

Ridder, Ad; Schwartz, Adam

Working Paper

Large Deviations without Principle: Join the Shortest Queue

Tinbergen Institute Discussion Paper, No. 05-003/4

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Ridder, Ad; Schwartz, Adam (2004) : Large Deviations without Principle: Join the Shortest Queue, Tinbergen Institute Discussion Paper, No. 05-003/4, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/86466>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



TI 2005-003/4

Tinbergen Institute Discussion Paper

Large Deviations without Principle

*Ad Ridder*¹

*Adam Shwartz*²

¹ Faculty of Economics and Business Administration, Vrije Universiteit Amsterdam, and Tinbergen Institute,

² Electrical Engineering, Technion – Israel Institute of Technology.

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Amsterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Please send questions and/or remarks of non-scientific nature to driessen@tinbergen.nl.

Most TI discussion papers can be downloaded at <http://www.tinbergen.nl>.

Large deviations without principle: Join the shortest queue

Ad Ridder

Econometrics

Vrije Universiteit Amsterdam*

Adam Shwartz

Electrical Engineering

Technion—Israel Institute of Technology†

January 5, 2005

Abstract

We develop a methodology for studying “large deviations type” questions. Our approach does not require that the large deviations principle holds, and is thus applicable to a large class of systems. We study a system of queues with exponential servers, which share an arrival stream. Arrivals are routed to the (weighted) shortest queue. It is not known whether the large deviations principle holds for this system. Using the tools developed here we derive large deviations type estimates for the most likely behavior, the most likely path to overflow and the probability of overflow. The analysis applies to any finite number of queues. We show via a counterexample that this system may exhibit unexpected behavior.

1 Introduction

The theory of large deviation is an important tool in the analysis of performance of computer communications networks. Much effort is devoted to the development of a theory of large deviations which will apply to a large class of models, namely to a large class of queueing systems, and providing a useful expression for the rate function so that the theory can be applied. However, this proves to be a difficult task. For example, the relatively broad study [5] does not cover our simple system, while the general methods of identifying the rate function [1] does not cover any priority systems, and requires checking rather complex hypothesis. See also their references for related studies. Sources of relevant theory include [4,

*Work of the first author was performed in part while visiting the Technion.

†Work of the second author was performed in part while visiting the Vrije Universiteit, Amsterdam, and was supported in part by Fund for the promotion of research at the Technion.

5, 6, 10, 17, 19, 24]. For a different point of view, bypassing the sample-path arguments, see [9] and references therein. At the other extreme, there are attempts to apply the theory to specific systems and questions. This is done by establishing first the validity of the large deviations principle for the specific system under investigation, and then calculating the probabilities and “optimal paths” of interest: see for example [2, 3, 8, 15, 23] and the applications in [17]. In some papers, the authors assume the large deviations principle to apply, and draw conclusions: see for example [11, 12, 13, 18]. For some cases, direct arguments allow to bypass the theory: see e.g. [21, 20] and the references in [9].

The “Join the shortest queue” (JSQ) system consists of K queues with exponential servers. Arrivals are routed to the “shortest weighted” queue. Previous work on this system covers the two dimensional system. For the symmetric system (where service rates at the queues are equal and both weights equal 1), the analysis in [17, Chapter 15.10] establishes the large deviations principle, and computes the most likely path as well as the probability of overflow. A variation of this system includes dedicated arrival streams for each queue. This variant is analyzed in [1], where the large deviations principle is established. This is not a generalization of the JSQ, since all analyses of this variant system do not allow the dedicated arrivals to have zero rate. Turner [22] studies the two dimensional variant system with equal weights, as well as with weights that correspond to joining the shortest expected waiting time. Using the large deviations principle and ingenious coupling arguments, the most likely path to overflow and the probability of overflow are computed. McDonald et. al [7, 14] use a more detailed analysis to study the variant JSQ with two queues and derive a more precise description of the distribution on the way to overflow.

The purpose of this paper is to show that some basic techniques of large deviations can be useful even if it is not known whether the large deviations principle holds for the system of interest. We do this through the analysis of a simple queueing system: Join the Shortest Queue. Our analysis covers the original JSQ system with any finite number of queues and any weights. Our major tools are coupling arguments, and large deviations results for the Poisson process and the M/M/1 queue. In Section 2 we give a formal description of the model. We then restrict our attention to the two dimensional system, so as to expose the ideas in a relatively simple setting. In Section 3 we describe the most likely behavior of this system, starting at any point. The way overflow occurs and the probability of overflow are given in Section 4. We show that, starting with an empty system, overflow always occurs by following the “weighted diagonal” at constant speed, possibly lingering at 0 before starting the excursion. For this case, we give an exact expression for the rate function. However, even in the two-dimensional case we show that, if the starting point is not the empty state, then it is possible that overflow occurs by emptying one queue (but not the other!) and then proceeding towards the “weighted diagonal.” We conclude in Section 5 with a description of the most likely behavior in any dimension, and comments on the overflow problem in higher dimensions.

2 The Join the Shortest Queue model

The two dimensional Join the Shorter Queue (JSQ) system consists of two infinite queues, each with its own server. Service times are exponential, with parameter μ_1 and μ_2

respectively. The total service rate is $\mu := \mu_1 + \mu_2$. There are Poisson arrivals with rate λ . Let x_i denote the number of customers in queue i . An arrival is routed to one of the queues according to a control policy of the following type. The control is characterised by two positive numbers r_1, r_2 such that whenever $x_1/r_1 \leq x_2/r_2$, the arriving customer is routed to queue 1, and otherwise to queue 2. To describe the control geometrically we imagine the line $r \triangleq (r_1, r_2) \cdot t, t \in \mathbb{R}_+$. When the state $x = (x_1, x_2)$ of the system is on or above the line r , arrivals join queue 1, and when the state is below the line, arrivals join queue 2. Let us call this line the control diagonal, or simply the diagonal. In accordance with the traditional JSQ terminology we say that arrivals join the shorter queue, although they actually join the “shortest weighted” queue.




Figure 1: Jump directions and rates in the JSQ system

Generalizing to more than two queues, service rates are $\mu_i, i = 1, \dots, K$, and the control is determined by the (strictly) positive numbers $r_i, i = 1, \dots, K$. If the vector of queue sizes is $x = (x_1, \dots, x_K)$ then an arrival is routed to the queue with the smallest value of x_i/r_i . We assume a fixed rule to break ties. We let i_x denote the queue to which an arrival will be sent, if the queue sizes are given by x . The particular choice is of no consequence to our analysis: we only assume that $i_x = i_{\alpha x}$ for all $\alpha > 0$, that is, the choice depends only on the relative sizes of the queues. The case $r_i = r_j$ corresponds to the traditional JSQ and then r is indeed the diagonal. Another special JSQ control is obtained by setting $r_i = \mu_i$: in this case the arriving customers join the queue with the smallest expected waiting time.

Denote by e_i the unit vector in direction $i = 1, \dots, K$, that is, $e_1 = (1, 0, \dots, 0)$ etc. We also denote $e_{-i} = -e_i$. Vectors are viewed as row vectors. We use the notation $|x| = x_1 + \dots + x_K$. We use x to denote points in the state space \mathbb{Z}_+^K , the vectors with positive integer entries. Here $x = (x_1, \dots, x_K)$ is the vector of queue sizes. However, since we are dealing with scaled processes, we shall use the space \mathbb{R}_+^K , the positive K -dimensional orthant. Our scaled process is denoted by z_n : this is the process were all jumps are of size $1/n$, and all rates are multiplied by n . As before, $z_{1,n} = z_n \cdot e_1$ is the first coordinate of the process z_n .

To summarize our conventions, we describe the JSQ process in terms of jump directions

and jump rates [17]. They are

direction	rate	Event
e_1	$\lambda_1(x) = \lambda \mathbf{1}[x_1/r_1 \text{ is smallest}]$	Arrival to queue 1
e_{-1}	$\lambda_{-1}(x) = \mu_1$	Departure from queue 1
\vdots	\vdots	\vdots
e_K	$\lambda_K(x) = \lambda \mathbf{1}[x_K/r_K \text{ is smallest}]$	Arrival to queue K
e_{-K}	$\lambda_{-K}(x) = \mu_K$	Departure from queue K .

We are interested in the most likely behavior of this system, and in its transient overflow probabilities. That is, given a starting point and a time T , we are interested in the probability that at time nT the process reaches a given point nx (or actually, reaches a small neighborhood of that point which, in our case, turns out to be the same) for $n \rightarrow \infty$. More specifically, we study the large deviations limit of this probability, and we find the so-called optimal behavior of the system until overflow.

In our analysis we couple the process to several other processes, which we now define. Fix a point x and recall that i_x denotes the queue to which an arrival will be routed. The *local process* z_n^l is obtained by routing all arrivals to queue i_x , regardless of changes in queue sizes. The *reduced process* z_n^r is a scaled one dimensional M/M/1 process with arrival rate $n\lambda$, service rate $n\mu = n\mu_1 + \dots + n\mu_K$, jumps of size $1/n$ and starting point $|x|$. Note that the reduced process is stable (in the sense that it reaches 0 with probability one) if and only if $\mu > \lambda$. As we shall show, this is the stability condition for the JSQ system as well. Finally, the scaled *component process* z_n^c is the $K + 1$ -dimensional process where coordinate i is a (scaled) Poisson process with rate $n\mu_i$, the $K + 1$ st coordinate is a (scaled) Poisson process with rate $n\lambda$ and all jumps are of size $1/n$. All four processes are coupled by using the same Poisson processes as the arrival and (potential) departure processes. Thus, every arrival entails a jump of size $1/n$ in (a component of) all four processes, etc.

3 Most likely behavior

In this section we describe the most likely behavior of the process that starts at a point x . We provide a detailed analysis of the two-dimensional system. The generalization to K dimensions is given in Section 5. The behavior described below is the path that the process follows, with probability nearly equal 1, in the sense of Kurtz's Theorem [17, Theorem 5.3].

Definition 1 Fix $T > 0$. We call z_∞ the most likely behavior for a sequence z_n of processes over $[0, T]$ if the following hold. Given ε small enough, there is a constant $C_1 > 0$ and a function $C_2(\varepsilon) > 0$ so that

$$\mathbb{P}_x \left(\sup_{0 \leq t \leq T} |z_n(t) - z_\infty(t)| \geq \varepsilon \right) \leq C_1 e^{-nC_2(\varepsilon)} \quad \text{all } n > 0. \quad (1)$$

Note that the definition implies that, necessarily, $z_\infty(0) = x$.

3.1 Most likely behavior in dimension 2

If $x_1 > 0$, $x_2 > 0$ and x is not on the “diagonal,” then at least locally one coordinate of the process is a Poisson process, and the other is the difference of two Poisson processes. Thus Kurtz’s theorem applies [17], and we have

Lemma 1 *For x above r and off the boundary $0 < x_1/r_1 < x_2/r_2$ define*

$$z_\infty(t) = x + (\lambda - \mu_1)e_1 \cdot t - \mu_2 e_2 \cdot t. \quad (2)$$

Let T_x be the first time $z_\infty(t)$ hits either the line $r(t)$ or the vertical line $(0, y)$. Then the most likely path until time T_x is to follow z_∞ . Moreover, C_2 is quadratic near $\varepsilon = 0$. Finally, for any given $\alpha > 0$, the estimate (1) is uniform in $\{x : T_x > \alpha\}$. The analogue conclusion holds if x is below r .

Proof. This is a direct application of [17, Theorem 5.3]. Consider the local process: it is identical to the JSQ process over the region of interest, and the local process satisfies the conditions of [17, Theorem 5.3]. ■

Note that, depending on the parameters, the most likely behavior may drive the process closer to r : in this case, depending on x , the process may hit r first or the vertical line first. The behavior depends both on the parameters and on the starting point. In general, if we start above r then, since arrivals all join queue 1, $z_{2,\infty}$ can only decrease, so that T_x is finite. However, if the system is stable then (if we start above r) the slope of $z_\infty(t)$ is in the cone between $(e_1, -e_2)$ clockwise to $(-e_1, 0)$. More precisely,

Corollary 2 *Suppose x is above r and off the boundary. Then the most likely path moves down with rate μ_2 .*

- (i) *If $\lambda > \mu_1$ then the most likely path moves to the right. If $\lambda = \mu_1$ then it moves straight down. In both cases it moves closer to the line r .*
- (ii) *If $\mu_1 > \lambda$ and $\mu_2/(\mu_1 - \lambda) \geq r_2/r_1$ then the most likely path moves to the left. If the last inequality is strict, it draws closer to r . If it is an equality, it moves parallel to r .*
- (iii) *If $\mu_1 > \lambda$ and $\mu_2/(\mu_1 - \lambda) < r_2/r_1$ then the most likely path moves to the left, and away from r .*

Suppose now x is below r and off the boundary. Then the most likely path moves to the left with rate μ_1 .

- (iv) *If $\lambda > \mu_2$ then the most likely path moves up. If $\lambda = \mu_2$ then it moves straight to the left. In both cases it moves closer to the line r .*
- (v) *If $\mu_2 > \lambda$ and $\mu_1/(\mu_2 - \lambda) \geq r_1/r_2$ then the most likely path moves down. If the last inequality is strict, it draws closer to r . If it is an equality, it moves parallel to r .*
- (vi) *If $\mu_2 > \lambda$ and $\mu_1/(\mu_2 - \lambda) < r_1/r_2$ then the most likely path moves down, and away from r .*
- (vii) *There is no combination of coefficients so that the most likely path moves away from r both when below and when above the line r .*

Proof. This is a direct consequence of Lemma 1. Using (2), (i)-(vi) are established by elementary geometry. The last conclusion follows from (iii) and (vi) by simple algebra. ■




Figure 2: Drift condition (vi)

Lemma 3 Suppose the starting point $x \neq 0$ is on r and that the condition of Corollary 2(iii) holds. Let $T_x \triangleq x_1(\mu_1 - \lambda)^{-1}$. Fix $T < T_x$ and define z_∞ through (2). Then z_∞ is the most likely path until time T , C_2 is quadratic near $\varepsilon = 0$ and for any $\alpha > 0$, this estimate is uniform in $\{x : T_x > \alpha\}$.

Proof. Let z_n^l denote the local process where arrivals join queue 1. If we couple the two queueing systems (that is, drive them with the same arrival and departure processes) then, at least until one process hits the boundary, $z_{1,n}(t) \leq z_{1,n}^l(t)$ and $z_{2,n}(t) \geq z_{2,n}^l(t)$. This holds since some arrivals to queue 1 for the process z_n^l are sent to queue 2 in the JSQ system (when below r). However, z_n^l satisfies the conditions of Kurtz theorem and so (1) holds. This implies that the scaled JSQ process moves away from r . But once it does, its increments are exactly those of z_n^l , and another application of Kurtz theorem establishes the Lemma. ■

Lemma 4 Suppose the starting point $x \neq 0$ is on r . Suppose further that neither of the conditions of Corollary 2(iii) or (vi) hold. Let

$$z_\infty^r(t) = x_1 + x_2 + (\lambda - \mu_1 - \mu_2) \cdot t.$$

If the system is stable, define T_x as the first time $z_\infty^r(t) = 0$. Otherwise, choose an arbitrary finite T_x . Fix $T < T_x$ and define z_∞ through

$$z_{2,\infty}(t)/r_2 = z_{1,\infty}(t)/r_1 \quad \text{and} \quad z_{1,\infty}(t) + z_{2,\infty}(t) = z_\infty^r(t). \quad (3)$$

Then the most likely path until time T is z_∞ . Moreover, C_2 is quadratic near $\varepsilon = 0$. Finally, for any $\alpha > 0$, the estimate (1) is uniform in $\{x : T_x > \alpha\}$.

Proof. The proof uses standard ideas which will be elaborated for the large deviations calculation, and so will only be described briefly here. First note that these two conditions define a unique path. By coupling the JSQ process with the reduced process we immediately have that the second condition of (3) must hold: that is, in the sense of Kurtz theorem, the scaled process must satisfy this condition. Consider now an initial short period of time. Since the component process satisfies the conditions of Kurtz theorem, we know that the process (with high probability) cannot move far. If it stays near r , we repeat the argument, to show that it stays near r until T . Once it moves away from r , we can apply Corollary 2 to conclude that it will move in the direction of r . Thus the process stays near r until T (in the sense of Kurtz theorem). Thus the first condition of (3) holds as well. ■

If x is on the boundary, there are two cases.

Lemma 5 Suppose $x_1 = 0$ and $x_2 > 0$.

- (i) If $\mu_1 \geq \lambda$ then (1) holds with $z_\infty(t) = x - \mu_2 e_2 \cdot t$ for $T < T_x = \frac{x_2}{\mu_2}$.
- (ii) If $\mu_1 < \lambda$ then (1) holds where $z_\infty(t)$ is given by (2), until z_∞ meets r .

Proof. Over the region of interest, the JSQ process is identical to the local process. Until the process reaches either r or the point 0, the two coordinates are statistically independent. The first coordinate is an M/M/1 queue, starting at 0. In case (i) it is stable, and so it stays near 0. In case (ii) it is not stable, so it follows its drift: see [17, Chapter 11]. The second coordinate is a Poisson process, so it follows its mean drift. ■

Note that we have now established our stability claim: the JSQ system is stable if $\mu_1 + \mu_2 > \lambda$. We conclude the analysis of the most likely path by describing the behavior starting at 0.

Lemma 6 Suppose $x = 0$. If the JSQ is stable then the most likely behavior is to stay near 0. If it is unstable, then the most likely behavior is to follow r , so that $z_{1,\infty}(t) + z_{2,\infty}(t) = (\lambda - \mu_1 - \mu_2) \cdot t$.

Proof. If the JSQ is stable then whenever it moves away from 0, the previous results show that it must move towards 0.

In the unstable case, the reduced system follows a straight line away from 0 with speed at least $\lambda - \mu_1 + \mu_2 > 0$. This will be precisely the speed if we are away from the boundaries. But as soon as we move away from 0, our previous results imply that the process must move towards r , and in particular away from the boundaries. Therefor the process moves along r with speed $\lambda - \mu_1 - \mu_2$. ■

4 Most likely path to overflow: dimension 2

We are interested in the probability and the most likely path to buffer overflow. We provide an explicit formula for the case that the system starts empty. Let $C = nc$ be some high level for queue sizes. Buffer overflow occurs when the system reaches any state where one or both queues exceed the high level: $\{(x_1, x_2) : x_1 \geq nc \text{ or } x_2 \geq nc\}$. Unless we start very close to the set of overflow states, this set is hit at the unique system state $x = (r_1, r_2)n\beta$ where the diagonal crosses the boundary of the overflow set. For example, if $r_1 > r_2$ then $\beta = c/r_1$. In this section we assume explicitly that the JSQ system is stable, for otherwise the event of reaching overflow is not rare, and the most likely behavior brings the process to the mentioned overflow state (with probability equal nearly 1, as demonstrated in the previous section).

The concept of most likely (or optimal) path is related to the usual scaling of the system and taking large deviations limits, as in [17]. The scaled JSQ process z_n is defined in Section 2. This is a piecewise constant process. It is often more convenient to work with a process with continuous paths, which is obtained by a linearly interpolating the original process between jump points. We do not distinguish between the piecewise-constant jump process and the piecewise-linear version, since the two are exponentially equivalent [17].

A path ϕ is a continuous function $[0, T] \rightarrow \mathbb{R}_+^2$ mapping a finite time interval into the state space of the scaled JSQ process. The open ε -ball around a path ϕ is the collection of paths

$$B_\varepsilon(\phi) = \left\{ \psi : \sup_{0 \leq t \leq T} |\psi(t) - \phi(t)| < \varepsilon \right\}.$$

Although we neither assume nor prove that the large deviations principle holds for our system, we shall make use of the following definitions.

Definition 2 *A sequence $\{y_n\}$ satisfies the large deviations principle with rate function $I(\cdot)$ if I is non negative, lower semicontinuous and*

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_x(y_n \in C) &\leq - \inf_{\phi \in C} I(\phi) \quad \text{for all closed } C, \text{ and} \\ \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_x(y_n \in G) &\geq - \inf_{\phi \in G} I(\phi) \quad \text{for all open } G. \end{aligned}$$

We call ℓ the local rate function if for any absolutely continuous ϕ on $[t_0, t_1]$,

$$I(\phi) = \int_{t_0}^{t_1} \ell(\phi(t), \dot{\phi}(t)) dt, \quad (4)$$

and $I(\phi) = \infty$ if ϕ is not absolutely continuous.

Let x^0 and x^T be two points and ϕ be a path with $\phi(0) = x^0$, and $\phi(T) = x^T$. We are concerned with the probability that z_n stays close to ϕ , given that $\lim_{n \rightarrow \infty} z_n(0) = x^0 = \phi(0)$. We denote this by $\mathbb{P}_{x^0}(z_n \in B_\varepsilon(\phi))$.

Definition 3 *We say that a path ϕ is more likely than the path ψ if*

$$\lim_{\varepsilon \downarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{x^0}(z_n \in B_\varepsilon(\psi)) \leq \lim_{\varepsilon \downarrow 0} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{x^0}(z_n \in B_\varepsilon(\phi)).$$

We call ϕ the optimal path from x^0 to x^T if this holds for all $\psi \neq \phi$.

Note that this relation is transitive, that is, if ϕ is more likely than ψ and ψ is more likely than η then ϕ is more likely than η .

We would like to know what these optimal paths look like, given any starting and ending point. The most interesting paths are those leading to overflow. We have already indicated that overflow is reached via the diagonal point $(r_1, r_2)\beta$, hence the optimal path to overflow always ends (at time T) at this point. Furthermore, since the first queue cannot grow below the diagonal and the second queue cannot grow above the diagonal, the only way to reach the overflow point $(r_1, r_2)\beta$ is by reaching the diagonal and then following it in an upward direction (both coordinates increase). A path or a piece of path where both coordinates increase (strictly!) is called *increasing*. We can split our search of optimal path into two:

- Given a starting point x^0 and a time T , how is the diagonal reached?
- Given a starting point on the diagonal and a time T , what is the optimal path that increases along the diagonal?

We deal with the second question: it will be clear that the same methods apply to the first. So, consider paths which start on the diagonal, $r^0 := (r_1, r_2)\alpha$ for some $0 < \alpha < \beta$, increase along the diagonal, and end at the overflow point $r^T := (r_1, r_2)\beta$ at time T . Notice that we restrict—for the moment—the starting point to avoid the point 0. The reason is that if you stay near the diagonal starting at r^0 , then both servers are kept busy. This makes it easier to treat the process.

Lemma 7 *Let ψ be path from r^0 to r^T , increasing along the diagonal. Let*

$$\phi(t) := r^0 + (r_1, r_2) \frac{\beta - \alpha}{T} t, \quad 0 \leq t \leq T.$$

Then ϕ is more likely than ψ .

Proof. Recall $|x| := x_1 + x_2$. We show that the probability for the JSQ process to stay near ϕ is (asymptotically) the same as the probability for the reduced process to stay near $|\phi|$ and, in addition, it is larger than the probability to stay near any other path ψ that stays on the diagonal.

Recall the definition of the coupled reduced process z_n^r (an M/M/1 queue with arrival rate λ and service rate $\mu = \mu_1 + \mu_2$). For $\varepsilon < \min\{r_1^0, r_2^0\}$, since ψ is increasing and $|\psi(0)| = |r^0| \neq 0$, we have that $|\psi(t)| \geq \varepsilon > 0$ and

$$\{z_n \in B_\varepsilon(\psi)\} \subset \{z_n^r \in B_\varepsilon(|\psi|)\} \subset \{z_n^r > 0\}. \quad (5)$$

The inclusions (5) hold since if the JSQ process stays near ψ then the coupled reduced process satisfies these conditions. However, for the reduced process we have an LDP with a convex local rate function. Hence straight lines are more likely than other increasing paths [17, Lemma 5.16]. Here there is only one straight line from $|r^0|$ to $|r^T|$, viz. $|\phi|$. Hence, we have

$$\begin{aligned} \lim_{\varepsilon \downarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{r^0}(z_n \in B_\varepsilon(\psi)) \\ &\leq \lim_{\varepsilon \downarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{|r^0|}(z_n^r \in B_\varepsilon(|\psi|)) \\ &\leq \lim_{\varepsilon \downarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{|r^0|}(z_n^r \in B_\varepsilon(|\phi|)) \\ &= \lim_{\varepsilon \downarrow 0} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{|r^0|}(z_n^r \in B_\varepsilon(|\phi|)). \end{aligned} \quad (6)$$

The last equality holds since the M/M/1 process satisfies the LDP.

We now claim that, in fact, both departure processes as well as the arrival process are nearly straight lines. This follows from [17, Lemma 7.25] as follows. Consider the augmented M/M/1 queue where we look at a 4-dimensional process, consisting of the 3 independent Poisson processes of the component process z^c , and the M/M/1 process whose arrival process is the Poisson λ process, and whose service process consists of (the superposition of) the other two independent processes. This M/M/1 process has exactly the desired distribution, and by [17, Lemma 7.25], each of the coordinate processes follows a straight line (note that in

our case, the rates and hence ℓ in [17, Eq. 7.25] do not depend on x , and r is linear so that y there is fixed, so that $\theta(s)$ and $\lambda_j(r(s))$ do not depend on s , and so wr^i there is a straight line). But if arrivals and departures follow a straight line, and the M/M/1 queue follows $|\phi|$, then by the definition of the JSQ, it follows ϕ . Therefore, we have for all $\varepsilon > 0$

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{|r^0|} (z_n^r \in B_\varepsilon(|\phi|)) \\ \leq \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{r^0} (z_n \in B_\varepsilon(\phi)) . \end{aligned} \quad (7)$$

Take $\varepsilon \downarrow 0$ in (7) and use (6) to conclude that ϕ is more likely than ψ . \blacksquare

We now extend Lemma 7 to include the starting point 0.

Lemma 8 *Let $r^0 = (0, 0)$, $r^T = (r_1, r_2)\beta$. Define $\phi(t) := (r_1, r_2)\beta t/T$, $0 \leq t \leq T$. Let ψ be any other path from 0 to r^T on $[0, T]$ which increases along the diagonal. Then ϕ is more likely than ψ .*

Proof. The difficulty here is that if one of the queues is empty but the other is not, then z_n does not behave like z_n^r . Let $t(\varepsilon)$ be the first time that $\psi_i(t) > \varepsilon$ for all i . Then after $t(\varepsilon)$, if z_n is in $B_\varepsilon(\psi)$ then none of the queues empties again. So let ϕ_ε be the direct path from $\psi(t(\varepsilon))$ to r^T . We can now repeat the arguments of Lemma 7, with two minor changes. First, here the starting point of $z^r(t(\varepsilon))$ is $|\psi(t(\varepsilon))|$ which is not necessarily equal to $|z_n(t(\varepsilon))|$: however their distance for the paths of z_n that stay in $B_\varepsilon(\psi)$ is at most ε . As $\varepsilon \rightarrow 0$, this does not change the proof. Finally,

$$\begin{aligned} \mathbb{P}_0 (z_n \in B_\varepsilon(\psi)) &\leq \mathbb{P}_0 (z_n \in B_\varepsilon(\psi), t \geq t(\varepsilon)) \\ &\leq \mathbb{P}_0 (z_n^r \in B_\varepsilon(|\phi|), t \geq t(\varepsilon)) \end{aligned}$$

so that the argument of (6) applies once we note that

$$|\mathbb{P}_0 (z_n^r \in B_\varepsilon(|\phi|)) - \mathbb{P}_0 (z_n^r \in B_\varepsilon(|\phi|) \text{ for } t \geq t(\varepsilon))| \leq e^{-\eta(\varepsilon)n}$$

where $\eta(\varepsilon) \rightarrow 0$ as $\varepsilon \downarrow 0$. With these changes the previous proof applies. \blacksquare

Remark The fact that paths with constant speed are most likely among all increasing paths along the diagonal also holds for decreasing paths. \blacksquare

Lemma 8 compares ϕ only to other strictly increasing paths. We show below that for large T the most likely path stays near 0 and then follows ϕ .

We call a straight line path with constant speed a *direct path*. From the two preceding Lemmas we deduce

Corollary 9 *Let ϕ be a direct path, increasing along the diagonal. Then*

$$I(\phi) := -\lim_{\varepsilon \downarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{x^0} (z_n \in B_\varepsilon(\phi)) \quad (8)$$

exists. Let T be the duration of the path ϕ , and $\ell(\cdot)$ the local rate function of the (non empty) $M(\lambda)/M(\mu)/1$ queue [17, Expression (7.17)]. Then

$$I(\phi) = J(T; |\phi(T) - \phi(0)|) := T\ell\left(\frac{|\phi(T) - \phi(0)|}{T}\right), \quad (9)$$

where I is given in (4) and J is strictly convex in T . If ϕ is decreasing than the result holds with a minus sign for $|\phi(T) - \phi(0)|$ in (9).

Proof. That the limit exists follows directly from the proof of Lemma 7 (and of Lemma 8 in case the starting point is 0), particularly by inequalities (6) and (7). Also these inequalities say that the limit equals the Large Deviations cost rate going from $|\phi(0)|$ to $|\phi(T)|$ in T time units by the $M(\lambda)/M(\mu)/1$ queue. The cost function is as stated in (9). Convexity follows by differentiation since ℓ is strictly convex. ■

Remark A similar result (limit exists and exact expression) holds for any direct path ϕ which does not cross the diagonal, for the following reason. Consider a (vector) process where the coordinates are statistically independent. Suppose an equation of the type (8) holds for each coordinate separately, with rate $I_i(\phi_i)$ for the i th coordinate. Then it is easy to show that (8) also holds for the full process, and $I(\phi) = \sum I_i(\phi_i)$. More explicitly,

$$\begin{aligned} \{z_{1,n} \in B_{\varepsilon/2}(\phi_1) \text{ and } z_{2,n} \in B_{\varepsilon/2}(\phi_2)\} &\subset \{z_n \in B_\varepsilon(\phi)\} \\ &\subset \{z_{1,n} \in B_\varepsilon(\phi_1) \text{ and } z_{2,n} \in B_\varepsilon(\phi_2)\} \end{aligned}$$

so that the desired limit exists, and the rate functions add. Now a direct path that does not cross the diagonal is of one of the following types.

- Except possibly for the starting and/or ending point, the path lies entirely in one of the two interiors of the JSQ, i.e., either below or above the diagonal, and away from the boundaries. In these areas the system is equivalent to the local process, for which one coordinate is a Poisson process and the other an independent $M/M/1$ queue, so that we have an LDP (end points are easily dealt with in case they lie on a boundary, cf. proof of Lemma 8).
- The path goes along one of the boundaries. Here again the process is equivalent to the local process, comprising of a Poisson and an independent $M/M/1$ queue. Since the LDP holds for each of the independent coordinates, it also holds for the process, and the rate function is simply the sum of the two rates, one for each coordinate. ■

Remark We call the expression $I(\phi)$ the *rate* or the *cost* of the path ϕ . (9) holds for any direct path in the quadrant which does not cross any boundary. Clearly for increasing paths (both coordinates grow) other than along the diagonal, the cost is ∞ . ■

Remark We have not yet settled the optimal paths! Lemmas 7 and 8 say only that direct paths along the diagonal are the most likely increasing paths. The optimal path may decrease from the start until some time epoch and then increase. This was our first question earlier in this section. ■

First we consider paths starting at 0. When the time to overflow T is large enough, the optimal path remains at 0 and then moves straight with constant speed $\mu_1 + \mu_2 - \lambda$ to the overflow point. Otherwise, it leaves 0 immediately and moves with constant speed.

Theorem 10 Let $T_0^* := |r^T|(\mu_1 + \mu_2 - \lambda)^{-1}$. The most likely way for the stable JSQ to reach the overflow point $r^T = (r_1, r_2)\beta$ starting at 0, is:

- (i) If $T \leq T_0^*$ then the optimal path is the direct path between 0 and r^T , i.e., $\phi(t) = (r_1, r_2)\beta t / T$, $t \in [0, T]$, with cost rate $I(\phi) = J(T, |r^T|)$.
- (ii) If $T > T_0^*$ then the optimal path is to stay in 0 until $T - T_0^*$, and then proceed along the diagonal with speed $\mu_1 + \mu_2 - \lambda$, i.e.

$$\phi(t) = (r_1, r_2)(r_1 + r_2)^{-1}(\mu_1 + \mu_2 - \lambda)(t - (T - T_0^*)), \quad t \in [T - T_0^*, T],$$

and $\phi(t) = 0$, $t \in [0, T - T_0^*]$, with cost rate

$$I(\phi) = J(T_0^*, |r^T|) = (r_1^T + r_2^T) \log((\mu_1 + \mu_2)\lambda^{-1}).$$

Proof. Let ψ a path containing a detour. That is, there are $t_1 < t_2 < t_3$ with $\psi(0) = 0$, $\psi(T) = r^T$, $\psi(t_1)$ and $\psi(t_2)$ are on the diagonal and $\psi(t)$ is not on the diagonal for $t_1 < t < t_2$. Since $\min\{z_{1,n}(t)/r_1, z_{2,n}(t)/r_2\}$ can only increase on the diagonal, then if z_n is near ψ with positive probability, then necessarily $\psi(t_2)$ is below and to the left of $\psi(t_1)$. Since $\psi(T) = r^T$ and since increase only occurs on the diagonal, there must be a t_3 , the smallest time after t_2 so that $\psi(t_3) = \psi(t_1)$. Define $\Delta = t_3 - t_1$ and consider

$$\bar{\psi}(t) = \begin{cases} 0 & 0 \leq t \leq \Delta \\ \psi(t - \Delta) & \Delta \leq t \leq t_3 \\ \psi(t) & t_3 \leq t. \end{cases} \quad (10)$$

We now use coupling to show that the probability to stay near $\bar{\psi}$ is larger than that of ψ . To do this, consider the scaled component process z_n^c . Construct the coupled process \bar{z}_n^c by interchanging its segments as in (10). Since Poisson processes are memoryless, the distribution of the two processes is the same. By this coupling,

$$\begin{aligned} \mathbb{P}_0(z_n \in B_\varepsilon(\psi)) &\leq \mathbb{P}_0(\sup |z_n(t) - \psi(t)| < \varepsilon, \text{ for } t \in [0, t_1] \cup [t_3, T]) \\ &= \mathbb{P}_0(\sup |\bar{z}_n(t) - \bar{\psi}(t)| < \varepsilon, \text{ for } t \in [\Delta, T]). \end{aligned}$$

Since the JSQ system is stable,

$$\begin{aligned} \lim_{\varepsilon \downarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0(\sup |\bar{z}_n(t) - \bar{\psi}(t)| < \varepsilon, \text{ for } t \in [\Delta, T]) \\ = \lim_{\varepsilon \downarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_0(\bar{z}_n \in B_\varepsilon(\bar{\psi})). \end{aligned}$$

Thus paths without detours are more likely. But then Corollary 9 applies, and the cost is that of the reduced M/M/1 queue, which gives (i) and (ii). \blacksquare

4.1 General starting point

The methods leading to Theorem 10 can be used to derive the most likely path to overflow, starting at any point. Below we provide some results in this direction: for proofs we refer to [16]. We need the following fact.

Lemma 11 *Let ϕ be a concatenation of K direct paths. Then the limit*

$I(\phi) = -\lim_{\varepsilon \downarrow 0} \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_{x^0}(z_n \in B_\varepsilon(\phi))$ exists and $I(\phi) = \sum_{k=1}^K I(\phi_k)$. If ϕ_k does not cross the diagonal, then $I(\phi_k)$ has an explicit representation (e.g. (9)).

The proof of this lemma uses standard continuity arguments, and in particular continuity of the rate function with respect to shifts, and homogeneity of the process away from the diagonal and the boundaries.

Optimal paths to overflow cannot be too complicated.

Theorem 12 *The optimal path to overflow has the following properties. It is composed of a finite number of direct paths. The direction or speed do not change in the interiors (i.e. they only change either on the diagonal or when at least one coordinate equals 0). If it spends time below the diagonal, then it does not spend time above the diagonal and conversely. If $\phi_i(t) = 0$ for some $t > 0$, then $\phi_i(t) = 0$ in some interval containing t .*

The proofs of these statements use arguments as in Lemma 7.

More can be obtained when the starting point is on the diagonal. Recall the definition of the reduced process, and the related cost function (9). The latter is the cost rate for an M/M/1 queue, which is a strictly convex, unimodal function. Its minimum is easily determined:

$$T_d := \arg \min_{T > 0} I_d(T) = \frac{|r^T - r^0|}{\mu_1 + \mu_2 - \lambda}.$$

Theorem 13 *Suppose the starting point is on the diagonal.*

(i) The direct path is optimal if $T \leq T_d$.

(ii) There is a T^ such that for $T > T^*$, the optimal path to overflow follows the most likely behavior to 0, stays there until $T - T^*$ and then follows a direct path with speed $\mu_1 + \mu_2 - \lambda$ to overflow. The cost rate equals*

$$(r_1^T + r_2^T) \log \frac{\mu_1 + \mu_2}{\lambda}.$$

It may be surprising at first glance that this is not the general case.

Theorem 14 *Suppose the parameters are such that the most likely path moves down and away from the diagonal. Fix a non-zero starting point r^0 on the diagonal. Then there is a T such that the most likely path to overflow in time T moves down until queue 2 empties, then moves left and then up to the diagonal, and finally through a direct path to overflow. Throughout the path, its first coordinate (queue 1) stays away from 0.*

The basic argument here is continuity. For details and an explicit example see [16].

The conclusions we can draw for the two dimensional system are as follows. The most likely behavior can be computed explicitly, starting at any point. The most likely path to overflow can be found by solving a finite dimensional optimization problem: we need to compute the rate function over a finite collection of candidate paths, each composed of several direct path segments. For each such segment, the rate can be computed explicitly.

However, it is quite clear that the complexity of the problem grows quite fast as the dimension of the system increases. In the next section we comment on the extension.

5 Extension to higher dimensions

In this section we describe how the results for two dimensions can be extended to the more general model. As we shall see, the results for the most likely behavior extend naturally, and the complexity grows but is quite manageable. For the overflow problem the situation is more complex.

5.1 Most likely behavior in higher dimensions

The most likely behavior can be analyzed much in the same way as the two dimensional system, and so we only comment on this.

(i) First note that if the JSQ system is unstable, namely $\sum_{i=1}^K \mu_i < \lambda$ then, by the previous arguments, the most likely behavior is to approach the line $r(t)$ and then follow that line, where the reduced system satisfies

$$z_{\infty}^r(t) = x^r + \left(\lambda - \sum_{i=1}^K \mu_i \right) \cdot t.$$

(ii) Suppose that some queues are empty, and some are not. Without loss of generality, assume that $x_i = 0$, $i = 1, \dots, i_x$ and $x_i > 0$, $i > i_x$. Suppose moreover that the subsystem consisting of i_x queues is unstable, that is $\sum_{i=1}^{i_x} \mu_i < \lambda$. Then queues $1, \dots, i_x$ are statistically independent of the rest of the queues until one of the other queues reaches a value of $z_{i,n}(t)/r_i$ that is equal or lower than that of one of the first i_x queues. Therefore, for our local (short-time) analysis, we need not consider those queues: their behavior is simply to follow their drift $-\mu_i e_i$. We therefor ignore those queues, which amounts to setting $K = i_x$. Now by definition, the reduced system is unstable, and moreover, it is coupled to our system whenever all queues are non empty. But by definition, none of the ratios $z_{i,n}(t)/r_i$ will be larger than 1 unless all are at least 1. Thus we see that the total size of the queues $1, \dots, i_x$ increases at the rate $\lambda - \sum_{i=1}^{i_x} \mu_i$ of the reduced system, and as analyzed before the increase is towards r .

(iii) If the subsystem is stable, than the same reasoning shows that these queues will remain empty. The queues we ignored recieve no arrivals, and therefore grow smaller. As another queue empties, the stable system which now consists of one more queue is obviously stable (more departures with the same arrival rate), and so eventually all queues will empty.

(iv) Suppose now that there are no empty queues, but that a subset of the queues have nearly the same product x_i/r_i , and further, that this product is smaller than for other queues. Again, without loss of generality, assume $x_i/r_i < x_j/r_j$ for all $i \leq i_x < j$. As before, we can consider the first group of queues, ignoring the rest. If the subsystem is unstable, that is $\sum_{i=1}^{i_x} \mu_i < \lambda$ then, as before, the most likely path is for this set of queues to grow along r , until $z_{i,n}(t)/r_i \geq z_{j,n}(t)/r_j$ for some $j > i_x$. This will eventually happen, as the other queues all decrease while the first queues all grow. At that point, if with the additional queue the system is still unstable, the present analysis holds, with new growth rates.

(v) Finally, if the subsystem is stable, it may decrease along r . To find out if this is the case, consider the queue (among the first i_x queues), say j , with the smallest ratio of μ_i/r_i . This queue is our candidate to empty slower than the rest of the queues, thus staying “above

the diagonal" r . To find out if this happens, we compute the rate at which the remaining queues empty. This is done by considering the reduced system with $i_x - 1$ queues, then calculating the rate at which, say, $z_{1,n}(t)/r_1$ decreases, and comparing this to μ_j/r_j . If the latter is larger, then we know that queue j empties as fast as the remaining group, and the most likely behavior is for these i_x queues to decrease together along r . However, if μ_j/r_j is smaller, then indeed it will stay larger even though no arrivals join queue j . We then repeat the procedure, to check if another queue should be excluded.

Once this procedure is complete, The remaining group of queues empty along r , until they are all empty. This concludes the last case in our analysis.

We note that, unlike other systems, here the boundaries do not hamper our analysis. This is the case since, for our purposes, the coordinates decouple except at 0 and on r . In the first case, a one-dimensional analysis provides the picture. In the second case, the boundaries play no role, and the analysis near r can be completed.

5.2 Most likely path to overflow in higher dimensions

We note that the methods of this paper apply to the higher dimensional system. These methods show that the rate function is the correct (asymptotic-logarithmic) value of the probabilities of overflow, even though the general large deviations principle may not hold. The main results, namely Theorem 10 as well as the results of subsection 4.1 hold as stated. Consequently, the infinite-dimensional optimization problem (finding the optimal path) is reduced to a finite dimensional problem: finding the parameters of the optimal path, which is constructed from a finite collection of straight lines with fixed speed. Moreover, the rate function for each such line possesses an explicit expression, since it is the rate for a collection of independent processes, where one process is an M/M/1 queue and the rest are Poisson processes. However, it is also clear that the complexity of the problem increases very fast with the dimension of the system, since the number of possibilities increases rapidly, and complicated paths, such as in Theorem 14 need to be considered. Fortunately, the behavior of this system is relatively simple since increase may only happen along a one-dimensional surface—the diagonal.

References

- [1] Atar, R. and P. Dupuis (1999) Large deviations and queueing networks: methods for rate identification. *Stoch. Proc. Appl.* 84: 255–296.
- [2] Bertsimas D., L. Paschalides and J.N. Tsitsiklis (1999) Large deviations analysis of the generalized processor sharing policy. *Queueing Systems* 32: 319–349.
- [3] Chang, C-S. (1995) Sample path large deviations and intree networks. *Queueing Systems* 20: 7–36.
- [4] Dembo, A. and O. Zeitouni (1996) *Large deviations techniques and applications*, Second edition. Springer Verlag.

- [5] Dupuis, P. and R.E. Ellis (1995) The large deviations principle for a general class of queueing systems I. *Trans. American Math. Soc.* 347: 2689–2751.
- [6] Dupuis, P. and R.E. Ellis (1995) *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley.
- [7] Foley, R.D. and D. McDonald (2001) Join the shortest queue: stability and exact asymptotics. *Ann. Appl. Prob.* 11: 569–607.
- [8] Ganesh, A.J. and N. O’Connell (2002) A large deviation principle with queueing applications. *Stoch. and Stoch. Reports* 73: 25–35.
- [9] Ganesh, A.J., N. O’Connell and D. Wischik (2004) *Big Queues*. Springer.
- [10] I. Ignatiouk-Robert (2004) Large deviations for processes with discontinuous statistics. *Ann. Prob.* to appear.
- [11] Mandjes, M. (1999) Rare event analysis of the state frequencies of a large number of markov chains. *Stoch. Models* 15: 577–592.
- [12] Mandjes, M. and A. Ridder (1999) Optimal trajectory in a queue fed by a large number of sources. *Queueing Systems* 31: 137–170.
- [13] Mandjes, M. and A. Weiss (2003) Sample path large deviations of a multiple time-scale queueing model. Preprint.
- [14] McDonald, D. (1996) Overloading parallel servers when arrivals join the shorest queue. pp. 169–196 In: Glasserman, P., K. Sigman and D. Yao (ed.) *Stochastic Networks: stability and rare events*. Springer Verlag.
- [15] Puhalskii, A.A. and W. Whitt (1998) Functional large deviation principles for waiting and departure processes. *Prob. Engin. Info. Sci.* 12: 479–507.
- [16] Ridder, A. and A. Schwartz (2004) Large deviations methods for the join-the-shortest-queue model. Report.
- [17] Schwartz, A. and A. Weiss (1995) *Large deviations for performance analysis: queues, communication and computing*. Chapman Hall.
- [18] Schwartz, A. and A. Weiss (1999) Multiple time scales in Markovian ATM models I. Formal calculations. CC PUB 267, Elec. Engin. Technion.
- [19] Schwartz, A. and A. Weiss (2005) Large Deviations with diminishing rates. to appear, *Math. of Oper. Res.*.
- [20] Stolyar, A.L. (2003) Control of End-to-End Delay Tails in a Multiclass Network: LWDF Discipline Optimality. *Ann. Appl. Prob.* 13: 1151–1206.
- [21] Stolyar A.L. and K. Ramanan (2001) Largest Weighted Delay First Scheduling: Large Deviations and Optimality. *Ann. Appl. Prob.* 11: 1–48.

- [22] Turner, S. (2000) Large deviations for join the shorter queue. In: McDonald, D. and S. Turner (ed) *Analysis of Communication Networks: Call Centres, Traffic and Performance*. The Fields Institute.
- [23] Wischik D.J. (2001) Sample path large deviations for queues with many inputs. *Ann. Appl. Prob.* 11: 379–404.
- [24] Weiss, A. (1986), A new technique for analyzing large traffic systems. *Adv. Appl. Prob.* 18: 506–532.