CCIT Report #441 August 2003

TREELIKE PARALLEL SERVER STATIONS IN HEAVY TRAFFIC*

Rami Atar
Department of Electrical Engineering
Technion–Israel Institute of Technology
Haifa 32000, Israel

August 6, 2003

Abstract

This paper studies a diffusion control problem arising as the formal limit of a queueing system scheduling problem in the asymptotic heavy traffic regime of Halfin and Whitt. The queueing system consists of several customer classes and many exponential servers working in parallel, grouped in several stations according to their type. Different types of servers offer service to customers of a given class at possibly different (and possibly zero) rates. The diffusion control problem does not seem to have explicit solutions even for simple cost criteria and therefore a characterization of optimal solutions via Hamilton-Jacobi-Bellman (HJB) equations is addressed. Denote by $\mathcal{G}$ the graph having a node for each class, a node for each type, and an edge joining a class and a type if the corresponding service rate is nonzero. The problem turns out to be different in nature depending on whether $\mathcal{G}$ is a tree or not; here we assume $\mathcal{G}$ is a tree. Our main result is the existence and uniqueness for solutions of the HJB equation. Since the cost per unit time is not assumed to be bounded, the analysis requires developing polynomial moment estimates on the state processes. In establishing these estimates, a key role is played by an integral formula relating queue length and idle time processes, that may be of independent interest. Our results cover three classes of problems. (i) Service rates are either class- or type-dependent (with general trees and costs); (ii) Trees satisfying $\mathrm{diam}(\mathcal{G}) \leq 3$ (with general service rates and costs); (iii) Cost per unit time is, in an appropriate sense, comparable to the system's state (and trees and service rates are general).

Keywords: Multiclass queueing networks, scheduling control, heavy traffic regime of Halfin and Whitt, buffer-station trees, optimal control of diffusions, Hamilton-Jacobi-Bellman equations, polynomial moment estimates.

*Research supported in part by the Israel Science Foundation (grant no. 126/02), the US–Israel Binational Science Foundation (grant no. 1999179) and the fund for promotion of research at the Technion

Contents

1 Introduction

We consider optimal scheduling control for a class of queueing networks operating in heavy traffic, in the sense that the load on the system is nearly equal to its capacity. As often occurs, exact analysis of the control problem is unavailable and an asymptotic approach is taken, where a parametrization of the problem is introduced and a diffusion control problem is obtained in the limit. The parametrization that has been more common in research papers on related problems (referred to here as 'conventional' heavy traffic) is one where arrival and service rates are both scaled up in a way that the system operates near full capacity. Recently, several papers have studied a different parametrization, proposed by Halfin and Whitt [10], where increase of arrivals is balanced by scaling up the *number* of (identical) servers in each service station, while keeping the service time distribution of the individual servers fixed. In the limit as the parameter grows without bound, conventional heavy traffic typically gives rise to diffusion processes in the orthant with reflection on the boundary, whereas under the Halfin-Whitt (HW) regime one obtains diffusion processes taking values in the Euclidean space. The subject of this paper is the optimal control of diffusions that correspond to the heavy traffic limit of a family of multiclass networks in the HW regime.

The queueing network considered has $\bar{\imath}$ customer classes and many exponential servers grouped in $\bar{\jmath}$ stations according to their type. Service to a class-i customer by a type-j server is performed at rate depending only on i and j (see Figure 1). The rate can be zero, in which case we say that type-j servers cannot serve class-i customers. Every customer requires service only once. Scheduling of jobs in the system is regarded as control. The cost is an expected discounted cumulative function of performance criteria such as queue lengths, number of idle servers, or number of customers of each class present at each station. The parameterization is such that the arrival rates and the number of servers at each station are nearly proportional to a large parameter n, while service rates are nearly constants as functions of n. In the case of a single class and a single type the number of servers is subtracted from the number of customers in the system and normalized by $\sqrt{n}$. The positive and, respectively, negative parts of this process represent the normalized queue length and number of idle servers. In the limit as $n \to \infty$ one obtains a diffusion on $\mathbb{R}$ with state dependent drift [10]. When multiple classes and types are considered, the processes representing the number of customers of each class present in the system are recentered about quantities arising from a corresponding static fluid model. Identifying these quantities requires solving a linear program associated with a static allocation problem studied in [12]. For motivation on the model and on this asymptotic regime see [4] and [13].

Depending on the application, scheduling in the queueing system can be modelled as nonpreemptive or preemptive. Under nonpreemptive scheduling policies every customer completes service with the server it is first assigned. Preemptive policies allow service to a customer to be stopped and resumed at a later time, possibly in a different station, and so customers can be moved instantaneously between the queue and service stations (among stations that offer service to the corresponding class). In a Markovian queueing network one observes that the dimension of the state space under nonpreemptive scheduling is greater than that under preemptive scheduling (cf. Section 2), and as a result, a similar statement holds for the corresponding controlled diffusions (even if the original queueing system is not Markovian). However, [4] shows that in the case of a single station it is possible to associate with a nonpreemptive scheduling model a controlled diffusion corresponding to a preemptive scheduling model. Thus for nonpreemptive scheduling, the asymptotic model lies in a lower dimension than the original model. Let $\mathcal{G}$ denote the graph having a node for each class, a node for each type, and an edge joining a class and a type if and only if the corresponding rate at the queueing system is zero. Reasoning of Mandelbaum and Reiman [16] along with the results of Harrison and López [12] indicate that for general $\bar{\imath}$ and $\bar{\jmath}$ a dimension reduction as referred to above requires that the graph $\mathcal{G}$ be a tree (cf. Section 4). Also necessary for this reduction is a certain condition on the static fluid model, namely the complete resource pooling condition of [12]. Both these conditions are assumed in this paper. It was proved in [12] that under the complete resource pooling condition a great simplification holds in the conventional heavy traffic regime: Servers in all stations work as if they were a single server with large capacity, and the diffusion becomes one dimensional. However, this condition does not lead to a one dimensional problem in the HW regime.

A related but rather different form of cooperation between the service stations is present in

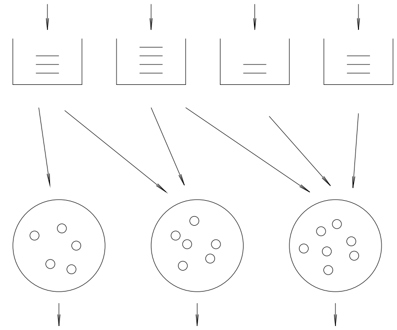

Figure 1: Parallel station system with four customer classes (buffers) and three server types (service stations)

our setting. It is a strong form of a work conservation assumption that we refer to as *joint work conservation* and regard as a constraint on the scheduling control. Under this condition, servers are not allowed to idle when there are any customers in the queue. To maintain joint work conservation in the queueing system and prevent servers at a station from idling when customers are present in the queue, possibly belonging to a class that this station cannot serve, it must be possible to appropriately rearrange customers in the stations. Considering a condition analogous to joint work conservation in the diffusion model, one can show that this condition can be maintained, basically because in a queueing system under preemptive policy such a rearrangement is possible with probability approaching one as the parameter becomes large. We present an argument toward showing that joint work conservation is in fact optimal in a class of examples where the cost is a monotone function of the length of each queue; however, this optimality question is not the main focus of the current paper. In fact, strong work conservation is not optimal for arbitrary cost criteria, but it is believed that the argument can be greatly generalized, and this will be the subject of future research. Technically, this condition simplifies the analysis in that it compactifies the control space. It is assumed throughout.

The main result of this paper is the characterization of the diffusion control problem's value as the solution to the associated HJB equation, uniqueness for this partial differential equation (PDE), and existence of optimal Markov controls. Such issues are well understood for diffusion control problems in bounded domains, or in unbounded domains but where the cost is bounded [5, 7]. Also, in case of unbounded domain and cost, but bounded drift coefficient, existence and uniqueness for the PDE directly follow from [14]. The difficulty in the current model stems from the fact that the domain, the cost and the drift coefficient are unbounded. The question then requires deeper understanding of the model and in particular, developing estimates on moments of the state process that are sub-exponential as a function of time. Our results cover three families of problems. (i) Service rates are either class- or type-dependent (with general trees and costs); (ii) Trees satisfying $\mathrm{diam}(\mathcal{G}) \leq 3$ (with general service rates and costs); (iii) Cost per unit time comparable to $\|X\|^m$ where X denotes the state process and $m \geq 1$ (with general trees and service rates).

Recent results on the HW regime include the following. Puhalskii and Reiman [19] extend the work of Halfin and Whitt to multiple customer classes, priorities and phase-type distribution. Mandelbaum, Massey and Reiman [17] establish functional law of large numbers and central limit theorems for a wide class of Markovian networks in the HW regime. Armony and Maglaras [1] model and analyze rational customers in equilibrium, and Garnett, Mandelbaum and Reiman [8] study models with abandonments from the queue. Papers where a control theoretic approach was taken to study queueing networks in this regime are few. The diffusion control problem associated with scheduling jobs in a system with multiple customer classes and a single type of servers was analyzed in Harrison and Zeevi [13]. In a similar setting, Atar, Mandelbaum and Reiman [4] establish asymptotic optimality of scheduling policies for the queueing network derived from the HJB equation associated with the diffusion problem. A special case where explicit, pathwise solutions to the diffusion control problem are available appeared in [3]. Finally, with regard to the analysis of parallel server systems in conventional heavy traffic, we mention again Harrison and López [12], where the corresponding Brownian control problem is identified and solved, Williams [20], where a dynamic threshold scheduling policy is proposed for the queueing system, conjectured to be asymptotically optimal, and Mandelbaum and Stolyar [18] where asymptotic optimality of a simple scheduling policy is proved for convex delay costs.

Here is an informal description of the diffusion model identified and studied in this paper. $\mathcal{I}$ and $\mathcal{J}$ are index sets for classes and, respectively, types. The controlled diffusion processes $X_i, i \in \mathcal{I}$ represent the number of customers of class i present in the system, appropriately recentered, having initial condition x_i and taking values in $\mathbb{R}$. The processes Y_i, $i \in \mathcal{I}$ and Z_j, $j \in \mathcal{J}$ represent queue length of class i and, respectively, the number of idle servers at station j. For $i \in \mathcal{I}$, $j \in \mathcal{J}$, Ψ_{ij} is a recentered version of the number of class-i customers at station j. Finally, $\tilde{W}_i$ are Brownian motions representing the effect of fluctuations in arrival and service times. The dynamics are described by equations (1)–(3). The constraint (4) comes from the definition of these processes and the constraint (5) expresses joint work conservation.

$$X_i(t) = x_i + \tilde{W}_i(t) - \sum_j \mu_{ij} \int_0^t \Psi_{ij}(s)ds, \quad i \in \mathcal{I} \tag{1}$$

$$\sum_j \Psi_{ij} = X_i - Y_i, \quad i \in \mathcal{I}, \tag{2}$$

$$\sum_i \Psi_{ij} = -Z_j, \quad j \in \mathcal{J}, \tag{3}$$

$$Y_i \geq 0, Z_j \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}. \tag{4}$$

$$\sum_j Z(t) > 0 \quad \text{implies} \quad \sum_i Y_i(t) = 0, \quad t \geq 0. \tag{5}$$

It is possible to encode equations (1)–(5) in a single equation of the form $X(t) = x + rW(t) + \int_0^t b(X(s), U(s))ds$ where U is a control process taking values in a compact space and W is a simple Brownian motion. This formulation is useful, but in the introduction we only consider the form (1)–(5).

The three cases of the main result alluded to above are treated by three different strategies. One of our basic tools is an integral formula that expresses a relation directly between Y, Z and $\tilde{W}$, not involving X and Ψ. Denoting $If = \int_0^\cdot f$ and by I_n the nth power of the operator I, we show

$$\sum_{i\in\mathcal{I}}(x_i + \tilde{W}_i - Y_i) + \sum_{j\in\mathcal{J}} Z_j + \sum_{i\in\mathcal{I}}\sum_{n=1}^{m_i} a_{i,n} I_n(x_i + \tilde{W}_i - Y_i) + \sum_{j\in\mathcal{J}}\sum_{n=1}^{m_j} b_{j,n} I_n Z_j = 0, \tag{6}$$

where m_i, m_j, $a_{i,n}$ and $b_{j,n}$ are positive constants. The special form this equation takes when μ_{ij} depend only on i or only on j makes it possible to get estimates on the moments of X that are polynomial in x and t in case (i) of the main result. The formula (6) may be of independent interest since it can be seen to express a relation between the data $\tilde{W}$, the control process U and the one dimensional process $\sum_i X_i$ alone (cf. Section 7).

This formula is used also in case (ii) of the main result, along with a certain property of the system (1)–(5). This property, that we call the *non-idling property*, should be understood as one of the system of equations rather than the stochastic processes, as it does not regard $\tilde{W}$ as a Brownian motion:

If the system starts with $x_i > 0$ and $t \mapsto \tilde{W}_i(t)$ are strictly increasing, then $Z_j(t) = 0$, $j \in \mathcal{J}$, $t \geq 0$.

For trees of diameter three or less we can show that the property holds and that it implies polynomial moment estimates on the state process. It is conceivable that the argument proving the estimates for trees of diameter three can be extended if this property is proved in greater generality. Since it can be seen that if x and $\tilde{W}$ vanish then so do X, Y and Z, heuristically the property expresses a very weak form of the intuitively sensible statement: "the more work there is, the less the servers are idle". Therefore how generally the non-idling property holds appears to be a basic question on the model, and one would like to understand it irrespective of the goals of the current work.

In case (iii) of the main result we assume that the cost per unit time is comparable to a power of the norm of the state process X. This assumption is not the most natural in the context of queueing, since for example cost functions depending on queue lengths cannot be treated. However, it can be useful if one is interested in stabilizing the dynamical system about a nominal model, since the cost penalizes deviations from the static fluid model. On the technical side, penalizing deviations from a nominal model simplifies the problem in that moment estimates are required under one particular control rather than under all controls, and as a result we can treat the model at full generality, as far as the tree and the service rates are concerned.

Apart from their own contribution the results of this paper are a first step toward identifying scheduling policies for the queueing network that are, in an appropriate sense, asymptotically optimal [2]. As in [4], such policies can be derived from the solution to the diffusion control problem. In addition, such asymptotic analysis would fully justify the relation between the queueing model and the diffusion control model (1)–(5), a relation that is only obtained here by means of formal limits.

The organization of the paper is as follows. Section 2 describes the queueing system model. Section 3 introduces the rescaled parameters and processes. Section 4 explains how our central assumption that the graph $\mathcal{G}$ is a tree is necessary for the dimensionality reduction referred to above. Section 5 introduces the controlled diffusion and states the main result. Section 6 reduces the problem of estimating the state process to that of estimating a one dimensional process, namely $|\sum_i X_i|$. Section 7 develops the integral formula (6) for Y and Z and establishes moment estimates in case (i) of the main result. Section 8 studies the non-idling property for trees of diameter three and establishes moment estimates in case (ii) of the main result. Section 9 treats case (iii) and summarizes the estimates in all cases. Section 10 discusses optimality of joint work conservation. Finally, in the appendix the main result is proved based on the moment estimates.

Notation.

Vectors are considered as column vectors. For a vector x let $\|x\| = \sum|x_i|$. For two column vectors v, u, $v \cdot u$ denotes their scalar product. The symbols e_i denote the unit coordinate vectors and $e = (1, \ldots, 1)'$. The dimension of e may change from one expression to another, and for example $e \cdot a + e \cdot b = \sum_i a_i + \sum_j b_j$ even if a and b are of different dimension. Denote by $B(m, r)$ the open Euclidean ball of radius r about m. $C^{m,\varepsilon}(D)$ [resp., $C^m(D)$] denotes the class of functions on $D \subset \mathbb{R}^{\bar{\imath}}$ for which all derivatives up to order m are Hölder continuous uniformly on compact subsets of D [resp., continuous on D]. $C_{\rm pol}(\mathbb{R}^{\bar{\imath}})$ denotes the class of continuous functions f on $\mathbb{R}^{\bar{\imath}}$, satisfying a polynomial growth condition: there are constants c and r such that $|f(x)| \le c(1 + \|x\|^r)$, $x \in \mathbb{R}^{\bar{\imath}}$. Let $C^{m,\varepsilon}_{\rm pol} = C_{\rm pol} \cap C^{m,\varepsilon}$, $C^m_{\rm pol} = C_{\rm pol} \cap C^m$, and let $C^m_{{\rm pol},+}$ be the class of nonnegative functions in $C^m_{\rm pol}$. Let $\mathbb{R}_+ = [0, \infty)$. If X is a process or a function on $\mathbb{R}_+$, $\|X\|^*_t = \sup_{0 \le s \le t}\|X(s)\|$, and if X takes real values, $|X|^*_t = \sup_{0 \le s \le t}|X(s)|$. $X(t)$ and X_t are used interchangeably. For a locally integrable function $f : \mathbb{R}_+ \to \mathbb{R}$ denote $If = \int_0^{\cdot} f(s)ds$. In case that f is vector- or matrix-valued, If is understood elementwise. The symbol c denotes a deterministic positive constant whose value may change from line to line.

2 The parallel station model

Although the results of this paper do not address the queueing system but only the associated diffusion processes, the stochastic model of the queueing system will be described in full detail. The diffusion model will then be obtained from it as a formal limit.

Consider a queueing system with $\bar{\imath}$ customer classes and $\bar{\jmath}$ service stations (see Figure 1). At each service station there are many independent servers of the same type, in the sense that they are statistically identical. Each customer requires service only once and can be served indifferently by any servers at the same station, but possibly at different rates at different stations. Only some stations can offer service to each class. When referring to the physical location of customers we say that they are in the buffer, or in the queue if they are not being served, and we say that they are in a certain station if they are served by a server of the corresponding type. There is one buffer per each customer class and one station per each server type.

A complete probability space (Ω, F, P) is given, supporting all stochastic processes defined below. Expectation with respect to P is denoted by E. There are $\bar{\imath}$ customer classes and $\bar{\jmath}$ server types. Since the set of all classes and all types constitutes the vertex set of certain graphs introduced below, it is convenient to label the classes as $1, \ldots, \bar{\imath}$ and the types as $\bar{\imath}+1, \ldots, \bar{\imath}+\bar{\jmath}$:

$$\mathcal{I} = \{1, \ldots, \bar{\imath}\}, \quad \mathcal{J} = \{\bar{\imath}+1, \ldots, \bar{\imath}+\bar{\jmath}\}.$$

The buffers [resp., stations] are labeled as the corresponding classes [resp., types]. For $j \in \mathcal{J}$ let N_j^n be the number of servers at station j. Thus the total number of servers is $e \cdot N^n$.

Let $X_i^n(t)$ denote the total number of class-i customers in the system at time t. Let $Y_i^n(t)$ denote the number of class-i customers in the queue at time t. Let $Z_j^n(t)$ denote the number of idle servers in station j at time t. And let $\Psi_{ij}^n(t)$ denote the number of class-i customers in station j at time t. We have the vector valued processes $X^n = (X_i^n)_{i\in\mathcal{I}}$, $Y^n = (Y_i^n)_{i\in\mathcal{I}}$, $Z^n = (Z_j^n)_{j\in\mathcal{J}}$, and the matrix valued process $\Psi^n = (\Psi_{ij}^n)_{i\in\mathcal{I}, i\in\mathcal{J}}$, $\bar{\imath} \times \bar{\jmath}$. Straightforward relations are expressed by the following equations:

$$Y_i^n + \sum_{j\in\mathcal{J}} \Psi_{ij}^n = X_i^n, \quad i \in \mathcal{I}, \tag{7}$$

$$Z_j^n + \sum_{i\in\mathcal{I}} \Psi_{ij}^n = N_j^n, \quad j \in \mathcal{J}, \tag{8}$$

$$Y_i^n(t), Z_j^n(t) \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}, t \geq 0. \tag{9}$$

The further constraints that $X_i^n \geq 0$ and $\Psi_{ij}^n \geq 0$ will turn out to be insignificant due to recentering and rescaling applied in Section 3.

To define arrival processes, let, for each $i \in \mathcal{I}$, $\{\check{U}_i(k), k \in \mathbb{N}\}$ be a sequence of strictly positive i.i.d. random variables with mean $E\check{U}_i(1) = 1$ and squared coefficient of variation

$$\mathrm{Var}(\check{U}_i(1))/(E\check{U}_i(1))^2 = C_{U,i}^2 \in [0, \infty).$$

Assume also that the sequences are independent. Let

$$U_i^n(k) = \frac{1}{\lambda_i^n} \check{U}_i(k),$$

where $\lambda_i^n > 0$. With $\sum_1^0 = 0$, define

$$A_i^n(t) = \sup\{l \geq 0 : \sum_{k=1}^{l} U_i^n(k) \leq t\}, \quad t \geq 0.$$

The renewal processes A_i^n are used to model arrivals: The number of arrivals up to time t is equal to $A_i^n(t)$. Note that the first class-i customer arrives at $U_i^n(1)$, and the time between the $(k-1)$st and kth arrival of class-i customers is $U_i^n(k)$.

The service times are independent exponential random variables, (more general renewal model for the service processes leads to a far more complicated diffusion limit) with mean depending on the customer class and server type. More precisely, let S_{ij}^n, $i \in \mathcal{I}, j \in \mathcal{J}$ be Poisson processes with rate $\mu_{ij}^n \in [0, \infty)$ (where a zero rate Poisson process is the zero process). These processes are assumed to be mutually independent, and independent of the arrival processes A_i^n. Let $T_{ij}^n(t)$ denote the time up to t devoted to a class-i customer by a server, summed over all type-j servers. Then

$$T_{ij}^n(t) = \int_0^t \Psi_{ij}^n(s)ds, \quad i \in \mathcal{I}, j \in \mathcal{J}, t \geq 0.$$

The number of service completions of class-i customers by all type-j servers by time t is $S_{ij}^n(T_{ij}^n(t))$.

The four principal quantities in the system now satisfy the following relation, with $X_i^{0,n} := X_i^n(0)$:

$$X_i^n(t) = X_i^{0,n} + A_i^n(t) - \sum_j S_{ij}^n\left(\int_0^t \Psi_{ij}^n(s)ds\right), \quad i \in \mathcal{I}, t \geq 0. \tag{10}$$

The description of the system dynamics is not complete unless full information is given on the process Ψ^n. In a stochastic control framework, Ψ^n can be regarded as a control process (as commented below, this makes sense for preemptive policies only) and is therefore determined as the solution to a control problem. Considering a cost of the form

$$C^n = E\int_0^\infty e^{-\gamma t} L(X^n(s), \Psi^n(s))ds,$$

one attempts selecting Ψ^n as the process that minimizes the cost among an appropriate family of processes. Since characterizing optimal or nearly optimal controls for such a problem can be very complicated, we adopt an asymptotic approach where n-dependent scaling is applied so as to obtain a diffusion control problem in the limit.

As mentioned in the introduction, there is a major difference between the models with and without preemption in setting up the queueing control problem. For example, in the special case where the arrivals are Poisson, the queueing control problems become Markov decision problems, and while the state process for the model with preemption is X^n, the state space for the model without preemption must include also Ψ^n. Therefore considering Ψ^n as a control is appropriate only for the model with preemption. However, under the assumptions of this paper queueing models with and without preemption are both asymptotically described by the same diffusion model, as argued in Section 4.

3 Fluid scaling and diffusion scaling

A sequence of systems as described above is considered, of which the symbol n is the index. As first proposed by Halfin and Whitt [10], the parameters scale as follows: The arrival rates scale up by n; the service rates of the individual servers converge; the total number of servers scales up by n. We also assume that the proportion of number of servers in each station converges to a positive number.

Scaling of parameters: There are constants $\lambda_i \in (0,\infty)$, $\mu_{ij} \in [0,\infty)$ and $\nu_j \in (0,\infty)$, $i \in \mathcal{I}$, $j \in \mathcal{J}$ such that

$$n^{-1}\lambda_i^n \to \lambda_i, \quad \mu_{ij}^n \to \mu_{ij} \tag{11}$$

and

$$n^{-1}N_j^n \to \nu_j.$$

It may be convenient to let n be the total number of servers $e \cdot N^n$ in the nth system, in which case ν_j sum to one, but we do not insist on this convention in this paper. while μ_{ij} stands for the limit service rate for an individual type-j server working on class i, we need a notation also for the total (normalized) service rate of class-i customers in station j:

$$\bar{\mu}_{ij} = \nu_j \mu_{ij}.$$

As will be seen, the rates $\bar{\mu}_{ij}$ are significant in the fluid model, while μ_{ij} are meaningful in the diffusion model.

In [11], [12] a static optimization problem is defined and studied for parallel server systems in the conventional heavy traffic setting. The same problem plays a key role in the current setting as well. It introduces a central assumption on the limit parameters λ_i, $\bar{\mu}_{ij}$ indicating that the sequence of systems is asymptotically critically loaded (see [18] for a different formulation of this condition).

Linear program: Minimize ρ subject to

$$\sum_{j\in\mathcal{J}} \bar{\mu}_{ij}\xi_{ij} = \lambda_i, \quad i \in \mathcal{I}, \tag{12}$$

$$\sum_{i\in\mathcal{I}} \xi_{ij} \le \rho, \quad j \in \mathcal{J}, \tag{13}$$

$$\xi_{ij} \ge 0, \quad i \in \mathcal{I}, j \in \mathcal{J}. \tag{14}$$

Heavy traffic condition: There exists a unique optimal solution (ξ^*, ρ^*) to the linear program. Moreover, $\rho^* = 1$, and $\sum_{i\in\mathcal{I}} \xi_{ij}^* = 1$ for all $j \in \mathcal{J}$.

In the rest of this paper, ξ^*_{ij} denotes the quantities from the above condition. Let also $x^* = (x^*_i, i \in \mathcal{I})$, $x^*_i = \sum_j \xi^*_{ij}\nu_j$.

In a fluid model where class-i customers are served at station j at rate $\bar{\mu}_{ij}$, the quantities $\sum_j \xi^*_{ij}\bar{\mu}_{ij}$ represent the actual service rate of class-i customers when each station j allocates a fraction ξ^*_{ij} of its servers to class-i jobs. The quantities x^*_i represent the total mass of class-i customers in the system. The heavy traffic condition indeed asserts that the system is critically loaded in the following sense. On one hand, equation (12) states that customers arriving at rates λ_i can be processed by the system. On the other hand, under this condition it is impossible to have a $\tilde{\xi}$ satisfying (13) (with $\rho \le 1$) and (14), and a $\tilde{\lambda} \neq \lambda$, $\tilde{\lambda}_i \ge \lambda_i$, $i \in \mathcal{I}$, such that (12) holds for $\bar{\mu}, \tilde{\lambda}$ and $\tilde{\xi}$ (since this would violate the uniqueness of ξ^*). Thus, the system cannot process arrivals at rates greater than λ_i.

Finer information on the scaling is as follows.

Scaling of parameters, continued: Letting $\lambda_i, \nu_i \in (0,\infty), \mu_{ij}, \in [0,\infty)$ be as above, there are constants $\hat{\lambda}_i, \hat{\mu}_{ij} \in \mathbb{R}$, $i \in \mathcal{I}, j \in \mathcal{J}$, such that

$$n^{1/2}(n^{-1}\lambda^n_i - \lambda_i) \to \hat{\lambda}_i, \qquad n^{1/2}(\mu^n_{ij} - \mu_{ij}) \to \hat{\mu}_{ij}, \tag{15}$$

$$n^{1/2}(n^{-1}N^n_j - \nu_j) \to 0. \tag{16}$$

Scaling of initial conditions: There are constants x_i, y_i, z_j, ψ_{ij} such that the deterministic initial conditions satisfy

$$\begin{aligned} \hat{X}^{0,n}_i &:= n^{-1/2}(X^{0,n}_i - nx^*_i) \to x_i, \qquad (17)\\ \hat{Y}^{0,n}_i &:= n^{-1/2}Y^{0,n}_i \to y_i,\\ \hat{Z}^{0,n}_j &:= n^{-1/2}Z^{0,n}_j \to z_j,\\ \hat{\Psi}^{0,n}_{ij} &:= n^{-1/2}(\Psi^{0,n}_{ij} - \xi^*_{ij}\nu_j n) \to \psi_{ij}, \end{aligned}$$

where, due to (7), (8) and (9) one must assume $y_i + \sum_j \psi_{ij} = x_i$, $z_j + \sum_i \psi_{ij} = 0$, $y_i \ge 0$, and $z_j \ge 0$, $i \in \mathcal{I}$, $j \in \mathcal{J}$.

We consider two levels of scaling for the processes involved. The processes rescaled at the fluid level are defined as

$$\begin{aligned} \bar{X}^n_i(t) &= n^{-1}X^n_i(t),\\ \bar{Y}^n_i(t) &= n^{-1}Y^n_i(t),\\ \bar{Z}^n_j(t) &= n^{-1}Z^n_j(t),\\ \bar{\Psi}^n_{ij}(t) &= n^{-1}\Psi^n_i(t). \end{aligned}$$

The processes recentered (about their formal fluid limit) and rescaled at the diffusion level are as follows

$$\hat{A}_i^n(t) = n^{-1/2}(A_i^n(t) - \lambda_i^n t), \qquad \hat{S}_{ij}^n(t) = n^{-1/2}(S_{ij}^n(nt) - n\mu_{ij}^n t),$$

$$\hat{X}_i^n(t) = n^{-1/2}(X_i^n(t) - nx_i^*), \tag{18}$$

$$\hat{Y}_i^n(t) = n^{-1/2}Y_i^n(t), \qquad \hat{Z}_j^n(t) = n^{-1/2}Z_j^n(t), \tag{19}$$

$$\hat{\Psi}_{ij}^n(t) = n^{-1/2}(\Psi_{ij}^n - \xi_{ij}^*\nu_j n).$$

The relations (7), (8) and (9) take the new form

$$\hat{Y}_i^n + \sum_j \hat{\Psi}_{ij}^n = \hat{X}_i^n, \quad i \in \mathcal{I}, \tag{20}$$

$$\hat{Z}_j^n + \sum_i \hat{\Psi}_{ij}^n = 0, \quad j \in \mathcal{J}, \tag{21}$$

$$\hat{Y}_i^n, \hat{Z}_j^n \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}. \tag{22}$$

Using the definitions above of the rescaled processes and the relation $\lambda = \mu(\xi^*)$, one finds that (10) takes the form

$$\hat{X}_i^n(t) = \hat{X}_i^{0,n} + r_i\hat{W}_i^n(t) + \ell_i^n t - \sum_j \mu_{ij}^n \int_0^t \hat{\Psi}_{ij}^n(s)ds \tag{23}$$

where

$$r_i\hat{W}_i^n(t) = \hat{A}_i^n(t) - \sum_j \hat{S}_{ij}^n\Big(\int_0^t \bar{\Psi}_{ij}^n(s)ds\Big),$$

$$\ell_i^n = n^{1/2}(n^{-1}\lambda_i^n - \lambda_i) - \sum_i n^{1/2}(\mu_{ij}^n - \mu_{ij})\xi_{ij}^*\nu_j.$$

With (15) we have

$$\lim_n \ell_i^n = \lambda_i - \sum_j \mu_{ij}\xi_{ij}^*\nu_j. \tag{24}$$

One is free to choose the values of r_i, and it is convenient to choose them so that with the formal substitution $\bar{\Psi}_{i,j}^n = \xi_{i,j}^*\nu_j$ one has

$$\lim_n E[(\hat{W}_i^n(1))^2] = 1. \tag{25}$$

Namely,

$$r_i = (\lambda_i C_{U,i}^2 + \lambda_i)^{1/2}.$$

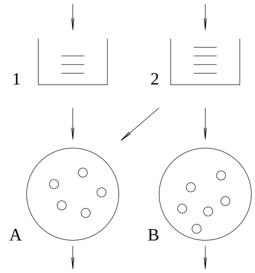

Figure 2: An example with three activities

4 Buffer-station tree

We use terminology from [12] to differentiate buffer-station pairs $(i,j) \in \mathcal{I} \times \mathcal{J}$ according to whether $\bar{\mu}_{ij} > 0$ (equivalently, $\mu_{ij} > 0$) and whether $\xi^*_{ij} > 0$. We somewhat deviate from [12] in notation but borrow their concepts. A pair (i,j) is said to be an *activity* if $\bar{\mu}_{ij} > 0$. This notion indicates that class-i arrivals can be processed in the fluid model by station j. If (i,j) is an activity, it is said to be *basic* if $\xi^*_{ij} > 0$, and *nonbasic* otherwise. This indicates that under the allocation matrix ξ^*, class-i arrivals are actually processed by station j in the fluid model. To emphasize that activities depend only on the fluid model, we sometimes refer to [basic] activities as *[basic] activities for the fluid model.* A similar notion for the prelimit model will also be useful. A pair (i,j) is said to be an *activity for the queueing network* if $\mu^n_{ij} > 0$ for some n.

The association of buffers with stations according to activities and basic activities can be encoded in graphs (see [20]). Let $\mathcal{G}_a$ be a graph with vertex set $\{1, 2, \ldots, \bar{\imath} + \bar{\jmath}\} = I \cup J$: A node is associated with each buffer and a node is associated with each station. An undirected edge is associated with each activity, namely every edge in the graph is between a buffer node (some $i \in \mathcal{I}$) and a station node (some $j \in \mathcal{J}$), and there is an edge between i and j if and only if (i,j) is an *activity.* Let $\mathcal{G}_b$ be a graph having the same nodes as $\mathcal{G}_a$ and having an edge between a node $i \in \mathcal{I}$ and a node $j \in \mathcal{J}$ if and only if (i,j) is a *basic activity.* Finally, we encode the activities for the queueing network in a graph $\mathcal{G}$ having the same nodes as $\mathcal{G}_a$ and $\mathcal{G}_b$, and an edge between a node $i \in \mathcal{I}$ and a node $j \in \mathcal{J}$ if and only if (i,j) is an *activity for the queueing network.*

As stated in the introduction, we are interested in generalizing the results of [4] where the asymptotic analysis of a queueing control problem for a system without preemption can be reduced to that of a system with preemption. As we show in an example, if there are nonbasic activities in the fluid model, the limit behavior is expected to be different for both. The example is explained heuristically and cannot replace a rigorous proof, but is nevertheless convincing (the observation and example are due to Mandelbaum and Reiman [16]).

Example 1 Consider a system with two customer classes and two server types as depicted in Figure 2. Servers of type A can serve both classes and servers of type B can only serve class-2 customers (in the examples we refer to stations as $A, B,...$ rather than as numbers in $\mathcal{J}$). Consider first a case where all activities are basic. For example, let $\nu_1 = \nu_2 = 1$, $\mu_{1A} = \mu_{2A} = \mu_{2B} = 1$, and $\lambda_1 = 1/2$, $\lambda_2 = 3/2$. Then there is a unique optimal solution to the linear program: $(\xi^*_{1A}, \xi^*_{2A}, \xi^*_{2B}) = (1/2, 1/2, 1)$. In a fluid model this means that half of the servers at station A serve class 1 and all other servers serve class 2. In the queueing model with large n, under any policy that keeps the deviations from the fluid model at the desired level $O(\sqrt{n})$, the number of class-1 and the number of class-2 customers in station A are $\Psi^n_{1A}, \Psi^n_{2A} = n/2 + O(\sqrt{n})$, while the number of class-2 customers in station B is $\Psi^n_{2B} = n + O(\sqrt{n})$. In a system with preemption the policy dynamically selects the values of these processes Ψ^n. In the diffusion model the control corresponds to the centered processes, namely to the $O(\sqrt{n})$ deviations referred to above. For a system without preemption to be governed by the same diffusion model, it should have the property that Ψ^n can track closely the corresponding diffusion control process Ψ. This means that it should be possible for the population in any activity to change by $O(\sqrt{n})$ in $o(1)$ units of time. It is proved in [4] that this is indeed possible in the case of a single station. This is also expected to be the case here, as we explain for activity $(2, A)$. Recall that the population in this activity is about $n/2$. If routing to this activity is stopped then the population in this activity will decrease due to service completions by $O(\sqrt{n})$ in $o(1)$ units of time. On the other hand, if routing to this activity is maintained and routing to the other activities is stopped, then the population in this activity will build up by $O(\sqrt{n})$ in $o(1)$ units of time. It is therefore possible to decrease or increase the population in this activity very rapidly, and tracking a desired control process is feasible. Consider now a case where μ and ν are as before, but $\lambda_1 = \lambda_2 = 1$. The unique solution to the linear program is now $(\xi^*_{1A}, \xi^*_{2A}, \xi^*_{2B}) = (1, 0, 1)$, and therefore activity $(2, A)$ is nonbasic. The queueing model will have only $O(\sqrt{n})$ customers in this activity. As a result it is impossible to reduce the population in this activity in $o(1)$ units of time. This stands in contrast with the case where preemption is allowed and customers can be moved in zero time between the buffer and any station that offers them service.

In view of the above example, since we are interested in problems in which nonpreemptive models possess the relatively simple diffusion limit that preemptive models have, we shall assume that *all activities are basic.*

Let now (i, j) be an activity for the queueing network but not an activity (namely $\mu_{ij} = 0$). By (15), $\mu^n_{ij} = \mu_{ij} + O(n^{-1/2}) = O(n^{-1/2})$. If the population in (i, j) ever gets to $O(\sqrt{n})$ (or more), then again one could not control it so as to change by $O(\sqrt{n})$ in $o(1)$ units of time, for reasons similar to those explained in the above example.

We therefore impose the assumption that *all activities for the queueing network are activities for the fluid model.*

As a result, the graphs $\mathcal{G}$, $\mathcal{G}_a$ and $\mathcal{G}_b$ must be the same. Finally, since in cases where the graph $\mathcal{G}$ is not connected the problem can be decomposed into subproblems on the components, we shall assume that $\mathcal{G}$ is connected. Here is a summary of the above assumptions.

One has $\mathcal{G} = \mathcal{G}_a = \mathcal{G}_b$, and each of these graphs is connected.

Lemma 1 *Let the heavy traffic condition hold. If $\mathcal{G}_b$ is connected then it is a tree.*

Proof: This result is due to [12]. See [20], Theorem 5.3 and Corollary 5.4. □

The condition that the graph $\mathcal{G}_b$ is connected is equivalent to the *complete resource pooling condition* of [12]. In view of the lemma, our assumption above is equivalent to the following.

Assumption 1 *The heavy traffic condition holds. The graph $\mathcal{G}$ is a tree. Moreover, every activity for the queueing network is a basic activity for the fluid model, i.e., $\mathcal{G}_a = \mathcal{G}_b = \mathcal{G}$.*

In what follows, $\mathcal{T}$ denotes the tree $\mathcal{T} := \mathcal{G} = \mathcal{G}_a = \mathcal{G}_b$. Write $i \sim j$ if (i, j) is an edge of $\mathcal{T}$ (equivalently, an activity) and $i \not\sim j$ otherwise. The edge set for the tree $\mathcal{T}$ is denoted by

$$\mathcal{E} = \{(i, j) \in \mathcal{I} \times \mathcal{J} : i \sim j\}.$$

Note that under Assumption 1, $\mu^n_{ij} = 0$, $n \geq 1$, whenever $i \not\sim j$. As a result, for every such (i, j) one has $\mu_{ij} = \hat{\mu}_{ij} = 0$ and $\Psi^n_{ij}, \bar{\Psi}^n_{ij}, \hat{\Psi}^n_{ij} = 0$.

<u>*Joint work conservation*</u>

A policy is said to be *work conserving* if it does not allow for a server to idle while a customer that it can serve is in the queue. In the current context one can consider a stronger condition for preemptive policies. Recall that if preemption is allowed, customers of each class can be moved between the queue and the various stations that offer service to them. Let t be given, and recall that the components of $X^n(t)$ denote the number of customers of each class present in the system at time t. Since t is fixed and its value will be immaterial in the following discussion, we omit it from the notation and write e.g., X^n for $X^n(t)$, Ψ^n for $\Psi^n(t)$. In a preemptive policy, the set of controls Ψ^n that can be applied at time t correspond to different rearrangements of the customers X^n in the stations and buffers. Let $\mathcal{X}^n$ denote the set of all possible values of X^n for which there is a rearrangement of customers with the property:

either there are no customers in the queue, or no server in the system idles. (26)

We shall say that a preemptive policy is *jointly work conserving* if it is work conserving and, in addition, for every s, if $X^n(s) \in \mathcal{X}^n$ then customers are arranged according to (26).

In the heavy traffic regime considered here, it is anticipated that it is nearly always the case that a rearrangement of customers according to (26) is possible. To explain this point, we show that there exists a constant $c_0 > 0$ such that the condition

$$\|X^n - nx^*\| \leq c_0 n \tag{27}$$

is sufficient for the existence of a rearrangement satisfying (26) for a given X^n, and provided that n is large enough. If one shows that the processes $\hat{X}^n(t) = n^{-1/2}(X^n(t) - nx^*)$ (see (18)) are tight (we have no such attempt in this paper; see [2]), then indeed (27) typically holds for large n. To show that under (27) it is always possible to meet (26), the following result will be useful.

Proposition 1 *Let Assumption 1 hold. Then given $\alpha_i, \beta_j \in \mathbb{R}$, $i \in \mathcal{I}, j \in \mathcal{J}$ satisfying $\sum \alpha_i = \sum \beta_j$ there exists a unique solution ψ_{ij} to the set of equations*

$$\sum_j \psi_{ij} = \alpha_i, \quad i \in \mathcal{I}, \tag{28}$$

$$\sum_i \psi_{ij} = \beta_j, \quad j \in \mathcal{J}, \tag{29}$$

where $\psi_{i,j} = 0$ for $i \not\sim j$.

Proof: See the appendix.

In view of Proposition 1 there is a map, denoted throughout by $G : \{(\alpha, \beta) \in \mathbb{R}^{\bar{\imath}+\bar{\jmath}} : \sum \alpha_i = \sum \beta_j\} \to \mathbb{R}^{\bar{\imath}+\bar{\jmath}-1}$, $\psi = G(\alpha, \beta)$. Clearly the map is linear. By (16) one has that $N^n = n\nu + o(n^{1/2})$. To simplify the explanation, we shall ignore in the discussion below the term $o(n^{1/2})$ and assume that $N^n = n\nu$. One then checks by substitution that, with $\psi = G(nx^*, N^n)$,

$$\psi_{ij} = n\xi^*_{ij}\nu_j, \quad i \sim j. \tag{30}$$

Given X^n satisfying (27) (with c_0 to be determined below) we now show that one can find Ψ^n, Y^n and Z^n that meet $\Psi^n_{ij} \geq 0$, (7), (8) and (9), as well as the condition that a jointly work conserving policy imposes, namely that $e \cdot Y^n$ and $e \cdot Z^n$ are not both positive (the symbol e is used here in different dimensions). First suppose X^n is such that $e \cdot X^n > e \cdot N^n$ (more customers than servers in the system at the given time). Set $Z^n = 0$ and let Y^n be any nonnegative vector satisfying $e \cdot Y^n = e \cdot X^n - e \cdot N^n$. Note that

$$\|Y^n\| = e \cdot Y^n = e \cdot X^n - e \cdot N^n = e \cdot (X^n - nx^*) \leq \|X^n - nx^*\|. \tag{31}$$

Then by Proposition 1, the relations (7) and (8) meet upon setting $\Psi^n = G(X^n - Y^n, N^n)$. Also, for $i \sim j$, by linearity of the map G and by (30), (31) and (27),

$$\begin{aligned}
\Psi^n_{ij} &= G(X^n - Y^n, N^n)_{ij} \\
&= G(nx^*, N^n)_{ij} + G(X^n - nx^* - Y^n, 0)_{ij} \\
&\geq n\xi^*_{ij}\nu_j - c_1(\|X^n - nx^*\| + \|Y^n\|) \\
&\geq n\xi^*_{ij}\nu_j - 2c_1(\|X^n - nx^*\|) \\
&\geq n\xi^*_{ij}\nu_j - 2c_1c_0n \\
&\geq 0,
\end{aligned}$$

where the last inequality follows from positivity of $\xi^*_{ij}\nu_j$ and an appropriate choice of the constant $c_0 > 0$. The case when $e \cdot X^n(t_0) \le e \cdot N^n$ can be treated analogously.

Based on the heuristic that $\hat{X}^n$ are expected to be tight stochastic processes, we have argued above that under preemptive policies it is nearly always possible to rearrange customers in the system so as to meet (26), and thus that joint work conservation is essentially equivalent to having (26) always hold. Under nonpreemptive policies, the definition of joint work conservation does not even make sense, since it is impossible to rearrange customers in the system. However, since we show that preemptive and nonpreemptive policies can be treated asymptotically under the same diffusion model [2], one expects that, in an appropriate sense, one can nearly achieve joint work conservation under nonpreemptive policies for large values of n. As far as the diffusion model introduced below is concerned, there will be no difficulty associating with it a property analogous to joint work conservation (cf. (36) below).

5 The controlled diffusion and statement of main results

We take formal limits as $n \to \infty$ in (19)–(23). The process $\hat{W}^n$ converges to a standard Brownian motion (cf. [4], proof of Lemma 4). Denoting the formal weak limit of $(\hat{X}^n, \hat{Y}^n, \hat{Z}^n, \hat{W}^n, \hat{\Psi}^n)$ by (X, Y, Z, W, Ψ), we get by (17), (24), (11) and the relation $\nu_j\mu_{ij} = \bar{\mu}_{ij}$:

$$X_i(t) = x_i + \tilde{W}_i(t) - \sum_j \mu_{ij} \int_0^t \Psi_{ij}(s)ds, \quad i \in \mathcal{I} \tag{32}$$

where $\tilde{W}_i(t) = r_i W_i(t) + \ell_i t$, W is a standard Brownian motion, and

$$\sum_j \Psi_{ij} = X_i - Y_i, \quad i \in \mathcal{I}, \tag{33}$$

$$\sum_i \Psi_{ij} = -Z_j, \quad j \in \mathcal{J}, \tag{34}$$

$$Y_i \ge 0, Z_j \ge 0, \quad i \in \mathcal{I}, j \in \mathcal{J}. \tag{35}$$

Also, the condition that joint work conservation is maintained for every n will dictate that a similar property holds for the diffusion model, namely

$$e \cdot Z(t) > 0 \quad \text{implies} \quad e \cdot Y(t) = 0, \quad t \ge 0, \tag{36}$$

or equivalently,

$$e \cdot Y(t) > 0 \quad \text{implies} \quad e \cdot Z(t) = 0, \quad t \ge 0.$$

One can formulate a diffusion control problem where (32)–(36) define the dynamics and constraints, with X the state process and Ψ the control. In such a formulation (33) and (34) define Y and Z

and (35) and (36) give the constraints. However, a formulation in the spirit of [13] (see also [4]) is possible, where controls lie in a compact space and in lower dimension. To this end, note that by (33) and (34),

$$e \cdot X = e \cdot Y - e \cdot Z,$$

and thus by (35) and (36),

$$e \cdot Y = (e \cdot X)^+, \quad e \cdot Z = (e \cdot X)^-.$$

Hence Y can be represented as

$$Y_i(t) = (e \cdot X(t))^+ u_i(t), \tag{37}$$

where

$$u_i(t) \geq 0, \quad e \cdot u(t) = 1.$$

Similarly,

$$Z_j(t) = (e \cdot X(t))^- v_j(t), \tag{38}$$

$$v_j(t) \geq 0, \quad e \cdot v(t) = 1.$$

Thus consider $U := (u, v)$ as a control process taking values in

$$\mathbb{U} := \{(u, v) \in \mathbb{R}^{\bar{\imath}+\bar{\jmath}} : u_i, v_j \geq 0, \sum u_i = \sum v_j = 1\}.$$

By Proposition 1 Ψ is given as $\Psi = G(X - Y, -Z)$ hence now

$$\Psi = \hat{G}(X, U) := G(X - (e \cdot X)^+ u, -(e \cdot X)^- v).$$

Letting $\ell = (\ell_1, \ldots, \ell_{\bar{\imath}})'$, $r = \mathrm{diag}(r_i)$, and

$$b = -\mu \circ \hat{G} + \ell, \tag{39}$$

we see that (32) can be written as

$$X(t) = x + rW(t) + \int_0^t b(X(s), U(s))ds.$$

Admissible systems and controlled processes

Definition 1 *i. A* complete filtered probability space *is a complete probability space* (Ω, F, P)*, equipped with a right continuous filtration* $(F_t)_{t \in \mathbb{R}_+}$ *of sub-*σ*-fields of* F*, such that* F_0 *contains all* P*-null sets of* F*. A stochastic process* X *on* (Ω, F, P) *is* adapted *to a filtration* (F_t)*, if for each* $t \geq 0$*,* X_t *is* F_t*-measurable. A stochastic process* U *is* F_t*-progressively measurable if for each* $t \geq 0$*, the mapping* $(s, \omega) \mapsto U_s(\omega)$ *is* $\mathcal{B}([0, t]) \otimes F_t$*-measurable, where* $\mathcal{B}([0, t])$ *is the Borel sigma-field on* $[0, t]$*.*

ii. We call

$$\pi = (\Omega, F, (F_t), P, U, W)$$

an admissible system *if*

1. $(\Omega, F, (F_t), P)$ *is a complete filtered probability space,*

2. U *is a* $\mathbb{U}$*-valued,* F*-measurable,* (F_t)*-progressively measurable process, and* W *is a standard* $\bar{\imath}$*-dimensional* (F_t)*-Brownian motion.*

The process U *is said to be a* control *associated with* π*.*

iii. We say that X *is a* controlled process *associated with initial data* $x \in \mathbb{R}^{\bar{\imath}}$ *and an admissible system* $\pi = (\Omega, F, (F_t), P, U, W)$*, if*

1. X *is a continuous process on* (Ω, F, P)*,* F*-measurable,* (F_t)*-adapted,*

2. $\int_0^t |b(X(s), U(s))| ds < \infty$ *for every* $t \geq 0$*,* P*-a.s.,*

3.
$$X(t) = x + rW(t) + \int_0^t b(X(s), U(s)) ds, \quad 0 \leq t < \infty, \tag{40}$$

   holds P*-a.s.*

The proposition below shows that there is a unique controlled process X associated with any x and π. With an abuse of notation we sometimes denote the dependence on x and π by writing P_x^π in place of P and E_x^π in place of E. Denote by Π the class of all admissible systems.

Proposition 2 *Let initial data* $x \in \mathbb{R}^{\bar{\imath}}$ *and an admissible system* $\pi \in \Pi$ *be given. Then there exists a controlled process* X *associated with* x *and* π*. Moreover, if* X *and* $\bar{X}$ *are controlled processes associated with* x *and* π*, then* $X(t) = \bar{X}(t)$*,* $t \geq 0$*,* P*-a.s.*

Proof: Note that $(x, U) \mapsto b(x, U)$ is continuous and $x \mapsto b(x, U)$ is Lipschitz uniformly in U. Consider b_m, a function that agrees with b on the ball $B(0, m)$, uniformly Lipschitz and bounded. Then strong existence and uniqueness for

$$X_m(t) = x + rW(t) + \int_0^t b_m(X_m(s), U(s)) ds, \quad 0 \leq t < \infty$$

holds by Theorem I.1.1 of [5]. Since $\|X_m(t)\| \leq \|x\| + c\|W(t)\| + c\int_0^t \|X_m(s)\| ds$, one has $\|X_m(t)\| \leq (\|x\| + c\|W\|_t^*)(1 + e^{ct})$ by Gronwall's lemma. Thus letting $\tau_m = \inf\{t : \|X_m(t)\| \geq m\}$, one has $\tau_m \to \infty$ a.s. Therefore $X(t) = \lim_m X_m(t)$ for all t defines a process that solves the equation (a strong solution). If X and $\bar{X}$ are both strong solutions, then for every m they both agree with X_m on $[0, \tau_m]$. Therefore they agree on $[0, \infty)$ a.s. □

Control problem and HJB equation

Let a constant $\gamma > 0$ and a function L be given and consider the cost

$$C(x,\pi) = E_x^\pi \int_0^\infty e^{-\gamma t} L(X(t),U(t))dt, \quad x \in \mathbb{R}^{\bar{\imath}}, \pi \in \Pi.$$

Our assumption on L is as follows.

Assumption 2 *i.* $L(x,U) \geq 0$, $(x,U) \in \mathbb{R}^{\bar{\imath}} \times \mathbb{U}$.

ii. The mapping $(x,U) \mapsto L(x,U)$ *is continuous.*

iii. There is $\varrho \in (0,1)$ *such that for any compact* $A \subset \mathbb{R}^{\bar{\imath}}$,

$$|L(x,U) - L(y,U)| \leq c\|x-y\|^\varrho$$

holds for $U \in \mathbb{U}$ *and* $x, y \in A$, *where* c *depends only on* A.

iv. There are constants $c > 0$ *and* $m_L \geq 1$ *such that* $L(x,U) \leq c(1 + \|x\|^{m_L})$, $U \in \mathbb{U}$, $x \in \mathbb{R}^{\bar{\imath}}$.

Define the value function as

$$V(x) = \inf_{\pi \in \Pi} C(x,\pi).$$

The HJB equation for the problem is (cf. [7])

$$\mathcal{L}f + H(x, Df) - \gamma f = 0, \tag{41}$$

where $\mathcal{L} = (1/2)\sum_i r_i^2 \partial^2/\partial x_i^2$, and

$$H(x,p) = \inf_{U \in \mathbb{U}} [b(x,U) \cdot p + L(x,U)].$$

The equation is considered on $\mathbb{R}^{\bar{\imath}}$ with the growth condition

$$\exists C, m, \quad |f(x)| \leq C(1 + \|x\|^m), \quad x \in \mathbb{R}^{\bar{\imath}}. \tag{42}$$

We say that f is a solution to (41) if it is of class C^2, and the equation is satisfied everywhere in $\mathbb{R}^{\bar{\imath}}$.

Definition 2 *Let* $x \in \mathbb{R}^{\bar{\imath}}$ *be given. We say that a measurable function* $h : \mathbb{R}^{\bar{\imath}} \to \mathbb{U}$ *is a* Markov control policy *if there is an admissible system* π *and a controlled process* X *corresponding to* x *and* π, *such that* $U_s = h(X_s)$, $s \geq 0$, P*-a.s. We say that an admissible system* π *is* optimal *for* x, *if* $V(x) = C(x,\pi)$. *We say that a Markov control policy is optimal for* x *if the corresponding admissible system is.*

Our main result is the following.

Theorem 1 *Let Assumptions 1 and 2 hold. In addition, let one of the following conditions hold.*

i. For $(i,j) \in \mathcal{E}$, μ_{ij} depends only on i; or for $(i,j) \in \mathcal{E}$, μ_{ij} depends only on j.

ii. The tree $\mathcal{T}$ is of diameter 3 at most.

iii. There are $c, C > 0$ such that $L(x,U) \geq c\|x\|^{m_L}$ for all $\|x\| > C$ and all $U \in \mathbb{U}$ (where m_L is as in Assumption 2).

Then there exists a classical nonnegative solution $f \in C^{2,\varrho}_{\text{pol}}(\mathbb{R}^{\bar{\imath}})$ to (41) (42). In cases (i) and (ii) this solution is unique in $C^2_{\text{pol}}(\mathbb{R}^{\bar{\imath}})$ and in case (iii) it is unique in $C^2_{\text{pol},+}(\mathbb{R}^{\bar{\imath}})$. Moreover, the value V is equal to f. Furthermore, there exists a Markov control policy that is optimal for all $x \in \mathbb{R}^{\bar{\imath}}$.

Deterministic setting

Equations (32), (33) and (34) can be considered in a deterministic setting. Use w in place of $x + \tilde{W}$ (where as in (32) x is the initial condition for X), and $x = x(t)$ [resp., y, z, ψ] in place of X [resp., Y, Z, Ψ]:

$$x_i(t) = w_i(t) - \sum_j \mu_{ij} \int_0^t \psi_{ij}(s)ds, \quad i \in \mathcal{I}, \tag{43}$$

$$\sum_j \psi_{ij} = x_i - y_i, \quad i \in \mathcal{I}, \tag{44}$$

$$\sum_i \psi_{ij} = -z_j, \quad j \in \mathcal{J}, \tag{45}$$

where μ_{ij} and ψ_{ij} vanish whenever $i \not\sim j$ (a convention that holds throughout). The constraints (35), (36) now have the form

$$y_i, z_j \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}, \tag{46}$$

$$\sum_i y_i(t) > 0 \text{ implies } \sum_j z_j(t) = 0, \quad t \geq 0. \tag{47}$$

A key step in the proof will be to show that (43)–(47) imply a uniform estimate on the state, of the following form:

$$\|x\|_t^* \leq c(1+t)^m (1 + \|w\|_t^*), \tag{48}$$

where c, m do not depend on t, w, y, z, ψ. This estimate, that we establish in the first two cases of Theorem 1, depends on both parts of Assumption 1: the heavy traffic condition and the tree structure. Thus we end this section with two counterexamples to (48), one where the uniqueness assumption that is a part of the heavy traffic condition does not hold, and one where the heavy traffic condition holds, but a nonbasic activity is used.

Example 2 Consider a two-class two-type system where at station A there are twice as many servers as in station B, and where each server in station B works twice as fast as each server in station A: $\nu_A = 2, \nu_B = 1$, $\mu_{1A} = \mu_{2A} = 1$, $\mu_{1B} = \mu_{2B} = 2$. Then $\bar{\mu}_{1A} = \bar{\mu}_{2A} = \bar{\mu}_{1B} = \bar{\mu}_{2B} = 2$, and with $\lambda_1 = \lambda_2 = 2$ it is easily seen that there are many optimal solutions to the linear program (12)–(14). For example, both

$$\xi = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \xi = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

are solutions. To show that (48) does not hold, consider $w = 0$ and note that (43) has the form:

$$\dot{x}_1 = -\psi_{1A} - 2\psi_{1B}, \quad \dot{x}_2 = -\psi_{2A} - 2\psi_{2B}.$$

Let $\psi_{1A} = -\psi_{2A} = k$ and let $\psi_{1B} = -\psi_{2B} = -k(1 + e^{-2t})/2$. One checks that (43)–(47) are satisfied with $x_1 = -x_2 = k(1 - e^{-2t})/2$, $y = 0$, $z = 0$, for any $k > 0$. Since k is arbitrary, (48) cannot hold.

Example 3 Consider a two-class two-type system as follows. $\nu_A = \nu_B = 1$, $\lambda = (1.5, 1.5)'$ and $\mu_{1A} = 3, \mu_{1B} = 1, \mu_{2A} = 1$ and $\mu_{2B} = 1$. Then the heavy traffic condition holds with $\xi^*_{1A} = 0.5$, $\xi^*_{1B} = 0$, $\xi_{2A} = 0.5$ and $\xi^*_{2B} = 1$. Taking $w = 0$, one checks that $\psi_{1A}(t) = -k$, $\psi_{1B}(t) = k$, $\psi_{2A}(t) = k$, $\psi_{2B}(t) = -k$ with $x_1(t) = y_1(t) = 2kt$, $x_2(t) = y_2(t) = 0$ and $z_A = z_B = 0$, satisfy all of (43)–(47) for any $k > 0$. Thus (48) does not hold.

6 Estimating the state X in terms of $e \cdot X$

While the relation $\|Y(t)\| + \|Z(t)\| \le c\|X(t)\|$ is immediate from (37), (38), the following result shows that in a weaker sense Y, Z and $x + \tilde{W}$ dominate X (or in the deterministic notation, y, z and w dominate x). The result only uses the relations (43)–(45) and not the further constraints (46), (47).

Proposition 3 *Let equations* (43)–(45) *hold. Then there is a constant* c_0 *not depending on* ψ, w, x, y, z *or* t*, such that*

$$\|I\psi\|^*_t + \|x\|^*_t \le c_0(\|w\|^*_t + \|Iy\|^*_t + \|Iz\|^*_t), \quad t \ge 0,$$

where $\|I\psi\|^*_t = \sum_{ij} |I\psi_{ij}|^*_t$.

Note that if, in addition, (46) and (47) are assumed, then $\|Iy\|^*_t = I(e \cdot x)^+(t)$ and $\|Iz\|^*_t = I(e \cdot x)^-(t)$. As a result of Proposition 3, the state x is dominated by w and $e \cdot x$ in the sense

$$\|x\|^*_t \le c_0(\|w\|^*_t + I|e \cdot x|(t)).$$

Lemma 2 *Let w be a measurable, locally bounded function and assume $x = w - \mu\int_0^\cdot x(s)ds$. Then*

$$x(t) = w(t) - \mu \int_0^t w(s)e^{-\mu(t-s)}ds, \tag{49}$$

and in particular, if $\mu > 0$ then $|x(t)| \leq 2|w|_t^*$, $t \geq 0$.

Proof: : Uniqueness of solutions x follows from Gronwall's lemma [6], and (49) is verified by substitution. □

Proof of Proposition 3: A node in a tree is said to be a leaf if there is exactly one edge joining it. Recall that the tree $\mathcal{T}$ has $\kappa = \bar{\imath} + \bar{\jmath}$ nodes and $\kappa - 1$ edges. Let $\mathcal{T}_1, \mathcal{T}_2 \dots, \mathcal{T}_{\kappa-1} = \mathcal{T}$ be an increasing sequence of trees as follows. For $n = 1, \dots, k-2$, $\mathcal{T}_n$ is obtained from $\mathcal{T}_{n+1}$ by deleting a leaf and the edge joining it. Note that $\mathcal{T}_1$ has exactly two nodes, one in $\mathcal{I}$, one in $\mathcal{J}$. Let $\mathcal{V}_n$ denote the vertex set of $\mathcal{T}_n$. Let $v_{n+1} = \mathcal{V}_{n+1} \setminus \mathcal{V}_n$ denote the node in $\mathcal{V}_{n+1}$ that does not belong to $\mathcal{V}_n$.

Denote $\mathcal{I}_n = \mathcal{I} \cap \mathcal{V}_n$ and $\mathcal{J}_n = \mathcal{J} \cap \mathcal{V}_n$. We shall show that for $n = 1, \dots, \kappa - 1$, if

$$\begin{aligned} x_i &= w_i - \sum_{j\in\mathcal{J}_n} \mu_{ij} I\psi_{ij}, && i \in \mathcal{I}_n, \\ \sum_{j\in\mathcal{J}_n} \psi_{ij} &= x_i + \alpha_i, && i \in \mathcal{I}_n, \\ \sum_{i\in\mathcal{I}_n} \psi_{ij} &= \beta_j, && j \in \mathcal{J}_n, \end{aligned} \tag{50}$$

then

$$\sum_{i\in\mathcal{I}_n, j\in\mathcal{J}_n} |I\psi_{ij}|_t^* + \sum_{i\in\mathcal{I}_n} |x_i|_t^* \leq c_n \left(\sum_{i\in\mathcal{I}_n} (|w_i|_t^* + |I\alpha_i|_t^*) + \sum_{j\in\mathcal{J}_n} |I\beta_j|_t^* \right). \tag{51}$$

The implication (50)⇒(51) is proved by induction on n.

<u>*Induction base:*</u> $n = 1$. $\mathcal{T}_1$ has exactly two nodes, say $i \in \mathcal{I}$ and $j \in \mathcal{J}$. By (50), $x_i = w_i - \mu_{ij} I\psi_{ij}$ and $I\psi_{ij} = I\beta_j$. Hence (51) holds.

<u>*Induction step:*</u> Assume that (50)⇒(51) holds for $n \in [1, \kappa - 2]$. Let (50) hold for $n + 1$. We show that (51) holds for $n + 1$ in the two cases.

Case 1: *The leaf node v_{n+1} is in $\mathcal{I}$.* Denote $i_0 = v_{n+1}$ and let j_0 denote the unique node $j \sim i_0$ in $\mathcal{T}_{n+1}$. The validity of (50) for $n+1$ implies the following regarding $i \in \mathcal{I}_n$, $j \in \mathcal{J}_n$:

$$\begin{aligned} x_i &= w_i - \sum_{j \in \mathcal{J}_n} \mu_{ij} I\psi_{ij}, && i \in \mathcal{I}_n, \\ \sum_{j \in \mathcal{J}_n} \psi_{ij} &= x_i + \alpha_i, && i \in \mathcal{I}_n, \\ \sum_{i \in \mathcal{I}_n} \psi_{ij} &= \beta_j, && j \in \mathcal{J}_n \setminus \{j_0\}, \\ \sum_{i \in \mathcal{I}_n} \psi_{ij_0} &= \beta_{j_0} - \psi_{i_0 j_0}, \end{aligned} \tag{52}$$

and the following regarding the node i_0:

$$\begin{aligned} x_{i_0} &= w_{i_0} - \mu_{i_0 j_0} I\psi_{i_0 j_0}, \\ \psi_{i_0 j_0} &= x_{i_0} + \alpha_{i_0}. \end{aligned} \tag{53}$$

By (52) and the induction assumption,

$$\sum_{i \in \mathcal{I}_n, j \in \mathcal{J}_n} |I\psi_{ij}|_t^* + \sum_{i \in \mathcal{I}_n} |x_i|_t^* \le c_n \left(\sum_{i \in \mathcal{I}_n} (|w_i|_t^* + |I\alpha_i|_t^*) + \sum_{j \in \mathcal{J}_n} |I\beta_j|_t^* + |I\psi_{i_0 j_0}|_t^* \right). \tag{54}$$

By (53),

$$x_{i_0} = w_{i_0} - \mu_{i_0 j_0} I\alpha_{i_0} - \mu_{i_0 j_0} I x_{i_0}. \tag{55}$$

Applying Lemma 2 to (55) and using again (53) shows

$$|x_{i_0}(t)| + |I\psi_{i_0 j_0}(t)| \le c(|w_{i_0}|_t^* + |I\alpha_{i_0}|_t^*). \tag{56}$$

Combining (54) and (56) establishes the validity of (51) for $n+1$.

Case 2: *The leaf node v_{n+1} is in $\mathcal{J}$.* Denote $j_0 = v_{n+1}$ and let i_0 denote the unique node $i \sim j_0$ in $\mathcal{T}_{n+1}$. Assuming (50) for $n+1$ implies

$$\begin{aligned} x_i &= w_i - \sum_{j \in \mathcal{J}_n} \mu_{ij} I\psi_{ij}, && i \in \mathcal{I}_n \setminus \{i_0\}, \\ \sum_{j \in \mathcal{J}_n} \psi_{ij} &= x_i + \alpha_i, && i \in \mathcal{I}_n \setminus \{i_0\}, \\ x_{i_0} &= (w_{i_0} - \mu_{i_0 j_0} I\psi_{i_0 j_0}) - \sum_{j \in \mathcal{J}_n} \mu_{i_0 j} I\psi_{i_0 j}, \\ \sum_{j \in \mathcal{J}_n} \psi_{i_0 j} &= x_{i_0} + (\alpha_{i_0} - \psi_{i_0 j_0}), \\ \sum_{i \in \mathcal{I}_n} \psi_{ij} &= \beta_j, && j \in \mathcal{J}_n, \end{aligned} \tag{57}$$

and

$$\psi_{i_0 j_0} = \beta_{j_0}. \tag{58}$$

By (57) and the induction assumption, (54) follows. Combining (54) with (58) gives (51) for $n+1$.

This completes the proof that (50) implies (51) for $n \in [1, \kappa - 1]$. The result follows on taking $n = \kappa - 1$ and substituting $\alpha = -y$ and $\beta = -z$. □

7 An integral formula for Y and Z

Equations (43)–(45) were used in the previous section as substitutes for (32)–(34). They express a relation between the quantities x, y, z, w and ψ. In this section we extract a relation between y, z and w alone. This relation, in the form of an integral equation, is a key element in treating parts (i) and (ii) of Theorem 1.

For $\alpha \in \mathbb{R}$ denote

$$T_\alpha f = f + \alpha I f.$$

Lemma 2 shows that T_α is invertible. It is easy to see the operators T_α, T_β commute. If $A = (\alpha_1, \dots, \alpha_k)$ is a finite real-valued sequence, denote

$$T_A = T_{\alpha_1} \circ \cdots \circ T_{\alpha_k}.$$

Then T_A does not depend on the order of the elements of A, but it depends on the multiplicity of each element. Let $\emptyset$ stand for the empty sequence and set $T_\emptyset$ as the identity map. Equations (43)–(45) imply

$$\sum_j T_{\mu_{ij}} \psi_{ij} = w_i - y_i, \quad i \in \mathcal{I}, \tag{59}$$

$$\sum_i \psi_{ij} = -z_j, \quad j \in \mathcal{J}. \tag{60}$$

Theorem 2 *Let* (59) *and* (60) *hold.*

i. y and z solve the following integral equation

$$\sum_{i \in \mathcal{I}} T_{A_i}(w_i - y_i) + \sum_{j \in \mathcal{J}} T_{B_j} z_j = 0, \tag{61}$$

where A_i and B_j are finite (possibly empty) sequences with values in $\{\mu_{ij} : (i,j) \in \mathcal{E}\}$.

ii. In the special case where $\mu_{ij} = \mu_j$ for $(i,j) \in \mathcal{E}$ equation (61) takes the form:

$$\sum_{i\in\mathcal{I}}(w_i - y_i) + \sum_{j\in\mathcal{J}} T_{\mu_j} z_j = 0. \tag{62}$$

iii. In the special case where $\mu_{ij} = \mu_i$ for $(i,j) \in \mathcal{E}$ equation (61) takes the form:

$$\sum_{i\in\mathcal{I}} T_{M_i}(w_i - y_i) + T_M(e \cdot z) = 0, \tag{63}$$

where $M_i = (\mu_{i'})_{i'\in\mathcal{I}, i'\neq i}$, $M = (\mu_{i'})_{i'\in\mathcal{I}}$.

Remark 1 **(a)** Writing I_n for the n-power $I \circ \ldots \circ I$ of the operator I, it is useful to note that the integral equation (61) can be written as

$$e\cdot w - e\cdot y + e\cdot z + \sum_{i\in\mathcal{I}}\sum_{n=1}^{m_i} a_{i,n} I_n(w_i - y_i) + \sum_{j\in\mathcal{J}}\sum_{n=1}^{m_j} b_{j,n} I_n z_j = 0. \tag{64}$$

Here, m_i, m_j, $a_{i,n}, b_{j,n}$ are positive constants that we do not give in explicit form.

(b) Recall that under (46) and (47) one has $y = (e\cdot x)^+ u$ and $z = (e\cdot x)^- v$. As a result, (64) expresses a relation between the data w, the controls u and v, and the quantity $e \cdot x$ alone.

We need some notation regarding the tree $\mathcal{T}$ to be used in this and the following sections. Fix one of the class nodes, i_0, as a root. (Analogous notation applies if we fix a station node, some j_0 as a root). For $k = 0, 1, \ldots$, let level k, denoted by l_k be the set of nodes of $\mathcal{T}$ at distance k from the root i_0 along the edges of $\mathcal{T}$ (see Figure 3). Note that $l_0 = \{i_0\}$ and that l_k is empty for all k large. Let also

$$L_k = l_0 \cup l_1 \ldots \cup l_k$$

be the set of nodes at distance at most k from the root, and

$$L_k^{\mathcal{I}} = L_k \cap \mathcal{I}, \quad L_k^{\mathcal{J}} = L_k \cap \mathcal{J}.$$

Note that the elements of $L_k^{\mathcal{I}}$ [resp., $L_k^{\mathcal{J}}$] are at even [resp., odd] distance from the root, not exceeding k. Let K be the largest k such that l_k is nonempty. For a node v at level k let $B(v)$ (B for 'below') be the set of nodes $v' \sim v$ at level $k+1$. For a node v at level $k \in [1, K]$ let $a(v)$ (a for 'above') be the unique node $v' \sim v$ at level $k-1$.

Proof of Theorem 2: Part i. We show the following. For $k \geq 1$ one has

$$\sum_{i\in L_{2k-2}^{\mathcal{I}}} T_{A_i^k}(w_i - y_i) + \sum_{j\in L_{2k-1}^{\mathcal{J}}} T_{B_j^k} z_j + \sum_{i\in l_{2k}} T_{C_i^k} \psi_{ia(i)} = 0 \tag{65}$$

where A_i^k, B_j^k, C_i^k are (possibly empty) sequences with values in $\{\mu_{ij} : (i,j) \in \mathcal{E}\}$, and summation over an empty set is regarded as zero. Part (i) will follow on taking k larger than K.

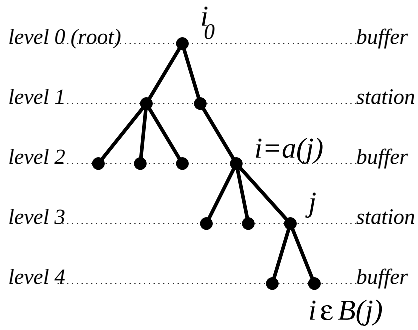

Figure 3: Buffer-station tree

We prove (65) by induction on k.

Induction base: $k=1$. By (59),

$$\sum_{j\in l_1} T_{\mu_{i_0 j}} \psi_{i_0 j} = w_{i_0} - y_{i_0},$$

and by (60)

$$\psi_{i_0 j} = -z_j - \sum_{i\in l_2} \psi_{ij}, \quad j \in l_1.$$

It follows that

$$w_{i_0} - y_{i_0} + \sum_{j\in l_1} T_{\mu_{i_0 j}} z_j + \sum_{i\in l_2} T_{\mu_{i_0 a(i)}} \psi_{ia(i)} = 0,$$

and (65) holds.

Induction step: Assume that (65) holds for k. Using (59) and then (60), for $i \in l_{2k}$ and $j = a(i)$ one has

$$\begin{aligned} T_{\mu_{ij}} \psi_{ij} &= - \sum_{j'\in B(i)} T_{\mu_{ij'}} \psi_{ij'} + w_i - y_i \\ &= \sum_{j'\in l_{2k+1}} \left\{ T_{\mu_{ij'}} z_{j'} + \sum_{i'\in l_{2k+2}} T_{\mu_{ij'}} \psi_{i'j'} \right\} + w_i - y_i. \end{aligned}$$

Apply $T_{C_i^k}$ on the above equation (where still $i \in l_{2k}$ and $j = a(i)$) to get

$$T_{C_i^k} T_{\mu_{ij}} \psi_{ij} = \sum_{j'\in l_{2k+1}} T_{C_i^k} T_{\mu_{ij'}} z_{j'} + \sum_{i'\in l_{2k+2}} T_{C_i^k} T_{\mu_{ia(i')}} \psi_{i'a(i')} + T_{C_i^k}(w_i - y_i). \tag{66}$$

Let $D_k = (\mu_{ia(i)})_{i\in l_{2k}}$. Apply T_{D_k} to (65) and use (66) to substitute for each summand in the third sum in (65) to get

$$0 = \sum_{i \in L^{\mathcal{I}}_{2k-2}} T_{D_k} T_{A^k_i}(w_i - y_i) + \sum_{j \in L^{\mathcal{J}}_{2k-1}} T_{D_k} T_{B^k_j} z_j$$

$$+ \sum_{i \in l_{2k}} \left\{ \sum_{j' \in l_{2k+1}} T_{D_k \setminus \{\mu_{ij'}\}} T_{\mu_{ij'}} T_{C^k_i} z_{j'} + \sum_{i' \in l_{2k+2}} T_{D_k \setminus \{\mu_{ij}\}} T_{\mu_{ia(i')}} T_{C^k_i} \psi_{i'a(i')} + T_{D_k \setminus \{\mu_{ij}\}} T_{C^k_i} (w_i - y_i) \right\}$$

where $D_k \setminus \{\mu\}$ denotes a sequence obtained from D_k by deleting from it one element of value μ. This proves that (65) holds for $k+1$ and completes the proof of Part i.

Part ii. In case that $\mu_{ij} = \mu_j$, by (59) and (60),

$$\sum_j T_{\mu_j} \psi_{ij} = w_i - y_i, \quad i \in \mathcal{I},$$

$$\sum_i T_{\mu_j} \psi_{ij} = -T_{\mu_j} z_j, \quad j \in \mathcal{J},$$

hence

$$\sum_i (w_i - y_i) + \sum_j T_{\mu_j} z_j = 0.$$

Part iii. In case that $\mu_{ij} = \mu_i$, by (59) and (60),

$$T_{\mu_i} \sum_j \psi_{ij} = w_i - y_i, \quad i \in \mathcal{I}, \tag{67}$$

$$\sum_i \psi_{ij} = -z_j, \quad j \in \mathcal{J}. \tag{68}$$

The result follows on applying $\prod_{i' \neq i} T_{\mu_{i'}}$ on (67) and $\prod_{i' \in \mathcal{I}} T_{\mu_{i'}}$ on (68). □

Proposition 4 *Let (59) and (60) hold. Assume also*

$$y_i \geq 0, z_j \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}, \tag{69}$$

$$e \cdot z(t) > 0 \quad \text{implies} \quad e \cdot y(t) = 0, \quad t \geq 0. \tag{70}$$

In cases (ii) and (iii) of Theorem 2 there are constants $c, m > 0$ such that

$$\|y(t)\| + \|z(t)\| \leq c(1+t)^m \|w\|^*_t, \quad t \geq 0. \tag{71}$$

Proof: In the case $\mu_{ij} = \mu_j$, (62) can be written as

$$e \cdot w - e \cdot y + e \cdot z + \sum_j \mu_j I z_j = 0.$$

By (69), z_j are positive and therefore

$$e \cdot w - e \cdot y + e \cdot z \leq 0.$$

Thus by (70)

$$0 \leq e \cdot z \leq (-e \cdot w)^+,$$

and therefore

$$e \cdot w \leq e \cdot y - e \cdot z \leq e \cdot w + t|e \cdot w|_t^*.$$

Since by (70)

$$\|y\| + \|z\| = |e \cdot y - e \cdot z|,$$

(71) follows.

In the case $\mu_{ij} = \mu_i$, applying $T_{\mu_i}^{-1}$ to (67) and by (68) (or by applying T_M^{-1} to (63)) we have

$$\sum_i T_{\mu_i}^{-1}(w_i - y_i) + e \cdot z = 0,$$

hence by Lemma 2 and positivity of y_i,

$$e \cdot w - e \cdot y + e \cdot z = \sum_i \mu_i \int_0^t (w_i(s) - y_i(s)) e^{-\mu_i(t-s)} ds \leq ct\|w\|_t^*.$$

By positivity of z_j and (70) we therefore have

$$\|z(t)\| \leq c(1+t)\|w\|_t^*. \tag{72}$$

By equation (64), the positivity of its coefficients and of y_i, and (72), there are constants $c, m > 0$ such that

$$\|y(t)\| \leq c(1+t)^m \|w\|_t^*. \tag{73}$$

Combining (72) and (73) establishes (71). □

By Propositions 3 and 4 we have the following.

Corollary 1 *Under the assumptions of Theorem 1(i), for any $m_1 \geq 1$, any initial condition $x \in \mathbb{R}^{\bar{I}}$ and any admissible system $\pi \in \Pi$,*

$$E_x^\pi \|X(t)\|^{m_1} \leq c(1 + \|x\|)^m (1+t)^m, \quad t \geq 0,$$

where c and m do not depend on x, π, t.

8 The non-idling property

As in the previous section, the following relations follow from (43)–(45):

$$\sum_j (\psi_{ij} + \mu_{ij} I\psi_{ij}) = w_i - y_i, \quad i \in \mathcal{I}, \tag{74}$$

$$\sum_i \psi_{ij} = -z_j, \quad j \in \mathcal{J}. \tag{75}$$

The constraints (46) and (47) were as follows:

$$y_i, z_j \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}, \tag{76}$$

$$\sum_i y_i(t) > 0 \text{ implies } \sum_j z_j(t) = 0, \quad t \geq 0. \tag{77}$$

Note that x (cf. equation (43)) is not a part of this system of equations, but can be obtained from it via

$$x_i = w_i - \sum_j \mu_{ij} I\psi_{ij}.$$

We say that *the system* (74)–(77) *incurs no idleness on* $[0, T]$ (compare with [20]) if $z_j(t) = 0$ for $t \in [0, T]$, $j \in \mathcal{J}$.

The non-idling property: *If the system starts with* $w_i(0) > 0$ *and* w_i*,* $i \in \mathcal{I}$*, are strictly increasing on* $[0, T]$*, then the system incurs no idleness on* $[0, T]$*.*

This section investigates a relation between the non-idling property and the uniform estimate

$$\|x\|_t^* \leq c(1+t)^m \|w\|_t^*, \tag{78}$$

where c, m do not depend on t, w, y, z, ψ. In particular, using the integral equation developed in the previous section, it shows that for trees of diameter not exceeding 3 the non-idling property implies the uniform estimate. A relatively simple argument then shows that the non-idling property holds for such trees, and the estimate follows.

A tree $\mathcal{T}$ of diameter 3 has the form depicted in Figure 4, where there are only two nodes $i_0 \in \mathcal{I}$ and $j_0 \in \mathcal{J}$ that are not leaves.

Theorem 3 *Let Assumption 1 hold and assume the diameter of the tree* $\mathcal{T}$ *is 3. Then the non-idling property implies the estimate* (78)*, where* c, m *do not depend on* t*,* w, y, z, ψ*.*

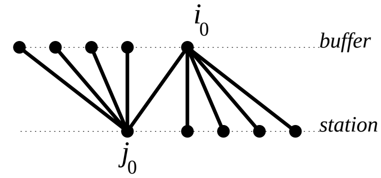

Figure 4: Tree of diameter 3

Lemma 3 *Let the assumptions of Theorem 3 hold. Fix $T > 0$ and let (ψ, y, z, w) be given, satisfying (74)–(77) on $[0, T]$. Let bounded measurable functions $f_i \geq 0$, $i \in \mathcal{I}$ be given. Then there is a constant c_0 that does not depend on T, f_i, ψ, y, z or w, and there exist $(\hat{\psi}, \hat{y}, \hat{z}, \hat{w})$ defined on $[0, T]$, satisfying (74)–(77) and moreover*

$$e \cdot \hat{y} \geq e \cdot y, \quad e \cdot \hat{z} \leq e \cdot z, \tag{79}$$

and

$$\hat{w}_i = w_i + f_i + \eta_i, \tag{80}$$

where η_i are nondecreasing and $0 \leq \eta_i \leq c_0 T \sum |f_i|^*_T$ on $[0, T]$.

Proof: It suffices to prove the lemma in the case where all but one of the functions f_i vanish, since the argument can then be repeated. Consider first the case where f_i vanish for all $i \neq i_0$ (and $f_{i_0} \geq 0$). Define $\hat{w}_i = w_i$, $\hat{y}_i = y_i$, $\hat{\psi}_{i,j} = \psi_{ij_0}$ for $i \neq i_0$. Define now $\hat{\psi}_{i_0 j}$ and $\hat{z}_j$ as follows. For $t \in \Theta := \{s \in [0, T] : e \cdot z(s) \geq f_{i_0}(s)\}$, let $j = \bar{\imath} + 1$, let $\Delta_j := z_j \wedge f_{i_0}$, and set $\hat{\psi}_{i_0 j} = \psi_{i_0 j} + \Delta_j$, $\hat{z}_j = z_j - \Delta_j$. Similarly, for $j \in [\bar{\imath}+2, \bar{\imath}+\bar{\jmath}]$, let $\Delta_j = z_j \wedge (f_{i_0} - \Delta_{\bar{\imath}+1} - \cdots - \Delta_{j-1})$, and $\hat{\psi}_{i_0 j} = \psi_{i_0 j} + \Delta_j$, $\hat{z}_j = z_j - \Delta_j$. This ensures that, for all j, $\hat{z}_j \geq 0$ and (75) is satisfied by the hat system, and, with $\delta(t) = 0$,

$$\sum_j \hat{\psi}_{i_0 j}(t) = \sum_j \psi_{i_0 j}(t) + f_{i_0}(t) - \delta(t). \tag{81}$$

For $t \in \Theta^c$, set $\hat{z}_j = 0$ and $\hat{\psi}_{i_0 j} = \psi_{i_0 j} + \Delta_j$ where now $\Delta_j = z_j$, $j \in \mathcal{J}$. Then again (75) holds for the hat system. Also (81) holds with $\delta(t) = f_{i_0}(t) - e \cdot z(t) \geq 0$. Now that $\hat{\psi}$ is defined on $[0, T]$, let

$$\hat{w}_{i_0} = w_{i_0} + \sum_j (\hat{\psi}_{i_0 j} - \psi_{i_0 j} + \mu_{i_0 j} I(\hat{\psi}_{i_0 j} - \psi_{i_0 j})) + \delta,$$

and $\hat{y}_{i_0} = y_{i_0} + \delta$. Then (74) holds on $[0, T]$. By (81), $\hat{w}_{i_0} = w_{i_0} + f_{i_0} + \eta_{i_0}$, where $\eta_{i_0} = \sum_j \mu_{i_0 j} I(\hat{\psi}_{i_0 j} - \psi_{i_0 j})$. Since $0 \leq \hat{\psi}_{i_0 j} - \psi_{i_0 j} \leq f_{i_0}$ by construction, the properties of η_{i_0} stated in the lemma hold. Inequalities (79) also hold by construction. Condition (77) holds for the hat system since on Θ, $\hat{y} = y$ and the set of times where $e \cdot z$ vanishes is a subset of that where $e \cdot \hat{z}$ vanishes; and on Θ^c, $e \cdot \hat{z} = 0$.

Next consider the case where for some $i_1 \neq i_0$, $f_i = 0$ on $[0, T]$ for all $i \neq i_1$. The argument is similar, but slightly more complicated because $i_1 \not\sim j$ for most j. We only indicate where the argument changes. The definition of $\hat{\psi}_{ij}$ and $\hat{z}_j$ differs as follows. $\hat{z}_j$ is defined as before, where i_0 is replaced by i_1. Recall that in the previous case, $\hat{\psi}$ is defined as $\hat{\psi} = \psi + \Delta$. In the current case, let $\hat{\psi}_{i_1 j_0} = \psi_{i_1 j_0} + \sum_j \Delta_j$, $\hat{\psi}_{i_0 j} = \psi_{i_0 j} + \Delta_j$ for $j \neq j_0$, and $\hat{\psi}_{i_0 j_0} = \psi_{i_0 j_0} - \sum_{j \neq j_0} \Delta_j$. Set

$$\hat{w}_{i_1} = w_{i_1} + \hat{\psi}_{i_1 j_0} - \psi_{i_1 j_0} + \mu_{i_1 j_0} I(\hat{\psi}_{i_1 j_0} - \psi_{i_1 j_0})) + \delta,$$

$$\hat{w}_{i_0} = w_{i_0} + \sum_{j \neq j_0} \mu_{i_0 j} I(\hat{\psi}_{i_0 j} - \psi_{i_0 j}).$$

Let $\hat{y}_{i_1}$ be defined analogously to the way $\hat{y}_{i_0}$ is defined in the previous case, and let $\hat{y}_i = y_i$ for all $i \neq i_1$. One checks that all properties (74)–(77) are satisfied by the hat system, and that the conclusions of the lemma all hold, except that $\hat{w}_{i_0} = w_{i_0} + f_{i_0} + \tilde{\eta}_{i_0}$ where there is no guarantee that $\tilde{\eta}_{i_0}$ is nondecreasing (in fact, we have assumed without loss that $f_{i_0} = 0$). But this is now corrected by applying a further modification as in the previous paragraph, where now one takes $f_{i_0} = \tilde{\eta}^-_{i_0}$ and $f_i = 0$ for all $i \neq i_0$. □

Proof of Theorem 3: Fix $T > 0$ and let (ψ, y, z, w) be given, satisfying (74)–(77) on $[0, T]$. Let f_i be defined on $[0, T]$ as

$$f_i(t) = \max(0, \sup_{0 \le s \le T} w_i(s)) - w_i(t) + t/T =: c_i - w_i(t) + t/T.$$

Let $(\hat{\psi}, \hat{y}, \hat{z}, \hat{w})$ be as in the conclusion of the lemma. Then $\hat{w}_i = c_i + t/T + \eta_i$, thus $\hat{w}_i(0) \ge 0$ and $\hat{w}_i$ is strictly increasing on $[0, T]$. Therefore by the non-idling property, $\hat{z}_j = 0$ on $[0, T]$, $j \in \mathcal{J}$. Hence by (61) and non-negativity of y_i,

$$\begin{aligned}\sum_i \hat{y}_i &\le \sum_i T_{A_i} \hat{w}_i \\ &\le c(1+t)^m \sum_i \hat{w}_i^*,\end{aligned}$$

where the last inequality uses the non-negativity of the coefficients of the integral operators T_{A_i}. Using (79) and (80),

$$e \cdot y(t) \le c(1+t)^m \sum_i (c_i + \eta_i(T)) \le c(1+t)^m T \|w\|_T^*.$$

Thus on $[0, T]$,

$$e \cdot y \le c(1+T)^{m+1} \|w\|_T^*.$$

Using again (61), now on w, y, z and equipped with the bound on y,

$$e \cdot z \le \sum_i T_{A_i}(y_i - w_i) \le c(1+t)^{m_1} \|w\|_T^*.$$

As a result,

$$\|y\|_T^* + \|z\|_T^* \le c(1+T)^{m_2} \|w\|_T^*,$$

where c, m_2 do not depend on T. The result now follows from Proposition 3. □

Lemma 4 *Let the assumptions of Theorem 3 hold. Let w_i be continuous functions. Then the non-idling property holds.*

Proof: It is more convenient here to work with the equivalent set of equations (43)–(45). For $i \neq i_0$, by (43), $x_i = w_i - \mu_{ij_0} \int_0^{\cdot} \psi_{ij_0}(s)ds$, and by (44), $\psi_{ij_0} = x_i - y_i \leq x_i$. Hence $x_i \geq \xi_i$, where ξ_i solves $\xi_i = w_i - \mu_{ij_0} \int_0^{\cdot} \xi_i(s)ds$. By Lemma 2 and since $w_i(0) \geq 0$, and w_i is nondecreasing, one has $\xi_i \geq 0$. As a result, $x_i \geq 0$.

Recall that $e \cdot z = (e \cdot x)^-$ and let $\tau = \inf\{t : e \cdot z(t) > 0\} = \inf\{t : e \cdot x(t) < 0\}$. Arguing by contradiction, assume that $\tau < T$. Since w is continuous, so is x, and since $e \cdot w(0) > 0$, $\tau > 0$. On $[0, \tau)$, $\psi_{ij_0} = z_{j_0} = 0$. Hence by the argument in the previous paragraph, applied for x_{i_0}, using the fact that w is strictly increasing, one has $x_{i_0}(\tau) > 0$. Thus by the result of the previous paragraph, $e \cdot x(\tau) > 0$. However, by the definition of τ and the continuity of $e \cdot x$ at τ, $e \cdot x(\tau) \leq 0$, a contradiction. As a result, $e \cdot x \geq 0$ on $[0, T)$, and by continuity of $e \cdot x$, the same holds on $[0, T]$. Therefore $e \cdot z = 0$, and by (76) $z_i(t) = 0$ for $i \in \mathcal{I}$ and $t \in [0, T]$. □

All results of this section remain valid for trees of diameter 2, as follows on applying them for trees of diameter 3 and choosing $\mu_{ij} = 0$ for appropriate (i, j) (nowhere in this section have we used the condition that $\mu_{ij} > 0$ for every (i, j) that is an edge of $\mathcal{T}$).

Theorem 3 and Lemma 4 imply the following.

Corollary 2 *Under the assumptions of Theorem 1(ii), for any $m_1 \geq 1$, any initial condition $x \in \mathbb{R}^{\bar{I}}$ and any admissible system $\pi \in \Pi$,*

$$E_x^\pi \|X(t)\|^{m_1} \leq c(1 + \|x\|)^m (1 + t)^m, \quad t \geq 0,$$

where c and m do not depend on π, x, t.

9 Case where cost is bounded below and summary of estimates

This section treats part (iii) of Theorem 1. First it shows that there exists an admissible system π_0 under which the state process satisfies a polynomial growth condition in time. Then it is shown that one can consider a subset of Π of admissible systems that, in a sense, 'switch' to π_0 after some time, without loosing optimality. The estimates for π_0 then remain valid in the 'switched' systems.

Proposition 5 *Let Assumption 1 hold and consider the system (43)–(45). Fix $i_0 \in \mathcal{I}$ and $j_0 \sim i_0$ and let $y = (e \cdot x)^+ u$, $z = (e \cdot x)^- v$, where $u(t) = e_{i_0}$ and $v(t) = e_{j_0}$ for all t. Then the estimate $\|x(t)\| \leq c(1 + t)^m \|w\|_t^*$ holds, where c and m do not depend on w, t.*

Proof: Let i_0 be the root. Then $j_0 \in B(i_0)$. It suffices to show that for $k = 0, 1, \dots$, if $i \in l_{2k}$ and $j \in l_{2k+1}$ then

$$|\psi_{a(j)j}(t)|, |\psi_{ia(i)}(t)|, |x_i(t)| \le c_k(1+t)^{m_k}\|w\|_t^*.$$

This claim is proved by backward induction on k.

Induction base: k is the smallest number n such that $l_{2n+2} = \emptyset$. If l_{2k+1} is not empty, let $j \in l_{2k+1}$. Then by (45), $\psi_{a(j)j} = -z_j = 0$. Let $i \in l_{2k}$. For $j \in B(i)$, $\psi_{ij} = 0$. Let $j = a(i)$. Then by (43), $x_i = w_i - \mu_{ij} I\psi_{ij}$ and by (44) and the assumptions, $\psi_{ij} = x_i - y_i = x_i$. Thus x_i solves $x_i = w_i - \mu_{ij} I x_i$ and by Lemma 2, $|x_i|, |\psi_{ij}| \le c|w_i|^*$.

Induction step: Assuming the claim holds for $k+1$ we show that it holds for k. Let $j \in l_{2k+1}$, $j \ne j_0$. Then by (45) and the assumption, $\psi_{a(j)j} = -\sum_{i \in B(j)} \psi_{ij}$. Thus by the induction assumption $|\psi_{a(j)j}| \le c(1+t)^{m_k}\|w\|^*$. Let $i \in l_{2k}$, $i \ne i_0$. Then by (43), (44), the bound on ψ_{ij}, $j \in B(i)$, just obtained, and the induction assumption,

$$x_i(t) = g(t) - \mu_{ia(i)} I x_i(t),$$

where $|g(t)| \le c(1+t)^m \|w\|_t^*$ with some constants c, m. Hence by Lemma 2, $|x_i(t)| \le c(1+t)^m \|w\|_t^*$. By (44) and the induction assumption, a similar bound then holds for $|\psi_{ia(i)}|$.

Consider now i_0 and j_0. Write x_0 for x_{i_0}, ψ_0 for $\psi_{i_0 j_0}$ and μ_0 for $\mu_{i_0 j_0}$. By (43)–(45) and the induction assumption, we have:

$$\begin{aligned} x_0 &= g_1 - \mu_0 I \psi_0, \\ \psi_0 + g_2 &= x_0 - y_0, \\ \psi_0 &= -z_0, \end{aligned}$$

where $|g_i(t)| \le c(1+t)^m \|w\|_t^*$ and c, m are some constants independent of t, w. Then $\psi_0 \le 0$ and therefore $x_0 \ge g_1$. Moreover, if $y_0(t) = 0$ then $\psi_0(t) \ge x_0(t) - g_2(t) \ge g_1(t) - g_2(t)$; and if $y_0(t) > 0$ then $y_0(t) \le e \cdot y(t) = e \cdot x(t)$, and $\psi_0(t) \ge x_0(t) - e \cdot x(t) - g_2(t) = -\sum_{i \ne i_0} x_i(t) - g_2(t)$. By the induction assumption thus $\psi_0(t) \ge -c(1+t)^m \|w\|_t^*$ holds for all t, with c, m some constants. Finally, $|x_0(t)| \le |g_1(t)| + \mu_0 t |\psi_0|_t^*$, and therefore a similar bound holds for x_0 as well. This completes the proof by induction. □

Throughout the rest of this section let i_0 and $j_0 \sim i_0$ be fixed. The following result shows that it suffices to consider only a certain subset of the set of admissible systems having the following property: There is a stopping time θ on the filtration associated with π, such that on $\{\theta < \infty\}$, $u(t) = e_{i_0}$ and $v(t) = e_{j_0}$ for all $t \ge \theta$.

Recall that on a given admissible system π, for every initial data x there is a corresponding controlled process (by Proposition 2).

Proposition 6 *Let the assumptions of Theorem 1(iii) hold. Let $x \in \mathbb{R}^{\bar{\imath}}$ be given. Then there is a set of admissible systems $\tilde{\Pi} \subset \Pi$ such that* *i. For every* $m_1 \geq 1$,

$$E_x^\pi \|X(t)\|^{m_1} \leq c(1+t)^m,$$

where c *and* m *depend on* x, π *and* m_1 *but not on* t*; and*

ii. $V(x) = \tilde{V}(x)$, *where*

$$\tilde{V}(x) = \inf_{\pi \in \tilde{\Pi}} C(x, \pi).$$

Proof: For any $\pi \in \Pi$ and any stopping time θ on π let π^θ be the admissible system obtained from π by setting $u(t) = e_{i_0}$ and $v(t) = e_{j_0}$ for $t \geq \theta$. Given x and $\pi \in \Pi$ such that $C(x,\pi) < \infty$ and for $\varepsilon > 0$ let

$$\sigma_\varepsilon = \inf\{t : E_x^\pi \int_t^\infty e^{-\gamma s} L(X(s), U(s)) ds \leq \varepsilon\}, \tag{82}$$

and define the stopping time

$$\theta_\varepsilon = \inf\{t \geq \sigma_\varepsilon : \|X(t)\| \geq \varepsilon^{-1}\}. \tag{83}$$

Note that σ_ε and θ_ε depend on x and π. Note also that

$$\theta_\varepsilon \geq \sigma_\varepsilon \to \infty, \qquad \text{as } \varepsilon \to 0. \tag{84}$$

For short, write π^ε for π^{θ_ε}. Define

$$\tilde{\Pi} = \{\pi^\varepsilon : \pi \in \Pi, \varepsilon \in (0,1)\}.$$

For Part (i), let $\pi \in \Pi$ and ε be given. Then on π^ε, and on the event $\{\theta_\varepsilon = \infty\}$, $\sup_{s \in [0,\infty)} \|X(s)\| \leq \varepsilon^{-1} \vee \|X\|^*_{\sigma_\varepsilon}$. Moreover, by Proposition 5, on the event $\{\theta_\varepsilon < \infty\}$, for $t \geq \sigma_\varepsilon$,

$$\|X(t)\| \leq c(1 + t - \sigma_\varepsilon)^{m_2} (\|X\|^*_{\sigma_\varepsilon} + \sup_{u \in [\sigma_\varepsilon, t]} \|\tilde{W}(u) - \tilde{W}(\sigma_\varepsilon)\|).$$

Recall that $\tilde{W}$ is a Brownian motion starting from the origin, and with constant drift. Hence $E_x^\pi(\sup_{u \in [\sigma_\varepsilon, t]} \|\tilde{W}(u) - \tilde{W}(\theta_\varepsilon)\|) \leq c(1+t-\sigma_\varepsilon)^{m_3}$ for some m_3 independent of ε and t, and the same holds for expectation under $E_x^{\pi^\varepsilon}$. Hence it suffices to prove that $E_x^{\pi^\varepsilon}[(\|X\|^*_{\sigma_\varepsilon})^{m_1}] = E_x^\pi[(\|X\|^*_{\sigma_\varepsilon})^{m_1}] < \infty$. This follows from an easy application of Gronwall's lemma, using the fact that $x \mapsto b(x, U)$ is Lipschitz uniformly in (x, U). Part (i) follows.

By Assumption 2(iv) and Proposition 5,

$$\begin{aligned}
&E_x^{\pi^\varepsilon}[1_{\theta_\varepsilon<\infty}\int_{\theta_\varepsilon}^\infty e^{-\gamma t}L(X(t),U(t))dt] \\
&\le cE_x^{\pi^\varepsilon}[1_{\theta_\varepsilon<\infty}\int_{\theta_\varepsilon}^\infty e^{-\gamma t}(1+t-\theta_\varepsilon)^m(1+\|X(\theta_\varepsilon)\|+\|\tilde W_t-\tilde W_{\theta_\varepsilon}\|)^{m_L}dt] \\
&\le c\int_0^\infty e^{-\gamma t}(1+t)^m dt\, E_x^{\pi^\varepsilon}[1_{\theta_\varepsilon<\infty}e^{-\gamma\theta_\varepsilon}(1+\|X(\theta_\varepsilon)\|)^{m_L}] \\
&\ + cE_x^{\pi^\varepsilon}[1_{\theta_\varepsilon<\infty}\int_{\theta_\varepsilon}^\infty e^{-\gamma t}(1+t-\theta_\varepsilon)^m\|\tilde W_t-\tilde W_{\theta_\varepsilon}\|)^{m_L}dt] \\
&:= cE_x^{\pi}[1_{\theta_\varepsilon<\infty}(1+\|X(\theta_\varepsilon)\|)^{m_L}e^{-\gamma\theta_\varepsilon}]+\alpha_2(\varepsilon) \\
&:= \alpha_1(\varepsilon)+\alpha_2(\varepsilon), \qquad (85)
\end{aligned}$$

where c does not depend on ε, and α_2 converges to zero as $\varepsilon\to 0$.

Next we show that $\alpha_1(\varepsilon)\to 0$ as $\varepsilon\to 0$. Below, we sometimes write θ for θ_ε. By definition of b (see (39)), $\|b(x,U)\|\le c(1+\|x\|)$, and $\|b(x,U)-b(y,U)\|\le c\|x-y\|$, where c does not depend on x, y and U. Thus by (40), for any $t\ge\theta$, one has on the event $\{\theta<\infty\}$

$$\begin{aligned}
\|X(t)-X(\theta)\| &\le c\|W(t)-W(\theta)\|+c\int_\theta^t[b(X(\theta),U(s))+(b(X(s),U(s))-b(X(\theta),U(s)))]ds \\
&\le c\|W(t)-W(\theta)\|+c(t-\theta)(1+\|X(\theta)\|)+c\int_\theta^t(X(s)-X(\theta))ds.
\end{aligned}$$

Hence by Gronwall's lemma

$$\|X(t)-X(\theta)\|\le c(\sup_{s\in[\theta,t]}\|W(s)-W(\theta)\|)(t-\theta)(1+\|X(\theta)\|)e^{c(t-\theta)}, \qquad (86)$$

where c is a deterministic constant not depending on t. Let τ be the stopping time defined as $\tau=\inf\{t>\theta:\|X(t)\|\le\|X(\theta)\|/2\}$ (and $\tau=\infty$ on the event $\{\theta=\infty\}$). By (83), $\|X(\theta)\|\ge\varepsilon^{-1}$. Hence by (86), on $\{\theta<\infty\}$,

$$\begin{aligned}
P(\tau-\theta<\varepsilon^{1/2}|F_\theta) &\le P(\sup_{t\in[\theta,\theta+\varepsilon^{1/2}]}\|X(t)-X(\theta)\|\ge(1+\|X(\theta)\|)/3|F_\theta) \\
&\le P(c\varepsilon^{1/2}e^{c\varepsilon^{1/2}}\sup_{s\in[\theta,\theta+\varepsilon^{1/2}]}\|W(s)-W(\theta)\|\ge 1/3|F_\theta) \\
&\le c\varepsilon^{1/2} \qquad (87)
\end{aligned}$$

for all $\varepsilon < 1$, where c does not depend on ε. Denote $\beta_\varepsilon = \tau_\varepsilon \wedge [\theta_\varepsilon + \varepsilon^{1/2}]$. Then by (82), (83), (87), and using the lower bound on $L(x, U)$, for all ε small,

$$\begin{aligned}
\varepsilon &\geq E_x^\pi[1_{\theta_\varepsilon<\infty} \int_{\theta_\varepsilon}^\infty e^{-\gamma s} L(X(s), U(s))ds] \\
&\geq cE_x^\pi[1_{\theta_\varepsilon<\infty} \int_{\theta_\varepsilon}^{\beta_\varepsilon} e^{-\gamma s}(1 + \|X(s)\|)^{m_L} ds] \\
&\geq cE_x^\pi[1_{\theta_\varepsilon<\infty}(1 + (1/2)\|X(\theta_\varepsilon)\|)^{m_L} e^{-\gamma\theta_\varepsilon}(\beta_\varepsilon - \theta_\varepsilon)] \\
&\geq cE_x^\pi\{1_{\theta_\varepsilon<\infty}(1 + (1/2)\|X(\theta_\varepsilon)\|)^{m_L} e^{-\gamma\theta_\varepsilon}\varepsilon^{1/2} P_x^\pi[\tau_\varepsilon - \theta_\varepsilon \geq \varepsilon^{1/2}|F_{\theta_\varepsilon}]\} \\
&\geq c\varepsilon^{1/2} E_x^\pi\{1_{\theta_\varepsilon<\infty}(1 + \|X(\theta_\varepsilon)\|)^{m_L} e^{-\gamma\theta_\varepsilon}\},
\end{aligned}$$

where $c > 0$ does not depend on ε. As a result, $\alpha_1(\varepsilon) \to 0$ as $\varepsilon \to 0$. Thus by (85), for every $\pi \in \Pi$ and $\varepsilon \in (0, 1)$ there is $\pi^\varepsilon \in \tilde{\Pi}$ such that

$$C(x, \pi^\varepsilon) \leq C(x, \pi) + \alpha(\varepsilon), \tag{88}$$

and $\alpha(\varepsilon) = \alpha_1(\varepsilon) + \alpha_2(\varepsilon) \to 0$ as $\varepsilon \to 0$. Hence $\tilde{V}(x) \leq V(x)$, and Part (ii) of the result follows. □

Here is a summary of our estimates in the three cases of Theorem 1.

Proposition 7 *Let the assumptions of Theorem 1 hold. In all cases (i)–(iii) of Theorem 1, for any x and any admissible system $\pi \in \Pi$ ($\pi \in \tilde{\Pi}$ in case (iii)), the following holds.*

i. For every $m_1 \geq 1$,

$$E_x^\pi\|X(t)\|^{m_1} \leq c(1 + t^m), \quad t \geq 0,$$

where the constants c and m do not depend on t (but may depend on x, π and m_1).

ii. There are constants c, m not depending on x such that $V(x) \leq c(1 + \|x\|)^m$.

iii. V is continuous on $\mathbb{R}^{\bar{\imath}}$.

iv. Let $D \subset \mathbb{R}^{\bar{\imath}}$ be a smooth domain. Let $g(t, x) = e^{-\gamma t}V(x)$, $(t, x) \in \mathbb{R}_+ \times \partial D$. Then $V = V_{D,g}$ in D, where $V_{D,g} = \inf_{\pi\in\Pi} C_{D,g}(x, \pi)$,

$$C_{D,g}(x, \pi) = E_x^\pi[\int_0^\tau e^{-\gamma t} L(X_t, U_t)dt + g(\tau, X_\tau)],$$

and $\tau = \inf\{t : X_t \notin D\}$.

Proof: Item (i) follows from Corollaries 1, 2 and Proposition 6. Item (ii) follows from Proposition 5 (in cases (i) and (ii) alternatively from Corollaries 1, 2). Items (iii) and (iv) are proved precisely as Proposition 4(ii,iii) of [4] (using Corollary 1 [resp., Corollary 2, Proposition 6] in cases (i) [resp., (ii), (iii)]). □

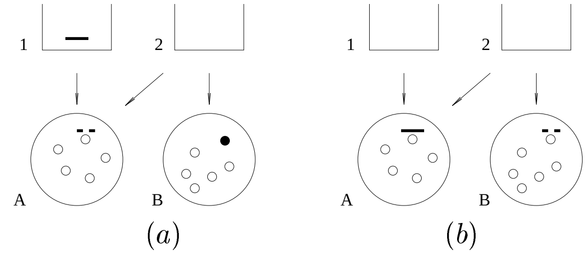

Figure 5: (a) A class-1 customer waiting in the queue, no free servers in station A, and one free server in station B. (b) A class-2 customer moved from station A to station B and the class-1 customer moved from the queue to station A

10 On joint work conservation

We argue that joint work conservation is in many cases a desired property for problems with preemption (for problems without preemption work conservation is typically not optimal, cf. [4], but since the distinction between the two is expected to disappear in the limit, it is expected that joint work conservation is, in a sense, asymptotically optimal). We only consider an example and do not attempt a rigorous treatment. We comment that even in the context of one station, proving that work conservation is optimal is not trivial, especially in a non Markovian system, as discussed in [4]. However, it is believed that the argument below can be greatly generalized, and this will be the subject of future work.

Example 4 The example consists of a Markovian system with two customer classes and two server types (see Figure 5). In (a) there is a customer of class 1 in the queue and a free server of type B. Although the free server cannot serve this customer, a rearrangement is possible so as to allow all customers to be served (b). We shall argue that 'good' preemptive policies prefer option (b) over (a). Clearly this question must be coupled with the cost criteria. For concreteness, consider a cost of the form $\sum c_i Y_i$ per unit time (weighted sum of queue lengths). Use the fact that optimality is obtained by feedback policies, that observe only the state: Number of customers at each class present in the system. Given a feedback policy π that leaves a customer in queue 1 when there is a free server in station B, we show there exists a policy $\hat{\pi}$ that pays a smaller cost on average. The argument is valid if (i) $\mu_{2A} \leq \max(\mu_{1A}, \mu_{2B})$ and otherwise if (ii) $c_1\gamma \geq c_2\mu_{2A}$. Arguing by coupling and assuming the system starts at the described state as time zero, let $\hat{\pi}$ move a class-2 customer, C_2, from station A to station B and let it move the class-1 customer waiting in the queue, C_1, into station A, until, at time τ, there is an arrival or a service in one of the systems (i.e., the system under π and that under $\hat{\pi}$). If the first to occur is an arrival or a service to a customer other than C_1 or C_2, $\hat{\pi}$ switches back to act like π for all times. In case (i) one can perform the coupling in such a way that service to C_1 under π always occurs later than the first between service to C_1 and service to C_2 under $\hat{\pi}$. If either service to C_1 or to C_2 under $\hat{\pi}$ is the first to occur, $\hat{\pi}$ then mimics π except for the single customer that under $\hat{\pi}$ is not present and may still be present for a while under π. In all cases the cost paid by $\hat{\pi}$ is not greater than that paid by π. In case (ii), if the first to occur is service to C_2 under π, an event we denote by η, then from that time on $\hat{\pi}$ keeps C_2 in the queue forever, and mimics π with regards to all other customers. On the event η the cost increases under $\hat{\pi}$ by at most $c_2 \int_0^\infty e^{-\gamma s} ds$, and the probability of η is μ_{2A}/Λ_0 where $\Lambda_0 = \lambda_1^n + n\lambda_2^n + \Psi_{1A}^{0,n}\mu_{1A} + \Psi_{2A}^{0,n}\mu_{2A} + \Psi_{2B}^{0,n}\mu_{2B}$. Hence $c_2\mu_{2A}\gamma^{-1}\Lambda_0^{-1}$ is lost. On the event η^c, the cost decreases under $\hat{\pi}$ by at least $c_1 \int_0^\tau e^{-\gamma s} ds$. Hence an amount of $c_1 E[1_{\eta^c} \int_0^\tau e^{-\gamma s} ds] \sim c_1 E\tau = c_1/\Lambda_0$ is gained. Therefore if $c_2\mu_{2A} \le c_1\gamma$, gain is greater than loss.

Appendix

Proof of Proposition 1: Let $i_0 \in \mathcal{I}$ be the root. We show the following claim by backward induction on k: For $k \in [1, K]$ even [resp., odd], if $i \in l_k$ [resp., $j \in l_k$] then $\psi_{ia(i)}$ [resp., $\psi_{ja(j)}$] is uniquely determined by (28) and (29).

Induction base: k is the largest n such that l_n is nonempty. If k is even, let $i \in l_k$. Then $B(i)$ is empty and (28) implies $\psi_{ia(i)} = \alpha_i$. The case k odd is similar. Induction step: Assume the claim holds for k. Consider the case where k is odd (the case k even is treated similarly). for $i \in l_{k-1}$, (28) shows $\psi_{ia(i)} = \alpha_i - \sum_{j\in B(i)} \psi_{ij}$, and since by the induction assumption ψ_{ij} are uniquely determined for $j \in B(i)$, so is $\psi_{ia(i)}$. This completes the proof by induction and the result follows. □

Proof of Theorem 1: Based on Proposition 7, the proof of Theorem 1 is similar to the proof of Theorem 2 of [4]. Since this is the main result of this paper, we have repeated it here with modifications, mainly to accommodate case (iii).

We first consider the equation (41) on a smooth open bounded connected domain $\tilde{D}$, satisfying an exterior sphere condition, with boundary conditions

$$f(x) = V(x), \quad x \in \partial\tilde{D}. \tag{89}$$

The key is a result from [9] regarding existence of classical solutions in bounded domains, with merely continuous boundary conditions. To use this result, we verify the following two conditions.

(i) $|H(x,p)| \le c(1 + \|p\|)$ for $x \in \tilde{D}$, where c does not depend on x or p.

(ii) $H(x,p) \in C^\varepsilon(\overline{\tilde{D}} \times \mathbb{R}^{\bar{I}})$, some $\varepsilon \in (0,1)$.

Item (i) is immediate from the local boundedness of $b(x,U)$ and $L(x,U)$. Next we show that item (ii) holds. For $\delta > 0$ let V be such that $H(y,q) \ge b(y,V)\cdot q + L(y,V) - \delta$. Write

$$H(x,p) - H(y,q) \le b(x,V)\cdot p + L(x,V) - b(y,V)\cdot q - L(y,V) + \delta.$$

Using the Hölder property of L in x uniformly for $(x,V) \in \overline{\tilde{D}} \times \mathbb{U}$, and the Lipschitz property of b in x, uniformly in (x,V),

$$H(x,p) - H(y,q) \le c\|p - q\| + c\|p\|\|x - y\| + c\|x - y\|^\varrho + \delta.$$

Since $\delta > 0$ is arbitrary, it can be dropped. This shows that H is Hölder continuous with exponent ϱ, uniformly over compact subsets of $\overline{\tilde{D}} \times \mathbb{R}^{\bar{\imath}}$. Hence (ii) holds.

Defining for $(x, z, p) \in \tilde{D} \times \mathbb{R} \times \mathbb{R}^{\bar{\imath}}$, $A(x, z, p) = (1/2)r^2 p$, $B(x, z, p) = H(x, p) - \gamma z$, one can write equation (41) in divergence form as

$$\operatorname{div} A(x, f, Df) + B(x, f, Df) = 0.$$

The hypotheses of Theorem 15.19 of [9] regarding the coefficients A and B hold in view of (i) and (ii). Indeed, B is Hölder continuous of exponent ϱ, uniformly on compact subsets of $\tilde{D} \times \mathbb{R} \times \mathbb{R}^{\bar{\imath}}$. Moreover, with $\tau = 0$, $\nu(z) = (1/2) \min_i r_i^2$, $\mu(z) = c(1 + \|z\|)$, $\alpha = 2$, $b_1 = 0$ and $a_1 = 0$, one checks that the conditions (15.59), (15.64), (15.66) and (10.23) of [9] are satisfied. Theorem 15.19 of [9] therefore applies. [We comment that there is a typo in the statement of the conditions of the theorem in [9]: the reference should be to condition (15.59) instead of (15.60).] It states that there exists a solution to (41) in $C^{2,\varrho}(\tilde{D}) \cap C(\overline{\tilde{D}})$, satisfying the continuous boundary condition (89). We denote this solution by f.

Let $x \in \tilde{D}$. Let π be any admissible system in Π, and let X be the controlled process associated with x and π. Let τ denote the first time X hits $\partial\tilde{D}$. Using Ito's formula for the $C^{1,2}(\mathbb{R}_+ \times \tilde{D})$ function $e^{-\gamma t} f(x)$, in conjunction with the inequality

$$\mathcal{L}f(y) + b(y, U) \cdot Df(y) + L(y, U) - \gamma f(y) \geq 0, \quad y \in \tilde{D}, U \in \mathbb{U},$$

satisfied by f, one obtains

$$f(x) \leq \int_0^{t\wedge\tau} e^{-\gamma s} L(X_s, U_s) ds + e^{-\gamma(t\wedge\tau)} f(X_{t\wedge\tau}) - \int_0^{t\wedge\tau} e^{-\gamma s} Df(X_s) \cdot r dW_s. \tag{90}$$

Taking expectation and then sending $t \to \infty$, using monotone convergence for the first term and bounded convergence for the second term, we have with $g(t, x) = e^{-\gamma t} V(x)$,

$$f(x) \leq E_x^\pi \Big[\int_0^\tau e^{-\gamma s} L(X_s, u_s) ds + e^{-\gamma\tau} V(X_\tau)\Big] = C_{\tilde{D},g}(x, \pi).$$

Taking infimum over $\pi \in \Pi$, we have

$$f(x) \leq V_{\tilde{D},g}(x) = V(x), \quad x \in \tilde{D},$$

where the last equality follows from Proposition 7 (iv).

In order to obtain the equality $f = V$ on $\tilde{D}$, we next show there exist optimal Markov control policies for the control problem on $\tilde{D}$. Let

$$\varphi(x, U) = b(x, U) \cdot Df(x) + L(x, U), \quad x \in \tilde{D}, U \in \mathbb{U}. \tag{91}$$

Note that φ is continuous on $\tilde{D} \times \mathbb{S}^k$. For each x, consider the set $\mathbf{U}_x \neq \emptyset$ of $U \in \mathbb{U}$ for which

$$\varphi(x, U) = \inf_{V \in \mathbb{U}} \varphi(x, V).$$

We show that there exists a measurable selection of $\mathbf{U}_x$, namely there is a measurable function h from $(\tilde{D}, \mathcal{B}(\tilde{D}))$ to $(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ with $h(x) \in \mathbf{U}_x$, $x \in \tilde{D}$.

Let $x_n \in \tilde{D}$ and assume $\lim_n x_n = x \in \tilde{D}$. Let U_n be any sequence such that $U_n \in \mathbf{U}_{x_n}$. We claim that any accumulation point of U_n is in $\mathbf{U}_x$. For, if this is not true, then by continuity of φ, there is a converging subsequence U_m, converging to $\tilde{U}$ and there is a $\hat{U}$ such that $\delta := \varphi(x, \tilde{U}) - \varphi(x, \hat{U}) > 0$. Hence for all m large, $\varphi(x_m, U_m) \geq \varphi(x, \hat{U}) + \delta/2 \geq \varphi(x_m, \hat{U}) + \delta/4$, contradicting $U_m \in \mathbf{U}_{x_m}$.

As a consequence, the assumptions of Corollary 10.3 in the appendix of [6] are satisfied, and it follows that there exists a measurable selection $h : \tilde{D} \to \mathbb{U}$ of $(\mathbf{U}_x, x \in \tilde{D})$.

We extend h to $\mathbb{R}^{\bar{\imath}}$ in a measurable way so that it takes values in $\mathbb{U}$ (but otherwise arbitrary). Clearly, $x \mapsto b(x, h(x))$ is measurable. Consider the autonomous SDE

$$X(t) = x + rW(t) + \int_0^t \hat{b}(X_s)ds, \tag{92}$$

where $\hat{b}(y)$ agrees with $b(y, h(y))$ on $\tilde{D}$, and is set to zero off $\tilde{D}$. Then $\hat{b}$ is measurable and bounded on $\mathbb{R}^{\bar{\imath}}$. By Proposition 5.3.6 of [15], there exists a weak solution to this equation. That is, there exists a complete filtered probability space on which X is adapted and W is a $\bar{\imath}$-dimensional Brownian motion, such that (92) holds for $t \geq 0$, a.s. On this probability space, consider the process $U_s = h(X_s)$. Since X has continuous sample paths and is adapted, it is progressively measurable (see Proposition 1.13 of [15]) and by measurability of h, so is U. Denote by π the admissible system thus constructed. Then for $s < \tau$, $U_s \in \mathbf{U}_{X_s}$, and

$$b(X_s, U_s) \cdot Df(X_s) + L(X_s, U_s) = H(X_s, Df(X_s)).$$

Hence

$$\mathcal{L}f(X) + b(X_s, U_s) \cdot Df(X_s) + L(X_s, U_s) - \gamma f(X_s) = 0, \quad s < \tau. \tag{93}$$

A use of Ito's formula and the convergence theorems just as before now shows that (90) holds with equality, and

$$f(x) = E_x^\pi[\int_0^\tau e^{-\gamma s} L(X_s, U_s)ds + e^{-\gamma\tau} V(X_\tau)] = C_{\tilde{D},g}(x, \pi), \quad x \in \tilde{D},$$

with g as above. This, together with Proposition 7(iv) shows that $f \geq V_{\tilde{D},g} = V$ on $\tilde{D}$. Summarizing, $f = V$ on $\tilde{D}$.

In particular, $V \in C^{2,\varrho}(\tilde{D})$ and is a classical solution to the HJB equation. $\tilde{D}$ can now be taken arbitrarily large, and this shows that $V \in C^{2,\varrho}(\mathbb{R}^{\bar{\imath}})$, and that it satisfies the HJB equation on $\mathbb{R}^{\bar{\imath}}$. In view of Proposition 7(ii), it also satisfies the polynomial growth condition. As a result, there exists a classical solution to (41) in $C^{2,\varrho}(\mathbb{R}^{\bar{\imath}})$, again denoted by f, satisfying (42), and moreover, $V = f$.

It remains to show uniqueness within $C^2_{\rm pol}(\mathbb{R}^{\bar\imath})$ and existence of optimal Markov control policies for the problem on $\mathbb{R}^{\bar\imath}$. Let $\bar f \in C^2_{\rm pol}(\mathbb{R}^{\bar\imath})$ be any solution to (41). Then analogously to (90) we obtain

$$\bar f(x) \le \int_0^t e^{-\gamma s} L(X_s, U_s) ds + e^{-\gamma t} \bar f(X_t) - \int_0^t e^{-\gamma s} D\bar f(X_s) \cdot r dW_s. \tag{94}$$

Taking expectation, sending $t \to \infty$, using the polynomial growth of $\bar f$ and the moment bounds on $\|X_t\|$ asserted in Proposition 7(i), one has that $\bar f(x) \le C(x, \pi)$, where $\pi \in \Pi$ ($\pi \in \tilde\Pi$ in case (iii)) is arbitrary. Consequently, $\bar f \le V$ on $\mathbb{R}^d$.

In cases (i) and (ii), the proof of existence of optimal Markov policies as well as the inequality $V \le \bar f$ on $\mathbb{R}^{\bar\imath}$ is completely analogous to that on $\tilde D$, where one replaces $\tilde D$ by $\mathbb{R}^{\bar\imath}$. The weak existence of solutions to (92) follows on noting that $\hat b$ satisfies a linear growth condition of the form $\|\hat b(y)\| \le x(1 + \|y\|)$, $y \in \mathbb{R}^{\bar\imath}$, and using again Proposition 5.3.6 of [15]. Then as before, (94) is satisfied with equality, and taking expectation and using the polynomial growth condition of $\bar f$ and the moment estimates on $\|X\|$ shows that $V = \bar f$ on $\mathbb{R}^{\bar\imath}$. We conclude that f is the unique solution in $C^2_{\rm pol}(\mathbb{R}^{\bar\imath})$, that $V = f$, and that there exists a Markov control policy, optimal for all $x \in \mathbb{R}^{\bar\imath}$.

In case (iii) there is no guarantee that the admissible system π constructed using (92) is in $\tilde\Pi$, and therefore the term $e^{-\gamma t} E^\pi_x \bar f(X_t)$ in (94) (that is satisfied with equality) may not tend to zero as $t \to \infty$. However, in this case we only claim uniqueness among nonnegative functions $\bar f$, and therefore using Ito's formula and (93) gives

$$\bar f(x) = E^\pi_x\Big[\int_0^t e^{-\gamma s} L(X_s, U_s) ds + e^{-\gamma s} \bar f(X_t)\Big] \ge E^\pi_x\Big[\int_0^t e^{-\gamma s} L(X_s, U_s) ds\Big],$$

and $\bar f(x) \ge C(x, \pi) \ge V(x)$. □

Acknowledgment The author wishes to thank Avi Mandelbaum and Marty Reiman for very valuable discussions.

References

[1] M. Armony and C. Maglaras. *Customer contact centers with multiple service channels,* Operations Research, to appear

[2] R. Atar. Asymptotically optimal scheduling for treelike parallel station systems. In preparation.

[3] R. Atar, A. Mandelbaum and M. Reiman. Brownian control problems for queueing systems in the Halfin-Whitt regime. *Systems and Control Letters,* to appear

[4] R. Atar, A. Mandelbaum and M. Reiman. Scheduling a multi-class queue with many exponential servers: asymptotic optimality in heavy-traffic. Preprint

[5] V. S. Borkar. *Optimal control of diffusion processes.* Longman, England 1989.

[6] S. N. Ethier and T. G. Kurtz. *Markov processes. Characterization and convergence.* Wiley, New York, 1986.

[7] W. H. Fleming and H. M. Soner. *Controlled Markov processes and viscosity solutions.* Springer-Verlag, New York, 1993.

[8] O. Garnett, A. Mandelbaum, and M. Reiman. *Designing a call center with impatient customers.* MSOM 4(3), 208-227, 2002.

[9] D. Gilbarg and N. S. Trudinger. *Elliptic partial differential equations of second order. Second edition.* Springer-Verlag, Berlin, 1983.

[10] S. Halfin and W. Whitt. Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* 29 (1981), no. 3, 567–588.

[11] J. M. Harrison. Brownian models of open processing networks: canonical representation of workload. *Ann. Appl. Probab.* 10 (2000), no. 1, 75–103.

[12] J.M. Harrison and M.J. López. Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, 33:339-368, 1999.

[13] J. M. Harrison and A. Zeevi. Dynamic scheduling of a multiclass queue in the Halfin-Whitt heavy traffic regime. *Operations Research,* to appear.

[14] H. Ishii. On uniqueness and existence of viscosity solutions of fully nonlinear second-order elliptic PDEs. *Comm. Pure Appl. Math.* 42 (1989), no. 1, 15–45.

[15] I. Karatzas and S. E. Shreve. *Brownian motion and stochastic calculus. Second edition.* Springer-Verlag, New York, 1991.

[16] A. Mandelbaum and M. I. Reiman. Private communication.

[17] A. Mandelbaum, W. A. Massey and M. I. Reiman. Strong approximations for Markovian service networks. *Queueing Systems* 30 (1998), no. 1-2, 149–201.

[18] A. Mandelbaum and A. L. Stolyar. Scheduling flexible servers with convex delay costs: heavy traffic optimality of the generalized $c\mu$ rule. Preprint

[19] A. A. Puhalskii and M. I. Reiman. The multiclass $GI/PH/N$ queue in the Halfin-Whitt regime. *Adv. in Appl. Probab.* 32 (2000), no. 2, 564–595.

[20] R. J. Williams. On dynamic scheduling of a parallel server system with complete resource pooling. *Analysis of communication networks: call centres, traffic and performance* (Toronto, ON, 1998), 49–71, Fields Inst. Commun., 28, Amer. Math. Soc., Providence, RI, 2000.