



Explicit Solution for a Network Control Problem in the Large Deviation Regime*

RAMI ATAR** and ADAM SHWARTZ*** {atar;adam}@ee.technion.ac.il
Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel

PAUL DUPUIS† dupuis@dam.brown.edu
Lefschetz Center for Dynamical Systems, Division of Applied Mathematics, Brown University, Providence, RI 02912, USA

Received 23 October 2002; Revised 14 August 2003

Abstract. We consider optimal control of a stochastic network, where service is controlled to prevent buffer overflow. We use a risk-sensitive escape time criterion, which in comparison to the ordinary escape time criteria heavily penalizes exits which occur on short time intervals. A limit as the buffer sizes tend to infinity is considered. In [2] we showed that, for a large class of networks, the limit of the normalized cost agrees with the value function of a differential game. In this game, one player controls the service discipline (who to serve and whether to serve), and the other player chooses arrival and service rates in the network. The game's value is characterized in [2] as the unique solution to a Hamilton–Jacobi–Bellman Partial Differential Equation (PDE). In the current paper we apply this general theory to the important case of a network of queues in tandem. Our main results are: (i) the construction of an explicit solution to the corresponding PDE, and (ii) drawing out the implications for optimal risk-sensitive and robust regulation of the network. In particular, the following general principle can be extracted. To avoid buffer overflow there is a natural competition between two tendencies. One may choose to serve a particular queue, since that will help prevent its own buffer from overflowing, or one may prefer to stop service, with the goal of preventing overflow of buffers further down the line. The solution to the PDE indicates the optimal choice between these two, specifying the parts of the state space where each queue must be served (so as not to lose optimality), and where it can idle. Referring to those queues which must be served as bottlenecks, one can use the solution to the PDE to explicitly calculate the bottleneck queues as a function of the system's state, in terms of a simple set of equations.

Keywords: large deviations, stochastic risk-sensitive control, differential game, HJB equation, tandem queue

AMS subject classification: 60F10, 60K25, 49N70, 93E20

* This research was supported in part by the United States–Israel Binational Science Foundation (BSF 1999179).

** Research of this author also supported in part by the fund for promotion of research at the Technion.

*** Research of this author also supported in part by INTAS grant 265, and in part by the fund for promotion of research at the Technion.

† Research of this author also supported in part by the National Science Foundation (NSF-DMS-0072004, NSF-ECS-9979250) and the Army Research Office (DAAD19-02-1-0425).

1. Introduction

In a previous work [2] we considered a stochastic control problem for a Markovian queueing network with deterministic routing, where the service stations may provide service to one or more queues, with each queue being limited by a finite buffer. The control refers to the service discipline at the service stations, and the cost involves the time till one of the queues first reaches its buffer limit. Such a problem can be regarded as the control of a Markov process up to the time it exits a domain, the domain being the rectangle associated with the buffer sizes. The cost is chosen so as to obtain a risk-sensitive (or rare-event) control problem: one considers $E_x e^{-c\sigma}$ as a criterion to be minimized, where $c > 0$ is fixed and σ denotes the exit time (the time when any one of the buffers first overflows). Such a criterion penalizes short exit times more heavily than ordinary escape time criteria (such as $E_x \sigma$, a criterion to be maximized). While the main result of [2] is the characterization of a limiting problem as the buffer sizes tend to infinity, the current work focuses on finding explicit solutions to this limit problem and on the interpretation thereof.

There are at least two motivations for the use of risk-sensitive criteria when designing policies for the control of a network. The first is that in many communication networks performance is measured in terms of the occurrence of rare but critical events. Buffer overflow is a principal example of such an event. The second motivation follows from the connection between risk-sensitive controls and robust controls. Indeed, as discussed in [8], the optimization of a single fixed stochastic network with respect to a risk-sensitive cost criteria automatically produces controls with specific and predictable robust properties. In particular, these controls give good performance for a family of perturbed network models (where the perturbation is around the design model and the size of the perturbation is measured by relative entropy), and with respect to a corresponding ordinary (i.e., not risk-sensitive) cost.

In many problems, one considers the limit of the risk-sensitive problem as a scaling parameter of the system converges, in the hope that the limit model is more tractable. This is the path followed in [2], in the asymptotic regime where time is accelerated and buffer lengths are enlarged by a factor n . The limit of the normalized cost was characterized both as the value function of a differential game, and as the solution to a corresponding nonlinear PDE. Both interpretations are important and useful. It is the interpretation as the value function of a differential game that is key in quantifying the robust aspects of the resulting control policy. However, because the PDE gives a necessary and sufficient characterization of the (a priori unknown) value function, it can provide qualitative information on the structure of the value function and the optimal controls. In particularly favorable circumstances one can go even further, and use the PDE for exact calculation and control policy synthesis. Indeed, if one can by any means guess a proper parametric form of the value function, then it is sometimes possible to use the PDE to verify that this is the correct form, and identify the unknown parameters in the representation for the value function. The instances in the control theory literature where this has been carried out are few and far between, especially when the state space of the

system has dimension greater than two. However, important information is obtained from these instances on the structure of optimal controls.

In this paper the focus is on applying the PDE obtained in [2] in the manner just described. In particular, we treat in detail the general case of a tandem queueing network, and construct an explicit solution to the PDE. In a system of queues in tandem, each server offers service to exactly one queue, and therefore the service control refers simply to whether each station should provide service or remain idle. Naturally, it is important to provide service to a queue in order to keep it from reaching the buffer limit. On the other hand, if the next buffer in line is nearly full, it might be necessary to idle the first queue in order to keep the second from overflowing. Besides factoring in how close all buffers are to their respective limits, one must consider the mean service rates and the likelihood of significant deviations from those mean service rates. Is it likely that the next queue down will stall and simply stop serving for a while? It is also possible that one will have to look even further ahead, and consider buffers further downstream.

Although the optimal control problem for the Markovian queueing system is fully described by a dynamic programming equation [2], such equations are typically solved numerically. It is hard to extract any global structural information from the exact equation, and even a numerical solution may not be feasible when buffers are large or the dimension of the state space is moderate. As discussed previously, the solution to the PDE for the limiting problem turns out to simplify things significantly. Roughly speaking, the PDE indicates the following structure of the asymptotic optimally controlled network. In the interior of the domain (i.e., when all buffers are away from their maximum capacity), and depending on the state of the system, service must be provided at certain service stations, while other stations may idle without causing loss of optimality. In that sense, the limit problem sharpens the control policy by emphasizing the importance of serving those “bottleneck” buffers (it is crucial to serve the bottleneck buffers, and completely unimportant to serve the others). The identification of the bottlenecks is nontrivial, and indeed queues with the smallest service rate are not necessarily bottlenecks. Instead, as hinted above, identification of the bottlenecks must include consideration of at least the following: (i) the relative closeness of all buffers to their maximum value, (ii) relative mean service rates, and (iii) relative uncertainties in the service rates. We identify a system of algebraic equations whose solution identifies all bottlenecks.

There is relatively little work on risk-sensitive and robust control of networks. Ball et al. have considered a robust formulation for network problems arising in vehicular traffic [4], where the cost structure is qualitatively different. Dupuis studies a robust control problem for networks in a deterministic setting and obtains explicit solutions for the value function [7]. The cost there is, in a sense, antipodal to the one considered in the current work, namely, the time till the system becomes empty (a criterion to be minimized). For other recent work on queueing control in a large deviation regime see [15], where a single server non-Markovian system (with quite general stationary increments input flows) is studied, and a particularly simple scheduling control policy is shown to be asymptotically optimal (see also an extension of the work in [14]).

The paper is organized as follows. Section 2 introduces the model and the PDE and states the main result. Section 3 contains discussion and interpretation. In section 4 we prove the main result.

Notation. The symbol \vee stands for maximum, while \wedge stands for minimum. Denote scalar product between two vectors as $x \cdot y$. For integers $i \leq j$, let $[i, j] \doteq \{1, \dots, j\}$.

2. Model and preliminaries

The queueing network control problem. We consider the following tandem network. There are J queues and J servers. Customers present at queue i at a certain time are said to be of class i at that time. Customers arrive to queue 1 according to a Poisson process of rate $\lambda \geq 0$. Service at queue i is provided by server i at exponential time with parameter $\mu_i > 0$, mutually independent and independent of the arrivals. After a customer is served by server i , it moves to queue $r(i)$, where $r(i) = i + 1, i = 1, \dots, J - 1, r(J) = 0$, and $i = 0$ is used to denote the “outside”. The state of the network is the vector of queue lengths, denoted by X . Let $\{e_i; i = 1, \dots, J\}$ denote the unit coordinate vectors, let $e_0 = 0$ and denote

$$\gamma_i = e_i - e_{r(i)}. \quad (1)$$

Note that $\gamma_i = e_i - e_{i+1}$ for $i = 1, \dots, J - 1, \gamma_J = e_J$, and that following service to queue i the state changes by $-\gamma_i$. The control is specified by the vector $u = (u_1, \dots, u_J)$, where $u_i = 1$ if customers in queue i are given service and $u_i = 0$ otherwise. We next consider the scaled process X^n under the scaling which accelerates time by a factor of n and down-scales space by the same factor. We are interested in a risk-sensitive cost functional that is associated with exit from a bounded set. Let G be the rectangle defined through

$$G = \{(x_1, \dots, x_J): 0 \leq x_1 < z_1; 0 \leq x_i \leq z_i, i = 2, \dots, J\}, \quad (2)$$

for some $z_i > 0, i = 1, \dots, J$. Note that G contains parts of, but not all of its boundary. Let

$$\sigma^n \doteq \inf\{t: X^n(t) \notin G\}.$$

The control problem is to minimize the cost $E_x e^{-nc\sigma^n}$, where E_x denotes expectation starting from x , and $c > 0$ is a constant. With this cost structure “risk-sensitivity” means that a typically short exit times are weighted heavily by the cost. As a result, even if exit within short time occurs with small probability, it may have a noticeable effect on the cost, and thus a “good” control will attempt to avoid such events as much as possible.

Although this paper do not directly treat the stochastic control problem but the corresponding PDE, we give here the precise formulation of the former, for completeness.

For each n , the state space is $G^n \doteq n^{-1}\mathbb{Z}_+^J \cap G$. The control (or action) space is given by

$$U \doteq \{(u_i), i = 1, \dots, J: 0 \leq u_i \leq 1, i = 1, \dots, J\}.$$

One considers state processes X^n that are similar to those defined for the original queueing network, except that time is accelerated (equivalently, arrival and service rates are multiplied) by a scaling parameter n , and space is scaled down by the same factor. More precisely, for $u \in U$ and $f: \mathbb{Z}_+^J \rightarrow \mathbb{R}$, let

$$\mathcal{L}^u f(x) = \lambda[f(x + e_1) - f(x)] + \sum_{i=1}^J u_i \mu_i \mathbb{1}_{\{x - \gamma_i \in \mathbb{Z}_+^J\}} [f(x - \gamma_i) - f(x)], \quad x \in \mathbb{Z}_+^J.$$

For $n \in \mathbb{N}$ let

$$\mathcal{L}^{n,u} f(x) = n\mathcal{L}^u g(nx),$$

where $f: n^{-1}\mathbb{Z}_+^J \rightarrow \mathbb{R}$ and $g(\cdot) = f(n^{-1}\cdot)$. A *controlled Markov process* starting from $x \in G^n$ will consist of a complete filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), P_x^{n,u})$, a state process X^n taking values in G^n and a control process u taking values in U , such that X^n is adapted to \mathcal{F}_t , u is measurable and adapted to \mathcal{F}_t , $P_x^{n,u}(X^n(0) = x) = 1$, and for every function $f: G^n \rightarrow \mathbb{R}$

$$f(X^n(t)) - \int_0^t \mathcal{L}^{n,u(s)} f(X^n(s)) ds$$

is an \mathcal{F}_t -martingale. $E_x^{n,u}$ denotes expectation with respect to $P_x^{n,u}$. For a parameter $c > 0$, the value function for the stochastic control problem is defined by

$$V^n(x) \doteq -\inf n^{-1} \log E_x^{u,n} e^{-nc\sigma_n}, \quad x \in G^n, \quad (3)$$

where the infimum is over all controlled Markov processes.

The domain and its boundary. It is possible for the controller to prevent any but the first queue from exceeding z_i , simply by turning off service to the preceding queue. However, the controller cannot prevent overflow of the first queue. Although it is in principle possible that the dynamics could exit through the portion of the boundary defined by queues $2, \dots, J$, it is always optimal for the controller to not allow this. Consider the simple two-class network illustrated in figure 1. The controller can prevent exit through the dashed portion of the boundary simply by stopping service at the first queue. As a consequence, there are, in general, three different types of boundary – the constraining boundary due to non-negativity constraints on queue length, the part of the boundary where exit can be blocked, and the remainder. These three types of boundary behavior result in the PDE in three types of boundary conditions. The three portions of the boundary are explicitly given as

$$\partial_c G = \{(x_1, \dots, x_J): 0 \leq x_1 < z_1, \text{ and } x_i = z_i \text{ for some } i > 1\},$$




Figure 1. A simple queueing network, a rectangular domain and three types of boundary. Full line: ∂_+G , dashed line: ∂_cG , and dotted line: ∂_oG .

$$\begin{aligned}\partial_oG &= \{(x_1, \dots, x_J): x_1 = z_1, \text{ and } 0 \leq x_i \leq z_i \text{ for all } i > 1\}, \\ \partial_+G &= \{(x_1, \dots, x_J): x_i < z_i \text{ for all } i, \text{ and } x_i = 0 \text{ for some } i\}.\end{aligned}$$

Note that ∂_cG , ∂_oG and ∂_+G partition the boundary ∂G of G . Also, $\partial_cG \subset G$ while $\partial_oG \cap G = \emptyset$. As usual, we will denote $G^\circ = G \setminus \partial G$ and $\overline{G} = G \cup \partial G$.

The Hamiltonian, PDE and viscosity solutions. It is a standard fact that the value functions V^n of the stochastic control problem considered above can be characterized by a dynamic programming equation. The results of [2] show that also the limit of V^n as $n \rightarrow \infty$ can be characterized by a certain equation, namely, a Hamilton–Jacobi–Bellman PDE, that can in fact be regarded as the dynamic programming equation for a certain deterministic game problem. For details on this PDE we need some notation (the reader is referred to [2] for details on the game). Let $l: \mathbb{R} \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ be defined as

$$l(x) \doteq \begin{cases} x \log x - x + 1, & x \geq 0, \\ +\infty, & x < 0, \end{cases}$$

where $0 \log 0 \doteq 0$. Define

$$M = \{m = (\bar{\lambda}, \bar{\mu}_1, \dots, \bar{\mu}_J): \bar{\lambda} \geq 0, \bar{\mu}_i \geq 0\}. \quad (4)$$

For $u \in U$ and $m \in M$ define

$$\begin{aligned}v(u, m) &\doteq \bar{\lambda} e_1 - \sum_{i=1}^J u_i \bar{\mu}_i \gamma_i, \\ \rho(u, m) &\doteq \lambda l\left(\frac{\bar{\lambda}}{\lambda}\right) + \sum_{i=1}^J u_i \mu_i l\left(\frac{\bar{\mu}_i}{\mu_i}\right),\end{aligned}$$

where γ_i are as in (1). Let

$$\begin{aligned}H(p, u, m) &= c + p \cdot v(u, m) + \rho(u, m) \\ &= c + \left[\bar{\lambda} p_1 + \lambda l\left(\frac{\bar{\lambda}}{\lambda}\right) \right] + \sum_{i=1}^J u_i \left[-\bar{\mu}_i \gamma_i \cdot p + \mu_i l\left(\frac{\bar{\mu}_i}{\mu_i}\right) \right]\end{aligned} \quad (5)$$

and let the Hamiltonian be defined as

$$H(p) = \sup_{u \in U} \inf_{m \in M} H(p, u, m) = \inf_{m \in M} \sup_{u \in U} H(p, u, m), \quad (6)$$

where the last identity, expressing the *Isaacs condition* is proved in [2]. The following simplification in the structure of the Hamiltonian will be useful. Using convexity and the fact that the slope of l at 0^+ is $-\infty$, the minimum over m is attained at $\bar{\lambda} = \lambda e^{-p_1}$, $\bar{\mu}_i = \mu_i e^{\gamma_i \cdot p}$. A straightforward calculation then shows that

$$\begin{aligned} H(p, u) &\doteq \inf_m H(p, u, m) = c + [(\lambda - \bar{\lambda}) + u_i(\mu_i - \bar{\mu}_i)] \\ &= c + \lambda(1 - e^{-p_1}) + \sum_{i=1}^J u_i \mu_i (1 - e^{\gamma_i \cdot p}). \end{aligned} \quad (7)$$

Define $I(x) = \{i: x_i = 0\}$. The PDE of interest is the following:

$$\begin{cases} H(DV(x)) = 0, & x \in G^\circ, \\ DV(x) \cdot \gamma_i = 0, & i \in I(x), x \in \partial_+ G, \\ V(x) = 0, & x \in \partial_0 G. \end{cases} \quad (8)$$

Since typically such equations do not possess classical solutions, the framework of viscosity solutions is useful (see [5]). This framework allows for functions that are merely continuous to be regarded as solutions. For $x \in G$, the set of superdifferentials $D^+V(x)$ and the set of subdifferentials $D^-V(x)$ are defined as

$$D^+V(x) = \left\{ p: \limsup_{y \rightarrow x} \frac{V(y) - V(x) - p \cdot (y - x)}{|y - x|} \leq 0 \right\}, \quad (9)$$

$$D^-V(x) = \left\{ p: \liminf_{y \rightarrow x} \frac{V(y) - V(x) - p \cdot (y - x)}{|y - x|} \geq 0 \right\}. \quad (10)$$

Definition 1. V is a viscosity solution to (8) if

$$H(p) \vee \max_{i \in I(x)} p \cdot \gamma_i \geq 0, \quad p \in D^+V(x), x \in G, \quad (11)$$

$$H(p) \wedge \min_{i \in I(x)} p \cdot \gamma_i \leq 0, \quad p \in D^-V(x), x \in G \setminus \partial_c G, \quad (12)$$

and $V(x) = 0$ for $x \in \partial_0 G$.

The following is a special case of [2, theorem 2].

Theorem 1. There exists a unique Lipschitz viscosity solution V to (8). Moreover, if $x_n \in G^n$, $n \in \mathbb{N}$ are such that $x_n \rightarrow x \in G$, then $\lim_{n \rightarrow \infty} V^n(x_n) = V(x)$.

For $i = 1, \dots, J$, let β_i denote the unique positive solution to

$$c + \lambda(1 - e^{\beta_i}) + \mu_i(1 - e^{-\beta_i}) = 0. \quad (13)$$

Set $b_i = \beta_i \sum_{j=1}^i e_j$. Our main result is the following:

Theorem 2. Assume $c > 0$. Then the viscosity solution to the PDE (8) is given by

$$V(x) = \min_{i=1,\dots,J} b_i \cdot (z - x). \quad (14)$$

3. Some remarks

3.1. Interpretation

As commented above, it is possible to characterize the value function V^n as the solution to a dynamic programming equation on the discrete space G^n . Moreover, if V^n is available in explicit form, one can use the dynamic programming equation to specify an optimal control policy for the problem. The results above provide an explicit expression only to the limit $V = \lim_n V^n$, and therefore we are unable to specify an optimal policy for the problem with finite n . Instead, we shall use the quantity V along with the HJB equation (8) to propose a policy for the queueing network problem (and the corresponding value V^n). It is plausible that the proposed policy is, in a sense, asymptotically optimal as $n \rightarrow \infty$ but we do not attempt to prove such a statement here. Instead, the discussion below could be regarded as a natural interpretation of the result. Moreover, since V can be considered as a good approximation to V^n (for large n), it can be used as an initial condition in a value iteration procedure that calculates V^n and an optimal policy.

The solution to the PDE stated here may be interpreted as follows. Let x be an interior point of G , and assume that V is differentiable at x . The form (14) implies that for some $j = j(x)$,

$$DV(x) = -b_j = -\beta_j \sum_{i=1}^j e_i. \quad (15)$$

Hence by (7), the equation $H(-b_j) = 0$ has the form

$$\sup_{u \in U} [c + \lambda(1 - e^{\beta_j}) + u_j \mu_j (1 - e^{-\beta_j})] = 0.$$

It is seen that the supremum is attained at any u for which $u_j = 1$ and $u_i \in [0, 1]$, $i \neq j$. The interpretation in terms of the service policy is that it is important to serve class j , while optimality does not depend on the service given to other classes. *The station $j = j(x)$ can therefore be regarded as a bottleneck:* it is crucial to serve station j when at state x . Since the PDE describes the limit of the stochastic control problem, one expects the bottleneck stations to have a similar property in the stochastic problem: Although it may be optimal for all servers to not idle when the system is near an interior point x , it is significant for the bottleneck server to work while if the other servers idle the cost is affected only by little.

A closer look at (15) reveals that all bottleneck stations belong to a certain set A' defined below. More precisely, given an interior point x where V is differentiable and denoting $y = z - x$, using (14), we see that (15) holds for j if and only if j satisfies

$$\beta_j(y_1 + \cdots + y_j) \leq \beta_i(y_1 + \cdots + y_i) \quad \forall i \neq j. \quad (16)$$

We claim that a necessary condition for j to satisfy (16) (for some y) is $j \in A'$, where

$$A' = \{k \in [1, J]: \mu_k \leq \mu_l, \text{ for all } l < k\},$$

and by convention, $1 \in A'$. To this end, write the explicit form of the solution to (13) as

$$e^{\beta_i} = (2\lambda)^{-1} [c + \lambda + \mu_i + (c^2 + \lambda^2 + \mu_i^2 + 2c(\lambda + \mu_i))^{1/2}],$$

and note that the positive solutions β_i to (13) are monotone in μ_i , in the sense that

$$\mu_i < \mu_j \iff \beta_i < \beta_j. \quad (17)$$

If (16) holds then $\beta_i \leq \beta_j$ for $j = 1, \dots, i$, and it follows from (17) that $j \in A'$.

It should be emphasized that the condition $j \in A'$ is only necessary for j to be a bottleneck, and being a bottleneck is really a function of the system's state. A more complicated necessary condition is given in section 4, which is of the form $j \in A(x)$. One interesting and perhaps counterintuitive phenomenon is that it is possible that a station is nearly empty while others are far from being empty, and still it is a bottleneck.

In the rest of this subsection we discuss a different queueing system for which a similar analysis is possible, including an explicit form for the limit of the value function and a clear interpretation of it regarding nearly optimal service policies for the queueing system (see [2] for proof of the results quoted below). Consider a single server that provides service to J classes (each customer requires service once). Service rate to queue i is exponential with rate μ_i , and the arrival process to class i is Poisson of parameter $\lambda_i > 0$, $i = 1, \dots, J$. All arrival and service processes are mutually independent. One considers the control space

$$U = \left\{ u: \sum_{i=1}^J u_i \leq 1, u_i \geq 0, i = 1, \dots, J \right\},$$

where u_i represents the fraction of service allocated to class i by the single server. Analogously to the problem discussed above, one defines a domain $G = \{(x_1, \dots, x_J): 0 \leq x_i < z_i, 1 = 1, \dots, J\}$, scaled queueing processes X^n and exit time σ^n , and attempts minimizing the cost $E_x e^{-nc\sigma^n}$ over an appropriate class of control processes. Let $V^n(x)$ denote the infimum. One can then show that $\lim_n V^n(x_n)$ exists whenever $x_n \rightarrow x \in G$. Moreover, the limit V is characterized in terms of a PDE of the form (8), where now $\gamma_i = e_i$, $i = 1, \dots, J$, and $H(p) = \sup_{u \in U} H(p, u)$. Here $H(p, u)$ has the form (analogous to (7)):

$$H(p, u) = c + \sum_{i=1}^J [\lambda_i (1 - e^{-p_i}) + u_i \mu_i (1 - e^{p_i})], \quad (18)$$




Figure 2. Priority to class i when the state is in G_i , $i = 1, 2$.

and the limit $V(x)$ can be explicitly calculated provided c is large enough as follows. For $i = 1, \dots, J$, define α_i as the unique positive solution to $c + \lambda_i(1 - e^{-\alpha_i}) + \mu_i(1 - e^{-\alpha_i}) = 0$, namely,

$$e^{\alpha_i} = (2\lambda)^{-1} [c + \lambda + \mu_i + ((c + \lambda + \mu_i)^2 - 4\lambda\mu_i)^{1/2}].$$

Then if c is large enough one has

$$V(x) = \min_{i=1, \dots, J} \alpha_i(z_i - x_i). \quad (19)$$

Let x be a point in the interior of G where the gradient of V is well defined. Then the PDE is satisfied in the classical sense at this point. Thus, $H(DV(x)) = H(-\alpha_j e_j) = 0$, where $j = j(x)$. Using (18), the equation $H(-\alpha_j e_j) = 0$ takes the form

$$\sup_{u \in U} [c + \lambda_j(1 - e^{-\alpha_j}) + u_j \mu_j(1 - e^{-\alpha_j})] = 0.$$

Clearly the supremum is attained at $u_i = 1_{i=j}$, $i = 1, \dots, J$. This means that it is optimal to serve class $j(x)$ at the state x . In the totally symmetric case, where $\mu_i = \mu$, $\lambda_i = \lambda$, $z_i = z$ for all i , the solution takes the form $V(x) = \alpha \min_i(z - x_i)$, and the optimal service discipline can be interpreted as “serve the longest queue.” An asymmetric two dimensional example is given in figure 2, where the domain G is divided into two subdomains G_1 and G_2 in accordance with the structure (19), and the optimal service discipline corresponds to giving priority to class i when the state is within G_i , $i = 1, 2$. This discipline gives priority to the queue with the (weighted) shortest free buffer space.

3.2. The perturbed rates

In the asymptotic analysis of rare events it is often the case that most of the probability mass of an event of interest is concentrated on the event that the stochastic process, say X^n , nearly follows a certain deterministic path, for large values of n . For example, the most likely way a stable M/M/1 queue overflows is by nearly following a deterministic path that is the solution to a fluid model in which arrival and service rates are reversed. This appears as a consequence of time-reversal arguments in [12]. For many other cases where one can compute deterministic paths about which most mass is concentrated see [13]. Similarly, in the current stochastic control problem, when the system

operates under the optimal control most contribution to the cost is obtained when the process X^n nearly follows a certain deterministic path. This path is now the solution to a fluid model with a control u and perturbed arrival and service rates $m = (\bar{\lambda}, \bar{\mu}_1, \dots, \bar{\mu}_J)$. At a point x where the viscosity solution V to (8) is differentiable, the correct values of u and m are those that achieve the max–min in $H(DV(x)) = \sup_u \inf_m H(DV(x), u, m)$ (cf. (6)).

We remark that in the example of tandem queues the optimal perturbed rates satisfy certain relations with the unperturbed rates, as shown in the two equations below. By (7), the relation $H(DV) = 0$ implies

$$\bar{\lambda} + \sum_{i=1}^J u_i \bar{\mu}_i = c + \lambda + \sum_{i=1}^J u_i \mu_i. \tag{20}$$

Moreover, by (7), by the fact that the minimum over m is attained at $\bar{\lambda} = \lambda e^{-p_1}$, $\bar{\mu}_i = \mu_i e^{\gamma_i \cdot p}$, and using $\sum_{i=1}^J \gamma_i = e_1$, we have

$$\bar{\lambda} \prod_{i=1}^J \bar{\mu}_i = \lambda \prod_{i=1}^J \mu_i. \tag{21}$$

Relation (20) was noticed by Avram [3]. The relation (21) appears to be new. A generalization of these equations to a more general network is possible and will appear elsewhere.

Remark. An interesting relation between the roots α_i and β_i and the busy cycle period was pointed out to us by Boxma [6]. For example, if B denotes the busy cycle period for an M/M/1 queue under the stationary distribution, then $E e^{cB} = e^\beta$, where

$$c + \lambda(1 - e^\beta) + \mu(1 - e^{-\beta}) = 0$$

(compare with (13)).

4. Proof of the main result

Recall that $I(x) = \{i \in [1, J]: x_i = 0\}$, and denote $B(x) = \{i \in [1, J]: x_i = z_i\}$. Note that $I(x)$ and $B(x)$ do not intersect. Since points x for which $x_1 = z_1$ are not in G , we have $1 \notin B(x)$ for all $x \in G$.

We show that the minimum over $i = 1, \dots, J$ in (14) can equivalently be performed over a smaller (state-dependent) set. As in section 3, let

$$A' = \{i \in [1, J]: \mu_i \leq \mu_j, \text{ for all } j < i\}, \tag{22}$$

and by convention let $1 \in A'$. Let also

$$A(x) = \{i \in A': \text{either } i + 1 \notin B(x); \text{ or } j > i \text{ and } [i + 1, j] \subset B(x) \text{ imply } \mu_i < \mu_j\}.$$

Lemma 1. The minimum in (14) is obtained over the indices in the set $A(x)$ defined above. More precisely,

$$V(x) = \min_{i=1,\dots,J} b_i \cdot (z - x) = \min_{i \in A(x)} b_i \cdot (z - x).$$

Proof. Fix $x \in G$ and recall that $z_1 - x_1 > 0$. If

$$b_i \cdot (z - x) \leq \min_{j \neq i} b_j \cdot (z - x), \quad (23)$$

then $\beta_i \leq \beta_j$ for $j \in [1, i]$, and it follows from (17) that $i \in A'$. Thus

$$V(x) = \min_{i=1,\dots,J} b_i \cdot (z - x) = \min_{i \in A'} b_i \cdot (z - x)$$

for all $x \in G$ (although not every $i \in A'$ is a minimizer). Fix now $i \in A'$, and suppose that

$$\text{there is some } j > i \text{ so that } [i + 1, j] \subset B(x) \text{ and } \mu_i \geq \mu_j. \quad (24)$$

Since $i \in A'$, if $i + 1 \notin A'$ then $\mu_{i+1} > \mu_i$. But then, if in addition $i + 2 \notin A'$ we have $\mu_{i+2} > \mu_i$. So, if $k \notin A'$ for all $i + 1 \leq k \leq j - 1$ then $\mu_k > \mu_i$, and together with $\mu_j \leq \mu_i$ we have that $j \in A'$. We conclude that under (24),

$$\text{there is } k \in [i + 1, j] \text{ such that } k \in A'. \quad (25)$$

By (24), $\beta_i \geq \beta_j$ and so

$$b_i \cdot (z - x) \geq b_j \cdot (z - x). \quad (26)$$

Therefore, for k as in (25) and under the assumptions in (24),

$$b_i \cdot (z - x) \geq b_k \cdot (z - x),$$

and since $k \in A'$, we need not consider i in the minimum. As a result, (14) is equivalently given as

$$V(x) = \min_{i \in A(x)} b_i \cdot (z - x), \quad x \in G,$$

where $A(x)$ is obtained from $A'(x)$ by deleting those $i \in A'(x)$ which are followed by empty queues, $[i + 1, j] \subset B(x)$, and such that $\mu_i \geq \mu_j$. This leaves in $A(x)$ an index i from $A'(x)$ only if either $z_{i+1} > x_{i+1}$, or $z_k = x_k$ for $i + 1 < k \leq j$ implies $\mu_i < \mu_j$. More formally,

$$\begin{aligned} A(x) &= A' \setminus \{i \in A': \exists j > i, [i + 1, j] \subset B(x), \mu_i \geq \mu_j\} \\ &= \{i \in A': \text{either } i + 1 \notin B(x); \text{ or } j > i \text{ and } [i + 1, j] \subset B(x) \text{ imply } \mu_i < \mu_j\}. \end{aligned} \quad (27)$$

□

The form of the proposed solution is the minimum of smooth functions. The set of superdifferentials at a point where the function is not smooth can be seen (using the

definition (9)) to consist of the convex hull of the gradients of the smooth functions defining it, at that point. Also, at the boundary, the fact that there are less constraints introduced by (9) on $p \in D^+V$ than there are when x is in the interior, has an effect of enlarging the set further. Thus, for example, the set of superdifferentials of the zero function from \mathbb{R}_+ to \mathbb{R} at zero is \mathbb{R}_+ . Using these considerations and the rectangular structure of the domain, one finds the general form for the superdifferential as follows. Let $x \in G$ be fixed. Set $A = A(x)$, $I = I(x)$, $B = B(x)$, and $O = \{1, \dots, J\} \setminus (I \cup B)$. Then any element $p \in D^+V(x)$ is given as

$$p = - \sum_{i \in A} v_i b_i + \delta,$$

for some $\{v_i\}$, where

$$\delta_i \geq 0, \quad i \in I, \quad \delta_i \leq 0, \quad i \in B, \quad \delta_i = 0, \quad i \in O, \quad (28)$$

and

$$v_i \geq 0, \quad i \in A, \quad v_i = 0, \quad i \notin A, \quad \sum_{i \in A} v_i = 1. \quad (29)$$

If we denote

$$\delta_{J+1} = 0, \quad (30)$$

then, using (7), the Hamiltonian is expressed as

$$\begin{aligned} H(p) &= c + \lambda(1 - e^{-p_1}) + \sum_{i=1}^J 0 \vee \mu_i(1 - e^{p \cdot \gamma_i}) \\ &= c + \lambda(1 - e^{-\sum_{i \in A} v_i \beta_i - \delta_1}) + \sum_{i=1}^J 0 \vee \mu_i(1 - e^{-v_i \beta_i + \delta_i - \delta_{i+1}}). \end{aligned} \quad (31)$$

Proof of theorem 2. Verifying the PDE for superdifferentials. To verify (11), it suffices to show that $H(p) \geq 0$ whenever $p \in DV^+(x)$ and $p \cdot \gamma_i < 0$ for all $i \in I$.

Step 1. We show first that $H(p) \geq 0$ whenever $p \in DV^+(x)$ and $p \cdot \gamma_i \leq 0$ for all $i = 1, \dots, J$. The forms of b_i and $\gamma_i = e_i - e_{i+1}$ imply the last inequality can be rewritten

$$-v_i \beta_i + \delta_i - \delta_{i+1} \leq 0, \quad i = 1, \dots, J. \quad (32)$$

In this case, (31) becomes

$$H(p) = h(v, \delta) \doteq c + \lambda(1 - e^{-\sum_{i \in A} v_i \beta_i - \delta_1}) + \sum_{i=1}^J \mu_i(1 - e^{-v_i \beta_i + \delta_i - \delta_{i+1}}). \quad (33)$$

The constraints we have put on (v, δ) define a convex set S , namely, the set of $(v, \delta) \in \mathbb{R}^J \times \mathbb{R}^{J+1}$ satisfying (28)–(30) and (32). The set S is bounded for the following reasons.

First, since ν is a probability vector (see (29)) we have $\nu_i \in [0, 1]$. Next, by (28) and $1 \notin B$, $\delta_1 \geq 0$. Finally, it follows from (32) (using the convention $\sum_s^t = 0$ if $t < s$) that

$$\delta_1 - \sum_{j=1}^{i-1} \nu_j \beta_j \leq \delta_i \leq \sum_{j=i}^J \nu_j \beta_j, \quad i = 1, \dots, J.$$

This shows the boundedness.

The function $(\nu, \delta) \mapsto h(\nu, \delta)$ is concave, and therefore to prove $h \geq 0$ on S it suffices to check the inequality on the set E of extreme points of S , which contains a finite number of points since the constraints are linear.

Lemma 2. All points $(\nu, \delta) \in E$ have the following form:

$$\begin{aligned} \nu_k &= 1 && \text{for some } k, \text{ and} \\ \delta_i^{(r,k)} &= \beta_k (\mathbb{1}_{i \geq r+1} - \mathbb{1}_{i \geq k+1}), && i = 1, \dots, J, \quad r = s, \dots, t-1, \end{aligned}$$

where t, s are defined below.

Proof. We will first obtain the general form of ν and then that of δ .

We claim that for any $(\nu, \delta) \in E$, ν is of the form $\nu = \mathbb{1}_k$ (short for $\nu_i = \mathbb{1}_{i=k}$, $i = 1, \dots, J$), for some $k \in A$. Assume this is false. Then there are $l, m \in A$, $l < m$, for which $\nu_l, \nu_m \in (0, 1)$. We will show that there is a vector Δ_m , such that replacing (ν, δ) by $(\nu', \delta') = (\nu - \varepsilon e_m, \delta + \varepsilon \Delta_m)$ maintains the relations (28), (30) and (32) for both $\varepsilon > 0$ and $\varepsilon < 0$ (provided $|\varepsilon|$ is small). A similar statement will hold also for $(\nu'', \delta'') = (\nu + \varepsilon e_l - \varepsilon e_m, \delta - \varepsilon \Delta_l + \varepsilon \Delta_m)$, and as a result all of (28)–(30) and (32) will hold for both $\varepsilon > 0$ and $\varepsilon < 0$, a contradiction to (ν, δ) being an extreme point of S . The construction of Δ_m based on (ν, δ) can be mimicked to construct Δ_l based on (ν', δ') (in particular, no use is made of the fact that ν_i sum to one, but only that some of its components are within $(0, 1)$), and therefore the latter construction is omitted.

Case 1. Inequality (32) holds as a strict inequality for $i = m$, i.e., $-\nu_m \beta_m + \delta_m - \delta_{m+1} < 0$. Then $-(\nu_m - \varepsilon) \beta_m + \delta_m - \delta_{m+1} \leq 0$ also holds (for both ε positive and negative), provided that $|\varepsilon|$ is small. Here we take $\Delta_m = 0$. We see that $(\nu - \varepsilon e_m, \delta)$ satisfies the requirements.

Case 2. Inequality (32) holds with equality for $i = m$, i.e., $-\nu_m \beta_m + \delta_m - \delta_{m+1} = 0$. Since $\nu_m \beta_m > 0$, either $\delta_m > 0$ or $\delta_{m+1} < 0$. Assume $\delta_m > 0$ (the case $\delta_{m+1} < 0$ can be treated analogously, and is therefore omitted). Let m' be the smallest $j \in [1, m]$ for which (32) holds with equality for all $i \in [j, m]$. For $i \in [m', m-1]$ (the set being empty and the statement void if $m' = m$) we have $-\nu_i \beta_i + \delta_i - \delta_{i+1} = 0$, hence if $\delta_{i+1} > 0$ then $\delta_i > 0$. Together with the fact that $\delta_m > 0$, this shows that

$$\delta_i > 0, \quad i \in [m', m]. \quad (34)$$

Moreover,

$$-\nu_{m'-1} \beta_{m'-1} + \delta_{m'-1} - \delta_{m'} < 0 \quad (35)$$

(the statement being void in case that $m' = 1$). Set $\Delta_m = -\beta_m \sum_{i=m'}^m e_i$, $v' = v + \varepsilon e_m$, and $\delta' = \delta + \varepsilon \Delta_m$. The perturbation $\varepsilon \Delta_m$ is chosen to cancel the change in $-v_m \beta_m$, and to preserve the left-hand side of (32) for all $i > m' - 1$. However, by (35) the inequality is maintained for $i = m' - 1$ as well, by taking $|\varepsilon|$ sufficiently small. Similarly, by (34), (28) and (30) are also maintained on taking $|\varepsilon|$ small.

This completes the construction of Δ_m . As described above, this leads to a contradiction, and we conclude that any extreme point $(v, \delta) \in E$ satisfies $v = \mathbb{1}_k$ for some $k \in A$.

When $v = \mathbb{1}_k$ for some $k \in A$, inequalities (32) can be rewritten as

$$\delta_i \leq \delta_{i+1}, \quad i \neq k, \quad \delta_k \leq \beta_k + \delta_{k+1}. \quad (32')$$

In particular, $0 \leq \delta_1 \leq \dots \leq \delta_k$. Let s denote the largest $j \leq k$ for which $j \in B \cup O$ (and $s = 0$ if there is no such j). The definitions of O and B imply $\delta_s = 0$, and thus

$$\delta_j = 0 \quad \text{for } 1 \leq j \leq s. \quad (36a)$$

Similarly, $\delta_{k+1} \leq \dots \leq \delta_J \leq 0$, and if t denotes the least $j \in [k+1, J]$ for which $j \in I \cup O$ (and $t = J+1$ if empty), then

$$\delta_j = 0, \quad t \leq j \leq J. \quad (36b)$$

Hence δ must satisfy

$$\begin{cases} 0 \leq \delta_{s+1} \leq \dots \leq \delta_k \leq \beta_k + \delta_{k+1}, & \text{when } s \leq k-1, \\ 0 \leq \delta_k, & \text{when } s = k. \end{cases} \quad (36c)$$

In addition,

$$\begin{cases} \delta_{k+1} \leq \dots \leq \delta_{t-1} \leq 0, & \text{when } t \geq k+2, \\ \delta_{k+1} \leq 0, & \text{when } t = k+1. \end{cases} \quad (36d)$$

We have just shown that (28), (30) and (32') imply (36a–d). On the other hand, clearly (36a–d) implies (32'). Moreover, as follows directly from the definition of s and t ,

$$i \in I \quad \text{for all } i \in [s+1, k] \quad \text{and} \quad i \in B \quad \text{for all } i \in [k+1, t-1]. \quad (37)$$

This shows that (36a–d) implies (28) and (30). Thus (28), (30) and (32') are equivalent to (36a–d). Now, the set of δ satisfying the constraints (36a–d) is easy to analyze. In particular, it is not hard to see that it has the following $t-s$ extreme points, indexed by k and $r \in [s, \dots, t-1]$, namely,

$$\delta_i^{(r,k)} = \beta_k (\mathbb{1}_{i \geq r+1} - \mathbb{1}_{i \geq k+1}), \quad i = 1, \dots, J, \quad r = s, \dots, t-1. \quad (38)$$

□

We now calculate h (cf. (33)) at each extreme point. To this end, note that if $(v, \delta) = (\mathbb{1}_k, \delta^{(r,k)})$ then by (38)

$$-v_i \beta_i + \delta_i - \delta_{i+1} = \begin{cases} -\beta_k, & i = r, \\ 0, & i \neq r, \end{cases} \quad i = 1, \dots, J,$$

and

$$\delta_1 = \begin{cases} \beta_k, & r = 0, \\ 0, & r > 0. \end{cases}$$

Substituting in (33), we have the following possibilities. If $r = 0$ then $\delta_i^{(r,k)} = \beta_k \mathbb{1}_{i \in [1,k]}$ thus $\sum v_i \beta_i - \delta_1 = 0$, and $-v_i \beta_i + \delta_i - \delta_{i+1} = 0$, $i = 1, \dots, J$. Hence $h(\mathbb{1}_k, \delta^{(r,k)}) = c > 0$. Otherwise,

$$h(\mathbb{1}_k, \delta^{(r,k)}) = c + \lambda(1 - e^{\beta_k}) + \mu_r(1 - e^{-\beta_k}). \quad (39)$$

In case that $r = k$, the right-hand side of (39) vanishes owing to the definition of β_k (see (13)). In case that $r < k$, recall that $k \in A$, and in particular, $k \in A'$. By (22) we therefore have $\mu_k \leq \mu_r$, hence by (13), $h(\mathbb{1}_k, \delta^{(r,k)}) \geq 0$. Finally, consider the case where $k < r \leq t - 1$. By (37), $[k + 1, r] \subset B$. Since $k \in A$, (27) implies that $\mu_k < \mu_r$, and we again conclude that $h(\mathbb{1}_k, \delta^{(r,k)}) > 0$. Having shown that $h(v, \delta) \geq 0$ for all extreme points of the set S , we conclude that the inequality holds for all $(v, \delta) \in S$.

Step 2. We now relax the condition (32). Let then (v, δ) satisfy (28)–(30), and assume that the inequality in (32) holds for all $i \in I$ (relaxing the assumption made in step 1, that it holds for all $1 \leq i \leq J$). Let $P = P(v, \delta)$ denote the set of i such that

$$-v_i \beta_i + \delta_i - \delta_{i+1} > 0. \quad (40)$$

If P is empty then the results of step 1 apply and $H(p) \geq 0$. Hence assume P is not empty. Let $j = j(v, \delta)$ be the least element in P .

Note that $J \notin P$. For if $J \in P$ then using (40) and (30) we find $\delta_J > 0$, and therefore $J \in I$. However, we get to assume (32) for all $i \in I$, which means that $-v_J \beta_J + \delta_J - \delta_{J+1} \leq 0$. This contradicts $J \in P$.

We proceed by backward induction on the value of $j(v, \delta)$. Note that the sets B and I depend on x . The argument below treats simultaneously all $x \in G$ by considering all possible sets B and I .

Induction step. Assumption: for all (I, B, v, δ) satisfying (28)–(30), such that $P(v, \delta) \cap I = \emptyset$ and $j(v, \delta) \in [i + 1, i + 2, \dots, J - 1]$, one has $H(p) \geq 0$. Let (I, B, v, δ) be such that $P(v, \delta) \cap I = \emptyset$ and $j(v, \delta) = i$ ($i \geq 1$). Then $-v_i \beta_i + \delta_i - \delta_{i+1} > 0$. Modify δ_{i+1} by increasing it so as to get equality, i.e., set $\delta'_{i+1} > \delta_{i+1}$ so that $-v_i \beta_i + \delta_i - \delta'_{i+1} = 0$. This modification does not change the value of $0 \vee \mu_i(1 - e^{-v_i \beta_i + \delta_i - \delta_{i+1}})$ (but keeps it zero), and it can only decrease (or leave unchanged) the value of $0 \vee \mu_{i+1}(1 - e^{-v_{i+1} \beta_{i+1} + \delta_{i+1} - \delta_{i+2}})$. Hence by (31), the value of $H(p)$ is only decreased. At the same time, $j(v, \delta') > j(v, \delta)$. Note that in modifying δ , (28) need not hold for B and I . However, clearly there are other sets, B' and I' with which it holds. For example,

$$B' = \{i \in [1, J]: \delta'_i < 0\}, \quad I' = \{i \in [1, J]: \delta'_i \geq 0\}, \quad O' = \emptyset. \quad (41)$$

Moreover, $1 \notin B'$, since we had $1 \notin B$ and δ_1 was not modified. By the induction assumption we therefore obtain that $H(p) \geq 0$.

Induction base. We show that for all (I, B, v, δ) satisfying (28)–(30), such that $P(v, \delta) \cap I = \emptyset$ and $j(v, \delta) = J - 1$, one has $H(p) \geq 0$. We have $-v_{J-1}\beta_{J-1} + \delta_{J-1} - \delta_J > 0$. Similar to before, we set $\delta'_J > \delta_J$ so that $-v_{J-1}\beta_{J-1} + \delta_{J-1} - \delta'_J = 0$. We claim that $P(v, \delta') = \emptyset$. Indeed, since $P(v, \delta) = \{J - 1\}$, clearly $P(v, \delta') \subset \{J\}$. However,

$$-v_J\beta_J + \delta'_J = -v_J\beta_J - v_{J-1}\beta_{J-1} + \delta_{J-1} \leq 0,$$

where we have used the fact that $J - 1 \in P(v, \delta)$ implies $J - 1 \notin I$, and hence by (28) $\delta_{J-1} \leq 0$. This shows that $P(v, \delta') = \emptyset$. As in the previous paragraph, because of the modification of δ , (28) need not hold for the I and B we started with, but there are other sets I' and B' (defined, e.g., as in (41)) with which it holds. The results of step 1 therefore apply, and therefore $H(p) \geq 0$.

This completes the argument by induction and establishes (11) for superdifferentials.

Verifying the PDE for subdifferentials. We are required to show that (12) holds for $p \in D^-V(x)$, where $x \in G \setminus \partial_c G$ (and hence $B(x) = \emptyset$). Unless $D^-V(x)$ is empty, the general form of $p \in D^-V(x)$, $x \in G$ is

$$p = -b_k + \delta,$$

where $k \in A$, and

$$\delta_i \leq 0, \quad i \in I, \quad \delta_i = 0, \quad i \notin I. \tag{42}$$

Here we have used the fact that V is the minimum of smooth functions, and therefore away from the boundary $D^-V(x)$ can have at most one element, equal to the gradient of any minimizing function. It suffices to show that $H(p) \leq 0$ whenever $p \cdot \gamma_i = -1_{i=k}\beta_k + \delta_i - \delta_{i+1} > 0$ for all $i \in I$. Using (31), (42) and (13), we find

$$\begin{aligned} H(p) &= c + \lambda(1 - e^{\beta_k - \delta_1}) + \sum_{i=1}^J 0 \vee \mu_i(1 - e^{-1_{i=k}\beta_k + \delta_i - \delta_{i+1}}) \\ &= c + \lambda(1 - e^{\beta_k - \delta_1}) + \sum_{i \notin I} \mu_i(1 - e^{-1_{i=k}\beta_k - \delta_{i+1}}) \\ &\leq c + \lambda(1 - e^{\beta_k}) + \sum_{i \notin I} \mu_i(1 - e^{-1_{i=k}\beta_k}) \\ &\leq c + \lambda(1 - e^{\beta_k}) + \mu_k(1 - e^{-\beta_k}) \\ &= 0. \end{aligned}$$

References

[1] R. Atar and P. Dupuis, A differential game with constrained dynamics, *Nonlinear Analysis* 51(7) (2002) 1105–1130.
 [2] R. Atar, P. Dupuis and A. Shwartz, An escape time criterion for queueing networks: Asymptotic risk-sensitive control via differential games, *Math. Oper. Res.* 28(4) (2003) 801–835.
 [3] F. Avram, On asymptotic optimal control problems for stochastic networks, Preprint (2002).

- [4] J. Ball, M. Day and P. Kachroo, Robust feedback control for a single server queueing system, *Math. Control Signals Systems* 12 (1999) 307–345.
- [5] M. Bardi and I. Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman Equations* (Birkhäuser, Boston, 1997).
- [6] O.J. Boxma, Personal communication.
- [7] P. Dupuis, Explicit solution to a robust queueing control problem, *SIAM J. Control Optim.* 42(5) (2003) 1854–1875.
- [8] P. Dupuis, M.R. James and I.R. Petersen, Robust properties of risk-sensitive control, *Math. Control Signals Systems* 13 (2000) 318–332.
- [9] P. Dupuis and H.J. Kushner, Minimizing escape probabilities: A large deviations approach, *SIAM J. Control Optim.* 27 (1989) 432–445.
- [10] R.J. Elliott and N.J. Kalton, The existence of value in differential games, in: *Memoirs of the American Mathematical Society*, Vol. 126 (Amer. Math. Soc., Providence, RI, 1972) pp. iv+67.
- [11] W.H. Fleming and P.E. Souganidis, PDE-viscosity solution approach to some problems of large deviations, *Ann. Scuola Norm. Sup. Pisa Cl. Sci.* 4 (1986) 171–192.
- [12] A. Shwartz and A. Weiss, Induced rare events: analysis via large deviations and time reversal, *J. Appl. Probab.* 25 (1993) 667–689.
- [13] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis: Queues, Communication and Computing* (Chapman and Hall, New York, 1995).
- [14] A. Stolyar, Control of end-to-end delay tails in a multiclass network: LWDF discipline optimality, *Ann. Appl. Probab.* 13(3) (2003) 1151–1206.
- [15] A. Stolyar and K. Ramanan, Largest weighted delay first scheduling: Large deviations and optimality, *Ann. Appl. Probab.* 11 (2001) 1–48.