_0)] < b \quad \text{for all } y : |y| \leq \zeta.$$

Lemma 7.9: If the conclusions of Theorem 7.8 hold, then the expected first return of Y to a ζ neighborhood of the origin that happens at least t_0 seconds after having started in that neighborhood, can be made close to t_0, if the constants ζ, b, and δ of Theorem 7.8 are chosen to be small.

Theorem 7.10: If the necessary conditions for Theorems 7.7 and 7.8 hold, then the long term average rates can be made arbitrarily close to the fluid model rates, by making n large enough.

Chapter 7

Fluid Limit Analysis

7.1 Preliminary Lemmas

In the this chapter, we will state and prove Theorem 7.5 that shows that given any sequence of initial conditions and system scales $\{\mathbf{n}_m\}$ with $\langle \mathbf{n}_m \rangle \to \infty$ there exists a subsequence $\{\mathbf{n}_j\} \subset \{\mathbf{n}_m\}$ for which the sequence of system trajectories with initial condition y_j, thresholds $n_j h$, and time and space scaled by $\langle \mathbf{n}_j \rangle$ converges to a fluid model trajectory satisfying a set of differential equations.

The proof of Theorem 7.5 depends on the lemmas presented in this section. The reader may either read this section first, or skip to the proof of Theorem 7.5 in Section 7.1 and turn back to this section as needed.

Lemmas 7.1, 7.2, and 7.3 all assume that we start with some sequence $\{\mathbf{n}_j\}$ that satisfies the following property concerning the residual arrival and service time processes in the fluid limit:

Property 1: $\{\mathbf{n}_j\}$ is a sequence of initial condition y_j and scale n_j pairs with $\langle \mathbf{n}_j \rangle \to \infty$ and

$$U^{\mathbf{n}_j}(0) \to \bar{U}(0) \quad V^{\mathbf{n}_j}(0) \to \bar{V}(0)$$

for some $\bar{U}(0)$, $\bar{V}(0)$.

This property ensures that there are well defined residual arrival and service times in the fluid limit. (When we eventually prove Theorem 7.5, when a sequence does not satisfy Property 1, we will find a subsequence for which it does.)

Lemma 7.1 is a form of the Functional Strong Law of Large numbers for renewal processes, and is taken from [32].

Lemma 7.2 is a new result showing that the thinned arrivals (the packets that make it beyond the policing point) converge to a fluid limit along a subsequence.

Lemma 7.3 is a result taken from [32] showing that the residual initial arrival and service times decline to zero at rate 1 in the fluid limit. The lemma also shows that the sequence of functions we use to take the fluid limit are uniformly integrable, which will later be used to show a sequence of expected values evaluated at a time t converges to 0 when the sequence of functions evaluated at time t converge to 0 almost surely.

We say that $f_j(t) \to f(t)$ *uniformly on compact sets* (u.o.c.) if for each $t \geq 0$

$$\lim_{j \to \infty} \sup_{0 \leq s \leq t} |f_j(s) - f(s)| = 0.$$

We also use the notation $\dot{f}(t) = \frac{d}{dt} f(t)$ where such a derivative exists. If a function $f(\cdot)$ is differentiable at t, we say that t is a *regular point*.

Lemma 7.1 (Dai, Lemma 4.2 of [32]). *Suppose that* $\{\mathbf{n}_j\}$ *is a sequence satisfying*

Property 1. Then for almost all ω,

$$\langle \mathbf{n}_j \rangle^{-1} \Phi^k(\rfloor \langle \mathbf{n}_j \rangle t \lfloor) \to P_k' t \ \textit{u.o.c.},$$

$$\langle \mathbf{n}_j \rangle^{-1} E_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \ \to \alpha_k (t - \bar{U}_k(0))^+ \ \textit{u.o.c.},$$

$$\langle \mathbf{n}_j \rangle^{-1} S_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \ \to \mu_k (t - \bar{V}_k(0))^+ \ \textit{u.o.c.}$$

Proof. See Lemma 4.2 of Dai [32]. The result is an instance of the Strong Law of Large Numbers for Renewal Processes [50]. □

Lemma 7.2 (Thinned Arrival Convergence). *Suppose that* $\{\mathbf{n}_v\}$ *is a sequence satisfying Property 1. Then for almost all* ω, *there exists a subsequence* $\{\mathbf{n}_j\} \subset \{\mathbf{n}_v\}$ *such that*

$$\langle \mathbf{n}_j \rangle^{-1} \Lambda^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{\Lambda}(t) \quad \textit{u.o.c.}$$

where $\bar{\Lambda}(t)$ *is some Lipschitz continuous process with for all regular* $t \geq 0$,

$$\dot{\bar{\Lambda}}_f(t) \leq \alpha_f \qquad \textit{for each flow } f. \tag{7.1}$$

Proof. By Lemma 7.1,

$$\langle \mathbf{n}_v \rangle^{-1} E_k^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) \to (\alpha_k t - V\bar{(0)}_k)^+ \ \text{ u.o.c.} \tag{7.2}$$

for each class k. For notational convenience in the development that follows, we define:

$$\bar{E}_k(t) \triangleq (\alpha_k t - \bar{V}(0)_k)^+$$

$$\Delta^v(t) \triangleq \langle \mathbf{n}_v \rangle^{-1} E^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) - \bar{E}(t). \tag{7.3}$$

Pick a compact time interval $[s_0, s_1]$. Because the number of packets allowed past the policing point cannot exceed the number of packets arriving to the policing point in any time interval (5.1), we have

$$\frac{1}{\langle \mathbf{n}_v \rangle}\left[\Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle (t+\varepsilon)) - \Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t)\right] \leq \frac{1}{\langle \mathbf{n}_v \rangle}\left[E^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle (t+\varepsilon)) - E^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t)\right] \quad (7.4)$$

for any positive $\varepsilon \leq s_1 - s_0$ and $t : s_0 \leq t \leq s_1 - \varepsilon$. Adding $-\Delta^v(t+\varepsilon)$ and $\Delta^v(t)$ to both sides and substituting (7.3) and (7.2), we have

$$\frac{1}{\langle \mathbf{n}_v \rangle}\Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle (t+\varepsilon)) - \Delta^v(t+\varepsilon) - \left[\frac{1}{\langle \mathbf{n}_v \rangle}\Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) - \Delta^v(t)\right] \leq \bar{E}(t+\varepsilon) - \bar{E}(t) \leq \varepsilon\alpha.$$

Define the family of functions:

$$\mathfrak{L}_v(s_0, t) := \sup_{s \in [s_0, t]} \left[\langle \mathbf{n}_v \rangle^{-1} \Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle s) - \Delta^v(s)\right]$$

Because the arguments of the sup functions are vectors, sup is taken component-wise. Note that for any (t, ε) with $t \in [s_0, s_1 - \varepsilon]$,

$$\begin{aligned} \mathfrak{L}_v(s_0, t+\varepsilon) &= \mathfrak{L}_v(s_0, t) \vee \mathfrak{L}_v(t, t+\varepsilon) \\ \text{and} \quad \mathfrak{L}_v(t, t+\varepsilon) &\leq \varepsilon\alpha + \mathfrak{L}_v(t, t) \\ &\leq \varepsilon\alpha + \mathfrak{L}_v(s_0, t). \end{aligned}$$

Thus $\mathfrak{L}_v(s_0, t+\varepsilon) - \mathfrak{L}_v(s_0, t) \leq \varepsilon\alpha$ and clearly $\mathfrak{L}_v(s_0, t+\varepsilon) - \mathfrak{L}_v(s_0, t) \geq 0$ because $\mathfrak{L}_v(s_0, \cdot)$ is monotone. Hence the functions $\mathfrak{L}_v(s_0, \cdot)$ are equicontinuous and individually Lipschitz continuous. Thus, by Arzela's theorem, there exists a further subsequence $\{\mathbf{n}_j\} \subseteq \{\mathbf{n}_v\}$ such that

$$\mathfrak{L}_j(s_0, t) \to \bar{\Lambda}(t)$$

uniformly on the compact set $t \in [s_0, s_1]$ for some monotone-nondecreasing, Lipschitz-continuous process $\bar{\Lambda}(t)$. But by (7.2), $\Delta^j(t) \to 0$ uniformly on compact sets. Because of this and the fact that $\langle \mathbf{n}_j \rangle^{-1} \Lambda^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s)$ is monotone, it follows that the maximizing values of each sup term approach $\langle \mathbf{n}_v \rangle^{-1} \Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t)$. Thus

$$\sup_{s \in [s_0, t]} \left[\frac{1}{\langle \mathbf{n}_v \rangle} \Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle s) - \Delta^v(s) \right] \to \frac{1}{\langle \mathbf{n}_v \rangle} \Lambda^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) \to \bar{\Lambda}(t).$$

Because the choice of $[s_0, s_1]$ was arbitrary, we have $\langle \mathbf{n}_j \rangle^{-1} \Lambda^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{\Lambda}(t)$ u.o.c. Furthermore, (7.2) and (7.4) imply that $\bar{\Lambda}(t)$ satisfies (7.1). □

Lemma 7.3 (Lemmas 4.3 & 4.5 of Dai [32]). *Suppose that* $\{\mathbf{n}_j\}$ *is a sequence satisfying Property 1. Then almost surely:*

$$\lim_{j \to \infty} \langle \mathbf{n}_j \rangle^{-1} U_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) = (\bar{U}_f(0) - t)^+ \text{ u.o.c.},$$

$$\lim_{j \to \infty} \langle \mathbf{n}_j \rangle^{-1} V_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) = (\bar{V}_k(0) - t)^+ \text{ u.o.c.}$$

Also, for each fixed $t \geq 0$, *the sets of functions:*

$$\left\{ \langle \mathbf{n}_j \rangle^{-1} U^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) : \langle \mathbf{n}_j \rangle \geq 1 \right\},$$

$$\left\{ \langle \mathbf{n}_j \rangle^{-1} V^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) : \langle \mathbf{n}_j \rangle \geq 1 \right\},$$

$$\left\{ \langle \mathbf{n}_j \rangle^{-1} Q^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) : \langle \mathbf{n}_j \rangle \geq 1 \right\}$$

are uniformly integrable.

Proof. See Lemmas 4.3 and 4.5 of Dai [32] □

The following lemma will be used later to show that because all of the systems we consider are work conserving, which we stated with the integral expression (5.8), the fluid limit must also be work conserving. In the lemma below, the notation $D_{\mathbb{R}}[0, \infty)$ denotes the space of right continuous functions on $\mathbb{R}_+$ having left limits on $(0, \infty)$, and endowed with the Skorohod topology [44]. $C_{\mathbb{R}}[0,\infty) \subset D_{\mathbb{R}}[0,\infty)$ is the subset of continuous paths.

Lemma 7.4 (Lemma 2.4 of Dai and Williams [45]). *Let* $\{(z_j, \chi_j)\}$ *be a sequence in* $D_{\mathbb{R}}[0,\infty) \times C_{\mathbb{R}}[0,\infty)$*. Assume that* χ_j *is nondecreasing and* (z_j, χ_j) *converges to* $(z,\chi) \in C_{\mathbb{R}}[0,\infty) \times C_{\mathbb{R}}[0,\infty)$ *u.o.c. Then for any bounded continuous function* f*,*

$$\int_0^t f(z_j(s))d\chi_j(s) \to \int_0^t f(z(s))d\chi(s) \;\; u.o.c.$$

Proof. See Lemma 2.4 of Dai and Williams [45]. □

7.2 Convergence to a Fluid Limit along a Subsequence

The following proof parallels the proof of Theorem 4.1 of Dai [32]. The proof here differs in that we deal with a sequence of systems and two different types of fluid limits, as discussed in Section 6.

Theorem 7.5. *Suppose one of the following cases hold for some constant* ζ*:*

TFL Case: $\{\mathbf{n}_m\} = \{(n_m, y_m)\}$ *is a sequence of (system scale, relative initial condition) pairs satisfying:*

$$|y_m| \leq \zeta \quad \text{and} \quad \langle \mathbf{n}_m \rangle_\zeta = n_m \to \infty.$$

JFL Case: $\{\mathbf{n}_m\} = \{(n_m, y_m)\}$ *is a sequence of (system scale, relative initial condition) pairs satisfying:*

$$|y_m| > \zeta \quad \text{and} \quad \langle \mathbf{n}_m \rangle_\zeta = \frac{1}{\zeta} n_m |y_m| \to \infty.$$

Then for almost all ω *there exists a subsequence* $\{\mathbf{n}_j\} \subseteq \{\mathbf{n}_m\}$ *for which*

$$\langle \mathbf{n}_j \rangle^{-1} \mathfrak{X}^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{\mathfrak{X}}(t) \qquad u.o.c.$$

for some fluid trajectory $\bar{\mathfrak{X}}(t)$ *with components*

$$\bar{\mathfrak{X}}(t) \triangleq [\bar{X}(t); \bar{T}(t); \bar{\Lambda}(t); \bar{h}]$$

where, in turn, the process $\bar{X}(t)$*has components*

$$\bar{X}(t) \triangleq [\bar{Q}(t); \bar{U}(t); \bar{V}(t); \bar{B}(t)]$$

where $\bar{B}(t) \equiv 0$*. The process* $\bar{\mathfrak{X}}(t)$ *may depend upon* ω *and the choice of subsequence* $\{\mathbf{n}_j\}$ *but must satisfy the following properties for all* $t \geq 0$*:*

$$\bar{U}_f(t) = (t - \bar{U}_f(0))^+, \quad \bar{V}_k(t) = (t - \bar{V}_k(0))^+, \tag{7.5}$$

$$\bar{T}_k(t) \text{ is nondecreasing and starts from zero,} \tag{7.6}$$

$$\bar{I}_i(t) := t - C_i \bar{T}(t) \text{ is nondecreasing ,} \tag{7.7}$$

$$\bar{D}_k(t) := \mu_{s(k)}(\bar{T}_k(t) - \bar{V}_k(0))^+, \tag{7.8}$$

$$\bar{A}(t) := R^\top \bar{D}(t) + \bar{\Lambda}(t), \tag{7.9}$$

$$\bar{Q}(t) := \bar{Q}(0) + \bar{A}(t) - \bar{D}(t), \tag{7.10}$$

$$\bar{Q}(t) \geq 0, \tag{7.11}$$

$$\int_0^\infty (C\bar{Q}(t)) d\bar{I}(t) = 0, \tag{7.12}$$

where (7.5), (7.6), and (7.8) hold for each flow f *and class* k*, while (7.7) holds for each station* i*. Assignments (7.7), (7.8), (7.9), and (7.10) define* $\bar{I}(t)$*,* $\bar{D}(t)$*,* $\bar{A}(t)$*, and* $\bar{Q}(t)$

respectively. Also, the following hold for each flow f for regular $t \geq 0$:

$$\dot{\bar{\Lambda}}_f(t) = 0 \qquad \text{whenever } \bar{Q}_k(t) > h \text{ for some } k \in \mathcal{C}(f), \tag{7.13}$$

$$\dot{\bar{\Lambda}}_f(t) = \alpha_f 1(t \geq \bar{U}_f(0)) \qquad \text{whenever } \bar{Q}_k(t) < h \text{ for all } k \in \mathcal{C}(f), \tag{7.14}$$

$$\dot{\bar{\Lambda}}_f(t) \leq \alpha_f. \tag{7.15}$$

Also, for station i and for any k, l such that $\{k, l\} \in C(i)$ the following are satisfied for all regular $t \geq 0$:

$$w_k^{-1} \dot{\bar{D}}_k(t) \geq w_l^{-1} \dot{\bar{D}}_l(t) \qquad \text{whenever } Q_k(t) > 0\ , \tag{7.16}$$

$$w_k^{-1} \dot{\bar{D}}_k(t) = w_l^{-1} \dot{\bar{D}}_l(t) \qquad \text{whenever } Q_k(t) > 0 \text{ and } Q_l(t) > 0. \tag{7.17}$$

In addition the following case specific conditions on $\bar{h}$, and $\bar{X}(0)$ hold:

$$\textbf{TFL Case:} \qquad \bar{h} = h, \qquad |\bar{X}(0) - \bar{h}\mathbf{e}| \leq \zeta, \tag{7.18}$$

$$\textbf{JFL Case:} \qquad 0 \leq \bar{h} \leq h, \qquad |\bar{X}(0) - \bar{h}\mathbf{e}| = \zeta. \tag{7.19}$$

Proof. We first consider the parts of the proof that require case specific analysis. For the TFL case we have

$$\frac{n_m}{\langle \mathbf{n}_m \rangle} h = \frac{n_m h}{n_m} = h,$$

thus

$$\frac{n_j}{\langle \mathbf{n}_j \rangle} h \to \bar{h} = h$$

along any subsequence $\{\mathbf{n}_j\} \subset \{\mathbf{n}_m\}$ giving us the first part of (7.18). For the JFL case we have

$$\frac{n_m}{\langle \mathbf{n}_m \rangle} h = \frac{n_m h}{\zeta^{-1} n_m |y_m|} = \frac{\zeta h}{|y_m|} < h. \tag{7.20}$$

Thus by the Bolzano-Weierstrass Theorem, there exists a further subsequence $\{\mathbf{n}_i\} \subseteq \{\mathbf{n}_m\}$ for which

$$\frac{n_i}{\langle \mathbf{n}_i \rangle} h \to \bar{h} \tag{7.21}$$

holds for some $\bar{h}$. The first part of (7.19) follows from (7.20). We also note that in the JFL case, (7.20) implies

$$\bar{h} = \lim_{i\to\infty} \zeta h |y_i|^{-1}. \tag{7.22}$$

For the TFL case, we use the definition of relative initial condition (5.2) and $\langle \mathbf{n} \rangle$ to write

$$\begin{aligned} \langle \mathbf{n}_i \rangle^{-1} |X^{\mathbf{n}_i}(0) - n_i h\mathbf{e}| &= \frac{1}{n_i}(|n_i y_i + nh\mathbf{e} - nh\mathbf{e}|) \\ &\le |y_i| \le \zeta. \end{aligned} \tag{7.23}$$

Similarly for the JFL case we have

$$\begin{aligned} \langle \mathbf{n}_i \rangle^{-1} |X^{\mathbf{n}_i}(0) - n_i h\mathbf{e}| &= \left| \frac{n_i(y_i + h\mathbf{e})}{n_i |y_i|} - \bar{h}\mathbf{e} \right| \\ &\le \zeta + (\zeta h |y_i|^{-1} - \bar{h})|\mathbf{e}| \qquad (7.24) \\ &\le \zeta + h|\mathbf{e}|. \qquad (7.25) \end{aligned}$$

Thus we may apply the Bolzano-Weierstrass theorem in either case to conclude that there is a subsequence $\{\mathbf{n}_r\} \subseteq \{\mathbf{n}_i\}$ for which $\langle \mathbf{n}_r \rangle^{-1} X^{\mathbf{n}_r}(0) \to \bar{X}(0)$ for some $\bar{X}(0)$. In addition $\langle \mathbf{n}_r \rangle^{-1} H^{\mathbf{n}_r}(\langle \mathbf{n}_r \rangle t) \to 0$ u.o.c. because $H^{\mathbf{n}_r}(\langle \mathbf{n}_r \rangle t)$ is bounded by a constant by its definition. Thus,

$$\langle \mathbf{n}_r \rangle^{-1} X^{\mathbf{n}_r}(0) \to [\bar{Q}(0); \bar{U}(0); \bar{V}(0); 0] \tag{7.26}$$

$$\langle \mathbf{n}_r \rangle^{-1} H^{\mathbf{n}_r}(t) \to 0 \quad \text{u.o.c.} \tag{7.27}$$

Property (7.26) allows us to use Lemma 7.3 to conclude

$$\langle \mathbf{n}_r \rangle^{-1} \begin{bmatrix} U^{\mathbf{n}_r}(\langle \mathbf{n}_r \rangle t) \\ \\ V^{\mathbf{n}_r}(\langle \mathbf{n}_r \rangle t) \end{bmatrix} \to \begin{bmatrix} \bar{U}(t) \\ \\ \bar{V}(t) \end{bmatrix} \quad \text{u.o.c.}$$

where $\bar{U}(t)$ and $\bar{V}(t)$ satisfy (7.5).

In the TFL case, the second part of (7.18) follows from (7.23). For the JFL case, (7.24) combined with (7.22) imply that $(\zeta h|y_r|^{-1} - \bar{h}) \to 0$ and thus $|\bar{X}(0) - \bar{h}\mathbf{e}| = \zeta$, giving us the second part of (7.19).

From this point on, the arguments apply for both cases. $T^{\mathbf{n}_i}$ satisfies

$$\langle \mathbf{n}_r \rangle^{-1} [T^{\mathbf{n}_i}(\langle \mathbf{n}_r \rangle t) - T^{\mathbf{n}_r}(\langle \mathbf{n}_r \rangle s)] \leq (t - s). \tag{7.28}$$

Thus by Arzela's theorem [46], there exists a further subsequence $\{\mathbf{n}_v\} \subseteq \{\mathbf{n}_r\}$ for which

$$\langle \mathbf{n}_v \rangle^{-1} T^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) \to \bar{T}(t).$$

Property (7.6) follows from (5.6). Property (5.7) implies

$$\langle \mathbf{n}_v \rangle^{-1} I^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) \to \bar{I}(t)$$

u.o.c. where $\bar{I}(t)$ satisfies (7.7).

By Lemma 7.1, $\langle \mathbf{n}_v \rangle^{-1} S_k^{\mathbf{n}_v}(\langle \mathbf{n}_v \rangle t) \to (\mu_k t - \bar{V}(0)_k)^+$ u.o.c. for each class k. This fact combined with (5.9) and (7.28) gives (7.8).

Property (7.26) allows us to use Lemma 7.2 to conclude that there is a subsequence $\{\mathbf{n}_j\} \subset \{\mathbf{n}_v\}$ where

$$\langle \mathbf{n}_v \rangle^{-1} \Lambda^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{\Lambda}(t) \quad \text{u.o.c.}$$

for some Lipschitz continuous process $\bar{\Lambda}(t)$ satisfying (7.15).

Lemma 7.1 combined with (5.3) gives us $\langle \mathbf{n}_j \rangle^{-1} A_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{A}_k(t)$ u.o.c. where $\bar{A}_k(t)$ is defined by (7.9). Furthermore, $\bar{A}_k(t)$ is Lipschitz continuous because it is equal to a linear combination of functions we have already shown to be Lipschitz continuous.

Thus using (5.4) we have that

$$\langle \mathbf{n}_j \rangle^{-1} Q^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \to \bar{Q}(t) \quad \text{u.o.c.} \tag{7.29}$$

where $\bar{Q}(t)$ is a Lipschitz continuous function given by (7.10). Property (7.11) follows easily from (5.5).

The next few arguments are similar to the proof of Proposition 4.2 in [33].

Suppose that $\bar{Q}_k(t) > \bar{h}$ for some $k \in \mathcal{C}(f)$. By Lipschitz continuity of $\bar{Q}_k(t)$, there exists some small $\tau > 0$ such that

$$\min_{t \le s \le t+\tau} \bar{Q}_k(s) > \bar{h}.$$

By the uniformity of the queue convergence in (7.29) and the threshold convergence in (7.21), there exists j^* such that for all $j > j^*$, $Q_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) > n_j h$ for all $s \in [t, t+\tau]$. Thus, by (5.1) one finds that

$$\Lambda_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) - \Lambda_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) = 0 \quad \forall s \in [t, t+\tau].$$

Therefore, it follows that

$$\bar{\Lambda}_f(s) - \bar{\Lambda}_f(t) = 0 \quad \forall s \in [t, t+\tau]$$

and consequently, $\dot{\bar{\Lambda}}_f(t) = 0$, which is (7.13).

Suppose that $\bar{Q}_k(t) < \bar{h}$ for all $k \in \mathcal{C}(f)$. First note that in this case $\bar{h} > 0$ and therefore (7.21) implies that $n_j \to \infty$. By the Lipschitz continuity of $\bar{Q}_k(t)$ for each k, there exists some small $\tau > 0$ such that

$$\max_{k \in \mathcal{C}(f)} \max_{s \in [t, t+\tau]} \bar{Q}_k(s) < \bar{h}.$$

Because $n_j \to \infty$, the uniformity of the convergence in (7.29), and the convergence in (7.21), there exists j' such that for all $j > j'$, $Q_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) < n_j h$. Furthermore there exists a $j^* \geq j'$ such that for all $j > j^*$ and $k \in \mathcal{C}(f)$, $Q_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) < n_j h - \sqrt{n_j} h \varsigma$. Thus by (5.1):

$$\Lambda_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) - \Lambda_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) = E_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) - E_f^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t) \qquad \forall s \in [t, t+\tau]$$

and consequently we have (7.14).

Suppose that for some class k, $\bar{Q}_k(t) > 0$. By the Lipschitz continuity of $\bar{Q}_k(t)$ there exists some small $\tau > 0$ such that

$$\min_{t \leq s \leq t+\tau} \bar{Q}_k(s) > 0.$$

Because of the uniformity of convergence in (7.29) there exists j^* such that for all $j > j^*$,

$$Q_k^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle s) > 0 \;\; \forall s \in [t, t+\tau].$$

By (5.10), for almost all ω, and all classes l we have

$$w_k^{-1}[D_k(\langle \mathbf{n}_j \rangle s) - D_k(\langle \mathbf{n}_j \rangle t)] \geq w_l^{-1}[D_l(\langle \mathbf{n}_j \rangle s) - D_l(\langle \mathbf{n}_j \rangle t)] \qquad \forall s \in [t, t+\tau]$$

and thus we have (7.16).

If $\bar{Q}_l(t) > 0$ and $\bar{Q}_k(t) > 0$ (7.16) is true as written or with the k and l and indices swapped. This implies (7.17).

We observe that (5.8) is equivalent to

$$\int_0^\infty f(\chi_j) dz_j = 0 \tag{7.30}$$

where

$$\chi_j := \langle \mathbf{n}_j \rangle^{-1} C_i Q^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t)$$

$$z_j := \langle \mathbf{n}_j \rangle^{-1} I_i^{\mathbf{n}_j}(\langle \mathbf{n}_j \rangle t)$$

$$f(\cdot) := (\cdot) \wedge 1. \tag{7.31}$$

Noting that χ_j and z_j meet the required conditions for Lemma 7.4 we have,

$$\int_0^\infty [C_i \bar{Q}(t)] \wedge 1 d\bar{I}_i(t) = 0 \tag{7.32}$$

which is equivalent to (7.12). □

7.2.1 Sliding Modes

We refer to equations (7.5) through (7.17) as the fluid model equations, while we call a solution to the fluid model equations a fluid trajectory. The fluid model equations are ordinary differential equations with respect to time, but the values of the derivatives are discontinuous functions of the state. For example (7.13) and (7.14), imply that $\dot{\bar{\Lambda}}_f(x)$ is discontinuous because the value changes abruptly when a queue in $\mathcal{C}(f)$ crosses the threshold $\bar{h}$. Differential equations with right-hand discontinuities, as they are called, may not have unique solutions for a given initial condition [47]. Fortunately, we will not have to show that the fluid model equations admit a unique solution. Instead, it will suffice to show that all solutions converge to an equilibrium.

Consider the boundary set on the state space $\mathcal{S}_f$, defined to be the states for which all queues in $\mathcal{C}(f)$ are at or below their thresholds and at least one such queue is exactly at its threshold. Relation (7.15) restricts

$$\dot{\bar{\Lambda}}_f(x) \in [0, \alpha_f] \text{ for } x \in \mathcal{S}_f \ .$$

Thus the fluid model equations are differential inclusions at the switching boundaries. Fillipov [47] gives a general treatment of discontinuous differential equations with differential inclusions at the switching boundaries. One type of solution to such equations is a sliding mode, where the trajectory "sticks" to the switching boundary.

For instance, the fluid model of a network with a single queue, served at rate μ, and single flow arriving at rate $\alpha > \mu$ will admit a sliding mode solution. The fluid model equations dictate that

$$\dot{\bar{Q}}(t) \begin{cases} = \alpha - \mu & \text{when } \bar{Q}(t) < \bar{h} \\ \in [0, \alpha] & \text{when } \bar{Q}(t) = \bar{h} \\ = -\mu & \text{when } \bar{Q}(t) > \bar{h} \end{cases}$$

for regular t. One can show that a solution, and in this case the only solution, of this system is a trajectory that goes from its initial condition to the sliding boundary in finite time, and then stays at the sliding boundary $\bar{Q}(\cdot) \equiv \bar{h}$. Once the trajectory sticks to the boundary, it must be that the policing is throttling the flow so that its thinned rate $\dot{\bar{\Lambda}}(t)$ is equal to μ. We informally call this a *sliding mode policing.* A sliding mode policing in the fluid model looks almost if the queue were giving a real number valued throttle signal to the policing point. However, in the original stochastic system this must correspond to a sequence of on and off policing events, with some amount of time between events at least as long as the service time of a single packet. This chattering gets smoothed into a continuous trajectory in the fluid scaling.

7.3 Upgrading Convergence along Subsequences to Convergence on Sequences

In the previous section, we showed that for both TFL and JFL sequences, we can extract a sample path dependent subsequence that converges to a fluid trajectory. The objective of this section is to upgrade this result to convergence along the original sequence, and not just the sample path dependent subsequence. In particular, we show in Lemma 7.6 which follows that if a functional $\mathfrak{F}$ of the fluid model trajectory goes to zero in a time proportional to the initial condition's distance from the fluid model equilibrium, then we get our desired convergence property. In later sections, we will invoke Lemma 7.6 choosing $\mathfrak{F}$ to extract the service rates from the fluid model, and later choosing $\mathfrak{F}$ to extract the distance from the fluid model equilibrium. Lemma 7.6 is a generalization of an argument used by Dai in the proof of Theorem 4.2 of [32].

Lemma 7.6. *Suppose that $\mathfrak{F}$ is a functional that maps $\mathbb{R}^r \times \mathbb{R}^+$ into $\mathbb{R}^s \times \mathbb{R}^+$ where r is the dimension of $\mathfrak{X}^{\mathbf{n}}(t)$ and s is arbitrary. Also suppose that $\mathfrak{F}$ is continuous on the topology of uniform convergence on compact sets.*

Recall that $\bar{h}$ is a component of $\bar{\mathfrak{X}}(t)$, and suppose that the following statement is true:

- *The fluid model equations (7.5) - (7.17) are such that any trajectory $\bar{\mathfrak{X}}(t)$ with $\bar{h} > 0$ that satisfies them, $\bar{\mathfrak{X}}(t)$ must also satisfy*

$$\mathfrak{F} \circ [\bar{\mathfrak{X}}(\,\cdot\,)](t) \equiv 0 \quad \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}|. \tag{7.33}$$

Then,

(TFL Case:) *For any sequence* $\{\mathbf{n}_m\} = \{(n_m, y_m)\}$ *that satisfies* $|y_m| \le \zeta$ *and* $n_m \to \infty$ *for some* $\zeta > 0$,

$$\left| \mathfrak{F} \circ \left[\frac{1}{\langle \mathbf{n}_m \rangle} \mathfrak{X}^{\mathbf{n}_m}(\langle \mathbf{n}_m \rangle \cdot\,) \right](t) \right| \to 0 \qquad a.s. \tag{7.34}$$

for each $t \ge \zeta t_0$*, and any initial condition* y.

(JFL Case:) *Suppose Condition (7.33) can be strengthened to hold for any trajectory* $\bar{\mathfrak{X}}$ *with* $\bar{h} \ge 0$*. Then for any sequence* $\{\mathbf{n}_m\} = \{(n_m, y_m)\}$ *that satisfies*

$$|y_m| > \zeta \text{ and } n_m|y_m| \to \infty$$

for some $\zeta > 0$*, conclusion (7.34) holds.*

Proof. By Theorem 7.5, for almost all sample paths ω, and for any subsequence $\{\mathbf{n}_j\} \subset \{\mathbf{n}_m\}$ there is a sample-path-dependent further-subsequence $\{\mathbf{n}_{l(\omega)}\} \subset \{\mathbf{n}_j\}$ for which

$$\langle \mathbf{n}_{j_i(\omega)} \rangle^{-1} \mathfrak{X}^{\mathbf{n}_{j_i(\omega)}}(\langle \mathbf{n}_{j_i(\omega)} \rangle t, \omega) \to \bar{\mathfrak{X}}(t, \omega) \quad \text{u.o.c.} \tag{7.35}$$

where $\bar{\mathfrak{X}}(t, \omega)$ satisfies (7.6) - (7.17). The notation $l(\omega)$ and $\bar{\mathfrak{X}}(t, \omega)$ emphasize that the further-subsequence and fluid trajectory depend on ω. Now fix an ω for which subsequences have convergent further subsequences as described. For the next few steps we suppress the ω arguments to simplify notation, but the reader should remember we are working with a particular ω. We also note that the $\bar{h}$ component of $\bar{\mathfrak{X}}$ satisfies $\bar{h} > 0$ in the TFL Case and $\bar{h} \ge 0$ in JFL case by conclusions (7.18) and (7.19) of Theorem 7.5 respectively. Consequently, condition (7.33) holds for the TFL Case, and the strengthened version of condition (7.33) holds in the JFL Case. Thus in both cases, we have

$$\mathfrak{F} \circ \left[\bar{\mathfrak{X}}(\,\cdot\,) \right](t) = 0 \quad \text{for each } t \ge |X(0) - h\mathbf{e}| t_0.$$

Furthermore $|X(0) - h\mathbf{e}| \le \zeta$ in both cases by (7.18) and (7.19), thus

$$\mathfrak{F} \circ \left[\bar{\mathfrak{X}}(\,\cdot\,)\right](t) = 0 \quad \text{for each } t \ge \zeta t_0. \tag{7.36}$$

Because $\mathfrak{F}$ is assumed to be continuous on the topology of uniform convergence on compact sets, (7.35) implies

$$\mathfrak{F} \circ \left[\langle \mathbf{n}_l \rangle^{-1} \mathfrak{X}^{\mathbf{n}_l}(\langle \mathbf{n}_l \rangle\,\cdot\,)\right](t) \to \mathfrak{F} \circ [\bar{\mathfrak{X}}(\cdot)](t) \quad \text{u.o.c.},$$

which combined with (7.36) yields

$$\left|\mathfrak{F} \circ \left[\langle \mathbf{n}_l \rangle^{-1} \mathfrak{X}^{\mathbf{n}_l}(\langle \mathbf{n}_l \rangle\,\cdot\,)\right](t)\right| \to 0 \tag{7.37}$$

for each $t \ge \zeta t_0$. So for this fixed ω, any subsequence $\{\mathbf{n}_j\} \subseteq \{\mathbf{n}_m\}$ has a further subsequence $\{\mathbf{n}_{l(\omega)}\} \subseteq \{\mathbf{n}_j\}$ for which (7.37) holds. Therefore the original sequence $\{\mathbf{n}_m\}$ converges for this fixed ω:

$$\left|\mathfrak{F} \circ \left[\langle \mathbf{n}_m \rangle^{-1} \mathfrak{X}(\langle \mathbf{n}_m \rangle\,\cdot\,, \omega)\right](t)\right| \to 0$$

for each $t \ge \zeta t_0$. But the same argument can be used to conclude that this holds for almost all ω. Thus, we have (7.34). □

7.4 Convergence to Fluid Model Rates on a Compact Time Interval

The objective of this section is to use Lemma 7.6 to conclude that the rates of the stochastic system are close to those of the fluid model. As a consequence of the convergence to the fluid limit in Theorem 8.9 being uniform on compact sets, and not uniform on $\mathbb{R}^+$ we will only be able to show that the rates are close over a compact time interval.

Theorem 7.7. *Suppose there exists t_0 such that*

$$M^{-1}\dot{\bar{T}}(t) \equiv R \quad \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}| \tag{7.38}$$

for any fluid model trajectory $\bar{\mathfrak{X}}(t)$ with limiting threshold $\bar{h} > 0$ satisfying (7.5) - (7.17). Then, for any positive $\gamma < 1$ and $\zeta < 1$, there exists $L_1(\zeta, \gamma)$ such that for all $n \geq L_1$,

$$\inf_{|y| \leq \zeta} \mathrm{E}\left[M^{-1}T^{n,y}(nt_0)\right] \geq R(1-\zeta)(1-\gamma)nt_0. \tag{7.39}$$

(Note that we write $T^{n,y}$ in place of $T^{\mathbf{n}}$ to give extra emphasis on the dependence on y.)

Proof. Let $\{\mathbf{n}_l\} = \{(n_l, y_l)\}$ be a sequence of system scale and relative initial condition pairs satisfying $n_l \to \infty$ and $|y_l| \leq \zeta$. We invoke the TFL case of Lemma 7.6 by picking $\mathfrak{F}$ so that

$$\mathfrak{F} \circ [\bar{\mathfrak{X}}(\cdot)](t) := \bar{T}(\zeta^{-1}t) - \bar{T}(t) - MR(\zeta^{-1} - 1)t.$$

$\mathfrak{F}$ is easily seen to be continuous on the topology of uniform convergence on compact sets. Also note that

$$\mathfrak{F} \circ [\bar{\mathfrak{X}}(\cdot)](t) = 0 \qquad \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}|$$

by (7.38), and thus Lemma 7.6 yields:

$$\lim_{l\to\infty} \left| \frac{T^{\mathbf{n}_l}(n_l t_0) - T^{\mathbf{n}_l}(\zeta n_l t_0)}{n_l(1-\zeta)t_0} - MR \right| = 0 \quad \text{a.s.}, \tag{7.40}$$

where we have used the fact that $\langle \mathbf{n}_l \rangle_\zeta = n_l$. The left hand side of (7.40) is bounded from above by a constant for all l, and thus by the dominated convergence theorem [50]:

$$\lim_{l\to\infty} \mathrm{E}\left| \frac{T^{\mathbf{n}_l}(n_l t_0) - T^{\mathbf{n}_l}(\zeta n_l t_0)}{n_l(1-\zeta)t_0} - MR \right| = 0. \tag{7.41}$$

Also note (7.41) holds for any sequence $\{\mathbf{n}_l\}$ with $\langle \mathbf{n}_l \rangle \to \infty$ and $|y_l| \leq \zeta$, because these were the only restrictions for our initial choice of sequence.

Now pick a positive constant $\gamma < 1$. Observe that there exists a constant $L_1(\gamma, \zeta)$ such that whenever $n > L_1$,

$$\inf_{|y| \leq \zeta} \frac{\mathrm{E}\left[T^{\mathbf{n}}(nt_0) - T^{\mathbf{n}}(n\zeta t_0)\right]}{n(1-\zeta)t_0} \geq MR(1-\gamma)$$

for if otherwise we could construct a sequence $\{\mathbf{n}_l\}$ that violates (7.41). By the monotonicity of $T^{\mathbf{n}_l}(\cdot)$, we have (7.39). □

7.5 Stochastic System Attracted to Fluid Equilibrium

The objective of this section is to show that the stochastic system is attracted to the fluid model equilibrium. In particular we show that the expected norm of the state declines geometrically for starting states outside a neighborhood of the equilibrium. We also show that the expected norm of the state is small, some fixed time after starting inside the neighborhood. The proof technique is similar that of Theorem 3.1 of Dai [32].

Before stating the theorem of this section, we define the following notation:

$$Y^{n,y}(t) \triangleq \frac{1}{n}(X^{\mathbf{n}}(nt) - nh\mathbf{e}). \tag{7.42}$$

$Y^{n,y}(t)$ is a "relative" state of system-n, that is it is re-centered so that the candidate equilibrium is at the origin and is scaled by $1/n$ in size, and by n in time. This transformation parallels our definition of relative initial condition y in (5.2).

Theorem 7.8. *Suppose that there exists t_0 such that*

$$\bar{X}(t) \equiv \bar{h}\mathbf{e} \;\; \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}| \tag{7.43}$$

for any fluid model trajectory $\mathfrak{X}(t)$ with limiting threshold $\bar{h} \geq 0$ satisfying equations (7.6) - (7.17). Then the following conclusions are true:

i) For any $\zeta > 0$, and any positive $\delta < 1$ there exists $L_2(\zeta, \delta)$ such that for all $n \geq \zeta^{-1} L_2$ and $|y| > \zeta$,

$$\mathrm{E}\,|Y^{n,y}(t_0|y|)| \leq \delta |y|, \tag{7.44}$$

ii) For any $\zeta > 0$, and any $b > 0$ there exits $L_3(\zeta, b)$ such that for all $n \geq L_3$ and all $|y| \leq \zeta$,

$$\mathrm{E}\,|Y^{n,y}(t_0)| \leq b. \tag{7.45}$$

Proof. We first prove conclusion (i). Pick any sequence of pairs $\{\mathbf{n}_l\} = \{(n_l, y_l)\}$ satisfying $n_l|y_l| \to \infty$ and $|y_l| > \zeta$ for some $\zeta > 0$ (A JFL sequence). We invoke Lemma 7.6, picking $\mathfrak{F}$ such that

$$\bar{F}(t) \triangleq \mathfrak{F} \circ \left[\bar{X}(\cdot); \bar{T}(\cdot); \bar{\Lambda}(\cdot); \bar{h}\right](t) := \bar{X}(t) - \bar{h}\mathbf{e}.$$

Note that $\bar{F}(|\bar{X}(0) - \bar{h}e|t) = 0\,\forall t \geq t_0$ by Assumption (7.43), and $\mathfrak{F}$ is easily seen to be continuous on the topology of uniform convergence on compact sets. Lemma 7.6 yields

$$\frac{1}{\langle \mathbf{n}_l \rangle} |X^{\mathbf{n}_l}(\langle \mathbf{n}_l \rangle t) - n_l h \mathbf{e}| \to 0 \text{ a.s.}$$

for each $t \geq t_0$. Noting that $\langle \mathbf{n}_l \rangle = n_l|y_l|$, and taking $t = t_0$ we have that

$$\frac{1}{n_l|y_l|} |X^{\mathbf{n}_l}(n_l|y_l|t_0) - n_l h \mathbf{e}| \to 0 \text{ a.s.}$$

By Lemma 7.3, $(n_l|y_l|)^{-1} X^{\mathbf{n}_l}(n_l|y_l|t_0)$ is uniformly integrable. Therefore

$$\lim_{l \to \infty} \frac{1}{n_l|y_l|} \mathrm{E}\,|X^{\mathbf{n}_l}(n_l|y_l|t_0) - n_l h \mathbf{e}| = 0.$$

Applying definition (7.42), and noting that our initial choice of sequence $\{\mathbf{n}_l\}$ was arbitrary, up to some constraints, we have that the following statement is true:

a) For any $\zeta > 0$, and any sequence $\{\mathbf{n}_l\}$ with $n_l|y_l| \to \infty$, and $|y_l| > \zeta$,

$$\lim_{l\to\infty} \frac{1}{|y_l|} \mathrm{E}\,|Y^{n_l,y_l}(t_0|y_l|)| = 0. \tag{7.46}$$

We claim that this fact implies the following statement is true:

b) For any $\zeta > 0$, and any positive $\delta < 1$ there exists $L_2(\zeta, \delta)$ such that for all $n|y| \geq L_2$ and $|y| > \zeta$,

$$\frac{1}{|y|} \mathrm{E}\,|Y^{n,y}(t_0|y|)| \leq \delta. \tag{7.47}$$

Suppose statement (b) were not true. Then for some $\zeta > 0$ and some positive δ, we would have that for any L_2 there would exist a pair (n, y) with $n|y| > L_2$ and $|y| > \zeta$ with $\frac{1}{|y|}\mathrm{E}\,|Y^{n,y}(t_0|y|)| > \delta$. We therefore could construct a sequence that violates statement (a), which is a contradiction. A special case of (b) is when $n > L_2\zeta^{-1}$ and $|y| > \zeta$. Hence we have conclusion (i) of the lemma.

We now turn to showing conclusion (ii). Pick an arbitrary sequence of pairs $\{\mathbf{n}_l\} = \{(n_l, y_l)\}$ satisfying $n_l \to \infty$ and $|y_l| \leq \zeta$ for some constant ζ (A TFL sequence). We again invoke Lemma 7.6 by taking $\mathfrak{F}$ to be the same functional as before,

$$\bar{F}(t) \triangleq \mathfrak{F} \circ \left[\bar{X}(\cdot); \bar{T}(\cdot); \bar{\Lambda}(\cdot); \bar{h}\right](t) := \bar{X}(t) - \bar{h}\mathbf{e}.$$

Using Lemma 7.6, and the fact that $\langle \mathbf{n}_l \rangle_\zeta = n_l$ when $|y_l| \leq \zeta$ we have

$$\frac{1}{n_l}|X^{\mathbf{n}_l}(n_l t) - n_l h\mathbf{e}| \to 0 \text{ a.s.}$$

for each $t \geq \zeta t_0$. Now take $t = t_0$,

$$\frac{1}{n_l}\left|X^{\mathbf{n}_l}(n_l t_0) - n_l h\mathbf{e}\right| \to 0 \text{ a.s.}$$

By Lemma 7.3, $(n_l)^{-1}X^{\mathbf{n}_l}(n_l t_0)$ is uniformly integrable. Therefore

$$\lim_{l\to\infty}\frac{1}{n_l}\mathrm{E}\left|X^{\mathbf{n}_l}(n_l t_0) - n_l h\mathbf{e}\right| = 0.$$

Applying definition (7.42),and noting that our initial choice of sequence $\{\mathbf{n}_l\}$ was arbitrary, up to some constraints, we have that the following statement is true:

c) For any $\zeta > 0$, and any sequence $\{\mathbf{n}_l\}$ with $n_l \to \infty$, and $|y_l| \leq \zeta$,

$$\lim_{l\to\infty}\mathrm{E}\left|Y^{n_l,y_l}(t_0)\right| = 0. \tag{7.48}$$

We claim that fact (c) implies conclusion (ii). Suppose (ii) were not true. Then for some choice ζ and b, we would have that for every constant L_3, there would exist an $n \geq L_3$ and $y \leq \zeta$ satisfying $\mathrm{E}\left|Y^{n,y}(t_0)\right| > b$. This would allow us to construct a sequence that violates statement (c), which is a contradiction. □

7.6 Hitting Times on a Neighborhood of the Fluid Equilibrium

The objective of this section is to show that the results of Theorem 7.8 imply that the expected return time of the ζ ball around the fluid equilibrium is small. Later on we will combine this with the results of Theorem 7.7 that show that the expected rates, when starting from within the ball are close to the fluid rates on a compact time interval, to show the long term rates are close to the fluid rates.

The proof of Lemma 7.9 is adapted from the proof of Theorem 2.1(ii) of [48], which was for a discrete time Markov chain.

In the proof of Lemma 7.9, and in subsequent proofs, when we want to express $Y^{n,y}(t)$ without specifying an initial condition we will write $Y^n(t)$ where the choice of initial condition is implicit by the choice of probability measure. We define P_y to be a probability measure for which

$$\mathrm{P}_y\{Y^n(0) = y\} = 1,$$

and thus,

$$Y^n(t) = Y^{n,y}(t) \quad P_y\text{-a.s.}, \qquad \text{and} \qquad \mathrm{E}_y[Y^n(t)] = \mathrm{E}[Y^{n,y}(t)].$$

Lemma 7.9. *If for n fixed, we have the following inequalities*

$$\mathrm{E}_y|Y^n(t_0|y|)| \leq \delta|y| \quad \textit{for all } |y| > \zeta, \tag{7.49}$$

$$\mathrm{E}_y|Y^n(t_0)| \quad \leq b \qquad \textit{for all } |y| \leq \zeta, \tag{7.50}$$

then Y^n is positive Harris recurrent and furthermore,

$$\sup_{y \in B} \mathrm{E}_y[\tau_B^n(t_0)] \leq t_0 \left[1 + \frac{\zeta + b}{1 - \delta}\right] \tag{7.51}$$

where $B \triangleq \{y : |y| \leq \zeta\}$ and $\tau_B^n(t_0)$ is defined by

$$\tau_B^n(t_0) \triangleq \inf\{t \geq t_0 : Y^n(t) \in B\}. \tag{7.52}$$

Proof. That Y^n is positive Harris recurrent follows directly from Theorem 3.1 of [32].

The rest of the argument that follows is adapted from the proof of Theorem 2.1(ii) of Meyn and Tweedie [48]. We will use the following Fact taken from Theorem 14.2.2 of [49]:

Fact 1: (Meyn and Tweedie [49]) Suppose a discrete time Markov chain $\hat{\boldsymbol{\Phi}} = \{\hat{\Phi}_k,\, k \in \mathbb{Z}^+\}$ is defined on a general state space X with transition kernel $\hat{\mathrm{P}}(x, A) = \mathrm{P}(\hat{\Phi}x \in A)$, where $A \in \mathfrak{B}(\mathrm{X})$, the Borel subsets of X. If V and f are nonnegative measurable functions satisfying

$$\int \hat{\mathrm{P}}(x, dy)V(y) \le V(x) - f(x) + \tilde{b}1_B(x), \qquad x \in \mathrm{X} \tag{7.53}$$

then

$$\mathrm{E}_x\left[\sum_{k=0}^{\hat{\tau}_B - 1} f(\Phi_k)\right] \le V(x) + \tilde{b}$$

where

$$\hat{\tau}_B = \inf\{k \ge 1 : \Phi_k \in B\}.$$

We now turn to setting up our problem to make use of Fact 1. We define the following two mappings, the first mapping each y to a time $n(y)$, and the second mapping each y to a integer valued Lyapunov function $V(y)$:

$$n(y) \triangleq \begin{cases} |y|t_0 & \text{if } |y| > \zeta \\ t_0 & \text{if } |y| \le \zeta \end{cases} \tag{7.54}$$

$$V(y) \triangleq \frac{t_0}{1-\delta}|y| \tag{7.55}$$

Substituting our assignment of $n(y)$ into relation (7.49), and adding a term to that relation's right hand side so that it holds for y both inside and outside B, we have

$$\begin{aligned} &\mathrm{E}_y|Y^n(n(y))| \\ &\qquad \le \delta|y| + \left(\sup_{y\in B} \mathrm{E}_y|Y^n(t_0)|\right) 1_B(y) \\ &\qquad \le |y| - \frac{1-\delta}{t_0} n(y) + (1 - \delta + b)\, 1_B(y). \end{aligned}$$

We multiply both sides by $t_0/(1-\delta)$ to get

$$\mathrm{E}_y|V\left(Y^n(n(y))\right)| \le V(y) - n(y) + \tilde{b}1_B(y) \tag{7.56}$$

$$\text{where} \qquad \tilde{b} = t_0 + \frac{t_0}{1-\delta}b. \tag{7.57}$$

The transition kernel P^t for the Markov process Y^n is defined by

$$\mathrm{P}^t(y, A) = \mathrm{P}_y(Y^n(t) \in A),$$

where A is any set in $\mathfrak{B}(\mathrm{Y})$, the Borel subsets of the state space Y. We define the discrete time "embedded" Markov chain $\hat{\boldsymbol{\Phi}} = \{\hat{\Phi}_k, k \in \mathbb{Z}_+\}$ with transition kernel $\hat{\mathrm{P}}$ given by

$$\hat{\mathrm{P}}(y, A) = \mathrm{P}^{n(y)}(y, A).$$

Note that

$$\int \hat{\mathrm{P}}(y, dz)V(z) = \int \mathrm{P}^{n(y)}(y, dz)V(z) = \mathrm{E}|V\left(Y(n(y))\right)|. \tag{7.58}$$

Combining (7.58) with (7.56) we have

$$\int \hat{\mathrm{P}}(y, dz)V(z) \le V(y) - n(y) + \tilde{b}1_B(y). \tag{7.59}$$

Recognizing this is the form of expression (7.53), we may use Fact 1 to conclude

$$\mathrm{E}_y\left[\sum_{k=0}^{\hat{\tau}_B - 1} n(\bar{\Phi}_k)\right] \le V(y) + \tilde{b}, \tag{7.60}$$

where $\hat{\tau}_B = \inf\{k \ge 1 : \Phi_k \in B\}$ is the first return time of the embedded discrete time chain $\boldsymbol{\Phi}$ to the set B. Fix a particular ω and initial condition y. If the embedded chain hits B in $\hat{\tau}_B$ discrete steps, then the original chain must also hit B in a time equal to the sum of the embedded times that those discrete steps correspond to. It is also possible that the original chain hits B earlier, in addition to hitting at a time equal to the sum of the embedded times. Thus the first hitting time of the original chain satisfies

$$\inf\{t \geq 0 : Y^n(t) \in B\} \leq \sum_{k=0}^{\hat{\tau}_B - 1} n(\hat{\Phi}_k) \ \mathrm{P}_y\text{-a.s.}$$

for each $y \in \mathrm{Y}$. Furthermore, whenever the initial condition $y \in B$, the first embedded time is t_0 seconds by (7.54). Consequently, the time of the first hitting of B after t_0 seconds expire satisfies

$$\inf\{t \geq t_0 : Y^n(t) \in B\} \leq \sum_{k=0}^{\hat{\tau}_B - 1} n(\hat{\Phi}_k) \quad \mathrm{P}_y\text{-a.s.}$$

for each $y \in B$. Substituting (7.52), taking the expectation, and using (7.60), we have

$$\mathrm{E}_y[\tau_B^n(t_0)] \leq V(y) + \tilde{b} \text{ for all } y \in B.$$

Taking the $\sup_{y \in B}$ of both sides, substituting (7.57) and (7.55) we have

$$\sup_{y \in B} \mathrm{E}_y[\tau_B^n(t_0)] \leq t_0 + \frac{t_0}{1-\delta} \left[\zeta + b\right],$$

which is (7.51). □

7.7 Convergence of Long Term Rates

The objective of this section is to tie together all of the preceding results to conclude that the long-term rates of the stochastic system are close to the fluid rates for large enough n. We pick n large enough so that the stochastic system's rates are close to the fluid rates for the first t_0 seconds after having started in a ζ neighborhood and that the expected time of first return to a ζ neighborhood, t_0 seconds after having started there is close to t_0. At this point, intuition suggests that the long term rates must be close to the fluid rates over the long term. We formalize that intuition with an argument based on stopping times and the strong Markov property.

Theorem 7.10. *Suppose both of the following are true:*

- *For any fluid model trajectory $\bar{\mathfrak{X}}(t)$ with limiting threshold $\bar{h} \geq 0$ that satisfies (7.6) - (7.16),*

$$\bar{X}(t) \equiv \bar{h}\mathbf{e} \qquad \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}|. \tag{7.61}$$

- *For any fluid model trajectory $\bar{\mathfrak{X}}(t)$ with limiting threshold $\bar{h} > 0$ satisfying (7.6) - (7.16):*

$$M^{-1}\dot{\bar{T}}(t) \equiv R \qquad \forall t \geq t_0|\bar{X}(0) - \bar{h}\mathbf{e}| \tag{7.62}$$

  where R is a constant K dimensional vector of flow rates.

Then for any $\epsilon > 0$, there exists a system-scale n_c such that for all system-scales $n \geq n_c$

$$\lim_{t\to\infty} \frac{D^{n,y}(t)}{t} \geq (1-\epsilon)R \quad a.s.$$

Proof. We observe that equations (7.62) and (7.61) are the necessary conditions to apply Theorems 7.7 and 7.8 respectively. Therefore, we may arbitrarily pick the constants ζ, δ, and b of Theorem 7.8 and the constants ζ and γ of Theorem 7.7 (using the same ζ value in Theorems 7.7 as we use when we apply Theorem 7.8), and then fix an n satisfying

$$n > \max[L_1(\zeta,\gamma),\ \zeta^{-1}L_2(\zeta,\delta),\ L_3(\zeta,b)] \tag{7.63}$$

so that the conclusions of both Theorems 7.8 and 7.7 hold.

In addition, conclusions (i) and (ii) of Theorem 7.8 allow us to invoke Lemma 7.9 to conclude

$$\sup_{y \in B} \mathrm{E}_y[\tau_B^n(t_0)] \leq t_0(1 + \frac{\zeta + b}{1 - \delta}) \tag{7.64}$$

where $\tau_B^n(t_0)$ is defined by (7.52). Because the constants ζ, b, δ could have been selected to be arbitrarily small, equations (7.64) and (7.63) imply that the expected first hitting time of B, t_0 seconds after having started in B, can be made to be arbitrarily close to t_0 by choosing n large enough. For convenience we collect some of the constants in (7.64) in the term t_0' defined by

$$t_0' = t_0 \left[1 + \frac{\zeta + b}{1 - \delta} \right]. \tag{7.65}$$

We have also chosen n large enough so that the following conclusion from Theorem 7.7 holds,

$$\inf_{|y| \leq \zeta} \mathrm{E}\left[T^{n,y}(nt_0) \right] \geq MR(1 - \zeta)(1 - \gamma)nt_0. \tag{7.66}$$

Define the stopping times

$$\begin{aligned} \sigma_0 &= 0, \\ \sigma_{i+1} &= \inf\{t \geq t_0 + \sigma_i : Y(t) \in B\}, \ \ \forall i \geq 0. \end{aligned} \tag{7.67}$$

Figure 7.1 illustrates how these stopping times are defined. Note that for any initial condition $y \in \mathrm{Y}$ (the state space of Y^n) and index $i \geq 1$,

$$\mathrm{E}_y[\sigma_{i+1} - \sigma_i] = \mathrm{E}_{Y^{n,y}(\sigma_i)}[\tau_B^n(t_0)] \leq \sup_{\tilde{y} \in B} \mathrm{E}_{\tilde{y}}[\tau_B^n(t_0)] \leq t_0'. \tag{7.68}$$

This follows from the fact that $Y^{n,y}(\sigma_i) \in B$, the strong Markov property, the stopping time definitions (7.52) & (7.67), and expressions (7.64) & (7.65). Also, Y^n is positive

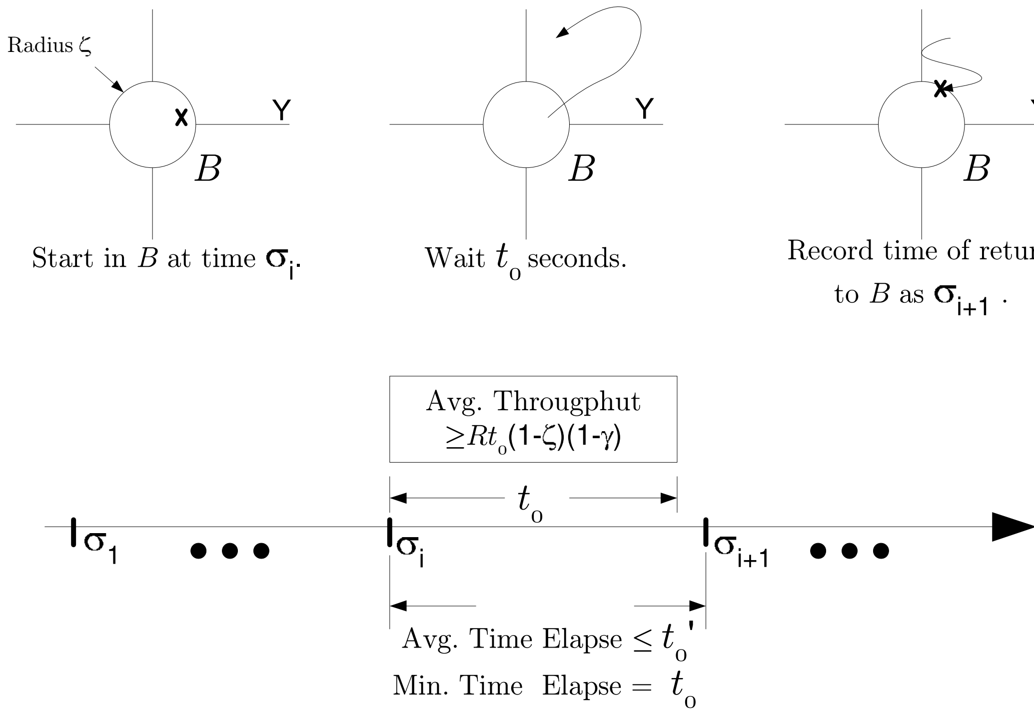

Figure 7.1: The top half of the figure illustrates the definition of the stopping times $\sigma_i, \sigma_{i+1}, ...$ The bottom half illustrates the intuition behind the proof of Theorem 7.10 by plotting the stopping times on a time line, and showing the bound on expected throughput between such stopping times. That the expected time elapse between σ_i and σ_{i+1} is upper-bounded by t_0' is a consequence of Lemma 7.9. The lower-bound on expected throughput between stopping times comes from Theorem 7.7.

Harris recurrent by Lemma 7.9 and therefore,

$$\mathrm{E}_y[\sigma_1] < \infty \tag{7.69}$$

for any $y \in \mathrm{Y}$. We define a counting process $N(t)$ for the stopping times σ_i as

$$N(t) = \inf\{i : \sigma_i \leq t\}.$$

Because Y^n is positive Harris recurrent, $\sigma_i < \infty$ almost surely, and therefore

$$N(t) \to \infty \qquad \text{a.s.} \tag{7.70}$$

We now turn to bounding the expected "arrival" rate of the stopping times σ_i. By (7.68) for each i,

$$\frac{\mathrm{E}_y[\sigma_i]}{i} = \frac{\sum_{j=1}^{i-1} \mathrm{E}_y[\sigma_{j+1} - \sigma_j] + \mathrm{E}_y\sigma_1}{i} \leq t_0'(1 + 1/i) + \frac{\mathrm{E}_y\sigma_1}{i} \tag{7.71}$$

Additionally, along any sample path

$$\frac{t}{N(t)} \leq \frac{\sigma_{N(t)+1}}{N(t)+1}\frac{N(t)+1}{N(t)}.$$

Thus by taking $\liminf_{t\to\infty} E_y(\cdot)$ of both sides, and using (7.70), and (7.71) we have

$$\liminf_{t\to\infty} \mathrm{E}_y\left[\frac{t}{N(t)}\right] \leq t_0'.$$

Also, by Fatou's Lemma

$$\mathrm{E}_y\left[\liminf_{t\to\infty} \frac{t}{N(t)}\right] \leq \liminf_{t\to\infty} \mathrm{E}_y\left[\frac{t}{N(t)}\right] \leq t_0'. \tag{7.72}$$

Recall that the process $T^{\mathbf{n}}(t) = T^{n,y}(t)$ is defined in terms of the time scale of X^n, and not that of Y^n, which we defined in expression (7.42) by scaling time by a factor of n. Therefore we define a re-scaled service process $\widetilde{T}^{n,y}$ with time re-scaled to match the time scale of the Process $Y(t)$ according to the definition

$$\widetilde{T}^{n,y}(a,b) = T^{n,y}(nb) - T^{n,y}(na). \tag{7.73}$$

We also define the random vectors ρ_i to track the throughput between stopping times σ_i as made precise by the definition

$$\rho_i = M^{-1}\widetilde{T}^{n}(\sigma_i, \sigma_i + \sigma_{i+1}), \tag{7.74}$$

where the y superscript in $\widetilde{T}^{n,y}$ is omitted because the initial condition will be specified implicitly by the choice of probability measure. Note that for $i \geq 1$ and each $y \in \mathrm{Y}$,

$$\begin{aligned} \mathrm{E}_y[\rho_i] &\geq \mathrm{E}_{Y^{n,y}(\sigma_i)}[M^{-1}\widetilde{T}^n(t_0)] \\ &\geq \inf_{\acute{y} \in B} \mathrm{E}_{\acute{y}}[M^{-1}\widetilde{T}^n(t_0)] \\ &\geq R t_0 (1-\zeta)(1-\gamma). \end{aligned} \tag{7.75}$$

This follows from the fact that $Y^{n,y}(\sigma_i) \in B$, the strong Markov property, the definition of σ_i (7.67), the definition of ρ_i (7.74), and relation (7.66). Figure 7.1 illustrates the fact that the throughput between stopping times σ_i and σ_{i+1} is lower-bounded according to relation (7.75).

Because, we have shown that Y^n is positive Harris recurrent, by [33] the following ergodic property holds for every measurable f on Y with $\pi(|f|) < \infty$,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t f(Y^n(s))ds = \pi(f) \;\; \mathrm{P}_y\text{-a.s. for each } y \in \mathrm{Y}$$

where π is the unique invariant distribution of Y^n. Assigning the function

$$f(y) := M^{-1}\dot{\widetilde{T}}^{n,y}(0)$$

to be the instantaneous service rates when the process is in state y, (Recall that we assumed the service rates are a function of the state in Section 5.7.2.) we have,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t f(Y^{n,y}(s))ds = \lim_{t\to\infty} \frac{1}{t} M^{-1}\widetilde{T}^{n,y}(t) = \mathcal{R} \text{ a.s.} \tag{7.76}$$

for some constant vector $\mathcal{R}$.

Consider the random variable

$$\mathcal{N} = \liminf_{t\to\infty} \frac{t}{N(t)}.$$

The random variable $\mathcal{N}$ is a P_π invariant random variable, and therefore is a constant. Hence by (7.72),

$$\mathcal{N} \text{ is a constant and } \mathcal{N} \leq t_0'. \tag{7.77}$$

A more detailed explanation of this argument is provided in the Appendix.

We observe that for any sample path the following inequalities hold,

$$\frac{t}{N(t)} \frac{M^{-1}\widetilde{T}^{n,y}(t)}{t} \leq \frac{\sum_{j=0}^{N(t)} \rho_j}{N(t)} \leq \frac{t}{N(t)} \frac{\sigma_{N(t)+1}}{t} \frac{M^{-1}\widetilde{T}^{n,y}(\sigma_{N(t)+1})}{\sigma_{N(t)+1}}. \tag{7.78}$$

Taking the $\liminf_{t\to\infty}$ of both sides, and using (7.76), (7.70), & (7.77) we have that

$$\liminf_{t\to\infty} \frac{\sum_{j=0}^{N(t)} \rho_j}{N(t)} = \mathcal{N}\mathcal{R} \text{ a.s.} \tag{7.79}$$

Now we need to apply the dominated convergence theorem. Before doing so, we note that

$$M^{-1}\frac{\widetilde{T}^{n,y}(\sigma_{N(t)+1})}{\sigma_{N(t)+1}} \leq M^{-1}\mathbf{1}$$

where $\mathbf{1}$ is a column vector of 1's of appropriate dimension. This fact combined with (7.78) and (7.77) yield that for each $i > 0$,

$$\inf_{k\geq i} \frac{\sum_{j=1}^{k} \rho_j}{i} \leq \liminf_{t\to\infty} \frac{t}{N(t)} M^{-1}\mathbf{1} \leq t_0(1+\phi)M^{-1}\mathbf{1}.$$

and thus the random variables

$$\left\{ \inf_{k\geq i} \frac{\sum_{j=1}^{k} \rho_j}{i} : i > 0 \right\}$$

are dominated by a constant. Consequently, the dominated convergence theorem applied to (7.79) yields,

$$\liminf_{i\to\infty} \mathrm{E}\left[\frac{\sum_{j=1}^{i} \rho_j}{i}\right] = \mathcal{N}\mathcal{R}.$$

Also for each $i > 0$ by (7.75),

$$\mathrm{E}\left[\frac{\sum_{j=1}^{i}\rho_j}{i}\right] \geq Rt_0(1-\gamma)(1-\zeta).$$

Thus,

$$\mathcal{N}\mathcal{R} \geq Rt_0(1-\gamma)(1-\zeta).$$

Substituting (7.72) we have that

$$\mathcal{R} \geq \frac{(1-\gamma)(1-\zeta)t_0}{t_0'}R.$$

Which by (7.65), (7.73), and (7.76) implies

$$\lim_{t\to\infty}\frac{1}{t}M^{-1}T^{n,y}(t) \geq \frac{(1-\gamma)(1-\zeta)}{1+\frac{\zeta+b}{1-\delta}}R \quad \text{a.s.}$$

Recall that the parameters γ, ζ, b, and δ may be chosen arbitrarily close to 0, and still have all of the preceding development hold by choosing a large enough n according to (7.63). Thus for a large enough n,

$$\lim_{t\to\infty}\frac{1}{t}M^{-1}T^{n,y}(t) \geq (1-\epsilon)R \quad \text{a.s.}$$

for any relative initial condition y. By the strong law of large numbers for renewal processes [50],

$$\lim_{t\to\infty}\frac{1}{t}S_k^{n,y}(t) \to m_k \quad \text{a.s.}$$

Thus by (5.9),

$$\lim_{t\to\infty}\frac{1}{t}D^{n,y}(t) \geq (1-\epsilon)R \text{ a.s.}$$

□

Chapter 8

Round Robin Network without Loops

In this section we specialize and consider networks whose flows traverse the network without splitting and without visiting the same class more than once. For such networks the routing matrix P is binary and nilpotent, because packets must leave the network in a bounded number of hops.

Our objective in this section is to show that, starting from any initial condition, the fluid model of such a network converges to an equilibrium state $\bar{h}\mathbf{e}$, in a time proportional to the distance of the initial condition from the equilibrium. Furthermore, we need to show that after reaching equilibrium, the departure rates for each flow are max-min fair. This will allow us to invoke Theorem 7.10 to conclude that the flows in the stochastic model of the network achieve close to the max-min fair share rates over the long term.

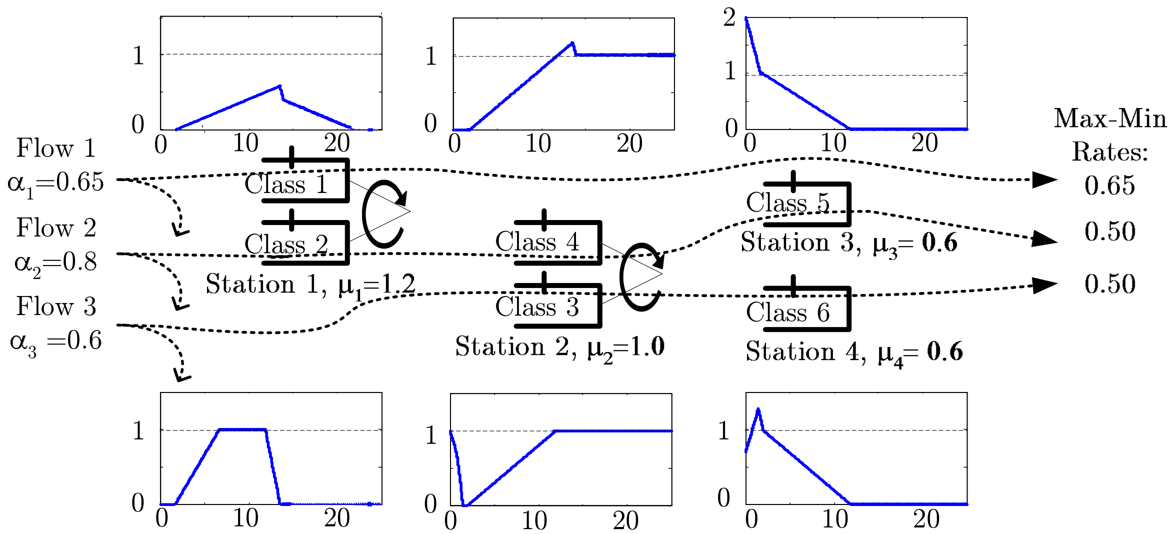

Figure 8.1: A fluid model trajectory of the example network introduced in Section 5.4.

8.1 Fluid Model Example.

To motivate the techniques we use in our proof, we consider the fluid model of the example network that we introduced in Section 5.4. We consider the TFL fluid limit of the example network, found by setting the base thresholds h to 1, and then scaling the thresholds, time, and space by $n \to \infty$. A trajectory of the resulting fluid model is illustrated in Figure 8.1. The illustrated trajectory has initial condition $\bar{Q}_2(0) = 1$, $\bar{Q}_5(0) = 2$, $\bar{Q}_6(0) = 0.7$, and all other queues starting from 0. This initial condition matches the initial conditions of the stochastic model trajectories pictured in Figure 5.1, but scaled in magnitude in proportion to the threshold size. Looking at the $h = 10$ and $h = 100$ cases of Figure 5.1, and then looking at the fluid model trajectory of Figure 8.1, we see that the trajectory of the stochastic system with larger thresholds looks more like the trajectory of the fluid model.

As we discussed in Section 5.4, we expect station 2 to be the primary bottleneck, and therefore its queues should fill to their thresholds. Afterwards, we should expect the fluid system to enter a *sliding mode*, (a term we introduced in Section 7.2.1) where the queues in the bottleneck station stick to their thresholds. Consequently the policing must throttle the arrivals so that the rates of the thinned arrival process match the service rate of the bottleneck station. Recall that in Section 7.2.1 we termed this phenomenon *sliding mode policing*.

Looking at figure 8.1, we indeed see that the queues at the station 2 do indeed fill up and eventually stick to their thresholds. After this happens, the rates of the fluid system exactly match the max-min fair share rates. However, the transient trajectory of the queues is rather complicated before the system reaches this final state. In order to apply Theorem 7.10, we need to be certain that the fluid model converges to the desired state from *any* initial condition, and in a time proportional to the distance of that initial condition from the equilibrium state. Therefore, we need a systematic argument that shows that for all initial conditions, and for all loop free round-robin networks, this convergence happens.

A natural approach is to find a Lyapunov function on the fluid model state space that declines to 0. To construct such a Lyapunov function, we need to find a quantity that declines monotonically along any fluid model trajectory. Intuition suggests that the queues downstream of the main bottleneck, station 2, should decline. However we see from the plots in Figure 8.1 that the class 6 queue actually increases for a time. This is explained as follows. The class 5 queue (queue 5 for short) starts above the policing threshold, and thus flow 2 starts off policed. Consequently, the rate of flow through queue 2 and queue 4 is zero in the beginning. As a result, station 2 devotes all of its service rate to queue 3, allowing queue 3 to drain into queue 6 at rate 1.0. This is faster than the service rate of queue 6, resulting in its level rising, and even causing flow 3 to be policed for a short time. This example shows that the queues downstream of the bottleneck do not monotonically decrease to 0.

However, we claim that there is a quantity that does decrease monotonically to zero, and that is the sum of all the queues downstream from any queue at the bottleneck station. In this case, that quantity is the sum of queues 5 and 6. To see this consider the rate of fluid being expelled from queue 5. Whenever queue 5 is nonempty, fluid leaves at rate 0.6. When queue 5 is empty, but an upstream queue is nonempty, fluid arrives to queue 5 at at least the bottle neck rate of 0.5, and as this is a fluid network, there is no delay in the fluid arriving from the upstream queue to queue 5. Because the service capacity of queue 5 exceeds 0.5, queue 5 expels fluid at a rate of at least 0.5 whenever an upstream queue is nonempty. Also, if queue 5 and all the other flow 2 queues are empty, flow 2 is not policed and thus its fluid arrives to the network at rate 0.8, a rate that exceeds the bottleneck rate. After passing through the bottleneck, the fluid arrival rate to queue 5 is at least 0.5, and this fluid arrives without delay. This exhausts all cases, so the rate of fluid leaving queue 5 is at least 0.5 at all times.

The same observations can be made about the fluid release from queue 6. We also see that the combined arrival rates to queue 5 and 6 can be no more than 1.0 because they are constricted by the service rate of station 2. We also note that because queue 5 and queue 6 can be served at a rate of 0.6 when they are nonempty fluid leaves them at a combined rate of at least 1.1 when at least one of them is nonempty. Thus their sum must fall monotonically to zero at a rate of at least 0.1. We will use this same strategy of considering the sum of the queues downstream of the bottleneck for the general proof we present later in the chapter.

We continue to look at the evolution of the example. Once the queues downstream of station 2 drop below their thresholds, fluid passes through the queue 3 and queue 4 at a rate of exactly 0.5. To see this, consider the possible cases. Either the queue 4 is nonempty, an upstream queue is nonempty, or flow 2 is not policed, or some combination of any of these. In all of these cases, there is either fluid in the queue 4 or fluid is arriving at faster than the bottleneck rate. The same can be said of queue 3. Because the queueing discipline is fair and work conserving, the only possibility is for each flow to be served at rate 0.5.

Once this has occurred, we should expect both queue 3 and queue 4 to rise monotonically towards their thresholds. Figure 8.1 shows that this indeed happens. When queue 4 reaches its threshold and activates policing, it continues to rise for a short time, before eventually falling back down to threshold. This is because the upstream queue has some fluid left in it. Eventually the fluid in the upstream queue empties, then queue 4 drains towards the threshold. When it reaches threshold, it enters a sliding mode. To see this, consider what would happen if the queue were to dip below threshold. There would be full release of the policing, causing a surge of fluid that would immediately push the queue back to its threshold. Thus, a sliding mode solution is the only possibility.

Similarly, queue 3 enters a sliding mode, but it does not "overshoot" the threshold as queue 4 does. This is because there are no queues upstream of queue 3. Once queues 3 and 4 settle to their equilibrium sliding modes, flows 2 and 3 have rates of 0.5, throughout the network. Consequently, station 1 has a remaining service capacity of $1.2 - 0.5 = 0.7$ to devote to the other flow it serves – flow 1. As flow 1's arrival rate is only 0.65, queue 1 drains to 0, and stays there. Flow 1's final rate is 0.65, and the governing constraint is its own rate of demand.

The example suggests the following strategy for proving the convergence of a general network:

- Identify the tightest bottleneck station in terms of rate per flow, or if the flows have different weights, the rate per unit weight of flow.

- Show that the sum of the queues downstream from the bottleneck empty, using the argument we gave for the example network.

- Once the downstream queues empty, the queues at the bottleneck station drain at constant rate. Furthermore, when the bottleneck queue is below threshold it must be rising monotonically. When it is above threshold, it may continue to rise as an upstream queue drains into it. However, the sum of the bottleneck queues and the upstream queues declines, because the flow is policed and new fluid is not being introduced. Using these two observations, we can construct a Lyapunov function of the form,

$$V = \sum [\text{Upstream Q's}] + \left([\text{Bottleneck Q}] - \bar{h}\right)^{+} + L\left([\text{Bottleneck Q}] - \bar{h}\right)^{-}.$$

Here L is a large enough number so that the when the bottleneck queue is below threshold, the effect of the bottleneck queue rising towards its threshold dominates any contribution of the other queues behavior. Once this Lyapunov function declines to 0 for each bottleneck queue at a bottleneck station, the bottleneck queues are all at their threshold $\bar{h}$. Also, the bottlenecked flows that use these queues all have been sliding-mode policed to rates that match the bottleneck rate.

- After the rates of the bottlenecked flows converge, they can be treated as constant rate, and can be removed from consideration by deducting their final rates from the capacity of the other stations through which the flows pass. Then one finds the most constrictive bottleneck from the "reduced network", and repeats the above analysis.

- At any point in this procedure, a flow's offered rate (its α) may be less then the most constrictive bottleneck station in the reduced network. In this case, one should proceed by considering the flow itself to be a bottleneck. One then shows that the queues that the flow passes through must drain to zero in finite time. One then removes the flow from consideration by deducting the flow's rate from the stations through which it passes, and then proceeds by finding the most constrictive bottleneck station or flow in the reduced network.

8.2 Pipeline Notation and Properties

In this section we will need to consider the queues that are either "upstream" or "downstream" of a particular queue in a systematic way. To that end, we define an ordered set of classes which we call a *Pipeline* . To make our definition precise, we define the following notation: We say that $k \prec l$ if flow $f = \mathrm{f}_{(k)}$ packets pass through the class k queue before passing through the class l queue, and $\mathbf{1}_j$ is a K dimensional column vector with a 1 at position j and zeros elsewhere.

Definition 1. *A* **Pipeline** *is an ordered set of classes* $\mathcal{P} = \{k_1 \prec k_2 \prec ... \prec k_m\}$ *such that*

$$P\mathbf{1}_{k_j} = \mathbf{1}_{k_{j+1}} \quad \forall j : 1 \leq j < m. \tag{8.1}$$

Because of the assumptions given in Section 5.1: that only class f receives exogenous arrivals for flow f, the routing matrix does not mix flows (i.e. $\mathrm{P}_{kl} = 0$ if $\mathrm{f}_{(k)} \neq \mathrm{f}_{(l)}$), the sum of any row of P is not more than 1; and because of the additional assumption that P is binary and nilpotent that we have made in this section, it is easy to verify that a pipeline $\mathcal{P}$ satisfies these additional properties:

$$[P\mathbf{1}_l]^\top \mathbf{1}_{k_j} = 0 \quad \forall l \notin \mathcal{P},\ \forall j : 2 \leq j \leq m \tag{8.2}$$

$$\alpha_{k_j} = 0 \quad \forall j : 2 \leq j \leq m. \tag{8.3}$$

It is also easy to verify that for any class k, there is a Pipeline $\mathcal{P}[k]$ that includes class k and all downstream queues. (i.e. $\mathcal{P} = \{k \prec k_2 \prec ... \prec k_m\}$ where $\mathrm{P}_{k_m,l} = 0$ for all l.) We call $\mathcal{P}[k]$ the pipeline *rooted* at k. Because flow f packets always enter as class f, $\mathcal{P}[f]$ is the pipeline containing all the classes that flow f passes through, and thus has all the elements of $K(f)$. In summary:

$$\mathcal{P}\,[k] = k \cup \text{Classes downstream of } k$$

$$\mathcal{P}\,[f] = \text{All classes of flow } f.$$

We also define:

$$\mathcal{H}[k] \triangleq \mathcal{P}[\mathrm{f}_{(k)}] \setminus \mathcal{P}[k], \qquad \text{Classes Upstream of } k$$

$$\mathcal{T}[k] \triangleq \mathcal{P}[k] \setminus k, \qquad \text{Classes Downstream of } k$$

where $\mathcal{H}$, $\mathcal{T}$ are mnemonics for Head and Tail respectively.

We state a series of lemmas to show additional properties of pipelines. In the lemmas, we use the notation

$$\mathcal{Q}(t, \mathcal{P}) \triangleq \sum_{k \in \mathcal{P}} \bar{Q}_k(t)$$

to represent the occupancy of the queues in $\mathcal{P}$. The following lemma is a law of conservation of flow for pipelines.

Lemma 8.1. *Consider a pipeline $\mathcal{P}$ with elements indexed as $\{k_1 \prec k_2 \prec ... \prec k_m\}$. Then for regular points t,*

$$\dot{\mathcal{Q}}(t, \mathcal{P}) = \dot{\bar{A}}_{k_1}(t) - \dot{\bar{D}}_{k_m}(t).$$

Proof. See the Appendix. □

The following lemma says that if fluid arrives to a pipeline at rate greater than r, and all the queues along that pipeline can offer at least rate r, then fluid passes through at a rate of at least r.

Lemma 8.2. *Suppose a pipeline $\mathcal{P}$ with elements indexed as $\{k_1 \prec k_2 \prec ... \prec k_m\}$ satisfies the following conditions for some r and τ, and for all regular points $t \geq \tau$:*

i) $\forall\, k_j \in \mathcal{P},\ \ \dot{\bar{D}}_{k_j}(t) \geq r$ *whenever* $\bar{Q}_{k_j}(t) > 0$.

ii) $\dot{\bar{A}}_{k_1}(t) \geq r$.

Then $\dot{\bar{D}}_{k_m}(t) \geq r$ *at regular points* $t \geq \tau$*.*

Proof. See the Appendix. □

The following lemma says that if all the queues along a pipeline have an available rate of r and one of those queues is nonempty, then fluid is expelled from the pipeline at a rate of at least r.

Lemma 8.3. *Suppose a pipeline* $\mathcal{P}$ *with elements indexed as* $\{k_1 \prec k_2 \prec ... \prec k_m\}$ *satisfies the following conditions for some rate* r*, and critical time* τ*:*

i) For all $k_j \in \mathcal{P}$ *and regular* $t \geq \tau$*,*

$$\dot{\bar{D}}_{k_j}(t) \geq r \quad \textit{whenever} \quad \bar{Q}_{k_j}(t) > 0.$$

ii) For some $k^* \in \mathcal{P}$ *and some time interval* $[s_0, s_1]$ *with* $s_0 \geq \tau$*,*

$$\bar{Q}_{k^*}(t) > 0 \quad \forall t \in [s_0, s_1].$$

Then $\dot{\bar{D}}_{k_m}(t) \geq r$ *at regular points* $t \in [s_0, s_1]$*.*

Proof. See the Appendix. □

8.3 Max-Min Fair Definitions

We say a vector of flow rates is $\mathbf{r}$ is *feasible* if

$$\sum_{f \in \mathbf{F}[i]} r_f \leq \mu_i$$

for each station i and $r_f \leq \alpha_f$ for each flow f. A vector of rates is $\mathbf{r}$ is *weighted max-min fair* if there is no flow rate r_f, which can be feasibly increased without decreasing the rate of some other flow f' for which $\frac{r_{f'}}{w_{f'}} \leq \frac{r_f}{w_f}$. In [28] it is shown that such a rate vector always exists.

Let $\mathbf{r}$ be a weighted max-min fair flow rate allocation. We say that the *normalized rate* of a flow f is the fair rate $\mathbf{r}_f$ divided by the weight w_f. We say a flow f *crosses* a station i if $f \in \mathbf{F}[i]$. A station i is *saturated* if its capacity constraint is tight in the max-min allocation (i.e. $\mu_i = \sum_{\mathbf{F}[i]} \mathbf{r}_f$). A *bottleneck station* of a flow f^* is any saturated station i that f^* crosses for which the normalized rate of f^* is greater or equal to the normalized rate of any other flow crossing that station. A flow has a *unique bottleneck* if it has only one bottleneck station.

If the normalized arrival rate α_{f^*} of flow f^* is less than or equal to the normalized rate any other flow f that crosses any station that f^* crosses, we say that flow f is *demand-limited*. Similarly, if the arrival rate α_{f^*} of flow f^* is strictly less than the normalized rate any other flow f that crosses any saturated station that f^* crosses, we say flow f is *strictly demand-limited*. It is straightforward to show [28] that a rate allocation vector is max-min fair if and only if every flow either has at least one bottleneck station or is demand limited.

In our proof that the long term rates of the stochastic network are close to the max-min fair rates for large thresholds, we will need to assume that each flow has either a unique bottleneck or is strictly demand limited. The following proposition shows that this is not a strong assumption.

Proposition 8.4. *Consider a network where the following are fixed: the arrival rate vector α, the flow weight vector w, and the routing. For almost all, with respect to*

Lesbesgue measure, choices of station capacity vector μ, each flow has either a unique bottleneck or is strictly demand limited in a max-min fair allocation.

Proof. See the Appendix. □

For the remainder of this chapter, we fix the K dimensional vector e by making the following assignment for each entry e_k,

$$e_k = \begin{cases} 1 & \text{if station } s(k) \text{ is a bottleneck of flow } \mathrm{f}_{(k)} \\ 0 & \text{otherwise.} \end{cases} \tag{8.4}$$

We also fix the full candidate equilibrium $\mathbf{e}$ by the assignment $\mathbf{e} := [e; 0_U; 0_V; 0_H]$ where 0_U, 0_V, 0_H are zero vectors of the same dimension as $U(t)$, $V(t)$, and $H(t)$ respectively.

8.4 Fluid Model Rate Lemmas

The following lemma supposes that some subset of the flows that cross a station i have been shown to converge to stable rates. The lemma concludes that the remaining flows of the station must share the remaining capacity in proportion to their weights.

Lemma 8.5. *Suppose that for some station i, some $\widetilde{\mathbf{C}}[i] \subseteq \mathbf{C}[i]$, and some $\widetilde{\mathbf{F}}[i] = \{f : f = f_{(k)}$ for some $k \in \widetilde{\mathbf{C}}[i]\}$, it has been shown that*

$$\dot{\bar{D}}_k(t) = \mu_i \dot{\bar{T}}_k(t) = r_{f_{(k)}} \;\; \forall k \in \mathbf{C}[i] - \widetilde{\mathbf{C}}[i]$$

for all $t \geq \tau$ for some critical time $\tau \geq \max(\bar{V}(0))$. Then for each $k^* \in \widetilde{\mathbf{C}}[i]$ and all regular $t \geq \tau$,

$$\dot{\bar{D}}_{k^*}(t) \geq w_{f_{(k^*)}} \frac{\tilde{\mu}_i}{\sum_{\widetilde{\mathbf{F}}[i]} w_f} \quad \textit{whenever } \bar{Q}_{k^*}(t) > 0.$$

where

$$\tilde{\mu}_i = \mu_i - \sum_{f \in \mathbf{F}[i] \setminus \widetilde{\mathbf{F}}[i]} r_f.$$

Proof. See Appendix. □

We will also need the following Lemma, concerning Lyapunov functions that decline at regular points t.

Lemma 8.6. *If $V(t) : \mathbb{R}^+ \to \mathbb{R}^+$ is an absolutely continuous function with $\dot{V}(t) < -a$ for almost all regular t for which $V(t) > 0$, then $V(s) \equiv 0$ for all $s \geq V(0)/a$.*

Proof. See the Appendix. □

8.5 Demand Limited Flow Analysis

Lemma 8.7. *Consider the max-min allocation vector $\mathbf{r}$, a strictly demand limited flow f^*, and all fluid model trajectories $\bar{\mathfrak{X}}(t)$ with $\bar{h} > 0$. Suppose that for each flow f satisfying $\frac{\mathbf{r}_f}{w_f} < \frac{\mathbf{r}_{f^*}}{w_{f^*}}$ it has been shown that for all initial conditions $\bar{X}(0)$ that*

$$\dot{\bar{D}}_k(t) \equiv \mathbf{r}_f \quad \forall k \in \mathcal{P}\left[f\right]$$

for all t greater than or equal to some critical time $\tau(\bar{X}(0)) \geq \max(\bar{U}(0)) \vee \max(\bar{V}(0))$.

Then for all $k \in \mathcal{P}\left[f^*\right]$ the following hold:

$$\dot{\bar{D}}_k(t) \equiv \mathbf{r}_{f^*}$$

$$Q_k(t) \equiv 0 = e_k$$

for all $t \geq \tau^*$ for some critical time $\tau^*(\bar{X}(0))$. Furthermore $\tau^*(\bar{X}(0)) \leq a|\bar{X}(0) - \bar{h}\mathbf{e}| + b\tau(\bar{X}(0))$ for some constants a and b.

Proof. Let $\widetilde{\mathcal{F}}$ be the set of all flows f for which $\mathbf{r}_f/w_f \geq r_{f^*}/w_{f^*}$. For each station i, define the reduced flow constituency set $\widetilde{\mathbf{F}}[i]$, the reduced class constituency set $\widetilde{\mathbf{C}}[i]$, and the reduced capacity $\tilde{\mu}_i$ by

$$\widetilde{\mathbf{F}}[i] \triangleq \mathbf{F}[i] \cap \widetilde{\mathcal{F}}, \qquad \widetilde{\mathbf{C}}[i] \triangleq \mathbf{C}[i] \cap \left[\bigcap_{\widetilde{\mathcal{F}}} \mathcal{P}\left[f\right] \right], \qquad \text{and}$$

$$\tilde{\mu}_i \triangleq \mu_i - \sum_{\mathbf{F}[i] \setminus \widetilde{\mathbf{F}}[i]} \mathbf{r}_f \tag{8.5}$$

respectively. Now consider any station i that flow f^* crosses. Also let s be the slackness of the capacity constraint at station i under the max-min allocation. In other words, $s := \mu_i - \sum_{\mathbf{F}[i]} \mathbf{r}_f$. Suppose $s = 0$. Then there must exist a flow $f \in \widetilde{\mathbf{F}}[i]$ for which

$$\frac{r_{f^*}}{w_{f^*}} < \frac{\mathbf{r}_f}{w_f}$$

by the following reasoning. If there did not exist such a flow, then f^* would have the highest normalized rate at a saturated station i, and thus station i would be a bottleneck for flow f^*. As we have assumed that f^* is strictly demand-limited, this is not possible.

Additionally, by the definition of the set $\widetilde{\mathbf{F}}[i]$, it must be that

$$\frac{\mathbf{r}_{f^*}}{w_{f^*}} \leq \frac{\mathbf{r}_f}{w_f} \quad \forall f \in \widetilde{\mathbf{F}}[i].$$

By listing the inequalities for each flow $f \in \widetilde{\mathbf{F}}[i]$ expressed by the above form and taking a convex combination of both sides of the inequalities in the list we deduce that

$$\frac{\mathbf{r}_{f^*}}{w_{f^*}} \leq \frac{\sum_{\widetilde{\mathbf{F}}[i]} \mathbf{r}_f}{\sum_{\widetilde{\mathbf{F}}[i]} w_f} = \frac{\mu_i - s - \sum_{\mathbf{F}[i^*] \setminus \widetilde{\mathbf{F}}[i^*]} \mathbf{r}_f}{\sum_{\widetilde{\mathbf{F}}[i^*]} w_f}.$$

Furthermore, if $s = 0$ the above inequality is strict because there would have been a strict inequality in the list we used to derive the above relation. Thus,

$$\frac{\mathbf{r}_{f^*}}{w_{f^*}} < \frac{\tilde{\mu}_i}{\sum_{f \in \widetilde{\mathbf{F}}[i]} w_f}$$

were the inequality is strict.

Let κ be the minimum slackness one finds by evaluating the above inequality across all the stations f^* crosses. More precisely,

$$\kappa := \min_{k \in \mathcal{P}[f^*]} \left[\frac{\tilde{\mu}_{s(k)}}{\sum_{f \in \widetilde{\mathbf{F}}[s(k)]} w_f} - \frac{\mathbf{r}_{f^*}}{w_{f^*}} \right].$$

Additionally it follows from Lemma 8.5 that for any $k \in \mathcal{P}[f^*]$,

$$\dot{\bar{D}}_k(t) \geq w_{f^*} \frac{\tilde{\mu}_{s(k)}}{\sum_{\widetilde{\mathbf{F}}[s(k)]} w_f} \qquad \text{whenever } \bar{Q}_k(t) > 0 \text{ for regular } t \geq \tau.$$

Because flow f^* is demand limited, $\alpha_{f^*} = r_{f*}$. Hence for all $k \in \mathcal{P}[f^*]$

$$\dot{\bar{D}}_k(t) \geq \alpha_{f^*} + \kappa \qquad \text{whenever } \bar{Q}_k(t) > 0 \text{ for regular } t \geq \tau.$$

While $\dot{\bar{A}}_{f^*}(t) \leq \alpha_{f^*}$ by fluid model equation (7.15). Thus by Lemma 8.3, which recall describes the departures from a pipeline that has a nonempty queue, and Lemma 8.1, the flow conservation rule for pipelines, we have

$$\dot{\mathcal{Q}}(t, \mathcal{P}[f^*]) \leq -\kappa$$

whenever $Q(t, \mathcal{P}[f^*]) > 0$ and $t \geq \tau$ regular. Thus by Lemma 8.6,

$$Q_k(t) \equiv 0 \qquad \forall k \in \mathcal{P}[f]$$

for all $t \geq \tau + |\mathcal{Q}(\tau, \mathcal{P}[f^*])|/\kappa$. Because the queues in $\mathcal{P}[f^*]$ cannot grow by more than $\tau\alpha_{f^*}$ in the first τ seconds, the same holds for all

$$t \geq \tau^* \triangleq \tau + (\alpha_{f^*}\tau + |\mathcal{Q}(0, \mathcal{P}[f^*])|)/\kappa.$$

Once the queues in the pipeline have stabilized to 0, the policing must be off and hence $\dot{\bar{A}}_f(t) = \alpha_f$ by (7.14). Because Lemma 8.1 says that flow is conserved through each pipeline $\{\mathcal{P}[f^*] - \mathcal{P}[k]\}$ for each $k \in \mathcal{P}[f^*]$, we may conclude

$$\dot{\bar{D}}_k(t) = \mathbf{r}_{f^*} \qquad \forall t \geq \tau^* \quad \forall k \in \mathcal{P}[f^*]. \tag{8.6}$$

□

8.6 Bottleneck Limited Flow Analysis

Lemma 8.8. *Consider the max-min allocation vector* $\mathbf{r}$*, a saturated station* i^**, and all valid fluid model trajectories* $\bar{\mathfrak{X}}(t)$ *with* $\bar{h} > 0$*. Let* $\mathcal{B}$ *be the maximum normalized rate of flows crossing* i^**. Equivalently let*

$$\mathcal{B} = \max_{f \in \mathbf{F}[i^*]} \frac{\mathbf{r}_f}{w_f},$$

and also make the following definitions:

$$\widetilde{\mathbf{F}}[i^*] := \arg \max_{f \in \mathbf{F}[i^*]} \frac{\mathbf{r}_f}{w_f}$$

$$\widetilde{\mathbf{C}}[i^*] := \{l \in \mathbf{C}[i] : \exists f \in \widetilde{\mathbf{F}}[i^*] \text{ with } f = f_{(l)}\}$$

Suppose that for every flow f *satisfying* $\mathbf{r}_f / w_f < \mathcal{B}$ *it has been shown that for all initial conditions* $\bar{X}(0)$ *that*

$$\dot{\bar{D}}_k(t) \equiv \mathbf{r}_f \quad \forall k \in \mathcal{P}[f]$$

for all t *greater than or equal to some critical time* $\tau(\bar{X}(0))$ *which satisfies* $\tau(\bar{X}(0)) \geq \max(\bar{U}(0)) \vee \max(\bar{V}(0))$.

Then for all $l \in \widetilde{\mathbf{C}}[i^*]$ *the following hold for all* $k \in \mathcal{P}\left[f_{(l)}\right]$

$$\dot{\bar{D}}_k(t) \equiv \mathbf{r}_{f_{(l)}}$$

$$Q_k(t) \equiv \bar{h} e_k$$

for all $t \geq \tau^*$ *for some critical time* $\tau^*(\bar{X}(0))$. *Furthermore* $\tau^*(\bar{X}(0)) \leq a\left|\bar{X}(0) - \bar{h}\mathbf{e}\right| + b\tau(\bar{X}(0))$ *for some constants* a *and* b.

Proof. Let $\widetilde{\mathcal{F}}$ be the set of all flows for which $\mathbf{r}_f/w_f \geq \mathcal{B}$. For each station i define the reduced flow constituency set $\widetilde{\mathbf{F}}[i]$, the reduced class constituency set $\widetilde{\mathbf{C}}[i]$, and the reduced capacity $\tilde{\mu}_i$, by the assignment rules (8.5). These assignments do not change the assignment we already made for $\widetilde{\mathbf{F}}[i^*]$ and $\widetilde{\mathbf{C}}[i^*]$ that we made in the statement of the lemma.

Consider a class $l \in \widetilde{\mathbf{C}}[i^*]$ and a station $i \neq i^*$ that flow $\mathrm{f}_{(l)}$ crosses. Also let s be the slackness of the capacity constraint at station i under the max-min allocation. Equivalently, $s := \mu_i - \sum_{\mathbf{F}[i]} \mathbf{r}_f$.

Suppose $s = 0$. Then there must exist a flow $f \in \widetilde{\mathbf{F}}[i]$ for which

$$\mathcal{B} = \frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} < \frac{\mathbf{r}_f}{w_f}$$

by the following reasoning. If there did not exist such a flow, then $\mathrm{f}_{(l)}$ would have the highest normalized rate at station i, and thus station i would be a bottleneck for flow $\mathrm{f}_{(l)}$. As we have assumed that flows may have only one bottleneck and station i^* is a bottleneck for flow $\mathrm{f}_{(l)}$, this is not possible.

Additionally, by the definition of $\widetilde{\mathbf{F}}[i]$ it must be that

$$\mathcal{B} = \frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} \leq \frac{\mathbf{r}_f}{w_f} \quad \forall f \in \widetilde{\mathbf{F}}[i].$$

By listing the inequalities for each flow $f \in \widetilde{\mathbf{F}}[i]$ expressed by the above form and taking a convex combination of both sides of the inequalities in the list we deduce that

$$\mathcal{B} = \frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} \leq \frac{\sum_{\widetilde{\mathbf{F}}[i]} \mathbf{r}_f}{\sum_{\widetilde{\mathbf{F}}[i]} w_f} = \frac{\mu_i - s - \sum_{\mathbf{F}[i^*] \setminus \widetilde{\mathbf{F}}[i^*]} \mathbf{r}_f}{\sum_{\widetilde{\mathbf{F}}[i^*]} w_f}.$$

Furthermore, if $s = 0$ the above inequality is strict because there would have been a strict inequality in the list we used to derive the above relation. Thus,

$$\frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} < \frac{\tilde{\mu}_i}{\sum_{\widetilde{\mathbf{F}}[i^*]} w_f}$$

where the inequality is strict. Let κ_1 be the minimum slackness one finds by evaluating the above inequality across all the stations that flows in $\widetilde{\mathbf{F}}[i^*]$ cross. More precisely,

$$\kappa_1 := \min_{l \in \widetilde{\mathbf{C}}[i^*]} \min_{k \in \mathcal{P}\left[\mathrm{f}_{(l)}\right] \setminus \{l\}} \left[\frac{\tilde{\mu}_{s(k)}}{\sum_{f \in \widetilde{\mathbf{F}}[s(k)]} w_f} - \frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} \right].$$

As flows in $\widetilde{\mathbf{F}}[i^*]$ are not demand limited, $\mathbf{r}_f / w_f < \alpha_f$ for each $f \in \widetilde{\mathbf{F}}[i^*]$. Let κ_2 be the minimum slackness in these inequalities, as made precise by the definition:

$$\kappa 2 := \min_{l \in \widetilde{\mathbf{C}}[i^*]} \left[\alpha_{\mathrm{f}_{(l)}} - \frac{\mathbf{r}_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}} \right]$$

Finally we set $\kappa = \kappa_1 \vee \kappa_2$.

Thus for each $l \in \widetilde{\mathbf{C}}[i^*]$ and all $k \in \mathcal{P}\left[\mathrm{f}_{(l)}\right] \setminus l$,

$$\frac{\tilde{\mu}_{s(k)}}{\sum_{\widetilde{\mathbf{F}}[s(k)]} w_f} \geq \mathcal{B} + \kappa. \tag{8.7}$$

By Lemma 8.5, the departure rates from each queue in $k \in \mathcal{P}\left[\mathrm{f}_{(l)}\right]$ satisfy,

$$\dot{\bar{D}}_k \geq w_{\mathrm{f}_{(l)}} \frac{\tilde{\mu}_{s(k)}}{\sum_{\widetilde{\mathbf{F}}[s(k)]} w_f} \quad \text{whenever } \bar{Q}_k(t) > 0 \text{ for regular } t \geq \tau. \tag{8.8}$$

Relation (8.7) together with property (8.8) imply that for any $l \in \widetilde{\mathbf{C}}[i^*]$, the departure rates of the queues upstream or downstream of l satisfy

$$\dot{\bar{D}}_k(t) \geq w_{\mathrm{f}_{(l)}}(\mathcal{B} + \kappa) \qquad \text{for each } k \in \left\{\mathcal{P}\left[\mathrm{f}_{(l)}\right] - l\right\} \text{ and any regular } t \geq \tau \text{ for which } \bar{Q}_k(t) > 0. \tag{8.9}$$

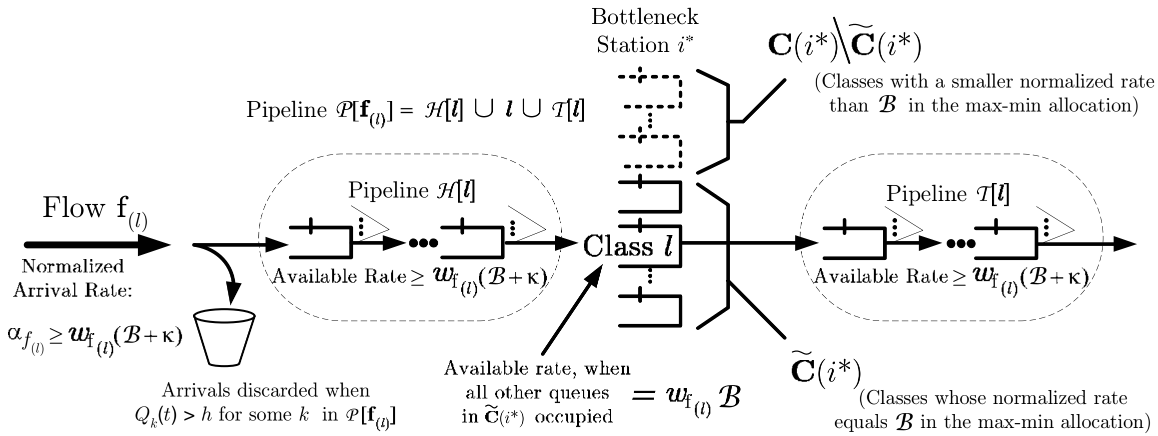

Figure 8.2: The class l queue at a bottleneck station i^*, and the available rates of the upstream and downstream queues.

Also, by the assumptions of this lemma

$$\frac{\tilde{\mu}_{i^*}}{\sum_{\widetilde{\mathbf{F}}[i^*]} w_f} = \mathcal{B}.$$

Substituting the above relation into (8.7), we have

$$\dot{\bar{D}}_l(t) \geq w_{f_{(l)}}\mathcal{B} \qquad \text{for each } l \in \widetilde{\mathbf{C}}[i^*] \text{ and any regular } t \geq \tau \text{ for which } \bar{Q}_l(t) > 0. \tag{8.10}$$

Properties (8.9) and (8.10) describe the available rate at each queue in a pipeline $\mathcal{P}\left[f_{(l)}\right]$ and we will refer to them often in the development that follows. An illustration of the situation described by these properties is shown in Figure 8.2.

Downstream Queues: We begin analyzing the occupancy of the queues downstream of each class in $\widetilde{\mathbf{C}}[i^*]$, with the goal of showing that they drain to zero in some

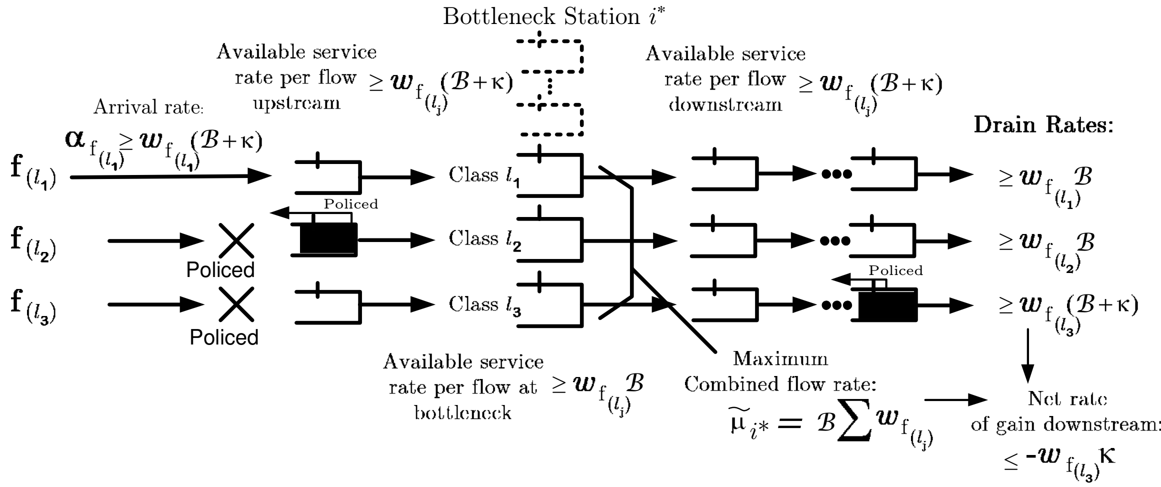

Figure 8.3: An illustration of why the combined queue depths of queues downstream of the bottleneck station i^* decline monotonically to 0. Note that class $l_1 \in \mathbf{PD}^c(t) \cap \mathbf{AZ}^c(t)$, class $l_2 \in \mathbf{PD}(t) \cap \mathbf{AZ}^c(t)$, and class $l_3 \in \mathbf{AZ}(t)$.

finite time. To help partition the possibilities, we define:

$$\mathbf{PD}(t) \triangleq \{l \in \widetilde{\mathbf{C}}[i^*] : \bar{Q}_k(t) \geq \bar{h} \text{ for some } k \in \mathcal{P}\left[\mathrm{f}_{(l)}\right]\}.$$

$$\mathbf{AZ}(t) \triangleq \{l \in \widetilde{\mathbf{C}}[i^*] : \bar{Q}_k(t) > 0 \text{ for some } k \in \mathcal{T}[l]\}.$$

PD, **AZ** are mnemonic for “Policed” set, and “Above Zero” set respectively. Also Recall $\mathcal{T}[l]$ is the ordered set of classes downstream of l.

Consider $l \in \mathbf{PD}^c(t) \cap \mathbf{AZ}^c(t)$, corresponding flow $\mathrm{f}_{(l)}$, and regular $t \geq \tau$. Class l_1 in Figure 8.3 is an example of such a class. Flow $\mathrm{f}_{(l)}$ is not policed at time t. Therefore $\dot{\bar{A}}_{\mathrm{f}_{(l)}}(t) = \alpha_{\mathrm{f}_{(l)}}$ by fluid model equations (7.14) and (7.9). By (8.7) $\alpha_{\mathrm{f}_{(l)}} \geq w_{\mathrm{f}_{(l)}}(\mathcal{B}+\kappa)$. This lower bound on the arrivals to the pipeline $\mathcal{P}\left[\mathrm{f}_{(l)}\right]$, combined with the properties (8.9) and (8.10) which describe the departure rates of queues in the pipeline $\mathcal{P}\left[\mathrm{f}_{(l)}\right]$, allow us to invoke Lemma 8.2 to conclude that $\dot{\bar{D}}_k(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B}$ where k is the last class in $\mathcal{P}\left[\mathrm{f}_{(l)}\right]$. Note that k is also the last class in the pipeline $\mathcal{T}[l]$, and class l empties into

$\mathcal{T}[l]$ at rate $\dot{\bar{D}}_l$. Thus by Lemma 8.1, which recall is the flow conservation lemma,

$$\dot{\mathcal{Q}}(t, \mathcal{T}[l]) \leq \dot{\bar{D}}_l - w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.11}$$

Consider $l \in \mathbf{PD}(t) \cap \mathbf{AZ}^c(t)$, for regular $t \geq \tau$ and corresponding flow $\mathrm{f}_{(l)}$. Class l_2 in Figure 8.3 is an example of such a class. By definition, flow $\mathrm{f}_{(l)}$ is policed at time t, and all queues for classes in $\mathcal{T}[l]$ are empty. Therefore there exists a class $l^* \in \mathcal{P}\left[\mathrm{f}_{(l)}\right] - \mathcal{T}[l]$ with $\bar{Q}_{l^*} > 0$. By (8.9), (8.10) and Lemma 8.3, $\dot{\bar{D}}_k(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B}$ where k is the last class in pipeline $\mathcal{P}\left[\mathrm{f}_{(l)}\right]$. Note that k is also the last class in the pipeline $\mathcal{T}[l]$, and class l empties into $\mathcal{T}[l]$ at rate $\dot{\bar{D}}_l$. Thus

$$\dot{\mathcal{Q}}(t, \mathcal{T}[l]) \leq \dot{\bar{D}}_l - w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.12}$$

Consider $l \in \mathbf{AZ}(t)$, and regular $t \geq \tau$. Class l_3 in Figure 8.3 is an example of such a class. By definition, $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{T}[l]$. By (8.9), Lemma 8.3,and Lemma 8.1

$$\dot{\mathcal{Q}}(t, \mathcal{T}[l]) \leq \dot{\bar{D}}_l - w_{\mathrm{f}_{(l)}}\mathcal{B} - \bar{w}\kappa. \tag{8.13}$$

Define the Lyapunov function

$$V(t) := \sum_{l \in \widetilde{\mathbf{C}}[i^*]} \mathcal{Q}(t, \mathcal{T}[l])$$

to be the sum of the queue occupancies downstream of bottleneck station i^*. Taking the derivative, and substituting (8.11), (8.12), and (8.13) we have for all regular $t \geq \tau$,

$$\begin{aligned}
\dot{V}(t) \leq & \sum_{l \in \mathbf{PD}^c(t) \cap \mathbf{AZ}^c(t)} \dot{\bar{D}}_l(t) - w_{\mathrm{f}_{(l)}}\mathcal{B} \\
& + \sum_{l \in \mathbf{PD}(t) \cap \mathbf{AZ}^c(t)} \dot{\bar{D}}_l(t) - w_{\mathrm{f}_{(l)}}\mathcal{B} \\
& + \sum_{l \in \mathbf{AZ}(t)} \dot{\bar{D}}_l(t) - w_{\mathrm{f}_{(l)}}(\mathcal{B} + \kappa).
\end{aligned}$$

By rearranging terms, we deduce that

$$\dot{V}(t) \leq \sum_{l \in \widetilde{\mathbf{C}}[i^*]} \left[\dot{\bar{D}}_l(t) - w_{\mathrm{f}_{(l)}} \mathcal{B} \right] - |\mathbf{AZ}(t)| \bar{w} \kappa \tag{8.14}$$

where $\bar{w} \triangleq \min_{f \in \mathcal{F}} w_f$. Using fluid model equations (7.7) and (7.12), as well as the definition of $\tilde{\mu}$ (8.5), we can bound the combined departure rates from the queues of station i^* according to

$$\begin{aligned} \sum_{\widetilde{\mathbf{C}}[i^*]} \mu_{i^*} \dot{\bar{T}}_l(t) + \sum_{\mathbf{C}[i^*] \setminus \widetilde{\mathbf{C}}[i^*]} \mu_{i^*} \dot{\bar{T}}_l(t) \leq \mu_{i^*} \\ \sum_{\widetilde{\mathbf{C}}[i^*]} \mu_{i^*} \dot{\bar{T}}_l(t) \leq \mu_{i^*} - \sum_{\mathbf{F}[i^*] \cap \widetilde{\mathcal{F}}} \mathbf{r}_f \\ \sum_{\widetilde{\mathbf{C}}[i^*]} \dot{\bar{D}}_l(t) \leq \tilde{\mu}_{i^*}. \end{aligned} \tag{8.15}$$

Substituting (8.15) into (8.14) we have,

$$\dot{V}(t) \leq \tilde{\mu}_i - \sum_{l \in \widetilde{\mathbf{C}}[i^*]} w_{\mathrm{f}_{(l)}} \mathcal{B} - |\mathbf{AZ}(t)| \bar{w} \kappa \qquad \forall t \geq \tau$$

which reduces to

$$\dot{V}(t) \leq -|\mathbf{AZ}(t)| \bar{w} \kappa \qquad \forall t \geq \tau.$$

Thus by Lemma 8.6, $V(t) \equiv 0$ for all $t \geq V(\tau)/(\bar{w}\kappa)$. Note that the downstream queues included in the sum that defines V can not collectively grow by more than $\mu_i \tau$ on the time interval $[0, \tau]$. Thus

$$Q_k(t) \equiv 0 \qquad \forall t \geq t^* \text{ and } k \in \mathbb{T} \tag{8.16}$$

where

$$\mathbb{T} := \bigcup_{l \in \widetilde{\mathbf{C}}[i]} \mathcal{T}[l] \quad \text{and} \quad t^* := (\bar{w}\kappa)^{-1} \left(\mu_i \tau + \sum_{k \in \mathbb{T}} Q_k(0) \right).$$

Upstream Queues: Having found a time for which all the queues downstream of the bottleneck drain to zero, we now consider the queues upstream of the bottleneck and in the bottleneck itself. Recall that $\mathcal{H}[l]$ is the ordered set of queues upstream from l, but not including l itself.

Consider regular $t \geq \tau + t^*$, $l \in \mathbf{PD}(t)$ and flow $\mathrm{f}_{(l)}$. By the definition of $\mathbf{PD}(t)$, $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{P}\left[\mathrm{f}_{(l)}\right]$. But we have established that $\bar{Q}_k(t) = 0$ for all $k \in \mathcal{T}[l]$. Thus either $\bar{Q}_l(t) > 0$ or $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{H}[l]$. Thus the drain from queue l satisfies $\dot{\bar{D}}_l(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B}$ by (8.9), (8.10), and Lemma 8.3.

Consider regular $t \geq \tau + t^*$, $l \in \mathbf{PD}^c(t)$ and flow $\mathrm{f}_{(l)}$. By (7.14), $\dot{\bar{A}}_{\mathrm{f}_{(l)}} = \alpha_{\mathrm{f}_{(l)}} > w_{\mathrm{f}_{(l)}}\mathcal{B}$. Thus by (8.9), (8.10), and Lemma 8.2, $\dot{\bar{D}}_l \geq w_{\mathrm{f}_{(l)}}\mathcal{B}$.

Thus, we have that for regular $t \geq \tau + t^*$ and all $l \in \widetilde{\mathbf{C}}[i^*]$, $\dot{\bar{D}}_l \geq w_{\mathrm{f}_{(l)}}\mathcal{B}$. This implies $\sum_{l \in \widetilde{\mathbf{C}}[i^*]} \dot{\bar{D}}_l(t) \geq \tilde{\mu}_i$ But by, (8.15), $\sum_{l \in \widetilde{\mathbf{C}}[i^*]} \dot{\bar{D}}_l(t) \leq \tilde{\mu}_i$. All of this is possible only if

$$\dot{\bar{D}}_l(t) \equiv w_{\mathrm{f}_{(l)}}\mathcal{B}, \;\; \forall t \geq \tau + t^*, \;\; \forall l \in \widetilde{\mathbf{C}}[i^*]. \tag{8.17}$$

Now consider the Lyapunov function

$$V_l(t) = \mathcal{Q}(t, \mathcal{H}[l]) + (\bar{Q}_l(t) - h)^+ + \frac{\alpha_{\mathrm{f}_{(l)}}}{w_{\mathrm{f}_{(l)}}\kappa}(\bar{Q}_l(t) - h)^- \tag{8.18}$$

which is equal to the sum of the queues upstream of l, plus the distance of $\bar{Q}_l$ from threshold $\bar{h}$ – where the distance is weighted more heavily when $\bar{Q}_l$ is below threshold. A level set of this Lyapunov function is illustrated in Figure 8.4. Suppose $l \in \widetilde{\mathbf{C}}[i]$ and t regular with $t \geq \tau + t^*$ We consider three cases, each of which is illustrated in Figure 8.5.

1. Suppose $\bar{Q}_l(t) < h$. Two subcases are possible. Subcase (1a): $\bar{Q}_k(t) = 0$ for all

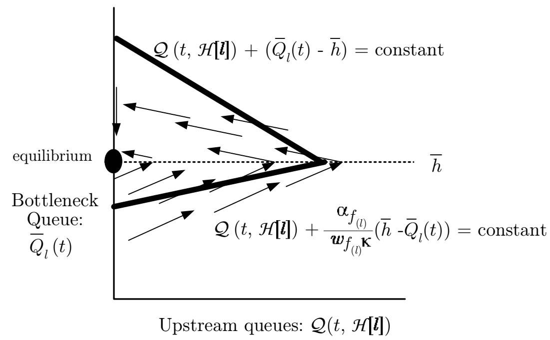

Figure 8.4: A level set of the Lyapunov function defined by (8.18). The level set is shown with a thick line. The arrows illustrate the possible state trajectories of the upstream and bottleneck queues.

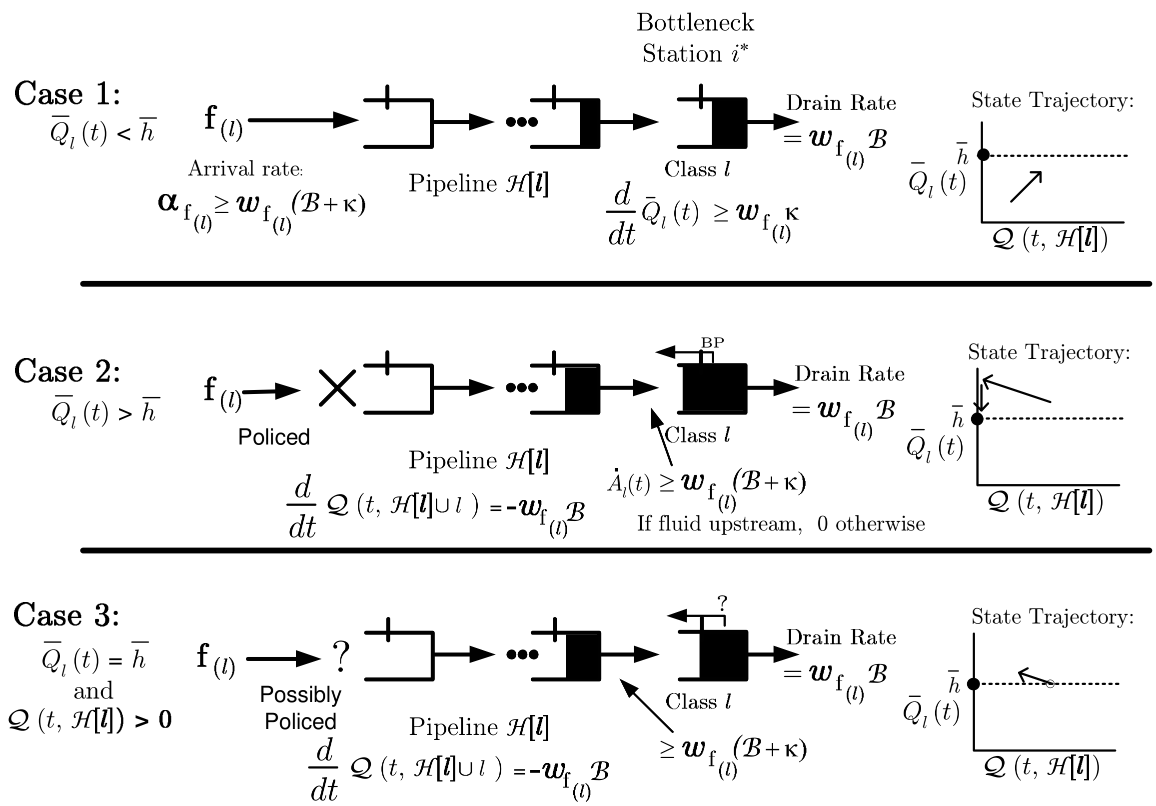

Figure 8.5: The three cases in analyzing the convergence of the upstream queues are shown. Also shown are the possible state trajectories for each of the cases. The question marks on the drawing of Case 3 reflect that it is indeterminate as to whether the back pressure is on when a queue level is exactly equal to the policing threshold.

$k \in \mathcal{H}[l]$ then the policing is off and thus

$$\dot{\bar{A}}_{\mathrm{f}_{(l)}}(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B} + w_{\mathrm{f}_{(l)}}\kappa.$$

Thus by (8.9) and Lemma 8.2,

$$\dot{\bar{A}}_l(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B} + w_{\mathrm{f}_{(l)}}\kappa.$$

Subcase (1b): $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{H}(l)$, then by (8.9) and Lemma 8.3, $\dot{\bar{A}}_l(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B} + w_{\mathrm{f}_{(l)}}\kappa$. Thus in either case,

$$\dot{\bar{A}}_l(t) \geq w_{\mathrm{f}_{(l)}}\mathcal{B} + w_{\mathrm{f}_{(l)}}\kappa \tag{8.19}$$

for regular $t \geq \tau + t^*$. This combined with relation (8.17), which describes the draining from class l, imply

$$\dot{\bar{Q}}_l(t) \geq w_{\mathrm{f}_{(l)}}\kappa. \tag{8.20}$$

Also by (8.19) and Lemma 8.1,

$$\dot{\mathcal{Q}}(t, \mathcal{H}[l]) < \alpha_{\mathrm{f}_{(l)}} - w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.21}$$

Differentiating (8.18), evaluating for the case where $\bar{Q}_l(t) < h$, and substituting (8.20) and (8.21) we have

$$\dot{V}_l(t) < -w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.22}$$

2. Suppose $\bar{Q}_l(t) > h$ Then the policing is on, and by (7.13). $\dot{\bar{A}}_l(t) = 0$. By Lemma 8.3, Lemma 8.1, and (8.9), and (8.10),

$$\dot{\mathcal{Q}}(t, \mathcal{H}[l] \cup l) \leq -w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.23}$$

Differentiating (8.18), evaluating for the case where $\bar{Q}_l(t) > h$, and substituting (8.23) we have

$$\dot{V}_l(t) \leq -w_{\mathrm{f}_{(l)}}\mathcal{B}. \tag{8.24}$$

3. Suppose $\bar{Q}_l(t) = h$, and $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{H}[l]$. Then $\bar{Q}_k(t) > \tilde{\varepsilon}$ for some $\tilde{\varepsilon} > 0$. Now consider any positive $\varepsilon \leq \tilde{\varepsilon}$ By Lemma 8.3, and (8.9), $\dot{\bar{A}}_l(s) \geq w_{\mathrm{f}_{(l)}}\mathcal{B}+\kappa$ for all regular $s \in [t, t+\varepsilon/\mu_{max}]$ where $\mu_{max} = \max_i \mu_i$. This combined with (8.17) implies that $\bar{Q}_l(s) > \bar{h}$ for any $s \in (t, \varepsilon/\mu_{max}]$. Thus for any positive $\delta < \varepsilon/\mu_{max}$, $\dot{\bar{A}}_{\mathrm{f}_{(l)}}(s) = 0$ for all $s \in (t+\delta, \varepsilon/\mu_{max}]$ because the policing is on in such a period by (7.13). Thus we have

$$\begin{aligned} V_l\left(t + \frac{\varepsilon}{\mu_{max}}\right) - V_l(t) &= \int_t^{t+\varepsilon/\mu_{max}} \dot{\bar{A}}_{\mathrm{f}_{(l)}}(s) - \dot{\bar{D}}_l(s) ds \\ &\leq \alpha_{\mathrm{f}_{(l)}}\delta - w_{\mathrm{f}_{(l)}}\mathcal{B}\varepsilon/\mu_{max}. \end{aligned}$$

Since δ is arbitrary, we have

$$V_l(t+\varepsilon/\mu_{max}) - V(t) \leq -w_{\mathrm{f}_{(l)}}\mathcal{B}\varepsilon/\mu_{max}. \tag{8.25}$$

Dividing (8.25) by ε and taking the limit as $\varepsilon \to 0$ we have

$$\dot{V}_l(t) \leq -w_{\mathrm{f}_{(l)}}\mathcal{B} \tag{8.26}$$

which therefore holds for any $l : \bar{Q}_l(t) = h$, regular $t \geq \tau + t^*$ with $\bar{Q}_k(t) > 0$ for some $k \in \mathcal{H}[l]$.

Thus by (8.22), (8.24), and (8.26) we have $\dot{V}_l(t) \leq -w_{\mathrm{f}_{(l)}}\mathcal{B}$ in all cases, for regular $t \geq \tau + t^*$.

To complete the argument that the upstream queues drain to zero in a time proportional to the initial condition, we define

$$\mathbb{H} := \bigcup_{l \in \widetilde{\mathbf{C}}[i^*]} \mathcal{H}[l],$$

$$\tau^* \triangleq \tau + t^* + \sum_{\mathbb{H} \cup \widetilde{\mathbf{C}}[i^*]} |\bar{Q}_k(0) - \bar{h}e_k| + c_1 c_2 (\tau + t^*) \tag{8.27}$$

where the constants c_1, c_2, c_3, and c_4 are defined by,

$$c_1 \triangleq 1 \vee \max_{k \in \{1,\ldots,K\}} \frac{\alpha_k}{\kappa w_{\mathrm{f}_{(k)}} \wedge 1}, \qquad c_2 \triangleq \max_{f \in \{1,\ldots,F\}} \alpha_f \vee \max_{i \in \{1,\ldots,d\}} \mu_i,$$

$$c_3 \triangleq \min \frac{\mu_i}{\sum_{\mathbf{F}[i]} w_f}, \qquad c_4 \triangleq c_1/c_3.$$

We have chosen these constants so that

$$\tau^* \leq \max_{l \in \widetilde{\mathbf{C}}[i^*]} \frac{V_l(\tau + t^*)}{w_{\mathrm{f}_{(l)}} \mathcal{B}} + \tau + t^*.$$

So by Lemma 8.6,

$$V_l(t) \equiv 0 \;\; \forall t \geq \tau_{j+1} \;\; \forall l \in \widetilde{\mathbf{C}}[i^*]$$

and thus

$$\bar{Q}_k(t) \equiv \bar{h}e_k = 0 \;\; \forall t \geq \tau^* \;\; \forall k \in \mathbb{H}. \tag{8.28}$$

$$\bar{Q}_l(t) \equiv \bar{h}e_l = 1 \;\; \forall t \geq \tau^* \;\; \forall l \in \widetilde{\mathbf{C}}[i^*]. \tag{8.29}$$

Relations (8.16), (8.17), (8.28), (8.29), and Lemma 8.1 together imply that the each Q_k in each pipeline $\mathcal{P}\left[\mathrm{f}_{(l)}\right]$ settle to e_k, and that the departure rates of each queue in each pipeline are $\mathcal{B}w_{\mathrm{f}_{(l)}}$ after time τ^*. Hence

$$\dot{D}_f(t) \equiv w_f \mathcal{B} = \mathbf{r}_f, \;\; \forall t \geq \tau^*, \;\; \forall f \in \widetilde{\mathbf{F}}[i^*] \tag{8.30}$$

and thus we have shown all of the conclusions of the lemma. $\square$

8.7 Final Result

Theorem 8.9. *Suppose that every flow has either a unique bottleneck station or is strictly demand-limited. Then for any ϵ, there exists a system-scale n_c, such that for all $n \geq n_c$:*

$$\lim_{t \to \infty} \frac{1}{t} D^{n,y}(t) \geq (1 - \epsilon) R \ a.s.$$

where $R_k = r_{f_{(k)}}$ for all $k \in \{1, ..., K\}$.

Proof. To prove the desired conclusion, we need to prove the necessary conditions to invoke Theorem 7.10. Recall that these conditions are that fluid model converges to its equilibrium in a time proportional to the initial condition's distance from the equilibrium, and that when $\bar{h} > 0$ the departure rates of the fluid model also converge to steady rates. We therefore divide the analysis into two cases: the case in which the fluid model threshold $\bar{h} = 0$, and the other case being the $\bar{h} > 0$ case.

$\bar{\mathbf{h}} = \mathbf{0}$ case

Suppose $\bar{h} = 0$. Consider any flow f and its pipeline $\mathcal{P}[f]$. By lemma 8.5, any queue $k \in \mathcal{P}[f]$ has a departure rate of $\dot{\bar{D}}_k(t) \geq w_f \mu_{s(k)} / [\sum_{\mathbf{F}[s(k)]} w_f]$ at regular $t \geq \max(\bar{V}(0))$. Also, whenever any queue $k \in \mathcal{P}[f]$ is nonempty, the policing is on and hence, $\dot{\bar{A}}_f(t) = 0$. Consequently by Lemma 8.1, $Q_k(t) \equiv 0$ for all $t \geq \max(\bar{V}(0)) + \mathcal{Q}(0, \mathcal{P}[f]) / r_{min}$ where $r_{min} := (\min_f w_f) \min_i (\mu_i / [\sum_{\mathbf{F}[i]} w_f])$. The same argument may be repeated for each flow f in the network. Hence,

$$\bar{Q}(t) \equiv 0 \qquad \forall t > a |\bar{X}(0) - \bar{h}\mathbf{e}|$$

for some constant a.

$\bar{\mathbf{h}} > \mathbf{0}$ case

Now suppose that $\bar{h} > 0$. We proceed by induction. We make the assignment $\widetilde{\mathcal{F}} := \{1, ..., g\}$ to represent the set of all flows we have not yet analyzed, and thus we will update $\widetilde{\mathcal{F}}$ as we proceed. We also define the function $\tau_1(\bar{X}(0)) = |\bar{U}(0)| + |\bar{V}(0)|$. As we proceed, we will define functions $\tau_2(\cdot)$, $\tau_3(\cdot)$,... in a similar fashion. Finally set $j := 1$. We will increment j as we proceed through the induction.

Induction Hypotheses: Suppose that the following hypothesis hold:

$$\tau_j(\bar{X}(0)) \leq a_j |\bar{X}(0) - \bar{h}\mathbf{e}| \qquad \text{for some constant } a_j \tag{8.31}$$

$$\bar{Q}_k(t) \equiv \bar{h} e_k \qquad \forall t \geq \tau_j \text{ and forall } k \in \bigcup_{f \in \widetilde{\mathcal{F}}^c} \mathcal{P}[f] \tag{8.32}$$

$$\dot{\bar{D}}_k(t) \mu_{s(k)} \equiv \mathbf{r}_f \qquad \forall t \geq \tau_j \text{ and for each } (f, k) \text{ satisfying}$$

$$f \in \widetilde{\mathcal{F}}^c \text{ and } k \in \mathcal{P}[f] \tag{8.33}$$

Note that all of these hypothesis hold for $j := 1$ and our initial choice of $\widetilde{\mathcal{F}}$ and $\tau_1(\cdot)$. Consider the following procedure:

1) Pick $f \in \arg\min_{\widetilde{\mathcal{F}}}$

- If f is strictly demand limited use Lemma 8.7. Take $\tau_{j+1}(\bar{X}(0)) := \tau^*(\bar{X}(0))$, where $\tau^*(\bar{X}(0))$ comes from the conclusion of Lemma 8.7. Set $\widetilde{\mathcal{F}} := \widetilde{\mathcal{F}} - f$, $j := j + 1$. The conclusions of Lemma 8.7 ensure that hypotheses (8.31), (8.32), and (8.33) remain true. Return to step 1 if $\widetilde{\mathcal{F}} \neq \emptyset$.

- If f has a unique bottleneck i^* use Lemma 8.8. Take $\tau_{j+1}(\bar{X}(0)) := \tau^*(\bar{X}(0))$, where $\tau^*(\bar{X}(0))$ comes from the conclusion of Lemma 8.8. Set $\widetilde{\mathcal{F}} := \widetilde{\mathcal{F}} - \widetilde{\mathbf{F}}[i^*]$ where $\widetilde{\mathbf{F}}[i^*]$ is defined as in Lemma 8.8, and set $j := j + 1$. The conclusions of Lemma 8.8

ensure that hypotheses (8.31), (8.32), and (8.33) remain true. Return to step 1 if $\widetilde{\mathcal{F}} \neq \emptyset$.

Because at least one flow is removed from $\widetilde{\mathcal{F}}$ at each step, the induction must terminate in a finite number of iterations. By carrying out this induction exhaustively, we will have verified all of the induction hypotheses for all flows in the network. Thus for all classes k

$$\dot{\bar{D}}_k(t) \equiv \mathbf{r}_{\mathrm{f}_{(k)}}, Q_k(t) \equiv \bar{h}e_k,$$

where both relations hold for all regular $t \geq b|\bar{X}(0) - \bar{h}e|$.

Concluding Step:

Combining the results from the $\bar{h} = 0$ and $\bar{h} > 0$ cases we have that in both cases,

$$\bar{Q}(|\bar{X}(0) - \bar{h}\mathbf{e}|t) \equiv 0 \qquad \forall t \geq t_0.$$

where $t_0 = a \vee b$. Because $t_0 \geq \tau_1 \geq \max[\bar{U}(0)] \vee \max[\bar{V}(0)]$, fluid model equation (7.5) ensures that the residual inter-arrival service times decline to 0 in the following way:

$$\begin{bmatrix} \bar{U}(|\bar{X}(0) - \bar{h}\mathbf{e}|t) \\ \bar{V}(|\bar{X}(0) - \bar{h}\mathbf{e}|t) \end{bmatrix} = 0 \qquad \forall t \geq t_0.$$

Thus

$$\bar{X}(|\bar{X}(0) - \bar{h}\mathbf{e}|t) \equiv 0 \qquad \forall t \geq t_0.$$

We have also shown that the departure rates of the fluid model converge to the fair rates when $\bar{h} > 0$, and we may express this fact as

$$M^{-1}\dot{\bar{T}}(|\bar{X}(0) - \bar{h}\mathbf{e}|t) \equiv R \qquad \forall t \geq t_0,$$

where $R_k := \mathbf{r}_{\mathrm{f}_{(k)}}$. Thus, we have the necessary conditions to invoke Theorem 7.10, to conclude that for any $\epsilon > 0$, there exists n such that

$$\lim_{t \to \infty} \frac{1}{t} D^{n,y}(t) \geq (1 - \epsilon) R \text{ a.s.}$$

□

8.8 Relaxing the Unique Bottleneck Requirement

Our final result that the flow rates of a round-robin network converge to max-min fair share rates depends on each flow having a unique bottleneck. We needed this condition to ensure that the bottleneck queues fill to their thresholds. If flows did have more than one bottleneck, one of the bottlenecks would still fill to the threshold in the fluid model, but which of these bottlenecks fills would depend on the initial condition. Consequently, our technique of showing that the stochastic system is attracted to the unique fluid model equilibrium would break down. Fortunately, the unique bottleneck condition can always be achieved by perturbing the weights and or the rates of the stations.

However, we feel it should be possible to eliminate the unique bottleneck requirement entirely. Consider the following intuitive conjectures. We conjecture that if a flow's weight were increased that its long term average rate cannot decrease, if we suppose all other parameters of the network are unchanged. Similarly, we conjecture that if a station that a flow crosses were increased in capacity, its long term rates cannot decrease. While these monotonicity conjectures as we call them seem like they should be true, they are not trivial to prove. If we could show the monotonicity conjectures, then we would be in a position to make pairwise long term rate comparisons of a network to a perturbed network where one parameter, a weight or a service capacity, has been reduced by an arbitrarily small amount. By making a string of such pairwise comparisons, it should be possible to show that the long-term rates of a network without unique bottlenecks can be made arbitrarily close to the long-term rates of a slightly perturbed network with unique bottlenecks.

There may be another pathway to relaxing the unique bottleneck requirement. It might be possible to show that the fluid model of a network without unique bottlenecks converges to an element in some "equilibrium set" of states – the set where at least one bottleneck of every flow is filled to its threshold. It may then be possible to show that the stochastic system is attracted to this same set of states, by showing that some distance metric between the current state and this equilibrium set contracts in expected value.

8.9 Uniformity of the Required Threshold for Different Demand Processes

In an application, a network like the round-robin network we have analyzed would be used to provide fair rates when the demand from the different flows is unknown and possibly changing over time. Such would be the case in an Internet packet switch for example. What we have shown is that there exists a large enough threshold scale factor n to achieve close to the fair rates for *each particular* choice of demand process. However, we have not shown that their exists a single scale factor which would work well for any choice of demand process.

For example, suppose the all the possible demand processes that the network might face are described by some vector of parameters v (i.e. the arrival rates α, and any other parameters needed describe the distribution of the inter-arrival times) that lives in some space $\mathcal{V}$. For each such vector v, our result shows that their exists a large enough scale factor $n(v, \epsilon)$ for which the long-term rates are within a factor $(1 - \epsilon)$ of the fair rates. However, we need a stronger result to show that $\sup_{v \in \mathcal{V}} n(v, \epsilon)$ is not infinite. We feel that it should be possible to derive such a result, but it is not trivial. One approach might be to show that the long term rate of each flow is monotone with respect to each component of v. If the space $\mathcal{V}$ were compact, then this monotonicity would allow us to upgrade our point-wise bound to a bound that holds uniformly over $\mathcal{V}$.

Chapter 9

Conclusion

In this work we have shown how the analysis of the flow rates of a stochastic network with a particular flow control scheme may be reduced to an analysis of a fluid model. While we have focused on a particular flow control scheme, we feel that same analysis could be carried out for many other control schemes. The key feature that enabled our approach was that our control scheme has a free parameter, n, which when increased makes the system look more and more like a deterministic fluid system.

In the future, we would like to make this notion more precise, and develop both a more systematic test to determine if a rate control scheme is amenable to this sort of fluid analysis, and a more systematic methodology to carry out the analysis.

It would also be desirable to find a way of quantifying how large n needs to be to have long-term rates close to the fair rates, rather than just knowing that such a threshold scale factor n exists. It might be possible to approximate the necessary n with a reflected Brownian motion model for example, but it is not clear that this approach would yield a value for n for which the long term rates are almost surely within $(1-\epsilon)$ of the fair rates.

We feel that in practice the actual size of the thresholds would be best determined by a simulation study. In this applied context, the methodology developed here would show whether or not it is possible to find large enough thresholds for a proposed flow control scheme to work, and therefore wether or not it is worthwhile to carry out a simulation study.