

Responding to Unexpected Overloads in Large-Scale Service Systems

Ohad Perry, Ward Whitt

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027-6699
{op2105, ww2040}@columbia.edu

We consider how two networked large-scale service systems that normally operate separately, such as call centers, can help each other when one encounters an unexpected overload and is unable to immediately increase its own staffing. Our proposed control activates serving some customers from the other system when a ratio of the two queue lengths (numbers of waiting customers) exceeds a threshold. Two thresholds, one for each direction of sharing, automatically detect the overload condition and prevent undesired sharing under normal loads. After a threshold has been exceeded, the control aims to keep the ratio of the two queue lengths at a specified value. To gain insight, we introduce an idealized stochastic model with two customer classes and two associated service pools containing large numbers of agents. To set the important queue-ratio parameters, we consider an approximating deterministic fluid model. We determine queue-ratio parameters that minimize convex costs for this fluid model. We perform simulation experiments to show that the control is effective for the original stochastic model. Indeed, the simulations show that the proposed queue-ratio control with thresholds outperforms the optimal fixed partition of the servers given known fixed arrival rates during the overload, even though the proposed control does not use information about the arrival rates.

Key words: service systems; call centers; overload controls; queue-ratio routing, many-server queues; deterministic fluid models

History: This paper was first submitted on August 27, 1928; this is the first revision.

1. Introduction

In a large-scale service system, such as a call center, under normal circumstances the arrival rates vary by time of day in a predictable way, and the staffing responds to that anticipated pattern, typically with fixed staffing levels over specified time intervals; see Aksin et al. (2007) and Gans et al. (2003) for background. However, occasionally, for various reasons, there may be unforeseen surges in demand, going significantly beyond the usual fluctuations, and lasting for a significant period of time. A demand surge might occur because of a catastrophic event in emergency response, a

system failure experienced by an alternative service provider, or an unanticipated intense television advertising campaign in retail. Such unexpected demand surges typically cause congestion that cannot be eliminated entirely. Since the demand surge is sudden and unexpected, it may not be possible to immediately change the staffing level.

Fortunately, there may be an opportunity to alleviate the congestion caused by the overload by getting help from another service system, which ordinarily operates independently. For example, with the reduction of telecommunication costs, it is more and more common to have networked call centers, often geographically dispersed, even on different continents. Such sharing is typically possible among different hospitals in a metropolitan area. It is often desirable to operate these service systems separately, but their connection provides opportunities, in particular, to provide assistance under overloads. In this paper we consider how that might be done and how to assess the costs and benefits.

An important consideration is that we typically do not want sharing under normal loads. One reason is that it is easier to manage the different facilities separately, e.g., by maintaining clear accountability. Another reason is that the agents in each service facility may be less effective and/or less efficient serving the customers from the other system, because each requires specialized skills not required for the other. We want to consider the case in which serving the other class is possible, but that there are penalties for doing so. We will assume that the service rates are slower for non-designated agents.

The proposed overload control applies directly to separate service systems run by a single organization, but could also be adopted by two different organizations by mutual agreement. Our analysis provides useful information about the likely consequences of any agreement, which should facilitate making the agreement. Current practice for call centers (that we are aware of) is limited to sharing within a single organization, and then only manually or on a regular basis under normal loading. Load-balancing schemes used in practice are described in §5.3 of Gans et al. (2003).

Thus, our goal is to develop a control to automatically detect when an overload has occurred (in either system, or in both) and, then, before the staffing levels can be changed, reduce the resulting

congestion by activating appropriate sharing from agents in the other system. We also want to prevent undesired sharing under normal loads. By focusing on this overload problem, we aim to contribute new insight into the longstanding question about the costs and benefits of resource pooling; see §4.2 of Aksin et al. (2007) and references cited therein. Here we focus on a situation where we want to turn on and off the pooling.

Organization of the paper. We start in §2 with a literature review. Next in §3 we introduce our proposed modelling approach. As an idealized model of two large-scale service systems, which ordinarily operate separately, but have the capability of serving customers from the other system, we consider the Markovian X call-center model having two homogeneous customer classes and two homogeneous agent pools, where all the agents are cross-trained but serve the other class inefficiently. For clarity, we provide a concrete example. We then introduce a cost framework in order to evaluate alternative controls. We indicate how we specify an overload incident and how we evaluate the performance consequence.

In §4 we introduce the proposed control, which is a variant of the queue-ratio controls introduced by Gurvich and Whitt (2007a,b). After reviewing that control (without thresholds), we show that it can perform very poorly for this unintended application, because it can induce inefficient sharing simultaneously in both directions. We then introduce our proposed alternative, which includes two thresholds, one for each direction of sharing.

In §5 we introduce a deterministic fluid model to approximate the overloaded system after the overload incident has occurred. We then introduce a convex cost structure and show how to select queue-ratio functions to minimize the long-run average cost in the overload incident for the fluid model. We then develop a numerical algorithm to compute the optimal queue-ratio functions for arbitrary convex cost functions. We exhibit explicit formulas for the optimal queue-ratio functions for special structured separable cost functions in §EC.4 in the e-companion.

In §6 we discuss how to set the threshold parameters. In §7 we conduct simulation experiments to show that the optimal control for the fluid model is effective for the stochastic X model. Finally, we state our conclusions in §8. Supporting material appears in the e-companion.

2. Literature Review

In this paper we contribute to the literature on overload (or congestion) control in queueing systems. There is a substantial literature studying controls that route (or assign) customers (or jobs) to servers, possibly exploiting thresholds, but many of these papers, like Bell and Williams (2005) and references therein, focus on single-server systems without customer abandonment, whereas we focus on many-server systems with customer abandonment; we only discuss the many-server literature. (The distinction is between routing to one of several servers, as opposed to routing to one of several pools of servers.) It is now understood that the presence of many servers changes the problem; e.g., see Gurvich and Whitt (2007b). One feature of many-server systems with customer abandonment we will exploit is the rate at which the transient distribution approaches its steady-state limit: It tends to be much faster for many-server queues. In particular, the systems we consider tend to reach steady-state in a few mean service times; we elaborate in §EC.1. Hence, in our analysis of performance during an overload incident, we approximate using the new steady state, determined by the new arrival rates (assumed constant). Customer abandonments ensure that the system remains stable.

Our paper can also be viewed as a contribution to the call-routing problem for multi-class and multi-site call centers with skill-based routing; see §5 of Gans et al. (2003) and §§2.3.3, 4.1, 4.2 of Aksin et al. (2007). Others have proposed responding to stochastic fluctuations and unexpected overloads by modulating demand in different ways: (i) admission control, (ii) making delay announcements that may induce customers to leave, use a different service channel (e.g., email instead of voice), or call back later, and (iii) acting to reduce service times, e.g., by curtailing cross-selling activities; see §3 of Aksin et al. (2007) and Armony and Gurvich (2006).

In contrast, our paper relates to the larger literature exploiting server flexibility (supply-side management). One approach is to have extra temporary servers available on short notice; see Bhandari et al. (2008) and references therein. Instead, we propose using servers that are already working; i.e., we propose a form of resource pooling, which exploits cross training; see §4.2 of Aksin et al. (2007) and §5.1 of Gans et al. (2003). As should be anticipated, though, our control tends

to be more effective in alleviating congestion (rather than just balancing the service degradation) when the less-loaded system actually has some slack. Our work draws on the queue-ratio control proposed in Gurvich and Whitt (2007a,b), which applies to very general network topologies. Here we consider the relatively difficult X model, allowing sharing in both directions (see Figure 1 below), but our approach makes the model behave more like the N model; see Tezcan and Dai (2006).

However, we make significant departures from the previous literature. First, we want resource sharing only in the presence of the unanticipated overload, and only in the proper direction, which depends on the nature of the overload. Hence, we turn on and off the sharing. Second, we regard the overload as a rare exceptional unanticipated event, rather than a stochastic fluctuation in demand. Thus, we think that it is inappropriate to perform a long-run steady-state analysis of system performance with alternating normal and overload periods (although that could be done). Instead, we focus on a single overload in isolation.

Since the system tends to be overloaded, even after sharing has been activated, system performance tends to be well approximated by deterministic fluid approximations, as in Whitt (2004). From a heavy-traffic perspective, the system operates in the so-called efficiency driven (ED) many-server heavy-traffic regime, instead of the quality-and-efficiency-driven (QED) regime; see Garnett et al. (2002), Gurvich and Whitt (2007a,b). Our paper also relates to the literature on arrival-rate uncertainty; see §4.4 of Gans et al. (2003) and §2.4 of Aksin et al. (2007). Arrival-rate uncertainty also tends to make deterministic fluid approximations remarkably accurate; e.g., see Whitt (2006), Bassamboo and Zeevi (2009) and their previous papers with Harrison.

In closing, we mention the large literature on detection outside queueing, such as control charts in statistical quality control, sequential analysis and change point problems, but we make no direct contact with it. However, detecting the new arrival rate is not the only issue: In simulations, our proposed control outperforms the optimal fixed partition of the servers given known arrival rates during the overload, even though our proposed control does not use direct information about the arrival rates; see §7.2.

3. The Modelling Approach

The X model. As an idealized model of two separate service systems with the capability of sharing, we consider the X model, depicted in Figure 1. The X model has two homogeneous customer classes and two homogeneous agent pools. We assume that each customer class has a service pool primarily dedicated to it, but all agents are cross-trained, so that they can handle calls from the other class, even though they may do so inefficiently or ineffectively. Under normal loading (at or near forecasted arrival rates), we want each class to be served only by its designated agents, without any help from cross-trained agents in the other service pool. We assume that staffing has been performed in standard ways, so that the number of agents in each pool is adequate to meet performance targets at forecasted arrival rates. However, we also want to automatically activate

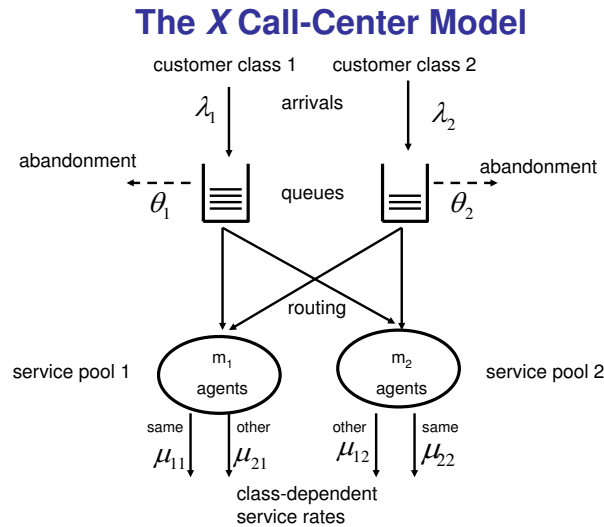


Figure 1 The X model

sharing when there are unexpected unbalanced overloads, either when only one class is overloaded or when both classes are overloaded but one is much more overloaded than the other.

More specifically, in this paper, we consider a fully Markovian model. Customers from the two classes arrive according to independent Poisson processes with arrival rates λ_1 and λ_2 . There is a queue for each customer class, with customers from each class entering service in order of arrival.

We assume that waiting customers have limited patience. A class- i customer will abandon if he does not start service before a random time that is exponentially distributed with mean $1/\theta_i$. There are two service pools, with pool j having m_j homogeneous servers working in parallel. The service times are mutually independent exponential random variables, but the mean may depend on both the customer class and the service pool. The mean service time for a class- i customer served by a type- j agent is $1/\mu_{i,j}$. Let the service times, abandonment times and arrival processes be mutually independent. Let $Q_i(t)$ be the number of class- i customers in queue and let $Z_{i,j}(t)$ be the number of type- j agents busy serving class- i customers, at time t . With the assumptions above, the stochastic process $(Q_i(t), Z_{i,j}(t); i = 1, 2; j = 1, 2)$ becomes a six-dimensional continuous-time Markov chain, given any routing policy that depends on this six-dimensional state.

In this context, under normal loading we want each class served only by agents from its own designated service pool; i.e., we want $Z_{1,2}(t) \approx Z_{2,1}(t) \approx 0$ for all t . One possible reason is that the value of service by agents from the other pool might be less, perhaps because they lack specialized skills. Another possible reason is that service by the cross-trained agents is less efficient; we might have the **strong inefficient-sharing condition**

$$\mu_{1,1} > \mu_{1,2} \quad \text{and} \quad \mu_{2,2} > \mu_{2,1}. \quad (1)$$

We examine the inefficient-sharing case. Throughout this paper, we assume the **basic inefficient-sharing condition**

$$\mu_{1,1}\mu_{2,2} \geq \mu_{1,2}\mu_{2,1}. \quad (2)$$

Clearly, condition (1) implies condition (2). These conditions play a role in §5.2.

In this X -model setting with inefficient sharing, we suppose that an unexpected overload occurs at some unanticipated time that changes the arrival rates. We assume that we are unable to immediately change the staffing levels in response to that unexpected overload. However, we do have the option of allowing some of the cross-trained agents from the less-loaded service pool serve customers from the more overloaded customer class. In addition, we do not know the new

arrival rates when the overload occurs. Thus we need to develop a control that depends on the system history; in some way we must discover that the arrival rates have indeed changed. That is challenging, because stochastic fluctuations under normal loading may make us think that the arrival rates have changed when in fact they have not. We illustrate with the following example.

EXAMPLE 1. To illustrate, consider a symmetric model with forecasted arrival rates $\lambda_1 = \lambda_2 = 90$ per unit of time, where the mean service time for customers served by designated agents is $\mu_{1,1}^{-1} = \mu_{2,2}^{-1} = 1.0$, while the mean service time for customers served by agents from the other pool is $\mu_{1,2}^{-1} = \mu_{2,1}^{-1} = 1.25$. We measure time in units of mean service times by designated agents, which for discussion we take to be 5 minutes. Notice that condition (1) holds here: For all agents, the mean time required to serve the other class is 25% greater than the mean time required to serve an agent's own class. Let customers abandon at rate $\theta_1 = \theta_2 = 0.4$.

Because serving the other class is less efficient, with these parameters it makes sense to operate the system as two separate systems. Following standard staffing methods for a single-class single-pool $M/M/m + M$ model, we may assign $m_1 = m_2 = 100$ agents to the two service pools. That makes the traffic intensities $\rho_1 \equiv \lambda_1/m_1\mu_{1,1} = \rho_2 = 0.90$, which we regard as normal loading. With this staffing, standard algorithms show that steady-state performance is quite good: 82% of the arrivals enter service immediately upon arrival without joining the queue, only 0.5% of the arrivals abandon, the average size of each queue is 1.1, and the expected conditional waiting time, given that the customer is served, is only 0.012 (about 3.6 seconds with a mean service time of 5 minutes).

Now suppose that, at some unanticipated time, the arrival rate for class 1 jumps to $\lambda_1 = 130$, while the arrival rate for class 2 remains at $\lambda_2 = 90$. If class 1 receives no help from pool 2, then class 1 experiences severe congestion. Assuming that the system reaches steady state after this shift in arrival rate (which does not take very long, approximately a few mean service times, as confirmed by simulations - see §EC.1), almost all class-1 customers must wait before starting service, 23% of the class-1 customers abandon, the average size of the class-1 queue becomes 75, the expected conditional waiting time given that a class-1 customer is served is 0.65 (3.25 minutes).

If, as system managers, we were able to recognize that the class-1 arrival rate had shifted to 130, then we might elect to reassign some of the class-2 agents. For example, we might let 25 of the pool-2 agents be devoted to serving class 1. That increases the total service rate responding to the class-1 arrival rate of 130 from 100 to $100 + (1/1.25)25 = 120$, while leaving a total service rate of $100 - 25 = 75$ to respond to the class-2 arrival rate of 90. Since sharing is inefficient, we must sacrifice 25 units of service rate for class 2 in order to gain 20 units of service rate for class 1.

Assuming that the two classes can be modelled as $M/M/m + M$ queues (which is only approximately correct for class 1 because its servers have become heterogenous), we can analyze the performance, e.g., by Whitt (2005). The pair of abandonment probabilities for the two classes changes from $(0.23, 0.005)$ to $(0.08, 0.17)$; the pair of mean queue lengths for the two classes changes from $(75, 1.1)$ to $(26, 38)$; and the pair of conditional expected waiting times given that the customer is served changes from $(0.65, 0.012)$ to $(0.205, 0.450)$ (1.03 minutes and 2.25 minutes, respectively). In this paper we develop a control that responds in a similar way, but does so automatically without having to know that the arrival rates made that specific shift, and without making a fixed partition of the agents.

Analysis with a cost function. The advantage of such sharing, or any other control that produces similar sharing by the inefficient cross-trained agents, depends upon the cost of the congestion experienced. To assess that cost, we will assume that there is a cost function C , with $C(Q_1(t), Q_2(t))$ representing the expected cost rate incurred at time t if the vector of queue lengths at time t is $(Q_1(t), Q_2(t))$. If the overload incident takes place over the time interval $[a, b]$, then the expected total cost would be

$$C_T \equiv E \left[\int_a^b C(Q_1(t), Q_2(t)) dt \right] = \int_a^b E[C(Q_1(t), Q_2(t))] dt. \quad (3)$$

We assume that the cost function C is convex and strictly increasing. The convexity explains why we might want to share when one class is much more overloaded than the other, no matter which class is overloaded.

In this context, our goal is to choose a routing policy, which may allow assignments to cross-trained agents, in order to achieve low (near-minimum) expected total cost for all possible overload incidents and resulting stochastic processes $(Q_1(t), Q_2(t))$, while producing only a negligible amount of sharing under normal loading. To define what we mean by an “overload incident,” We can first specify an interval $[a, c]$ over which the arrival-rate vector $(\lambda_1(t), \lambda_2(t))$ differs from the nominal vector. (We assume that the arrival process is a nonhomogeneous Poisson process with these new arrival rates.) However, we should also include an additional interval $[c, b]$ after time c to allow the vector queue-length $(Q_1(t), Q_2(t))$ to return to its nominal steady-state value. (Engineering judgement is required.) In our analysis, we simplify by restricting attention to scenarios, as in the example above, in which the pair of arrival rates (λ_1, λ_2) makes a sudden unexpected shift at some time, and remains at the new vector for a significant duration, so that the system reaches a new steady-state at the new arrival-rate vector. (Customer abandonment ensures that the system reaches steady state for any arrival-rate vector.) Our control applies more generally.

For such scenarios, we simplify by re-expressing our goal as minimizing the expected steady-state cost; i.e., we aim to minimize $C_T \equiv E[C(Q_1, Q_2)]$, where (Q_1, Q_2) is the vector of steady-state queue lengths associated with the new arrival-rate vector associated with the overload. We will use this steady-state overload framework to set the control parameters and demonstrate effectiveness, but the control applies to other overload scenarios. For this steady-state analysis to be effective, it is important that the system approaches the new steady state associated with the overload relatively quickly. As illustrated in the concrete example above, this tends to happen in a few mean service times. We discuss this important point further in §EC.1.

In the context of Example 1, we might have a shift in arrival rates lasting five hours. It might not be possible to change the staffing in response, because it is in the middle of the same day. The initial transient period might last 3 mean service times or 15 minutes, which is 5% of the total overload incident. There might then be a recovery period lasting about 5 mean service times or 25 minutes, after which the system returns to steady state. For such overloads, the steady-state is evidently reasonable, and it is essential for tractability. Even with this simplifying approximation, the control

problem for the stochastic system is very difficult. We will get an approximate solution only after exploiting a fluid approximation in addition to this steady-state analysis; see §5.2. Even with that approximation, the analysis with a general increasing convex cost function gets complicated; see §5.2. However, as a byproduct, there is a very nice simple story (explicit formulas for everything), provided that we assume a separable quadratic power cost function; see Proposition 5.

4. The Proposed Control

We start by briefly reviewing the *fixed-queue-ratio* (FQR) *routing rule* from Gurvich and Whitt (2007a) and then we show that the FQR rule without thresholds can perform poorly with inefficient sharing, where the conditions in the theorems of Gurvich and Whitt (2007a) are violated. Then we introduce our proposed modification of FQR in order to treat unexpected overloads. It involves general queue-ratio functions, as in Gurvich and Whitt (2007b), and thresholds, one of each for each direction of sharing.

4.1. FQR and its Difficulties with Inefficient Sharing

With two queues, FQR can be implemented by considering a (weighted) *queue-difference stochastic process* $D(t) \equiv Q_1(t) - rQ_2(t)$, $t \geq 0$, where r is a single target-ratio parameter that management can set. With FQR for the X model, a newly available agent in either service pool serves the customer at the head of the class-1 (class-2) queue if $D(t) > 0$ ($D(t) < 0$), and serves the customer at the head of its own queue if $D(t) = 0$. The goal of FQR is to maintain a nearly constant queue ratio: $Q_1(t)/Q_2(t) \approx r$ throughout time. When $r = 1$, FQR coincides with serving the longer queue.

Under regularity conditions, the FQR control has two very desirable features for large-scale service systems, which makes it possible to reduce the multi-class multi-pool staffing-and-routing problem to the well-understood single-class single-pool staffing problem. First, if the required conditions are satisfied, then FQR tends to produce *state-space collapse* (SSC); i.e., for the X model, the two-dimensional queue-length vector $(Q_1(t), Q_2(t))$ tends to evolve approximately as a one-dimensional process determined by the total queue length $Q_\Sigma(t) \equiv Q_1(t) + Q_2(t)$. In particular, $Q_i(t) \approx p_i Q_\Sigma(t)$ for $i = 1, 2$, where $p_1 = r/(1+r) = 1 - p_2$; e.g., see Figure EC.8 in §EC.2. Moreover,

it does so in a way such that all three stochastic processes - $Q_\Sigma(t)$, $Q_1(t)$ and $Q_2(t)$ - remain appropriately stable as $t \rightarrow \infty$. Indeed, Gurvich and Whitt (2007a) show that, under regularity conditions, FQR achieves SSC asymptotically in the quality-and-efficiency-driven (QED) many-server heavy-traffic limiting regime. Second, with FQR, it is possible to choose the ratio parameter r (or, equivalently, the queue proportions p_i) in order to determine the optimal level of staffing to achieve desired service-level differentiation; i.e., staffing costs are minimized subject to meeting class-dependent delay targets $P(W_i > T_i) = \alpha$; see EC.2 and Gurvich and Whitt (2007a). Gurvich and Whitt (2007b) also showed how to staff to minimize convex costs under normal loading. In that case, the asymptotically optimal control in the QED regime is not FQR, but a state-dependent generalization: the *queue-and-idleness-ratio* (QIR) control. Our optimal queue ratios for the fluid model under overloading with convex costs are of the same state-dependent form.

However, in our setting, where service provided by non-designated agents is inefficient, neither FQR nor QIR, without the extra thresholds, is appropriate in normal loading, because they induce undesired sharing. Because of the inefficient sharing, the system is *not* work-conserving; sharing causes the required workload to increase. Indeed, the conditions in the key theorems of Gurvich and Whitt (2007a,b) are violated. In fact, those conditions are actually needed to maintain stability. (However, for FQR without the thresholds, SSC is still achieved; the two queues explode together.)

EXAMPLE 2. To illustrate, consider the X model with parameters $m_1 = m_2 = 100$, $\mu_{1,1} = \mu_{2,2} = 1.0$, $\mu_{1,2} = \mu_{2,1} = 0.8$, $\lambda_1 = \lambda_2 = 0.99$ and $\theta_1 = \theta_2 = 0.0$ (no abandonment). Since the traffic intensities are $\rho_i = \lambda_i / m_i \mu_{i,i} = 0.99$, the two separate systems without sharing are stable (with mean queue length 85 and mean waiting time 0.85). However, if we use FQR with $r = 1$, then inefficient sharing is generated, so that a significant proportion of each agent pool is busy serving the other class. As a consequence, the arrival rate actually exceeds the service rate and the queue lengths diverge to infinity. Here, there still is SSC, but the two queue lengths diverge together.

This difficulty when FQR is applied inappropriately is illustrated by Figures 2 and 3. They show the sample paths of $Q_1(t)$ and $Z_{2,1}(t)$, starting empty, in one simulation run. After an initial transient period, the number of agents serving the other class fluctuates around $E[Z_{1,2}] = E[Z_{2,1}] \approx$

39, while the queue grows in an approximately linear rate; the simulation estimate is $E[Q_i(t)] \approx 6.8t$, $t \geq 0$. (These numerical values are estimated from multiple simulation runs. The confidence intervals are less than 1%. We develop analytical approximations to describe this behavior in a subsequent paper.)

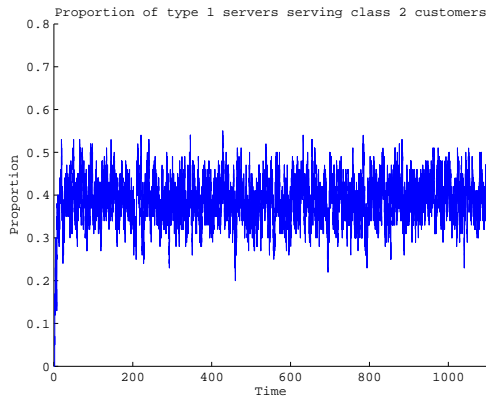


Figure 2 Sample path of $Z_{2,1}(t)$ for FQR




Figure 3 Sample path of $Q_1(t)$ for FQR

Customer abandonment necessarily prevents the queues from exploding. Even in the worst case, when all agents are dedicated to the wrong class, the system would be stable. However, there still is performance degradation, e.g., with $\theta_1 = \theta_2 = 0.2$ and $r = 1$ about 39% of the agents in each pool are busy serving customers from the other class which causes the queues to grow from 10, if there is no sharing, to 34. More details appear in §EC.2.

4.2. The Proposed Control: FQR-T

Here is the lesson from the previous subsection: If we are going to use a queue-ratio control, then we need to take extra measures to prevent sharing under normal loading. First, we want to prevent simultaneous inefficient sharing in both directions. Hence, we restrict the routing to *one-way sharing* at any time: We do not allow a newly available type-2 agent to serve a waiting class-1 customer if there are any type-1 agents busy serving class-2 customers. And similarly in the other direction. (However, over time, the direction of one-way sharing may change; we are not considering the so-called N model, which only allows one-way sharing in one fixed direction.)

From cost considerations, discussed in §5, we want to allow different ratio parameters $r_{1,2}$ and $r_{2,1}$ for the different ways we may share. (In general, we may need more complicated ratio functions or, equivalently, sharing regions; see §5.2, especially Figure 4.) In order to permit sharing only in the presence of unbalanced overloads, we suggest *fixed-queue-ratio routing with thresholds* (FQR-T). In addition to the two ratio parameters $r_{1,2}$ and $r_{2,1}$, we introduce two positive thresholds $\kappa_{1,2}$ and $\kappa_{2,1}$. We then define two queue-difference stochastic processes

$$D_{1,2}(t) \equiv Q_1(t) - r_{1,2}Q_2(t) \quad \text{and} \quad D_{2,1}(t) \equiv r_{2,1}Q_2(t) - Q_1(t). \quad (4)$$

As long as $D_{1,2}(t) < \kappa_{1,2}$ and $D_{2,1}(t) < \kappa_{2,1}$, we do not allow any sharing, i.e., we only let agents serve customers from their designated class.

However, available pool-2 agents are assigned to class-1 customers when $D_{1,2}(t) \geq \kappa_{1,2}$, provided that no pool-1 agents are still serving a class-2 customer. As soon as the first pool-2 agent is assigned to serve a class-1 customer, we drop the threshold $\kappa_{1,2}$, but keep the other threshold $\kappa_{2,1}$. (We could elect to add another threshold for the sharing; see §EC.4.) Once one-way sharing has been activated with pool 2 helping class 1, we use ordinary FQR with ratio parameter $r_{1,2}$. Upon service completion, a newly available type-2 agent serves the customer at the head of the class-1 queue (the class-1 customer who has waited the longest) if $D_{1,2}(t) > 0$; otherwise the agent serves a customer from his own class. In this phase, pool-1 agents only serve class-1 customers. Only one-way sharing in this direction will be allowed until either the class-1 queue becomes empty or the other difference process crosses the other threshold, i.e., when $D_{2,1}(t) \geq \kappa_{2,1}$. As soon as either of these events occurs, newly available pool-2 agents are only assigned to class 2 and the threshold $\kappa_{1,2}$ is reinstated.

We can initiate sharing in the opposite direction when first $D_{2,1}(t) \geq \kappa_{2,1}$ and there are no class-2 agents serving class-1 customers. At the first time both conditions are satisfied, we start sharing with a pool-2 agent serving a class-1 customer. When that first assignment takes place, we remove the threshold $\kappa_{2,1}$ and again use FQR with one-way sharing, but now with the ratio parameter $r_{2,1}$.

Upon arrival, a class- i customer is routed to pool i if there are idle servers; otherwise the arrival goes to the end of the class- i queue. An arrival might increase the queue to a point that sharing is activated. Then the first customer in queue is served by the other class (presumably the agent that has been idle the longest, but we do not focus on individual agents).

The queue-difference stochastic processes in (4) will never provide any instantaneous motivation to have agents of both types simultaneously inefficiently serving the other class if $r_{1,2} \geq r_{2,1}$. That property will be satisfied when we apply a cost function to specify the ratio parameters in §5.2.

To illustrate how FQR-T performs in normal loading (heavy load, but not overloaded), we again consider Example 2 with abandonments at rate $\theta_i = 0.2$. We let $r_{1,2} = r_{2,1} = 1$, so that there is no change from FQR above, but now we add thresholds $\kappa_{1,2} = \kappa_{2,1} = 10$. The performance is greatly improved with FQR-T compared to FQR without thresholds: $E[Z_{1,2}] = E[Z_{2,1}] \approx 2.0$ for FQR-T, while $E[Z_{1,2}] = E[Z_{2,1}] \approx 39$ for FQR. As a consequence, the performance for FQR-T is almost the same as without sharing. In particular, with FQR-T, the abandonment rate is slightly higher than without sharing (2.5% compared to 2.0%), but the average queue length is actually less (9.4 compared to 10.0). In fact, FQR-T can outperform no sharing with larger threshold values, because of the resource-pooling effect. For more details, see §EC.2.

5. The Fluid Approximation for the Steady State of the X Model

In order to obtain a tractable characterization of performance for FQR-T and find good queue-ratio parameters, we now introduce a deterministic fluid approximation. To describe the steady-state behavior of our model when there is no sharing, we first discuss the case of a single customer class served by a single service pool - the classical $M/M/m + M$ model, with arrival rate λ , individual service rate μ and abandonment rate θ . Afterwards we treat the more general X model.

5.1. One Class and One Pool

For the $M/M/m + M$ model, the approximating deterministic fluid model has been studied in Whitt (2004) via many-server heavy-traffic limits. Here we will derive the simple steady-state formulas directly. We assume that input and output (which we call fluid) occurs deterministically

at the specified rates. We think of the system as large and thus regard the number of customers and servers as continuous quantities as well. Thus, fluid arrives deterministically and continuously at constant rate λ . Fluid also is served and abandons deterministically and continuously at rates that are directly proportional to the number of busy servers and the queue length, respectively. If the “number” of busy servers is x , then fluid is served at rate $x\mu$; if the queue length is q , then fluid abandons at rate $q\theta$.

We say that the system is overloaded if the input rate exceeds the maximum possible total service rate. Given m servers, each working at rate μ , the maximum possible total service rate is $m\mu$. Thus the system is overloaded if $\lambda > m\mu$, and not overloaded otherwise. If the system is overloaded, then in steady state all servers will be busy and there will be a queue of waiting fluid, with content q , which can be determined simply by equating the rate in to the rate out, including customer abandonment: $rate\ in \equiv \lambda = m\mu + q\theta \equiv rate\ out$. As an immediate consequence, we get $q = (\lambda - m\mu)/\theta$. If the system is not overloaded, i.e., if $\lambda \leq m\mu$, then there will be no queue. Then we can describe the steady-state via the amount of spare service capacity (number of idle servers), s , which again can be determined by equating the rate in to the rate out: $rate\ in \equiv \lambda = (m - s)\mu \equiv rate\ out$. As an immediate consequence, we get $s = m - (\lambda/\mu)$. Without directly specifying whether or not the system is overloaded, we can write

$$q = \frac{(\lambda - m\mu)^+}{\theta} \quad \text{and} \quad s = \left(m - \frac{\lambda}{\mu}\right)^+, \quad (5)$$

where $(x)^+ \equiv \max\{x, 0\}$. We always have the *complementarity relation* $qs = 0$.

From the point of view of our analysis, we regard λ as an unknown parameter, but we consider the remaining parameters m , μ and θ as fixed and known. For any given λ , we can compute q and s as indicated above. With our overload control problem in mind, it is significant that we can recover λ from the pair (q, s) , because we want to learn about λ by observing (q, s) . If $q > 0$ and $s = 0$, then necessarily we are *overloaded*, and $\lambda = \theta q + m\mu$; if $q = 0$ and $s > 0$, then necessarily we are *underloaded* (which includes normally loaded), and $\lambda = (m - s)\mu$; if $q = 0$ and $s = 0$, then necessarily we are *critically loaded*, and $\lambda = m\mu$; we cannot have $q > 0$ and $s > 0$. For an overloaded

fluid queue, λ is an increasing linear function of q ; for an underloaded queue, λ is a decreasing linear function of s .

As discussed in Whitt (2004), we can also describe the transient behavior of the fluid model and determine other performance measures. For example, if the fluid model is overloaded, then the associated approximate potential steady-state waiting time (virtual waiting time for a customer with infinite patience) is $w = \log(\lambda/m\mu)/\theta_1 = \log(\rho)/\theta_1$, where $\rho \equiv \lambda/m\mu$ is the traffic intensity, satisfying $\rho > 1$; see (2.26) of Whitt (2004).

Note that an increasing convex function of w is an increasing convex function of λ for $\lambda \geq m\mu$. Since λ is a positive linear function of q under overloads, we see that an increasing convex function of w itself is a convex increasing function of q , as we have assumed in our optimization formulation. Similarly, the abandonment rate in the overloaded fluid model is $\theta q = \lambda - m\mu$, so the abandonment rate is an increasing linear function of q under overloads.

5.2. The Optimal Solution for the X Fluid Model

The X fluid model is a natural generalization of the single-class single-pool fluid model above. Now we have two deterministic arrival rates λ_1 and λ_2 , one for each class, with the additional parameters $\{m_j, \theta_i, \mu_{i,j}; i = 1, 2; j = 1, 2\}$. Closely paralleling the discussion above, we will be characterizing the steady-state performance in terms of the quantities (Q_1, Q_2, S_1, S_2) , where Q_i is the fluid content at the class- i queue, while S_j is the amount of spare capacity at pool j .

The steady-state behavior of the X fluid model depends on the number of agents from each pool assigned to (and actually busy serving customers from) each customer class, i.e., the deterministic vector $(Z_{1,1}, Z_{1,2}, Z_{2,1}, Z_{2,2})$, where $Z_{i,j}$ is the number of pool- j agents assigned to serve class- i customers, which is regarded as a continuous variable. To be legitimate assignments, we must have $Z_{i,j} \geq 0$ for all i and j with $Z_{1,1} + Z_{2,1} \leq m_1$ and $Z_{1,2} + Z_{2,2} \leq m_2$. Since these agents are actually busy serving customers, we must also have $\lambda_1 \geq Z_{1,1}\mu_{1,1} + Z_{1,2}\mu_{1,2}$ and $\lambda_2 \geq Z_{2,1}\mu_{2,1} + Z_{2,2}\mu_{2,2}$. Once we assign values to these variables $Z_{i,j}$, we reduce the X model to two single-class single-pool models. The arrival rate for class i is λ_i , while the service rate for class i is $Z_{i,1}\mu_{i,1} + Z_{i,2}\mu_{i,2}$. Class i

is then overloaded if and only if $\lambda_i > Z_{i,1}\mu_{i,1} + Z_{i,2}\mu_{i,2}$, in which case the steady-state fluid content in the class- i is

$$Q_i = \frac{\lambda_i - Z_{i,1}\mu_{i,1} - Z_{i,2}\mu_{i,2}}{\theta_i}. \quad (6)$$

If class i is not overloaded, then $Q_i = 0$. The spare capacity in pool j in steady state is $S_j = m_j - Z_{1,j} - Z_{2,j} \geq 0$, $j = 1, 2$.

In this X fluid model setting, for known arrival rates, our initial goal is to determine the minimum cost $C^*(\lambda_1, \lambda_2)$, which is the minimum of $C(Q_1(Z_{1,1}, Z_{1,2}), Q_2(Z_{2,1}, Z_{2,2}))$ for specified arrival-rate vector (λ_1, λ_2) , which we denote simply by $C(Z_{1,1}, Z_{1,2}, Z_{2,1}, Z_{2,2})$, over all feasible fixed assignment vectors $(Z_{1,1}, Z_{1,2}, Z_{2,1}, Z_{2,2})$ in \mathbb{R}^4 with $Q_i \equiv Q_i(Z_{i,1}, Z_{i,2})$ defined in (6). We let the asterisk denote the optimal solution. (We do not consider more general controls.) We will apply the optimal solution to find the optimal state-dependent queue-ratio functions.

Let q_i be the queue length of class i and let s_i be the spare capacity in pool i when there is no sharing. They can be expressed as in (5), with formulas depending on i . In the fluid model, we regard the system as being in normal loading if neither queue is overloaded without sharing, i.e., if $q_1 = q_2 = 0$, but the amount of spare capacity is not too large. Since the cost function is increasing and convex, under normal loading we achieve the minimum cost by letting $Z_{1,2} = Z_{2,1} = 0$ (no sharing) to obtain $Q_i = 0$ for $i = 1, 2$. The unexpected overload means that either $q_1 > 0$ or $q_2 > 0$, or both. Henceforth we assume that to be the case.

The natural model state is (λ_1, λ_2) , but an equivalent representation is (q_1, s_1, q_2, s_2) , where we always have the complementarity relation $q_1 s_1 = q_2 s_2 = 0$. If $q_i > 0$, then $\lambda_i = m_i \mu_{i,i} + q_i \theta_i$; if $s_i > 0$, then $\lambda_i = (m_i - s_i) \mu_{i,i}$. This alternative representation implies that, for the X fluid model, we can determine the arrival rates by observing the queue lengths and spare capacities.

Let $Z_{i,j}^*$ be the optimal value of the variable $Z_{i,j}$. We start by stating some basic propositions, which serve to simplify our X -fluid-model optimization problem. We first reduce the number of variables from four to two. The following is immediate.

PROPOSITION 1. (no idle agents) *If we do not have $Q_1^* = Q_2^* = 0$, then there should be no idle agents, i.e., $S_j^* = 0$ or, equivalently, $Z_{1,j}^* + Z_{2,j}^* = m_j$ for $j = 1, 2$.*

As a consequence of Proposition 1, if $q_1 > 0$, $q_2 = 0$ and $s_2 > 0$, then necessarily $Z_{1,2}^* > 0$. Moreover, either $Z_{1,2}^* \geq s_2$ or $Q_1^* = Q_2^* = 0$.

We next show that inefficient sharing implies no two-way sharing.

PROPOSITION 2. (one-way sharing) *Since the service rates satisfy the inefficient-sharing condition $\mu_{1,1}\mu_{2,2} \geq \mu_{1,2}\mu_{2,1}$ in (2), it suffices to consider one-way sharing; i.e., $Z_{1,2}^*Z_{2,1}^* = 0$.*

Proof. Suppose that $Z_{1,2} > 0$ and $Z_{2,1} > 0$, so that we have sharing in both directions. It suffices to assume that $Q_1 > 0$ and $Q_2 > 0$. We will show that, for appropriate positive variables $x_{1,2}$ and $x_{2,1}$, if we replace $(Z_{1,2}, Z_{2,1})$ by $(Z_{1,2} - x_{1,2}, Z_{2,1} - x_{2,1})$, then both queue lengths will decrease until one of the variables $Z_{1,2} - x_{1,2}$ or $Z_{2,1} - x_{2,1}$ becomes 0 or both queues become empty. We define $x_{2,1}$ as an appropriate constant multiple of $x_{1,2}$, so that we have a single real variable. To do so, let $\gamma_i \equiv \lambda_i - Z_{i,1}\mu_{i,1} - Z_{i,2}\mu_{i,2} > 0$ for $i = 1, 2$. Then let $x_{2,1} \equiv \beta x_{1,2}$, where $\beta \equiv (\gamma_2\mu_{1,2} + \gamma_1\mu_{2,2})/(\gamma_2\mu_{1,1} + \gamma_1\mu_{2,1})$. Then we consider what happens as we increase $x_{1,2}$, assuming that β remains constant. Let $\Delta_i \equiv \theta_i(Q_i(0) - Q_i(x_{1,2}))$, where $Q_i(x_{1,2})$ denotes Q_i with the initial vector of sharing levels $(Z_{1,2}, Z_{2,1})$ replaced by $(Z_{1,2} - x_{1,2}, Z_{2,1} - \beta x_{1,2})$. Then

$$\Delta_1 = x_{1,2}\gamma_1 \left(\frac{\mu_{1,1}\mu_{2,2} - \mu_{1,2}\mu_{2,1}}{\gamma_2\mu_{1,1} + \gamma_1\mu_{2,1}} \right) \quad \text{and} \quad \Delta_2 = x_{1,2}\gamma_2 \left(\frac{\mu_{1,1}\mu_{2,2} - \mu_{1,2}\mu_{2,1}}{\gamma_2\mu_{1,1} + \gamma_1\mu_{2,1}} \right) \quad (7)$$

Clearly, $\Delta_i \geq 0$ for both i if and only if inequality (2) holds. Moreover, from (6) and (7), we see that both queues become empty at the same level of $x_{1,2}$. Hence, we can decrease both variables $Z_{1,2}$ and $Z_{2,1}$ by increasing $x_{1,2}$ until one of these variables becomes 0 or both queue lengths simultaneously become 0. ■

As a consequence of Proposition 2, we can re-express the basic optimization problem, first, in terms of two convex real-valued functions of a single real variable, $C_{1,2}$ and $C_{2,1}$, and second, in terms of a single combined convex function of a single real variable, C_c . Let 1_A be the indicator function of the set A ; i.e., $1_A(x) = 1$ if $x \in A$ and $1_A(x) = 0$ otherwise. We put the short proof required in §EC.3.1.

PROPOSITION 3. (single-variable functions) *Since the inefficient-sharing condition (2) holds, the optimal cost can be expressed as*

$$C^*(\lambda_1, \lambda_2) = C^*(q_1, s_1, q_2, s_2) = \min \{C_{1,2}(Z_{1,2}), C_{2,1}(Z_{2,1})\} = \min \{C_c(Z_{1,2} - Z_{2,1})\} \quad (8)$$

over $Z_{1,2}$ and $Z_{2,1}$ such that $0 \leq Z_{1,2} \leq m_2$, $0 \leq Z_{2,1} \leq m_1$ and $Z_{1,2}Z_{2,1} = 0$, where

$$\begin{aligned} C_{1,2}(Z_{1,2}) &\equiv C_{1,2}(Z_{1,2}; \lambda_1, \lambda_2) \equiv C \left(\frac{(\lambda_1 - m_1\mu_{1,1} - Z_{1,2}\mu_{1,2})^+}{\theta_1}, \frac{(\lambda_2 - (m_2 - Z_{1,2})\mu_{2,2})^+}{\theta_2} \right) \\ &\equiv C_{1,2}(Z_{1,2}; q_1, s_1 = 0, q_2, s_2) \\ &\equiv C \left(\frac{(q_1 - \mu_{1,2}Z_{1,2})^+}{\theta_1}, \frac{(q_2 - s_2\mu_{2,2} + \mu_{2,2}Z_{1,2})^+}{\theta_2} \right), \end{aligned} \quad (9)$$

$$\begin{aligned} C_c(Z_{1,2} - Z_{2,1}) &\equiv C_{1,2}(Z_{1,2} - Z_{2,1})\mathbf{1}_{\{Z_{1,2} - Z_{2,1} \geq 0\}} + C_{2,1}(-(Z_{1,2} - Z_{2,1}))\mathbf{1}_{\{Z_{1,2} - Z_{2,1} < 0\}} \\ &= C_{1,2}(Z_{1,2})\mathbf{1}_{\{Z_{1,2} > 0\}} + C_{2,1}(Z_{2,1})\mathbf{1}_{\{Z_{2,1} > 0\}} + C(q_1, q_2)\mathbf{1}_{\{Z_{1,2} = Z_{2,1} = 0\}}, \end{aligned} \quad (10)$$

with q_i and s_i defined in (5), satisfying $q_1s_2 = q_2s_1 = 0$, and $C_{2,1}(Z_{2,1})$ defined analogously to $C_{1,2}(Z_{1,2})$ in (9). The functions $C_{1,2}$ and $C_{2,1}$ are continuous strictly convex functions of the single real variables $Z_{1,2}$ and $Z_{2,1}$ over their domain. If, in addition, the stronger inefficient-sharing condition $\mu_{1,1} > \mu_{1,2}$ and $\mu_{2,2} > \mu_{2,1}$ in (1) holds, then C_c is also a continuous strictly convex function of the single real variable $Z_{1,2} - Z_{2,1}$ over the domain specified in Proposition 3.

COROLLARY 1. (three intervals) *If the stronger inefficient-sharing condition (1) holds, then for each pair of arrival rates (λ_1, λ_2) or initial state (q_1, s_1, q_2, s_2) (without sharing), there are two thresholds $\zeta_{1,2} \geq \zeta_{2,1}$ such that exactly one of the following occurs:*

- (i) $Z_{1,2}^* > 0$ and $Z_{2,1}^* = 0$ for $Z_{1,2} - Z_{2,1} > \zeta_{1,2}$,
 - (ii) $Z_{2,1}^* > 0$ and $Z_{1,2}^* = 0$ for $Z_{1,2} - Z_{2,1} < \zeta_{2,1}$
 - (iii) $Z_{1,2}^* = Z_{2,1}^* = 0$ for $\zeta_{2,1} \leq Z_{1,2} - Z_{2,1} \leq \zeta_{1,2}$.
- (11)

The value of Corollary 1 will be clear when we turn our attention to the queue ratio r below. We can apply Proposition 2 to get further simplification if there is initially spare capacity. Then, from

the beginning, we know that we can only have sharing with help provided by the pool with spare capacity; i.e., if $q_1 > 0 > s_2$, then $Z_{1,2}^* > 0$ and $Z_{2,1}^* = 0$, so that it suffices to minimize $C_{1,2}(Z_{1,2})$.

It is natural to have the cost function C be smooth, in which case the optimal solution can be found by simple calculus. Proposition EC.1 concludes that, if the optimal solution found by calculus falls outside the feasible set, then the actual optimum value is obtained at the nearest boundary point.

It is easy to see that there is a one-to-one correspondence between the queue ratio $r \equiv Q_1/Q_2$ and the real variable $Z_{1,2} - Z_{2,1}$ used to specify the optimization problem in Proposition 3. That implies that there is a one-to-one correspondence between the fixed-agent-allocation optimization problem (choosing $Z_{1,2}$ and $Z_{2,1}$) and the (fixed) queue-ratio control problem (choosing state-dependent queue-ratio functions $r_{1,2}$ and $r_{2,1}$) in the fluid-model context. We establish it formally in §EC.3.3.

Finally, we provide a basis for an efficient algorithm to determine the equivalent optimal controls. To do so, we effectively reduce the dimension from two to one by observing that special weighted sums of the queue lengths (and corresponding weighted sums of the arrival rates) are independent of the agent-assignment variables $Z_{1,2}$ and $Z_{2,1}$. We only state the result for $Z_{1,2}$; the corresponding result for $Z_{2,1}$ is stated in §EC.3.4. The proof is verification by direct computation, so we omit it. For understanding, it may be helpful to refer to Figure 4 in the next subsection.

PROPOSITION 4. (constant weighted queue lengths) *Let*

$$a_{1,2} \equiv \frac{\mu_{2,2}\theta_1}{\mu_{1,2}\theta_2} \quad \text{and} \quad \tilde{a}_{1,2} \equiv \frac{\mu_{1,2}}{\mu_{2,2}}. \quad (12)$$

Consider any initial state (λ_1, λ_2) , or equivalently (q_1, s_1, q_2, s_2) , with $s_1 = 0$. Then

$$\begin{aligned} w_{1,2} &\equiv a_{1,2} \left(\frac{\lambda_1 - m_1\mu_{1,1}}{\theta_1} \right) + \left(\frac{\lambda_2 - m_2\mu_{2,2}}{\theta_2} \right) = a_{1,2}q_1 + q_2 - \frac{s_2\mu_{2,2}}{\theta_2} \\ &= a_{1,2}Q_1(Z_{1,2}) + Q_2(Z_{1,2}) - \frac{S_2(Z_{1,2})\mu_{2,2}}{\theta_2} \end{aligned} \quad (13)$$

for all $Z_{1,2}$ with $0 \leq Z_{1,2} \leq m_2$.

Proposition 4 implies that the locus of all nonnegative queue-length vectors $(Q_1, Q_2) \equiv (Q_1(Z_{1,2}), Q_2(Z_{1,2}))$ associated with initial state (λ_1, λ_2) , or equivalently (q_1, s_1, q_2, s_2) , with $s_1 = 0$,

is on the line $\{(Q_1, Q_2) : a_{1,2}Q_1 + Q_2 = w_{1,2}\}$ in the nonnegative quadrant. Thus, for any nonnegative constant $w_{1,2}$, the optimal queue-length vector (Q_1^*, Q_2^*) and the optimal queue-ratio $r_{1,2}^* \equiv Q_1^*/Q_2^*$ restricted to one-way sharing ($Z_{2,1} = 0$) are the same for all initial states (q_1, s_1, q_2, s_2) with $s_1 = 0$ satisfying (13) provided that $q_1 \geq Q_1^*$. In that case, $a_{1,2}Q_1^* + Q_2^* = w_{1,2}$. That same optimal queue-length vector and optimal queue ratio holds for all arrival pairs (λ_1, λ_2) where $s_1 = 0$, $Z_{2,1} = 0$ and

$$\lambda_1 + \tilde{a}_{1,2}\lambda_2 = \tilde{w}_{1,2} \equiv \frac{\theta_1\theta_2w_{1,2} + a_{1,2}\theta_2m_1\mu_{1,1} + \theta_1m_2\mu_{2,2}}{a_{1,2}\theta_2}. \quad (14)$$

And similarly for sharing in the other direction; see §EC.3.4.

5.3. Computing the Optimal Queue-Ratio Functions

We now demonstrate how to numerically find the optimal state-dependent queue ratios $r_{1,2}^*$ and $r_{2,1}^*$ as functions of the fluid state (Q_1, S_1, Q_2, S_2) . With the thresholds, this gives us a state-dependent *queue-ratio control with thresholds* (QR-T). To illustrate, we consider a (nonseparable) quadratic cost function of the form

$$C(Q_1, Q_2) = 3Q_1^2 + 2Q_2^2 + Q_1Q_2 + 10Q_1 + 5Q_2. \quad (15)$$

For any vector of arrival rates (λ_1, λ_2) we can assign one, and only one, point in the (Q_1, Q_2) plane, which represents the queue lengths associated with these arrival rates, when there is no sharing. To represent spare capacity, we allow negative values; i.e., $-Q_i$ is shorthand for $-S_i\mu_{i,i}/\theta_i$. (We actually plot $(Q_1 - S_1\mu_{1,1}/\theta_1, Q_2 - S_2\mu_{2,2}/\theta_2)$ even though the axes are simply labelled Q_i .)

We apply Proposition 4 to find the optimal queue ratios. We first consider when pool 2 helps class 1. To treat that case, we let $\lambda_2 = m_2\mu_{2,2}$, so that class 2 has no queue before pool 2 helps class 1. We then assume that $\lambda_1 > m_1\mu_{1,1}$ so that class 1 is overloaded. We then choose a large set of positive weighted arrival sums $\{\tilde{w}_{1,2}^1, \dots, \tilde{w}_{1,2}^n\}$ and find the optimal queue ratio for each. In the first step, we let $\lambda_1 \equiv \tilde{w}_{1,2} - \tilde{a}_{1,2}\lambda_2$, using (14). We then write (15) as a function of $Z_{1,2}$, take its derivative and find the optimal $Z_{1,2}^*$. Plugging $Z_{1,2}^*$ in the queue equations gives us the optimal queue lengths (for the specific arrival rates), and the optimal queue ratio $r_{1,2}^*$. We repeat this for

every $\tilde{w}_{1,2}^i$ to get the curve $1/r_{1,2}^*$ depicted in Figure 4. To find the curve $1/r_{2,1}^*$ we go through essentially the same procedure for $Z_{2,1}^*$.

Figure 4 simultaneously depicts the three optimal sharing regions in the two-dimensional state space and the two curves of optimal queue ratios. It was generated using Matlab on a system with the following parameters: $m_1 = m_2 = 100$, $\mu_{1,1} = \mu_{2,2} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_1 = \theta_2 = 0.3$. In




Figure 4 Curves of the optimal queue ratios for an X model with parameters $\mu_{i,i} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$, $\theta_i = 0.3$, $m_i = 100$ and the quadratic cost function $C(Q_1, Q_2)$ from (15). Also shown are two possible initial states, depicted by stars, and the lines of constant weighted queue lengths that relate the initial states to (reciprocals of) the optimal queue ratios, depicted by circles. Negative queues stand for spare capacity.

addition, Figure 4 shows how to find the optimal queue ratio for two possible initial queue lengths denoted by stars. When the initial queue-length vector is $(Q_1, Q_2) = (150, 0)$ (equivalently, $\lambda_1 = 145$ and $\lambda_2 = 100$), then the optimal queue-length vector is $(Q_1^*, Q_2^*) = (76.5, 91.8)$ and the optimal queue ratio is $r_{1,2}^* \equiv Q_1^*/Q_2^* = 0.83$. This optimal queue ratio is the intersection of the curve $1/r_{1,2}$ with the line with slope $-a_{1,2}$ that passes through $(150, 0)$ (the circle on the $1/r_{1,2}$ curve). When

the initial queue-length vector is $(Q_1, Q_2) = (0, 150)$ (equivalently, $\lambda_1 = 100$ and $\lambda_2 = 145$), we get $r_{2,1}^* = 0.41$ and $(Q_1^*, Q_2^*) = (46.6, 112.8)$. The optimal queue ratio is also the intersection of the curve $1/r_{2,1}$ with the line with slope $-a_{2,1}$ that passes through $(0, 150)$ (the circle on the $1/r_{2,1}$ curve).

Both the $1/r_{1,2}^*$ and $1/r_{2,1}^*$ curves seem to be linear, although that is actually not quite the case; the $r_{i,j}^*$'s are not constants for this cost function. For example, we already noted that, for $(\lambda_1, \lambda_2) = (145, 100)$ the optimal queue ratio is $r_{1,2}^* = 0.83$. If we change λ_1 to 110 then the optimal ratio becomes 0.80. For the other sharing direction, if $(\lambda_1, \lambda_2) = (100, 145)$, then $r_{2,1}^* = 0.41$, but if we change λ_2 to 110, then the optimal ratio changes to 0.38.

The fact the the two optimal-ratio curves are nearly linear in Figure 4 suggests that we can approximate the optimal queue-ratio function by fixed queue ratios, depending only on the direction of sharing; i.e., we can use FQR-T with only two values: one for $r_{1,2}$ and the other for $r_{2,1}$. In our example we may choose to use $r_{1,2} = 0.8$ and $r_{2,1} = 0.4$. The cost for using a nearly optimal ratio is very small in the fluid approximation, and even smaller in the stochastic system.

To understand when the optimal queue-ratio functions are nearly linear, as in the example above, and what the structure should be more generally, we investigate structured cost functions in §EC.4. We obtain explicit analytical expressions in special cases. We focus on separable cost functions: $C(Q_1, Q_2) = C_1(Q_1) + C_2(Q_2)$, where each component cost function C_i is strictly convex, strictly increasing and twice differentiable. For example, we find that QR-T reduces to FQR-T exactly when $C_i(Q_i) = c_i Q_i^{n_i}$ with $n_1 = n_2$; the case $n_1 = n_2 = 2$ is close to (15).

PROPOSITION 5. (explicit solution) *When $C(Q_1, Q_2) = c_1 Q_1^2 + c_2 Q_2^2$, FQR-T is optimal for the X fluid model with*

$$\begin{aligned} r_{1,2}^* &\equiv \frac{a_{1,2}c_2}{c_1} = \frac{c_2\mu_{2,2}\theta_1}{c_1\mu_{1,2}\theta_2}, & r_{2,1}^* &\equiv \frac{a_{2,1}c_2}{c_1} = \frac{c_2\mu_{2,1}\theta_1}{c_1\mu_{1,1}\theta_2}, \\ Z_{1,2}^* &= \frac{(c_1\mu_{1,2}\theta_1)(q_1 - (s_1\mu_{1,1}/\theta_1)) - (c_2\mu_{2,2}\theta_2)(q_2 - (s_2\mu_{2,2}/\theta_2))}{c_1\mu_{1,2}^2/\theta_1 + c_2\mu_{2,2}^2/\theta_2}, \\ Z_{2,1}^* &= \frac{c_2\mu_{2,1}\theta_2(q_2 - (s_2\mu_{2,2}/\theta_2)) - c_1\mu_{1,1}\theta_1(q_1 - (s_1\mu_{1,1}/\theta_1))}{c_1\mu_{1,1}^2 + c_2\mu_{2,1}^2}. \end{aligned} \tag{16}$$

In Proposition 5, the cost is specified by a single parameter: The ratio c_1/c_2 specifies the relative importance of the two queues. (The remaining parameter is equivalent to choosing the monetary units.) Finally, we caution that other cases (e.g., linear costs) can be quite different; see §EC.4.

5.4. Application to the Stochastic Model

We can directly apply the QR-T control derived above to the stochastic model. Figure 4 identifies three sharing regions to apply to the stochastic process $(Q_1(t), S_1(t), Q_2(t), S_2(t))$ once sharing has been activated. There are two regions for each direction of sharing; e.g., if sharing has been activated with pool 2 helping class 1, available pool-2 agents serve class-1 customers when the queue-length vector falls in the lower right region, whereas there is no sharing in the other two regions. The way to share is described in §4.2.

6. Choosing the Thresholds

We now consider how to choose the thresholds $\kappa_{1,2}$ and $\kappa_{2,1}$. These thresholds have two important roles: First, they automatically detect when the system becomes overloaded and, second, they prevent unwanted sharing in normal loading. If the thresholds are too large, then the queues may not reach them during the overload. (Abandonments necessarily keep the queues from increasing without bound, even under overloads.) On the other hand, as discussed in §4.1, if the thresholds are too small, then sharing may be activated too often, so that we may get inefficient sharing.

Unfortunately, the fluid analysis cannot reveal the “right” size of the thresholds, since the fluid queues are empty under normal loading. We need to understand the extent of the stochastic fluctuations, something which is not captured by the fluid approximation. At this point, it is convenient to apply many-server heavy-traffic limits to gain additional insight. To understand the general idea, it suffices to refer to established limits for the basic $M/M/n + M$ model, as in Garnett et al. (2002) and Whitt (2004). There, both fluid models and refined diffusion process models are obtained as limits as the scale increases, where scale is measured by the number n of servers. What is unusual here, though, is that we are simultaneously interested in the quality-and-efficiency-driven (QED) regime and the efficiency-driven (ED) or overloaded regime.

The QED regime is appropriate to describe normal loading, which is what prevails before the overload occurs, while the ED regime is appropriate to describe the overloaded system. In both cases, the arrival rate is allowed to grow as $n \rightarrow \infty$, while the service rate and abandonment rate are held fixed. The important insight is that the queue lengths tend to be of order $O(\sqrt{n})$ in the QED regime, as depicted by the diffusion limit in the QED regime, while the queue lengths tend to be of order $O(n)$ in the ED regime, as depicted by the fluid limit in the ED regime.

Thus, to prevent unwanted sharing when the system is normally loaded, we should choose the thresholds to be of size bigger than $O(\sqrt{n})$. That ensures that the weighted-queue-difference processes, $D_{1,2}$ and $D_{2,1}$, will not move above the thresholds by random fluctuations. On the other hand, we should choose the thresholds to be $o(n)$, so that the thresholds will be asymptotically negligible compared to the $O(n)$ fluid content. Then, asymptotically, they will be exceeded instantaneously when the overload occurs and they will not significantly alter the queue ratios. From this simple reasoning, we see that it suffices to have $\kappa_{i,j}^{(n)} = O(n^p)$ as $n \rightarrow \infty$ for $1/2 < p < 1$. (Incidentally, that scaling also makes the thresholds out of reach in normal loading in the law-of-the-iterated logarithm scaling of $(n \log \log n)^{1/2}$.)

This asymptotic analysis shows that the thresholds chosen in this way are asymptotically optimal, both during normal loading and during overload incidents. Asymptotically, the thresholds will be exceeded negligibly often during normal loading; asymptotically, the thresholds do not alter the optimal average cost in the overload incidents. For the case of normal loading, we can apply the QED results in Garnett et al. (2002); for the overload incidents, the ED results in Whitt (2004) provide only heuristic support, because they apply only to the $M/M/n + M$ model. We intend to prove the ED fluid model for the X model with FQR-T in a subsequent paper.

Of course, we actually have a system with one fixed n . When we want to apply the theory to a real system, with a finite number of agents, it becomes hard to distinguish between $O(n)$ and $O(\sqrt{n})$. For example, if $n = 100$, then both $10 = 0.1n$ and $10 = \sqrt{n}$. Thus, as in any application of asymptotic results, we should numerically verify that the values chosen are appropriate, and refine them if necessary, for which we can use simulation. For example, for a system with 100 agents in

each pool (and abandonment rates less than service rates), we found that $\kappa_{i,j} = 10$ is effective. We found that the performance is not too sensitive to the choice of the thresholds, provided that they are neither too small nor too large. We present simulation results for FQR-T under normal loading in §EC.5.2, including a sensitivity analysis for the thresholds.

7. Simulation Experiments

Our analysis has been based on a fluid approximation of a stochastic system. It remains to show that the fluid approximation is suitably accurate for the stochastic system and that the optimal control for the fluid model works well in the stochastic system. For those purposes, we conduct simulation experiments.

7.1. Accuracy Of The Fluid Approximation

In this subsection we investigate the accuracy of the approximation. To show how the accuracy increases as the system becomes larger, we simulated three cases, each case represents an element in a sequence of queueing systems indexed by n , scaled to satisfy a many-server heavy-traffic limit in the ED regime as $n \rightarrow \infty$. We use the same fixed service and abandonment rates as before ($\mu_{i,i} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_i = 0.3$). We consider a fixed queue ratio $r = 1$. We let the arrival rates be $\lambda_1 = 1.3n$ and $\lambda_2 = n$, when there are n agents in each service pool. The three cases we consider are $n = 25, 100, 400$. We let the thresholds for these three values of n be $\kappa_{1,2} = \kappa_{2,1} = 3, 10, 30$, respectively. (The thresholds were dropped when exceeded.)

Table 1 shows the results. Each result is the average of 5 independent simulation runs having 300,000 arrivals in each run. The half-width of 95% confidence intervals, calculated using the t random variable with 4 degrees of freedom, are also given.

To show both the actual performance and the convergence to the fluid limit as n increases, we display both the direct values and the scaled values, dividing by n . Since the scaled values tend to be nearly independent of n , we witness the heavy-traffic fluid limit. We see that the approximations get better as n increases, but they are already not too bad when $n = 25$.

	n=25		n=100		n=400	
perf. meas.	approx.	sim.	approx.	sim.	approx.	sim.
Q_1	13.9	13.5 ± 0.4	55.6	52.8 ± 1.2	222.2	216.7 ± 7.0
Q_1/n	0.56	0.54 ± 0.02	0.56	0.53 ± 0.01	0.56	0.54 ± 0.02
Q_2	13.9	15.7 ± 0.5	55.6	58.4 ± 1.2	222.2	223.1 ± 7.0
Q_2/n	0.56	0.63 ± 0.02	0.56	0.58 ± 0.01	0.56	0.56 ± 0.02
ratio	1.0	0.98 ± 0.02	1.0	0.90 0.00	1.0	0.96 0.00
$Z_{1,2}$	4.2	4.8 ± 0.2	16.7	17.7 ± 0.3	66.7	66.4 ± 2.2
$Z_{1,2}/n$	0.17	0.19 ± 0.01	0.17	0.18 ± 0.00	0.17	0.17 ± 0.01
Cost (thousands)	1.37	1.79 ± 0.01	19.35	20.28 ± 0.81	299.6	299.8 ± 19.2

Table 1 A comparison of steady-state performance measures in the fluid approximation with corresponding simulation results for the Markovian X model. For each n , there are n agents in each pool, with $\lambda_1 = 1.3n$, $\lambda_2 = n$, $\mu_{i,i} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_i = 0.3$. The thresholds $\kappa_{1,2} = \kappa_{2,1}$ are 3, 10, 30 for $n = 25, 100, 400$, respectively.

7.2. Comparing The Two Controls

In the fluid analysis, choosing the number of agents in each pool that are helping customers from the other class is equivalent to choosing the queue ratio, $r_{i,j}$. However, that is not the case in the actual stochastic system. With specified numbers of agents serving customers from the other class, the queue ratio fluctuates randomly. With specified queue ratios, the numbers of agents helping the other class fluctuates randomly. Moreover, with specified numbers of agents serving customers from the other class, the two queue-length processes evolve independently. In sharp contrast, with specified queue ratios, the queue-length processes are strongly dependent, as in Figure EC.8. This suggests that there is a big difference between the two controls in a real, stochastic system. We thus expect the average cost under FQR-T to be different than the average cost when fixing $Z_{i,j}$. We conducted simulation experiments to compare the two controls.

To compare the two controls - FQR-T, and fixed $Z_{i,j}$ - we simulated a system with $m_i = 100$ agents in each pool, arrival rates $\lambda_1 = 130$ and $\lambda_2 = 100$, service rates $\mu_{1,1} = \mu_{2,2} = 1$ and $\mu_{1,2} = \mu_{2,1} = 0.8$, and abandonment rates $\theta_1 = \theta_2 = 0.3$. Since class 1 is overloaded, we took $\kappa_{2,1} = \kappa_{1,2} = 10$, but once

we go over the threshold $\kappa_{1,2}$, we drop it, so that it becomes $\kappa_{1,2} = 0$.

Figure 5 presents simulation results comparing the two average costs for five different cases: (1) $r_{1,2} = 1.2, Z_{1,2} = 15$, (2) $r_{1,2} = 1, Z_{1,2} = 17$, (3) $r_{1,2} = 0.83, Z_{1,2} = 19$ (optimal point), (4) $r_{1,2} = 0.6, Z_{1,2} = 22$ and (5) $r_{1,2} = 0.4, Z_{1,2} = 25$. For each point, we fixed the queue-ratio $r_{1,2}$, and used FQR-T with this ratio. For each such $r_{1,2}$, there is an equivalent $Z_{1,2}$ in the fluid equations. Since this $Z_{1,2}$ is not necessarily an integer, we rounded it up to the smallest integer larger than $Z_{1,2}$, i.e., we used $\lceil Z_{1,2} \rceil$ in each simulation of the fixed- $Z_{i,j}$ control. According to the fluid approximation, the optimal queue ratio is $r_{1,2} = 0.83$, and the respective optimal $Z_{1,2}$ is equal to 18.4, rounded up to 19 in the simulation experiments.

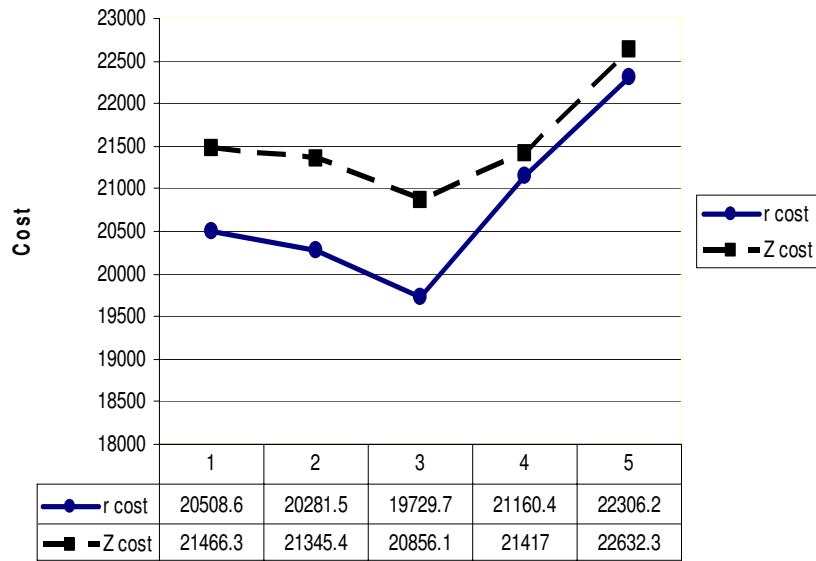


Figure 5 A Comparison of the cost of using FQR-T with different ratios and the cost of fixing the equivalent $Z_{1,2}$'s. In order, the five ratios considered are $r = 1.2, 1, 0.83, 0.6, 0.4$, where $r = 0.83$ is the optimal queue ratio according to the fluid model. Based on the fluid model, the equivalent $Z_{1,2}$'s are, respectively, 15, 17, 19, 22, 25, where each $Z_{1,2}$ has been rounded up to the smallest integer greater than its actual value. The average cost of 5 independent simulation runs is written beneath each point. The cost function is (15), and the system parameters are $m_i = 100$, $\lambda_1 = 130$, $\lambda_2 = 100$, $\mu_{i,i} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_i = 0.3$.

For each case, we conducted 5 independent simulation runs using FQR-T, and 5 independent simulation runs with a fixed $Z_{1,2}$, each run with 300,000 arrivals. The independent replications make

it possible to reliably estimate confidence intervals using the t statistic with 4 degrees of freedom. The large number of arrivals ensures that the transient behavior in the beginning of the simulation, before reaching steady state, does not affect the final simulation estimates of the steady-state averages. Additional simulation results are given in Table EC.1 in §EC.5, including the half-width of 95% confidence intervals and a comparison of the simulation to the fluid approximation.

There are several interesting observations to be made: First, the r -cost curve lies significantly below the Z -cost curve, which shows that FQR-T is a superior control. At the optimal point for FQR-T, the average cost under FQR-T is about 5.4% smaller than the average cost under the fixed- $Z_{1,2}$ control.

Secondly, FQR-T tends to be a more robust control. Small changes in r do not produce large changes in the cost. Note that the largest r value here (1.2) is 3 times as large as the smallest r value (0.4), whereas the largest Z value here (25) is only 1.6 times as large as the smallest Z value (15). Moreover, the average costs when using FQR-T with $r_{1,2} = 1.2$ and $r_{1,2} = 1$ are still smaller than the cost of fixing $Z_{1,2}$ at its optimal value. For further discussion, see §EC.5.

8. Conclusions

In this paper we studied ways to respond to unexpected overloads in large-scale service systems. We considered the Markovian X model with two customer classes and two service pools, assuming that agents are more effective serving customers from their own class than customers from the other class, as specified by the inefficient-sharing conditions in (1) and (2). Thus we want negligible sharing under normal loads, but we want to activate sharing when there is an unexpected overload at an unanticipated time, without knowing what the new arrival rates will be.

The main ideas for analyzing the performance and determining appropriate queue-ratio functions for the *queue-ratio with thresholds* (QT-R) and *fixed-queue ratio with thresholds* (FQR-T) controls we propose are: (i) to use steady-state analysis and (ii) to apply an approximating deterministic fluid model. The QR-T and FQR-T controls proposed for the actual stochastic system are direct applications of the corresponding optimal controls derived for the fluid model in §5.2. We developed

an algorithm to find the optimal queue-ratio curves for a general convex cost function in Proposition 4 and §5.3. The resulting QR-T control is easily understood as a partition of the state space into three sharing regions, as depicted in Figure 4, with two regions for each direction of sharing.

In Proposition 5 we also provided strong justification for FQR-T when the cost function has the form $C(Q_1, Q_2) = c_1 Q_1^2 + c_2 Q_2^2$ for some constants c_1 and c_2 . In that case, we proved that FQR-T is optimal for the fluid model (i.e., the optimal QR-T control reduces to an instance of FQR-T) and exhibited the explicit optimal queue-ratio parameters. Then the optimal queue-ratio parameters depend on the cost function C only via the single parameter c_1/c_2 , which succinctly captures the relative importance of the two queues. For other sharing regions, see §EC.4.

The main ideas for gaining insight into appropriate threshold values were to apply (i) many-server heavy-traffic asymptotics and (ii) simulation. Heuristically, we showed that the thresholds should be asymptotically optimal simultaneously for periods of normal loading and for periods of overload. Asymptotically, no tradeoff need be made. The requirement is that the thresholds should be of order $O(n^p)$ as $n \rightarrow \infty$, where $1/2 < p < 1$ and n is the system scale factor. We used simulation to verify that the thresholds work well for given finite n .

Our FQR-T (or QR-T) control is appealing for several reasons. First, it is automatic and simple; we need not directly discover the arrival rates in order to find out when overloads occur, and then decide what amount of sharing should be done. Instead, FQR-T automatically detects the time the system becomes overloaded, and then automatically enforces the optimal ratio, by observing only the size of the two queues. It is easier to use the information about the queues, which is readily available, than to use information about the arrival rates, which is not readily available. Moreover, simulation experiments indicate that FQR-T performs better (produces lower expected costs) than fixing $Z_{i,j}$ at their optimal values, even with known arrival rates; see Figure 5.

It remains to mathematically prove that the fluid model with QR-T or FQR-T arises as the many-server heavy-traffic limit of scaled overloaded queueing systems as the scale increases. It also remains to show that FQR-T is asymptotically optimal among all possible controls in such a

limit. It remains to develop similar controls for more complex multi-class multi-pool systems with non-exponential service-time and abandonment-time distributions.

References

- Aksin, Z., M. Armony, V. Mehrotra. 2007. The modern call center: a multi-disciplinary perspective on operations management research. *Production Oper. Management* **16** (6) 665–688.
- Armony, M., I. Gurvich. 2006. When promotions meet operations: cross-selling and its effect on call-center performance. Stern School of Business, NYU.
- Bassamboo, A., A. Zeevi. 2009. Near optimal data-driven staffing for large call centers. *Oper. Res.*, forthcoming.
- Bell, S. L., R. J. Williams. 2005. Dynamic scheduling of a parallel server system in heavy traffic with complete resource pooling: asymptotic optimality of a threshold policy. *Electronic J. Prob.* **10** 1044–1115.
- Bhandari, A., A. Scheller-Wolf, M. Harchol-Balter. 2008. An exact and efficient algorithm for the constrained dynamic operator staffing problem for call centers. *Management Sci.* **54** (2) 339–353.
- Gans, N., G. Koole, A. Mandelbaum. 2003. Telephone call centers: tutorial, review and research prospects. *Manufacturing Service Oper. Management* **5** (2) 79–141.
- Garnett, O., A. Mandelbaum, M. I. Reiman. 2002. Designing a call center with impatient customers. *Manufacturing Service Oper. Management* **4** (3) 208–227.
- Gurvich, I., W. Whitt. 2007a. Service-level differentiation in many-server service systems via queue-ratio routing. *Oper. Res.*, forthcoming. (<http://www.columbia.edu/~ww2040/recent.html>)
- Gurvich, I., W. Whitt. 2007. Scheduling Flexible Servers with Convex Delay Costs in Many-Server Service Systems. *Manufacturing Service Oper. Management*, forthcoming.
- Tezcan, T., J. G. Dai. 2006. Dynamic control of N -systems with many servers: asymptotic optimality of a static priority policy in heavy traffic. Georgia Institute of Technology.
- Whitt, W. 2004. Efficiency-driven heavy-traffic approximations for many-server queues with abandonments. *Management Sci.* **50** (10), 1449–1461.
- Whitt, W. 2005. Engineering solution of a basic call-center model. *Management Science* **51** (2) 221–235.
- Whitt, W. 2006b. Staffing a call center with uncertain arrival rate and absenteeism. *Production Oper. Management* **15** (1) 88–102.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

e-Companion

In this online e-companion we present additional material supplementing the main paper. The topics are ordered as they arise in the paper. In §EC.1 we discuss the way the transient distribution approaches its steady-state limit, both at the beginning and the end of an overload incident. In §EC.2 we provide additional discussion about the FQR and FQR-T controls, supplementing §4. In §EC.3 we present additional details about the optimal solution for the deterministic fluid model during the overload, supplementing §5. Finally, in §EC.5 we present additional simulation results about the performance of the control. In §EC.5.1 we present a table of detailed simulation results supporting Figure 5. In §EC.5.2 we present additional simulation results about the performance of FQR-T under normal loading. We perform a sensitivity analysis for the thresholds there.

EC.1. Time To Reach Steady State

An important aspect of our QR-T and FQR-T controls is the transient behavior of the system. When the overload incident occurs, the system must shift from steady state under normal loading to steady state under the overload. Afterwards, at the end of the overload period, there is a recovery period, during which the system shifts back to the original steady state. From analysis and extensive simulations, we conclude that these two transient periods do not dominate, so that it is possible to use steady-state analysis as a reasonable approximation. In this section, we provide some supporting simulation results and discuss the supporting mathematical results.

EC.1.1. Simulation Experiments

We start by doing a simulation experiment of an overload incident, including all five regimes: (i) steady state before the overload, (ii) transition to new steady state at the beginning of the overload, (iii) new steady state under the overload, (iv) recovery period and (v) original steady state again after the overload.

Our example is based on Example 1 in the main paper and the associated typical overload incident described at the end of §3. We assume that there is an overload incident that lasts 5 hours when the mean service times are 5 minutes. Given that we measure time in units of mean service

times, the overload incident lasts 60 time units. Thus, we simulate the system over the time interval $[0, 150]$, and have the overload begin at time 80 and end at time 140. Thus, the initial transient begins at time 80, while the recovery period begins at time 140.

We consider a large system with $n = 400$ agents in each pool. For the normal loading, we let $\lambda_1 = \lambda_2 = 380$; for the overload during $[80, 140]$, we let $\lambda_1 = 520$, while $\lambda_2 = 380$ as before. As in Example 1, we let the mean service time for customers served by designated agents be $\mu_{1,1}^{-1} = \mu_{2,2}^{-1} = 1.0$, while the mean service time for customers served by agents from the other pool is $\mu_{1,2}^{-1} = \mu_{2,1}^{-1} = 1.25$. We let customers abandon at rate $\theta_1 = \theta_2 = 0.4$.

Since class 1 experiences the overload, we will have pool 2 helping class 1 during the overload incident. Typical sample paths of the processes $Z_{1,2}(t)$ and $Q_1(t)$ generated by simulation are shown in Figures EC.1 and EC.2 below. A dotted horizontal line depicts the steady-state fluid approximation during the overload. We do not show the other processes. From corresponding plots of $Q_2(t)$ and $Q_1(t)$, it is evident that they move together during the overload, reflecting state-space collapse, but they move independently during normal loading. From the displayed sample path, we




Figure EC.1 Sample path of $Z_{1,2}(t)/400$ for FQR-T, with overload incident in $[80, 140]$ for $n = 400$.



Figure EC.2 Sample path of $Q_1(t)$ for FQR-T, with overload incident in $[80, 140]$ for $n = 400$.

see that the system indeed reaches a new steady state after a few mean service times, as claimed in §§2-3.

To elaborate, we also show corresponding sample paths in Figures EC.3 and EC.4 with $n = 100$

agents in each pool. One important observation is that in both systems ($n = 100$ and $n = 400$) it takes less than 3 time units for the queues to hit their fluid value, denoted by the dotted horizontal lines. The recovery time, after the overload incident has ended, is also very short, and is about 2 time units for the queues in both systems.

The story is different for the $Z_{1,2}(t)$ process. To make the connection between the two cases clear, we present the **proportion** of class-1 customers in pool 2 instead of the actual number, i.e., we show $Z_{1,2}(t)/n$ in Figures EC.1 and EC.3. First, when the overload begins at time 80, it takes some time until the queues hit the threshold $\kappa_{1,2}$. That is the reason why $Z_{1,2}(t)$ starts growing a bit after time 80. It is interesting to see how our choice of the thresholds influences this delay. Recall that we choose the thresholds to be of order of size less than $O(n)$ but greater than $O(\sqrt{n})$; see §6 for more details. In these simulations, we took $\kappa_{i,j} = 20$ for $n = 400$ and $\kappa_{i,j} = 10$ for $n = 100$. This explains why in the $n = 400$ system it takes less time for $Z_{1,2}(t)$ to start increasing than in the $n = 100$ system: The thresholds are relatively smaller for the bigger system.

We also observe a difference between the two systems after the arrival rates return to normal at time 140. At this time, the $Z_{1,2}(t)$ processes start decreasing immediately and in a very fast rate. But now, service-pool 2 stops serving class-1 customers faster in the small system. Let $T_{1,2}$ be the time it takes for pool 2 to stop serving all class-1 customers after the end of the overload incident (after 140 in our example). As an approximation, we have

$$E[T_{1,2}] \approx \sum_{j=1}^r \frac{1}{j \cdot \mu_{1,2}},$$

where $r \equiv Z_{1,2}(140)$. Hence, the larger $Z_{1,2}(140)$ is, the longer it takes $Z_{1,2}(t)$ (or equivalently, $Z_{1,2}(t)/n$) to reach zero after the arrival rates shift back to normal. Yet, in both cases $Z_{1,2}(t)/n$ drops below 0.1 in about 2 time units, so that the total service rate in service-pool 2 is greater than λ_2 in 2 time units after the shift. In summary, we see that the transient period is relatively short, and a steady-state analysis is reasonable to apply.



Figure EC.3 Sample path of $Z_{1,2}(t)/100$ for FQR-T, with overload incident in $[80, 140]$ for $n = 100$.




Figure EC.4 Sample path of $Q_1(t)$ for FQR-T, with overload incident in $[80, 140]$ for $n = 100$.

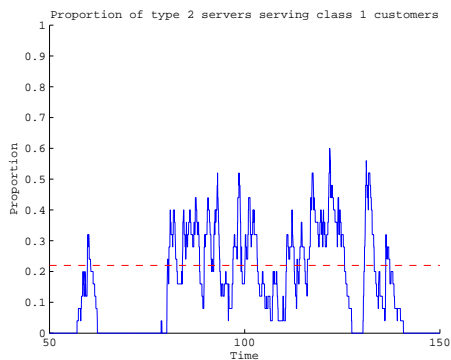


Figure EC.5 Sample path of $Z_{1,2}(t)/25$ for FQR-T, with overload incident in $[80, 140]$ for $n = 25$.




Figure EC.6 Sample path of $Q_1(t)$ for FQR-T, with overload incident in $[80, 140]$ for $n = 25$.

EC.1.2. Mathematical Analysis

We now provide further support. We first review mathematical analysis of the $M/M/n+M$ model; we next contrast with single-server models; afterwards we discuss implications for our X model

The $M/M/n+M$ model. Consider the $M/M/n+M$ model with arrival rate λ , service rate μ and abandonment rate θ . First, it is useful to consider the special case in which $\theta = \mu$; then the number in system is distributed the same as in an $M/M/\infty$ system with service rate $\theta = \mu$. Thus the number in system at time t has a Poisson distribution for each fixed initial state. An explicit expression for the mean $m(t)$ at time t , starting empty, is given in (20) of Eick et al. (1993). More generally, the mean $m(t)$ satisfies an ordinary differential equation (ODE); see Corollary 4

of Eick et al. (1993). These results show that $m(t)$ and the entire distribution reaches steady state approximately at time c/μ , some constant c times the mean service time $1/\mu$. The constant c depends on our criterion; the critical time constant is $1/\mu$, a mean service time.

For the more general overloaded $M/M/n+M$ model (without assuming that $\theta = \mu$), it is helpful to consider the deterministic fluid approximation in Whitt (2004). Formula (2.17) there shows that the fluid approximation for the number in queue, $q(t)$, starting with all the servers busy, again evolves as the $M/M/\infty$ ODE, but with arrival rate $\lambda - n\mu$ and service rate θ . That implies that the fluid queue content (approximating the number in queue), starting from all servers busy but no queue, reaches steady state approximately at time c/θ , some constant c times the mean abandonment time $1/\theta$. That too will be approximately c/μ provided that θ is not too different from μ . The critical time constant here is $1/\theta$, a mean time to abandon.

To illustrate this mathematical analysis, we do a simulation of the $M/M/n+M$ model. We base our example here on Example 1 in §1. In that example, the service rates in both pools are $\mu_{i,i} = 1$, the abandonment rates are $\theta_i = 0.4$ and the number of agents in each pool is 100. In this example the arrival rates changed at some instant from $(\lambda_1, \lambda_2) = (90, 90)$ to $(\lambda_1, \lambda_2) = (130, 90)$. We show what happens if class 1 receives no help from service-pool 2. Then the class-1 queue behaves like an $M/M/100 + M$ queue. Figure EC.7 depicts a simulated sample path of an $M/M/100 + M$ queue, when the system is initialized empty at time 0. The average steady-state queue length in the overload incident is about 75, and it can be seen that this steady-state value is reached within about 4 time units when the system is initialized empty. (Time is measured in units of mean service times). If we assume, as in our example above, that the system was operating before the arrival rates changed, then most of the agents were probably busy, and the time to reach the new steady state is about 2 time units (two mean service times).

Single-Server models. In §2 we stated that the number in system tends to approach steady state more quickly in many-server queues with abandonment than in single-server queues without abandonment. We should begin with a qualification: Slow approach to steady state occurs for single-server systems without abandonment when the system is heavily loaded. For single-server

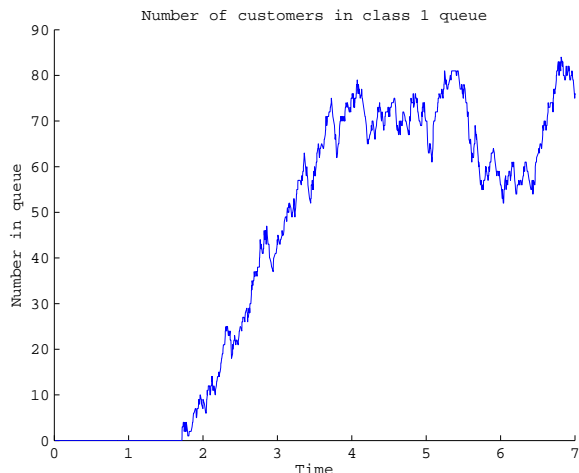


Figure EC.7 Time to reach steady state when $m_1 = m_2 = 100$, $\lambda_1 = 130$, $\lambda_2 = 90$, $\mu_{1,1} = \mu_{2,2} = 1$,
 $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_1 = \theta_2 = 0.4$.

queues, we refer to Section III.7.3 of Cohen (1982) on the relaxation time. Sections 4.6 and 5.1 of Whitt (1989) gives conventional heavy-traffic approximations (when $\rho \uparrow 1$ with n fixed, where $\rho \equiv \lambda/n\mu$ is the traffic intensity) for the time required for the mean number in system to reach steady state in the general $G/G/n$ model with fixed n and without customer abandonment. The time required to reach steady state is approximately $c/(1-\rho)^2$ mean service times, where c is a constant depending on the number of servers, n , the variability of the arrival and service processes (quantified explicitly) and again the criterion. Clearly the time to reach steady state can be quite long when ρ is high.

The X model. For our X model, there are two implications of the $M/M/n + M$ analysis above: First, when the overload incident begins, the queue length should be negligible, so that the fluid content in a newly overloaded queue will grow approximately linearly at rate $\lambda - n\mu$, because the opposite force $\theta q(t)$ will be small, since $q(t)$ is initially small. That means that the threshold will be quickly passed if there is a significant unbalanced overload.

For our more complicated X model with the QR-T control, after the threshold has been exceeded, the theoretical analysis for the $M/M/n + M$ model above provides a rough heuristic analysis indicating what should happen, but the actual evolution still depends on the state of the six-dimensional Markov chain $(Q_i(t), Z_{i,j}(t); i = 1, 2; j = 1, 2)$. Thus we rely on simulation to confirm

that the actual behavior is indeed similar to what occurs in these simple many-server models. We remark that the state-space collapse discussed in the next subsection indicates that $(Q_1(t), Q_2(t))$ should evolve approximately as a one-dimensional process, suggesting that the analysis above should not be too far off when the service rates $\mu_{i,j}$ do not differ greatly.

EC.2. More on FQR

In this section we present additional background on FQR; for more, see Gurvich and Whitt (2007a,b,c,d). We first illustrate the state-space collapse (SSC), which is the topic of Gurvich and Whitt (2007c). The conditions for SSC are satisfied if either the service rates only depend upon the customer class or the service rates only depend upon the agent pool. To illustrate, suppose that the service rates are independent of both class and pool, with $\mu_{1,1} = \mu_{1,2} = \mu_{2,1} = \mu_{2,2} = 1.0$. Figure EC.8 shows the plots of typical sample paths of the two queue-length processes when $\lambda_1 = \lambda_2 = m_1 = m_2 = 100$ and $\theta_1 = \theta_2 = 0.2$. From Figure EC.8, we can clearly see the SSC.




Figure EC.8 State-Space Collapse

We observed that, with FQR, it is possible to choose the ratio parameter r (or, equivalently, the queue proportions p_i) in order to determine the optimal level of staffing to achieve desired service-level differentiation. For example, under normal loading, our goal may be to choose staffing levels as small as possible subject to having 80% of class-1 customers wait less than 20 seconds, while 80% of class-2 customers wait less than 60 seconds. To see how this can be done with FQR, let T_i be

the class- i delay target (e.g., $T_1 = 0.033$ and $T_2 = 0.100$ for 20 seconds and 60 seconds if the mean service times are 10 minutes); let W_i be the class- i waiting time before starting service; let p_i be the queue proportion determined by r . As explained in Gurvich and Whitt (2007a), the following string of approximations show how the individual class- i performance targets $P(W_i > T_i) \leq \alpha$, for both i , can be reduced into a single-class single-pool performance target $P(W > T) \leq \alpha$ for an appropriate choice of the queue proportions p_i and the aggregate target T :

$$\begin{aligned} P(W_i > T_i) &\approx P(Q_i > \lambda_i T_i) \approx P(p_i Q_\Sigma > \lambda_i T_i) \approx P\left(Q_\Sigma > \sum_{k=1}^2 \lambda_k T_k\right) \\ &\approx P\left(\lambda W > \sum_{k=1}^2 \lambda_k T_k\right) \approx P(W > T) \leq \alpha, \end{aligned} \quad (\text{EC.1})$$

where we define $p_i \equiv \lambda_i T_i / (\lambda_1 T_1 + \lambda_2 T_2)$, $\lambda \equiv \lambda_1 + \lambda_2$ and $T \equiv (\lambda_1 T_1 + \lambda_2 T_2) / (\lambda_1 + \lambda_2)$. The first approximation in (EC.1) follows by a heavy-traffic generalization of Little's law, establishing that the steady-state queue-length and waiting-time random variables are related approximately by $Q_i \approx \lambda_i W_i$. The second approximation in (EC.1) is due to SSC: $Q_i \approx p_i Q_\Sigma$. The third approximation is obtained by choosing p_i as specified above. The fourth approximation in (EC.1) follows from the heavy-traffic generalization of Little's law once again, for the entire system: $Q_\Sigma \approx \lambda W$ for λ as defined above, where W is the waiting time for an arbitrary customer. The fifth and final approximation follows by the appropriate definition of the aggregate target T , as defined above. With this reduction, we can determine the overall staffing by using elementary established methods for the single-class single-pool model. That is, we choose the total number of agents, m , so that $P(W > T) \leq \alpha$ in the $M/M/m + M$ model. We then let $m_i = p_i m$. From (EC.1) and the fact that $r = p_1 / (1 - p_1)$, we see that the required ratio is

$$r = \frac{p_1}{1 - p_1} = \frac{\lambda_1 T_1}{\lambda_2 T_2}. \quad (\text{EC.2})$$

For the X model (and more generally), Theorem 4.1 of Gurvich and Whitt (2007d) shows that, if the service rates only depend on the service pool or the class (but not both), then FQR is asymptotically optimal to minimize linear staffing costs subject to service-level constraints, as above, in the QED many-server heavy-traffic regime.

As was shown in §4.1, with inefficient sharing FQR without the thresholds we add in FQR-T can cause the queues in the general X-model system to explode when there is no abandonment, because of the inefficient sharing. We now show that there is also serious performance degradation when we include customer abandonment. We use the same example as in §4.1, only adding abandonments with rates $\theta_1 = \theta_2 = 0.2$. As before, there are 100 agents in each pool. The arrival rates are $\lambda_1 = \lambda_2 = 99$ and the service rates are $\mu_{1,1} = \mu_{2,2} = 1$ and $\mu_{1,2} = \mu_{2,1} = 0.8$. To describe the performance degradation, we compare the performance to the no-sharing case. When there is no sharing, 2% of the customers abandon, the mean queue length is 10 and the mean conditional waiting time given that the customer is served is 0.10. On the other hand, for FQR with $r = 1$, again about 39% of the agents are busy serving customers from the other class. That reduces the effective service rate for each class from 100 to 92.2. As a consequence, about 7% of the customers abandon, the mean queue length is 34 and the average conditional waiting time given that the customer is served is 0.35.

Figures EC.9 and EC.10 show the sample paths of the number of agents in pool 1 helping class-2 customers, and the class-1 queue, respectively. Due to the symmetry of the system in our example, the $Z_{2,1}$ and Q_2 figures are very similar, and the fluid approximations for both queues and $Z_{i,j}$'s are equal.




Figure EC.9 Sample path of $Z_{2,1}(t)/100$ for FQR with $r = 1$ and abandonments at rate $\theta_1 = \theta_2 = 0.2$.




Figure EC.10 Sample path of $Q_1(t)$ for FQR with $r = 1$ and abandonments at rate $\theta_1 = \theta_2 = 0.2$.

In contrast, to illustrate how FQR-T performs, we consider the same example: Example 2 with abandonments at rate $\theta_i = 0.2$. We let $r_{1,2} = r_{2,1} = 1$, so that there is no change from FQR above, and we let the thresholds be $\kappa_{1,2} = \kappa_{2,1} = 10$. The results of a simulation experiment are shown in Figures EC.11 and EC.12. Numerical values were given in §4.2. The performance is greatly




Figure EC.11 Sample path of $Z_{2,1}(t)/100$ with FQR-T, $r = 1$.

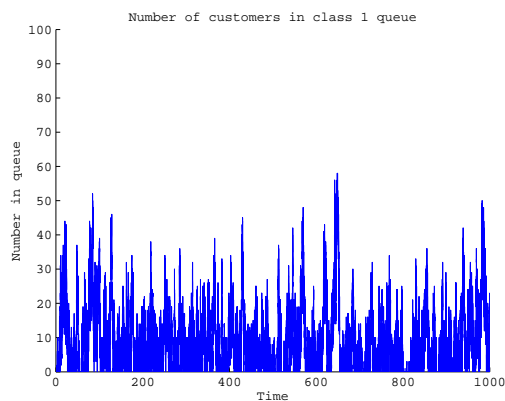


Figure EC.12 Sample path of $Q_1(t)$ with FQR-T, $r = 1$.

improved with FQR-T.

EC.3. Optimal Solution for the Fluid Model

In this section we provide additional material supplementing §5.

EC.3.1. Proof of Proposition 3.

The representation is an immediate consequence of Proposition 2. Since $C_{1,2}(Z_{1,2})$ is the composition of a strictly convex function and a linear function, it is a strictly convex function of $Z_{1,2}$; e.g., p. 38 of Rockafellar (1970); similarly for $C_{2,1}(Z_{2,1})$. To establish the convexity of C_c , first assume that C is differentiable. It suffices to show that the derivative of C_c with respect to $Z_{1,2} - Z_{2,1}$, denoted by C'_c , is nondecreasing. Existence and monotonicity of the derivative C'_c away from the boundary point $Z_{1,2} - Z_{2,1} = 0$ follows from the differentiability and convexity of $C_{1,2}$ and $C_{2,1}$, assuming that C is differentiable and convex. However, even if C is differentiable, the derivative of C_c need not exist at $Z_{1,2} - Z_{2,1} = 0$. It suffices to show that the left derivative is less than the right derivative at this point. The right derivative of C_c at 0, denoted by $C_c'^+(0)$, coincides

with the derivative $C'_{1,2}(0)$, while the left derivative of C_c at 0, denoted by $C_c'^-(0)$, coincides with $-C'_{2,1}(0)$. Let C'_i denote the partial derivative of C with respect to its i^{th} coordinate at the argument $(q_1 - (s_1\mu_{1,1}/\theta_1), q_2 - (s_2\mu_{2,2}/\theta_2))$, which is positive because C is increasing. Then observe that

$$C'_{1,2}(0) = -C'_1\left(\frac{\mu_{1,2}}{\theta_1}\right) + C'_2\left(\frac{\mu_{2,2}}{\theta_2}\right) \quad \text{and} \quad -C'_{2,1}(0) = -C'_1\left(\frac{\mu_{1,1}}{\theta_1}\right) + C'_2\left(\frac{\mu_{2,1}}{\theta_2}\right)$$

Hence, $C'_{1,2}(0) \geq -C'_{2,1}(0)$, so that $C_c'^+(0) > C_c'^-(0)$ if the two inequalities in (1) hold. These relations can be extended to non-differentiable functions C by working with left and right derivatives. ■

EC.3.2. Optimal Values Beyond the Boundaries

It is natural to have the cost function C be smooth, in which case the optimal solution can be found by simple calculus. The following result concludes that, if the optimal solution found by calculus falls outside the feasible set, then the actual optimum value is obtained at the nearest boundary point. Let $a \wedge b \equiv \min\{a, b\}$ and $a \vee b \equiv \max\{a, b\}$. We omit the proof, which is a standard convexity result.

PROPOSITION EC.1. (optimal values beyond the boundaries) *Let $\bar{Z}_{1,2}$ and $\bar{Z}_{2,1}$ be the values of $Z_{1,2}$ and $Z_{2,1}$ yielding minimum values of $C_{1,2}$ and $C_{2,1}$ in (9), and let $\hat{Z}_{1,2}$ and $\hat{Z}_{2,1}$ be the corresponding values yielding the minima ignoring the constraints in Proposition 3. Then $\bar{Z}_{1,2} = \hat{Z}_{1,2} \vee 0 \wedge m_2$, $\bar{Z}_{2,1} = \hat{Z}_{2,1} \vee 0 \wedge m_1$ and $(Z_{1,2}^*, Z_{2,1}^*)$ can assume only two possible values: $(\bar{Z}_{1,2}, 0)$ or $(0, \bar{Z}_{2,1})$.*

EC.3.3. The Relation between \mathbf{r} and \mathbf{Z}

In §5.2 we observed that there is a one-to-one correspondence between the queue ratio $r \equiv Q_1/Q_2$ and the real variable $Z_{1,2} - Z_{2,1}$ used to specify the optimization problem in Proposition 3. That implies that there is a one-to-one correspondence between the fixed-agent-allocation optimization problem (choosing $Z_{1,2}$ and $Z_{2,1}$) and the queue-ratio control problem (choosing a state-dependent queue-ratio r) in the fluid-model context.

PROPOSITION EC.2. (relating r and $Z_{1,2} - Z_{2,1}$) *For any given arrival-rate vector (λ_1, λ_2) or initial state (q_1, s_1, q_2, s_2) (without sharing), the queue ratio $r \equiv Q_1/Q_2$ is a strictly decreasing differentiable function of $Z_{1,2} - Z_{2,1}$, denoted by ϕ , as $Z_{1,2} - Z_{2,1}$ varies over its allowed domain in Proposition 3. Thus, the function ϕ has a unique inverse ϕ^{-1} and there exists a unique optimal $r^* \equiv r^*(q_1, s_1, q_2, s_2)$, which is characterized by*

$$r^* = \phi^{-1}(Z_{1,2}^* - Z_{2,1}^*), \quad (\text{EC.3})$$

where both r^* and $Z_{1,2}^* - Z_{2,1}^*$ are understood to be functions of the initial state (q_1, s_1, q_2, s_2) . Moreover, there are two thresholds $\eta_{1,2} > \eta_{2,1}$ such that we want one-way sharing with pool 2 helping class 1 if $r > \eta_{1,2}$, in which case we let $r_{1,2} = r^*$; we want one-way sharing with pool 1 helping class 2 if $r < \eta_{2,1}$, in which case we let $r_{2,1} = r^*$; and we want no sharing at all if $\eta_{2,1} \leq r \leq \eta_{1,2}$. The thresholds are obtained from the thresholds $\zeta_{1,2}$ and $\zeta_{2,1}$ in Corollary 1 by $\eta_{1,2} = \phi^{-1}(\zeta_{1,2})$ and $\eta_{2,1} = \phi^{-1}(\zeta_{2,1})$.

Proof. By (9), when pool 2 helps class 1, Q_1 is a strictly decreasing differentiable function of $Z_{1,2}$ and while Q_2 is a strictly increasing differentiable function of $Z_{1,2}$. On the other hand, when pool 1 helps class 2, Q_1 is a strictly increasing differentiable function of $Z_{2,1}$ and while Q_2 is a strictly decreasing differentiable function of $Z_{2,1}$. Thus $r \equiv Q_1/Q_2$ is a strictly decreasing differentiable function of $Z_{1,2} - Z_{2,1}$ over its domain, as claimed. ■

EC.3.4. Constant Weighted Queue Length

We now complete Proposition 4 by exhibiting the result for pool 1 helping class 2.

PROPOSITION EC.3. (constant weighted queue lengths with pool 1 helping class 2) *Let*

$$a_{2,1} \equiv \frac{\mu_{2,1}\theta_1}{\mu_{1,1}\theta_2} \quad \text{and} \quad \tilde{a}_{2,1} \equiv \frac{\mu_{2,1}}{\mu_{1,1}}. \quad (\text{EC.4})$$

Consider any initial state (λ_1, λ_2) , or equivalently (q_1, s_1, q_2, s_2) , with $s_2 = 0$. Let

$$w_{2,1} \equiv a_{2,1} \left(\frac{\lambda_1 - m_1\mu_{1,1}}{\theta_1} \right) + \left(\frac{\lambda_2 - m_2\mu_{2,2}}{\theta_2} \right) = a_{2,1} \left(q_1 - \frac{s_1\mu_{1,1}}{\theta_1} \right) + q_2. \quad (\text{EC.5})$$

Then

$$a_{2,1} \left(Q_1(Z_{1,2}) - \frac{S_1(Z_{2,1})\mu_{1,1}}{\theta_1} \right) + Q_2(Z_{2,1}) = w_{2,1} \quad (\text{EC.6})$$

for all $Z_{2,1}$ with $0 \leq Z_{2,1} \leq m_1$.

Just as with Proposition 4, Proposition EC.3 implies that the locus of all nonnegative queue-length vectors $(Q_1, Q_2) \equiv (Q_1(Z_{2,1}), Q_2(Z_{2,1}))$ associated with initial state (λ_1, λ_2) , or equivalently (q_1, s_1, q_2, s_2) , with $s_2 = 0$, is on the line $\{(Q_1, Q_2) : a_{2,1}Q_1 + Q_2 = w_{2,1}\}$ in the nonnegative quadrant. Thus, for any nonnegative constant $w_{2,1}$, the optimal queue-length vector (Q_1^*, Q_2^*) and the optimal queue-ratio $r_{2,1}^* \equiv Q_1^*/Q_2^*$ restricted to one-way sharing ($Z_{1,2} = 0$) are the same for all initial states (q_1, s_1, q_2, s_2) with $s_2 = 0$ satisfying (13) and $q_2 \geq Q_2^*$. Moreover, $a_{2,1}Q_1^* + Q_2^* = w_{2,1}$. That same optimal queue-length vector and optimal queue ratio holds for all arrival pairs (λ_1, λ_2) where $s_2 = 0$, $Z_{1,2} = 0$ and

$$\lambda_1 + \tilde{a}_{2,1}\lambda_2 = \tilde{w}_{2,1} \equiv \frac{\theta_1\theta_2w_{2,1} + a_{2,1}\theta_2m_1\mu_{1,1} + \theta_1m_2\mu_{2,2}}{a_{2,1}\theta_2}. \quad (\text{EC.7})$$

EC.4. Structured Separable Cost Functions

At the end of §5.3, we observed that we can obtain explicit analytical expressions for the optimal ratio control if we impose additional structure on our cost function. We give the main results in this section and provide supporting details in the next section.

EC.4.1. Main Results

We first assume that C is separable, i.e., $C(Q_1, Q_2) = C_1(Q_1) + C_2(Q_2)$, where each component cost function C_i is strictly convex, strictly increasing and twice differentiable. We start by assuming that the derivatives C'_i are strictly increasing, so that their inverses exist. Let $\Psi(Q_1) \equiv C'_1(Q_1)$ and let Ψ^{-1} be its inverse. Then one of the following relations between the queue lengths should hold, when we choose the one that minimizes the cost:

$$Q_1 = \Psi^{-1}(a_{1,2}C'_2(Q_2)) \quad \text{or} \quad Q_1 = \Psi^{-1}(a_{2,1}C'_2(Q_2)), \quad (\text{EC.8})$$

for $a_{1,2}$ defined in (12) and $a_{2,1}$ defined in §EC.3.4. If C'_1 is not strictly increasing, then we work with the left-continuous inverse of Ψ defined by $\Psi^{\leftarrow} \equiv \{x : \Psi(x) \geq y\}$.

Power functions. If the cost functions C_i are simple power functions, i.e., $C_i(Q_i) \equiv c_i Q_i^{n_i}$ for $i = 1, 2$, then we have that either $Q_1^* = r_{1,2}^* Q_2^{*(n_2-1)/(n_1-1)}$ or $Q_1^* = r_{2,1}^* Q_2^{*(n_1-1)/(n_2-1)}$, where

$$r_{1,2}^* \equiv \sqrt[n_1-1]{a_{1,2} c_2 n_2 / c_1 n_1} \quad \text{and} \quad r_{2,1}^* \equiv \sqrt[n_2-1]{a_{2,1} c_2 n_2 / c_1 n_1}. \quad (\text{EC.9})$$

When $n_1 = n_2$, Q_1^*/Q_2^* is a fixed queue ratio, either $r_{1,2}^*$ or $r_{2,1}^*$ for $r_{i,j}^*$ as in (EC.9). Thus, we need only to decide which way we should share, and then use FQR-T with the appropriate $r_{i,j}^*$; i.e., we are in the setting of Figure 4 with constant ratios for which we have explicit expressions.

Quadratic functions. In practice it may be difficult to actually specify an appropriate cost function. Thus, for practical application we suggest quadratic functions: $C_i(Q_i) \equiv c_i Q_i^2 + b_i Q_i + a_i$ for $i = 1, 2$. These functions might be obtained by performing an approximation (e.g., via Taylor series approximation to an analytical expression or least squares fit to data). In this case, we have either

$$Q_1^* - r_{1,2}^* Q_2^* = k_{1,2}^*, \quad \text{or} \quad Q_1^* - r_{2,1}^* Q_2^* = k_{2,1}^*, \quad (\text{EC.10})$$

for

$$r_{1,2}^* \equiv \frac{a_{1,2} c_2}{c_1} = \frac{c_2 \mu_{2,2} \theta_1}{c_1 \mu_{1,2} \theta_2}, \quad r_{2,1}^* \equiv \frac{a_{2,1} c_2}{c_1} = \frac{c_2 \mu_{2,1} \theta_1}{c_1 \mu_{1,1} \theta_2}. \quad (\text{EC.11})$$

and

$$k_{1,2}^* \equiv \frac{a_{1,2} b_2 - b_1}{2c_1}, \quad k_{2,1}^* \equiv \frac{a_{2,1} b_2 - b_1}{2c_1}. \quad (\text{EC.12})$$

In other words, we keep a fixed-queue ratio centered about a constant $k_{i,j}^*$ instead of zero. That is, we employ new thresholds after sharing has been activated. (The current thresholds $k_{i,j}^*$ are not to be confused with the thresholds $k_{i,j}$ used with the queue-difference processes in (4) to test for the occurrence of overloads. We use $k_{1,2}^*$ only after sharing with pool 2 helping class 1.) From the two formulas in (EC.11), we directly see how these ratio parameters and thresholds should depend on the model parameters. In particular, each ratio is either directly proportional or inversely proportional to each of six model parameters.

Quadratic and linear power functions. A natural simple cost function is the quadratic power function, which is a special case of the general power function with $n_1 = n_2 = 2$ and a special case

of the general quadratic function with $b_1 = b_2 = a_1 = a_2 = 0$. The optimal control then is precisely FQR-T ($k_1^* = k_2^* = 0$), as indicated in Proposition 5. In §EC.4.2 we also discuss the special cases $n_1 = 2, n_2 = 1$ and $n_1 = 1, n_2 = 1$.

FQR-T without cost functions. Clearly, FQR-T could be employed directly without specifying any cost function, using engineering judgment to set the parameters. Even if that is the case, the queue-ratio formulas in (EC.11) and possibly the centering formulas in (EC.12) provide important insight into how the control parameters should depend on the model parameters.

EC.4.2. Supporting Details About Structured Separable Cost Functions

We now supplement §EC.4 by providing more details about the fluid model with a separable cost function. As before, we assume that C is separable, and that each component cost-function C_i is strictly convex, strictly increasing and twice differentiable. We then relax the strictly-increasing assumption, and consider linear functions.

Let $C(Q_1, Q_2)$ be a separable function. We can write C as a function of one variable $Z_{1,2}$ or $Z_{2,1}$, depending on which way the sharing is done.

$$C(Q_1, Q_2) = C_1(Q_1) + C_2(Q_2) \equiv C_{1,(i,j)}(Z_{i,j}) + C_{2,(i,j)}(Z_{i,j}) \equiv C(Z_{i,j}), \quad (\text{EC.13})$$

where

$$C_{1,(1,2)}(Z_{1,2}) \equiv C_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} - \frac{Z_{1,2} \mu_{1,2}}{\theta_1} \right), \quad C_{2,(1,2)}(Z_{1,2}) \equiv C_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} + \frac{Z_{1,2} \mu_{2,2}}{\theta_2} \right).$$

and

$$C_{1,(2,1)}(Z_{2,1}) \equiv C_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} + \frac{Z_{2,1} \mu_{1,1}}{\theta_1} \right), \quad C_{2,(2,1)}(Z_{2,1}) \equiv C_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} - \frac{Z_{2,1} \mu_{2,1}}{\theta_2} \right).$$

Hence, the optimal $Z_{1,2}$ is achieved when

$$C'(Z_{1,2}) = -C'_{1,(1,2)}(Z_{1,2}) \left(\frac{\mu_{1,2}}{\theta_1} \right) + C'_{2,(1,2)}(Z_{1,2}) \left(\frac{\mu_{2,2}}{\theta_2} \right) = 0,$$

or equivalently,

$$C'_1(Q_1) = \frac{\mu_{2,2} \theta_1}{\mu_{1,2} \theta_2} C'_2(Q_2) \equiv a_{1,2} C'_2(Q_2). \quad (\text{EC.14})$$

Similarly, the optimal $Z_{2,1}$ is achieved when

$$C'_1(Q_1) = \frac{\mu_{2,1}\theta_1}{\mu_{1,1}\theta_2} C'_2(Q_2) \equiv a_{2,1} C'_2(Q_2). \quad (\text{EC.15})$$

The fact that C_i is strictly convex implies that $C''_i \geq 0$. If $C''_i > 0$ then C'_i is strictly increasing, and its inverse function exists. Let $\Psi(Q_1) \equiv C'_1(Q_1)$ and let Ψ^{-1} be its inverse. Then one of the following relations between the queues should hold:

$$\text{either } Q_1 = \Psi^{-1}(a_{1,2}C'_2(Q_2)) \quad \text{or} \quad Q_1 = \Psi^{-1}(a_{2,1}C'_2(Q_2)), \quad (\text{EC.16})$$

where we choose the relation that minimizes the cost-function $C(Q_1, Q_2)$.

If the inverse of Ψ does not exist, then we can work with the left-continuous inverse of Ψ defined by $\Psi^{\leftarrow}(y) \equiv \{x : \Psi(x) \geq y\}$.

We now consider separable cost functions of the form:

$$C(Q_1, Q_2) = c_1 Q_1^{n_1} + c_2 Q_2^{n_2}, \quad n_1, n_2 \in \mathbb{N}, \quad (\text{EC.17})$$

where each component is a power function. The optimal solution is given in the main paper. We observe that the compatible-ratio condition $r_{1,2} \geq r_{2,1}$ in S 4.2 holds, because

$$\frac{r_{1,2}^*}{r_{2,1}^*} = \frac{{}^{n_1-1}\sqrt{a_{1,2}}}{\sqrt{a_{2,1}}} = \frac{{}^{n_1-1}\sqrt{\mu_{1,1}\mu_{2,2}}}{\mu_{1,2}\mu_{2,1}} \geq 1,$$

under the inefficient-sharing condition (2). When $n_1 = n_2$ we get

$$\frac{Q_1^*}{Q_2^*} = r_{1,2}^* \quad \text{or} \quad \frac{Q_1^*}{Q_2^*} = r_{2,1}^*;$$

i.e., it is optimal to keep a fixed-queue ratio. Thus, we need only to decide on which way we should share, and then use FQR-T with the appropriate fixed queue ratio $r_{i,j}^*$.

These results explain why the optimal ratios in our numerical example with the cost function in (15) in §5.3 are almost constant. In the numerical example there are other terms, but the dominating ones are the quadratic terms. As the queues get larger, the influence of the smaller-power terms decreases, and the optimal ratios converge to fixed numbers. If the function is separable (as would

be the case if our example had not had the Q_1Q_2 term), then the convergence is to the same ratios as if the only terms are $c_1Q_1^n + c_2Q_2^n$. The mixed terms of power n change these numbers. For the cost function in (15), the Q_1Q_2 terms is also of power 2, and hence the optimal ratios converge to different numbers than (EC.9). But for that example, clearly the optimal ratios are nearly constant.

In §EC.4 we introduced the general separable quadratic cost function to provide a tractable approximation for a broad range of possible cost functions. We observed that the optimal queue-ratio function becomes a shifted version of FQR-T, which is just FQR-T centered at points $k_{1,2}^*$ and $k_{2,1}^*$ instead of centered at 0. We now illustrate the resulting control for a candidate cost function. In order to make the linear components have approximately equal weight to the quadratic components when the queue lengths are about 50, we divide the coefficients c_i for the quadratic terms by 10. We also omit the mixed term Q_1Q_2 , which violated the separability property. Instead of the cost function in (15), we now consider the cost function

$$C(Q_1, Q_2) \equiv 0.3Q_1^2 + 10Q_1 + 0.2Q_2^2 + 5Q_2. \quad (\text{EC.18})$$

The centering is depicted by the y-intercepts on the two lines in Figure EC.13.




Figure EC.13 The optimal queue ratios (shifted FQR) for an X model with parameters $\mu_{i,i} = 1$,

$\mu_{1,2} = \mu_{2,1} = 0.8$, $\theta_i = 0.3$, $m_i = 100$ and the separable quadratic cost function

$$C(Q_1, Q_2) \equiv 0.3Q_1^2 + 10Q_1 + 0.2Q_2^2 + 5Q_2 \text{ in (EC.18).}$$

We now consider the two linear cases. When one or more of the component cost functions C_i is

linear, we are led to modify our control. We indicate how our fluid-model analysis can be applied to generate alternative controls in these cases, but we do not examine their performance here.

$\mathbf{n}_1 = \mathbf{2}, \mathbf{n}_2 = \mathbf{1}$. The cost-function $C(Q_1, Q_2) = c_1 Q_1^2 + c_2 Q_2$ has one quadratic term and one linear term. The special structure of this function (C_2 not strictly convex) changes the control. Now, there is no longer dependence on the two queues, since Q_2 no longer comes into play. By (EC.16),

$$Q_1^* = \frac{a_{1,2}c_2}{2c_1} \equiv k_{1,2}^* \quad \text{or} \quad Q_1^* = \frac{a_{2,1}c_2}{2c_1} \equiv k_{2,1}^*.$$

Thus, we are no longer trying to keep a relation between the two queues, but instead we keep Q_1 not bigger than $k_{1,2}^*$ or $k_{2,1}^*$, depending which is optimal to use. To keep Q_1 at its optimal target, we modify our control: If class 1 is overloaded such that $q_1 > k_{1,2}^*$, then whenever $D_{1,2} \geq \max\{\kappa_{1,2}, k_{1,2}^*\}$ every newly available agent takes his next customer from the head of queue 1. Otherwise, every agent takes his next customer from the head of its own class queue.

If class 2 is overloaded, then we can have $Z_{2,1} > 0$ as long as we keep $Q_1 < k_{2,1}^*$. Hence, if $D_{2,1} < \kappa_{2,1}$ and $Q_1 < k_{2,1}^*$, then every newly available agent takes his next customer from the head of Q_2 . Otherwise, he will take his next customer from the head of his own class queue.

$\mathbf{n}_1 = \mathbf{1}, \mathbf{n}_2 = \mathbf{1}$. The purely-linear cost function $C(Q_1, Q_2) = c_1 Q_1 + c_2 Q_2$ is even more different than the functions we considered so far. However, it is well known that a linear function attains its minima on the boundaries of its domain. In our setting, this means that we either try to keep the queue that needs help at zero, or that we do not help it at all. When $Z_{1,2} > 0$, we have

$$\begin{aligned} C(Q_1, Q_2) &= C(Z_{1,2}) = c_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} - \frac{\mu_{1,2}}{\theta_1} Z_{1,2} \right) + c_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} + \frac{\mu_{2,2}}{\theta_2} Z_{1,2} \right) \\ &= c_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} \right) + c_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} \right) + \left(\frac{c_2 \mu_{2,2}}{\theta_2} - \frac{c_1 \mu_{1,2}}{\theta_1} \right) Z_{1,2}. \end{aligned}$$

Thus, if

$$\frac{c_1 \mu_{1,2}}{\theta_1} \leq \frac{c_2 \mu_{2,2}}{\theta_2} \tag{EC.19}$$

the function $C(Z_{1,2})$ is increasing, and its minima is attained when $Z_{1,2} = 0$. Otherwise, the function is decreasing, and it is optimal to take $Z_{1,2}$ as large as needed to ensure $Q_1 = 0$ (the simple calculus

gives us that $Z_{1,2}^* = m_2$, but of course class 1 may not have enough arrivals to fill both service pools). This means that we either share completely, or not share at all.

Similarly, for

$$\begin{aligned} C(Q_1, Q_2) = C(Z_{2,1}) &= c_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} - \frac{\mu_{1,1}}{\theta_1} Z_{2,1} \right) + c_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} + \frac{\mu_{2,1}}{\theta_2} Z_{2,1} \right) \\ &= c_1 \left(q_1 - \frac{s_1 \mu_{1,1}}{\theta_1} \right) + c_2 \left(q_2 - \frac{s_2 \mu_{2,2}}{\theta_2} \right) + \left(\frac{c_1 \mu_{1,1}}{\theta_1} - \frac{c_2 \mu_{2,1}}{\theta_2} \right) Z_{2,1}, \end{aligned}$$

if

$$\frac{c_2 \mu_{2,1}}{\theta_2} \leq \frac{c_1 \mu_{1,1}}{\theta_1} \quad (\text{EC.20})$$

then $C(Z_{2,1})$ is increasing, and its minima is attained at $Z_{2,1} = 0$. Otherwise, $C(Z_{2,1})$ is decreasing, and it is optimal to take $Z_{2,1}$ as large as needed (and possible) to make sure that $Q_2 = 0$.

Rewriting the inefficient-sharing condition (2), we get

$$\frac{\mu_{1,1}}{\mu_{2,1}} \geq \frac{\mu_{1,2}}{\mu_{2,2}}.$$

If the two inequalities (EC.19) and (EC.20) hold together, then

$$\frac{\theta_1 c_2}{\theta_2 c_1} \leq \frac{\mu_{1,2}}{\mu_{2,2}} \quad \text{and} \quad \frac{\mu_{1,1}}{\mu_{2,1}} \leq \frac{\theta_1 c_2}{\theta_2 c_1}, \quad (\text{EC.21})$$

but this contradicts the inefficient-sharing condition above, unless all the inequalities hold as equalities. Thus, we can have at most one of the inequalities, (EC.19) or (EC.20), hold under (2).

At first glance, it may seem from the discussion above that, when the holding cost is linear, we should not consider the system as an X model, but rather as an N model (sharing can be done in only one direction), if either (EC.19) or (EC.20) hold, or two independent I systems (no sharing at all), if none of these two inequalities hold. But that is not so. If there is spare capacity in one class, while the other class is overloaded, then it is always optimal to use the extra agents to help the overloaded class. Since we do not know what the overload incident will produce, we cannot restrict the model to an N model in advance.

Let us summarize what we have found: The cost analysis leads us to give priority to either queue 1 or queue 2. Suppose that it is optimal to give priority to queue 1. That leads us to set

the threshold for pool 2 helping class 1 at $\kappa_{1,2} = 0$. In the fluid model that will either produce the desired result $Q_1 = 0$ or $Q_1 > 0$ and $Z_{1,2} = m_2$, with pool 2 devoting all its effort to class 1. There remains another case: when pool 1 has spare capacity. In that case, within the fluid model, if pool 2 is overloaded, then pool 1 should devote all the required spare capacity to serving class 2. We should have $Z_{2,1} = s_1 \wedge q_2$. If $s_1 > q_2$, then the help pool 1 provides to class 2 makes both queues empty, and there is remaining spare capacity for pool 1. On the other hand, if $s_1 \leq q_2$, then we have exactly $Z_{2,1} = s_1$. Overall, there are three possible end results in the fluid model: (i) $Q_1 > 0$ and $Z_{1,2} = m_2$, (ii) $Q_1 = Q_2 = 0$, (iii) $Q_2 > 0$, $Q_1 = 0$ and $Z_{2,1} = s_1$.

We now must consider how to implement that control in the actual system. As indicated above, to give priority to queue 1 at all times, we can set $\kappa_{1,2} = 0$, and we always allow pool 2 to help class 1, even if $Z_{2,1} > 0$. The only difficulty is detecting whether or not pool 1 has spare capacity, so that we can have pool 1 helping class 2. For this purpose, we propose using a positive queue threshold for queue 2: We let an available agent in pool 1 help class 2 if, and only if, $Q_2 > \kappa_{2,1}$, $Q_1 = 0$ and $Z_{1,2} = 0$.

Since we allow pool 2 to serve class 1 all the time, we could possibly have simultaneous two-way sharing (both $Z_{1,2} > 0$ and $Z_{2,1} > 0$), but there should be only minimal simultaneous two-way sharing. It remains to further investigate this case.

EC.5. Additional Simulation Results

In this section we present additional simulation results.

EC.5.1. Comparing the Two Controls

We now supplement the comparison of the two controls (the fixed staffing levels versus QR-T) in §7.2 by presenting detailed simulation results. These are given in Table EC.1, including the half-width of 95% confidence intervals and a comparison of the simulation to the fluid approximation.

As stated before, for each case, we conducted 5 independent simulation runs using QR-T, and 5 independent simulation runs with a fixed $Z_{1,2}$, each run with 300,000 arrivals. The independent replications make it possible to reliably estimate confidence intervals using the t statistic with

4 degrees of freedom. The large number of arrivals ensures that the transient behavior in the beginning of the simulation, before reaching steady state, does not affect the final simulation estimates.

		Cost (in thousands)		actual ratio			actual $Z_{1,2}$	
policy		Approx.	Sim.	Approx.	Sim.	std.	Approx.	Sim.
FQR-T	$r = 1.20$	19.65	20.51 ± 0.64	1.20	1.07 ± 0.00	0.16 ± 0.01	15.0	15.9 ± 0.4
	$r = 1$	19.35	20.28 ± 0.81	1.00	0.90 ± 0.00	0.13 ± 0.01	16.7	17.7 ± 0.3
	$r = 0.83$	19.25	19.73 ± 0.64	0.83	0.76 ± 0.00	0.11 ± 0.00	18.4	18.9 0.5
	$r = 0.60$	19.56	21.16 ± 0.77	0.60	0.56 ± 0.00	0.08 ± 0.00	21.4	22.1 ± 0.3
	$r = 0.40$	20.75	22.31 ± 0.92	0.40	0.37 ± 0.00	0.06 ± 0.00	25.0	25.3 ± 0.3
fixed $Z_{1,2}$	$Z_{1,2} = 15$	19.65	21.47 ± 0.57	1.20	1.52 ± 0.08	1.93 ± 0.29	15.0	15.0 ± 0.0
	$Z_{1,2} = 17$	19.32	21.35 ± 0.46	0.96	1.13 ± 0.11	1.17 ± 0.45	17.0	17.0 ± 0.0
	$Z_{1,2} = 19$	19.26	20.86 ± 0.37	0.78	0.87 ± 0.07	0.75 ± 0.57	19.0	19.0 ± 0.0
	$Z_{1,2} = 22$	19.69	21.42 ± 0.60	0.56	0.61 ± 0.05	0.38 ± 0.04	22.0	22.0 ± 0.0
	$Z_{1,2} = 25$	20.75	22.63 ± 0.86	0.40	0.42 ± 0.01	0.33 ± 0.14	25.0	25.0 ± 0.0

Table EC.1 Full simulation results of Figure (5). The ‘approx’ columns show the anticipated results according to the fluid model, and the ‘sim.’ columns show the simulation results, together with half-width confidence interval.

We now provide additional observations about our simulation results for this example. Another important observation is that FQR-T is doing a better job in keeping the ratio between the two queues close to the desired ratio. The accuracy becomes even better when the system is larger (see the “ratio” row in Table 1 in the $n = 400$ columns). We have also included a column showing the simulated standard deviations of the ratios. Note how small the standard deviations are when using FQR-T, in comparison to the standard deviations when using the fixed- $Z_{1,2}$ control. Since FQR-T is working towards keeping the ratio between the two queues fixed throughout, the ratio between the two queues at any time point is approximately $r_{1,2}$. It also makes the two queues

strongly positively correlated, which reduces the overall variance. In contrast, under the fixed- $Z_{1,2}$ control, the two queues are independent with zero correlation.

The simulated ratio was calculated as a long-run average of the ratio between the two queues throughout the simulation time. We can compare it to Q_1/Q_2 from Table 1 which appears in §7.1 (in the $n = 100$ columns) which is also approximately 0.9. These agree closely because of state-space collapse, as in Figure EC.8 discussed in §EC.2.

Finally, Table EC.1 shows that the fluid approximation tends to underestimate the actual average cost in the stochastic model. That is understandable, because the fluid model ignores stochastic fluctuations, which will tend to increase the average costs with a convex cost function. However, note that the fluid approximation does do an excellent job in describing the relative costs. In particular, the fluid model succeeds in locating the correct minima for both controls.

EC.5.2. Performance of FQR-T Under Normal Loading

In §4.1 we saw that FQR-T performs well when $\kappa_{1,2} = \kappa_{2,1} = 10$ for Example 2 with $n = 100$ servers and $\lambda_i = 99$, showing that FQR without thresholds and one-way sharing can perform poorly. To supplement those simulation results and the simulation results in §7, in this section we present additional simulation results. As in Table 1, we consider three values of n : $n = 25$, $n = 100$ and $n = 400$. We let the arrival rates in both queues be $\lambda_i^{(n)} = 0.98n$. The service-rate and abandonment-rate parameters are fixed at $\mu_{i,i} = 1.0$, $\mu_{1,2} = \mu_{2,1} = 0.8$, $\theta_i = 0.2$. We let the thresholds be $\kappa_{1,2}^{(n)} = \kappa_{2,1}^{(n)} = 0.1n$, rounded up to 3.0 for $n = 25$. We compare the results of FQR-T to the $M/M/n + M$ model, which would prevail if there were absolutely no sharing at all. As before, we see that the mean queue lengths are actually slightly smaller with FQR-T. That shows that the little sharing that takes place with FQR-T is not so bad.

Due to the symmetry of the system under the parameters we chose, there is no difference between the steady-state values of both queues and service pools. Thus, we display only Q_1 and $Z_{1,2}$. We can see that as the system becomes larger, the sharing decreases, and the queue size gets closer to the queue length in an $M/M/n + M$ model.

	n=25		n=100		n=400	
perf. meas.	<i>I</i> model	sim.	<i>I</i> model	sim.	<i>I</i> model	sim.
$E[Q_1]$	5.1	4.8 ± 0.3	8.4	7.3 ± 1.0	11.3	10.5 ± 2.6
$E[Q_1/n]$	0.20	0.19 ± 0.01	0.08	0.07 ± 0.01	0.03	0.03 ± 0.01
$E[Z_{1,2}]$	–	1.3 ± 0.1	–	1.9 ± 0.2	–	1.3 ± 0.4
$E[Z_{1,2}/n]$	–	0.05 ± 0.01	–	0.02 ± 0.00	–	0.00 0.00

Table EC.2 A comparison of the exact *I*-model queues with simulation results for the steady-state performance measures of the *X* model in normal loading under FQR-T. The arrival rates are $\lambda_1^{(n)} = \lambda_2^{(n)} = 0.98n$ and the thresholds are $\kappa_{1,2}^{(n)} = \kappa_{2,1}^{(n)} = 0.1n$. Service rates are $\mu_{1,1} = \mu_{2,2} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and the abandonment rates are $\theta_1 = \theta_2 = 0.2$

EC.5.3. Sensitivity Analysis For the Thresholds

We now consider different values for the thresholds with and without one-way sharing. Our objective is to perform a sensitivity analysis for the thresholds for a finite system (in this case having 100 agents in each service pool), as a complimentary to the asymptotic line of reasoning in §6. The simulation results, displayed in Table EC.3, are for systems having $\lambda_i = 98$, $\mu_{1,1} = \mu_{2,2} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and $\theta_i = 0.2$, $i = 1, 2$. We vary the thresholds between $\kappa_{i,j} = 1$ and $\kappa_{i,j} = 30$, $i, j = 1, 2$. (Note that with $\kappa_{i,j} = 1$ FQR-T reduces to FQR.) For ease of exposition we take $r_{1,2} = r_{2,1} = 1$. The symmetry allows us to present the results for $E[Q_1]$ and $E[Z_{1,2}]$ only, and consider $\kappa_{1,2} = \kappa_{2,1}$. (If $r_{1,2} \neq r_{2,1}$ then the sensitivity analysis should be performed for each of the two thresholds separately.)

Table EC.3 clearly shows the benefits of using one-way sharing, since even with $\kappa_{i,j} = 1$ the performance is almost as good as when we add thresholds of size 15. However, recall that the thresholds play a vital role in our control: In addition to helping prevent unwanted sharing, they act as “overload detectors”: When $D_{i,j}(t)$ first crosses the threshold $\kappa_{i,j}$, we consider class-*i* queue to be overloaded, and sharing is activated with pool *j* helping queue *i*.

As discussed in §6, we do not want to have the thresholds too large, as they will fail to detect small overloads. Moreover, we see that it can actually be beneficial to share a little, even when the

system is not overloaded; Observe that the average queue length in the case $\kappa_{i,j} = 30$ is larger than when $\kappa_{i,j}$ is 10, 15 or 20. (See also the last paragraph in §4.)

Thus, in choosing the thresholds we need to make sure that under normal loadings they will not be crossed too often, but even small overloads will be detected. Here we see that any value in $\{10, \dots, 20\}$ is reasonable, both with one-way sharing and without. To ensure that even small overloads will be detected by the thresholds, it is probably best to take $10 \leq \kappa_{i,j} \leq 15$.

Insight From the Asymptotic Analysis of the Thresholds. The asymptotic analysis in §6 helps to find good candidates for the thresholds for larger systems. In our example, we can heuristically think of 30 as being of order $O(n)$, while 10 and 15 are of a smaller order, say $O(n^{0.6})$. Then $30 = 0.3n$, $15 \approx n^{0.6}$ and $10 \approx 2/3n^{0.6}$.

This line of reasoning hints at what the thresholds should be (approximately) for a larger system having the same service and abandonment parameters. For example, if $n = 1000$ then $\kappa_{i,j} = 0.3n = 300$ is too large, but $64 \approx n^{0.6} \leq \kappa_{i,j} \leq 2/3n^{0.6} \approx 42$ are good candidates for the thresholds. The threshold values can be determined using simulations, just as in Table EC.3.

	With One-Way Sharing		Without One-Way Sharing	
perf. meas.	$E[Q_1]$	$E[Z_{1,2}]$	$E[Q_1]$	$E[Z_{1,2}]$
$\kappa_{i,j} = 1$	8.4 ± 0.4	2.8 ± 0.3	29.9 ± 1.7	38.2 ± 0.6
$\kappa_{i,j} = 5$	8.3 ± 0.9	2.4 ± 0.3	8.6 ± 0.6	8.1 ± 0.1
$\kappa_{i,j} = 10$	7.5 ± 0.4	1.9 ± 0.2	7.4 ± 0.6	3.6 ± 0.2
$\kappa_{i,j} = 15$	7.2 ± 0.6	1.5 ± 0.2	7.1 ± 0.6	2.1 ± 0.1
$\kappa_{i,j} = 20$	7.5 ± 0.7	1.1 ± 0.2	7.3 ± 0.7	1.3 ± 0.2
$\kappa_{i,j} = 30$	8.2 ± 0.9	0.5 ± 0.1	8.2 ± 0.7	0.5 ± 0.1

Table EC.3 Sensitivity analysis of the effect of the thresholds in a system with 100 agents in each pool. The arrival rates are $\lambda_1 = \lambda_2 = 98$. Service rates are $\mu_{1,1} = \mu_{2,2} = 1$, $\mu_{1,2} = \mu_{2,1} = 0.8$ and the abandonment rates are $\theta_1 = \theta_2 = 0.2$. All the results are derived from five independent simulation runs.

References

- Cohen, J. W. 1982. *The Single Server Queue*, second ed., North-Holland, Amsterdam.
- Eick, S. G., W. A. Massey, W. Whitt. 1993. The physics of The Mt/G/infty queue. *Oper. Res.* **41** (4) 731–742.
- Gurvich, I., W. Whitt. 2007c. Queue-and-idleness-ratio controls in many-server service systems. *Math. Oper. Res.*, forthcoming. Available at: <http://www.columbia.edu/~ww2040/recent.html>
- Gurvich, I., W. Whitt. 2007d. Asymptotic optimality of queue-ratio routing for many-server service systems. working paper. Available at: <http://www.columbia.edu/~ww2040/recent.html>
- Rockafellar, T. R. 1970. *Convex Analysis*, Princeton University Press, Princeton, N.J.
- Whitt, W. 1989. Planning queueing simulations. *Management Sci.* **35** (11) 1341–1366.