



Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

On the Dynamic Control of Matching Queues

Itai Gurvich, Amy Ward

To cite this article:

Itai Gurvich, Amy Ward (2014) On the Dynamic Control of Matching Queues. *Stochastic Systems* 4(2):479-523.
<https://doi.org/10.1287/13-SSY097>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Stochastic Systems*. Copyright 2014 The Author(s). <https://doi.org/10.1287/13-SSY097>, used under a Creative Commons Attribution License: <http://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2014, The author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

ON THE DYNAMIC CONTROL OF MATCHING QUEUES

BY ITAI GURVICH AND AMY WARD

Northwestern University and University of Southern California

We consider the optimal control of matching queues with random arrivals. In this model, items arrive to dedicated queues, and wait to be matched with items from other (possibly multiple) queues. A match type corresponds to the set of item classes required for a match. Once a decision has been made to perform a match, the matching itself is instantaneous and the matched items depart from the system. We consider the problem of minimizing finite-horizon cumulative holding costs. The controller must decide which matchings to execute given multiple options. In principle, the controller may choose to wait until some “inventory” of items builds up to facilitate more profitable matches in the future.

We introduce a multi-dimensional *imbalance* process, that at each time t , is given by a linear function of the cumulative arrivals to each of the item classes. A non-zero value of the imbalance at time t means that no control could have matched all the items that arrived by time t . A lower bound based on the imbalance process can be specified, at each time point, by a solution to an optimization problem with linear constraints. While not achievable in general, this lower bound can be asymptotically approached under a dedicated item condition (an analogue of the local traffic condition in bandwidth sharing networks). We devise a myopic discrete-review matching control that asymptotically—as the arrival rates become large—achieves the imbalance-based lower bound.

1. Introduction. We consider the matching of items that arrive randomly over time. Items of different classes arrive sequentially and wait in their respective queues – a queue for each class. Items can leave the system only after being matched to items of other (possibly multiple) classes. Once matched, the items leave the system together. We refer to such systems as *matching queues* and are concerned with their optimal control. The items may be thought of as people with relevant needs and/or skills (such as in an online market setting) or as inanimate components that must be combined to form completed products (such as in a manufacturing setting).

In the example depicted in Figure 1, there are 4 classes of items, and items of class i arrive according to a time-varying Poisson process A_i having

Received January 2013.

MSC 2010 subject classifications: 60K25, 60J70, 60F17.

Keywords and phrases: Matching queues, optimal control, diffusion approximations.




FIG 1. A queueing network view of a system with four input streams and two matchings.

instantaneous rate $\lambda_i(t)$, $i = 1, 2, 3, 4$. Items of class 1 can be matched to items of class 2. Items of class 2 can be also matched with items of classes 3 and 4. This structure is reflected in the graph in Figure 1 where each rectangle corresponds to an item class and each of the circles A and B to matching types. When a class 1 item is matched with a class 2 item they both leave the system: matchings are instantaneous. An item of class 4 **must** be matched to both a class 3 and a class 2 item to depart. The matching-to-queue adjacency matrix (henceforth, the *matching matrix*) is given in this case by

$$M = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}.$$

The controller must decide when to perform a matching and which matching to perform given multiple options. If the decision is to perform d_A matchings of type A and d_B matchings of type B , $(Md)_i$ units are depleted from the class i queue, where $d = (d_A, d_B)$. The controller may choose to wait even if items are available. In Figure 1, suppose that there is a single item available in each of the class 1, class 2 and class 4 queues but none in the class 3 queue. The controller may be greedy and match the available class 1




FIG 2. The simplest matching network – one matching type and two item classes.

item with the available class 2 item, which would result in two items departing the system. It may be, however, preferable to “inventory” this class 2 item until a class 3 item arrives at which point it can be matched with one class 2 and one class 4 items – depleting three items.

A first (and central) step in optimizing these networks is to identify a good *state descriptor*. This choice is simple if one were to consider the simplest of matching networks – one with two item classes and one possible matching; see Figure 2. Assuming one performs matchings whenever there are items to be matched, either queue 1 or queue 2 must be empty. In turn, the *imbalance* between the arrivals

$$S(t) = A_1(t) - A_2(t),$$

is a sufficient state descriptor: the size of each of the queues at time t is determined by the value of the imbalance. In particular,

$$(1.1) \quad Q_1(t) = [A_1(t) - A_2(t)]^+ = [S(t)]^+ \text{ and } Q_2(t) = [A_1(t) - A_2(t)]^- = [S(t)]^-.$$

Given a strictly increasing non-negative cost function $C : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$, it is trivial now that under any control

$$(1.2) \quad C(Q_1(t), Q_2(t)) \geq C([S(t)]^+, [S(t)]^-), \quad t \geq 0.$$

If the imbalance at time t is non-zero, then matching all arriving items is infeasible and the instantaneous cost must be strictly positive.

More generally the definition of the imbalance process is not as simple. The conceptual implications, however, remain valid in great generality:

- If the imbalance process is non-zero at time t , then there is **no** control that could have matched all items that arrived by time $t > 0$.
- A *lower bound* on the holding costs under any matching control is given by a simple function of the imbalance process, as in (1.2).

In this paper, we show how to explicitly construct such an imbalance process. The imbalance S is constructed as a linear mapping of the arrival processes. If q_0 is the vector of queues at time $t = 0$, then

$$(1.3) \quad S(t) = Y^T(q_0 + A(t)),$$

where the *imbalance matrix* Y has dimension $\mathbf{I} \times \mathbf{J}$, with \mathbf{I} being the number of classes and \mathbf{J} the number of possible matchings; see Section 3. In the simple example of Figure 2, Y is the column vector $(1, -1)^T$.

Our objective is to propose a matching control that minimizes the finite horizon cumulative cost, $\int_0^u C(Q(t))dt$, where $u > 0$ is the time horizon and $C : \mathbb{R}_+^{\mathbf{I}} \rightarrow \mathbb{R}_+$ is the instantaneous cost function. Finite horizon (rather than long-run average) objectives are natural because the matching queues we consider here are inherently unstable: for the simple example in Figure 2, either the queue sizes blow up to ∞ (if $\lambda_1 \neq \lambda_2$) or the imbalance process is null recurrent (if $\lambda_1 = \lambda_2$).

A lower bound on the instantaneous holding cost based on the imbalance process in (1.3) is given by

$$(1.4a) \quad \min \quad C(q)$$

$$(1.4b) \quad \text{s.t.} \quad Y^T q = S(t)$$

$$(1.4c) \quad q \geq 0,$$

at every time $t > 0$. That is, one optimizes the queue sizes q subject to the constraint that the appropriately weighted (by Y^T) sum of the queues must equal the imbalance process. Obviously, there is no reason to expect that the lower bound in (1.4) is achievable as it ignores past actions; namely, the fact that once an item is matched, it can not be “unmatched”. In Figure 1, if the controller matches a class 2 item with a class 1 item when a class 4 item was waiting but a class 3 item arrives shortly after, the controller cannot “unmatch” the class 1 and 2 items in order to match instead the class 2, 3, and 4 items.

In this paper, we show that a simple discrete review matching control achieves the imbalance based lower bound for a large class of networks

(specifically, those that satisfy a dedicated item condition – see Assumption 1), when the arrival rates become large. In brief, our proposed control minimizes the costs given the backlog of items at specified decision epochs t_1, \dots, t_m : at time t_m , we determine the number of matches of each type d_m and the resulting queues q_m by solving:

$$\begin{aligned} (1.5a) \quad & \min && C(q_m) \\ (1.5b) \quad & \text{s.t.} && q_m = Q(t_m-) - Md_m, \\ (1.5c) \quad & && q_m, d_m \geq 0, \end{aligned}$$

where M is the matching matrix. The algorithm is fleshed out in detail in §4.

The discrete-review matching control (1.5) is (partly) greedy—in the case of Figure 1, our control will not “inventory” items of class 2 in anticipation of arrivals of classes 3 and 4. However, the term “greedy” must be used with care. Discrete review means that we do nothing between review periods, so that items accumulate. The inventory accumulated between review periods provides the flexibility to achieve the lower bound in (1.4) at each decision epoch. Decision epochs must be close enough so as not incur significant holding cost, but sufficiently spaced out so that enough items have accumulated between review periods to provide matching flexibility.

In summary, the main contribution of this paper is to introduce a model for matching queues, identify the imbalance process as a key concept, and prove that a simple control is asymptotically optimal for holding cost minimization.

Organization of the paper. We conclude this section with a literature review. We specify our model in §2, and explicitly construct the imbalance matrix Y and the imbalance process $S(t)$ in §3. We fully describe our proposed control in §4. We state our first asymptotic optimality result in §5. That first result assumes that, in fluid scale, all items can be matched. When there is no such fluid balance, our asymptotic optimality result requires additional conditions, and that setting is studied in §6. Proof essentials appear in §7 and numerical examples in §8. We make concluding remarks in §9. The proofs of propositions and theorems appear in the main body of the paper, and there is an [appendix](#) for the proofs of lemmas.

Literature review. There are two streams of literature that are relevant to our work: that on stochastic processing networks and that on assemble-to-order systems.

The stochastic processing networks (SPNs) literature. In Harrison’s SPN framework [6], a matching network can be viewed as a SPN with I classes of

items that are processed via J activities. For the matching network drawn in Figure 1, there are $I = 4$ classes and $J = 2$ activities that can be used to process these classes. Each of the J activities is undertaken by a resource and all processing is instantaneous (so that each resource has infinite capacity).

Our model can thus be viewed as an extension of assembly-queues. One early example is the model studied in [4], which has I classes, a single activity served by a finite capacity resource, and positive processing times. Although the matching requirement is similar, there are two important distinctions. The first is the fact that we have instantaneous processing. A consequence of this is that the imbalance process can be approximated by an unregulated diffusion. With positive processing times, the appropriate approximation is a regulated diffusion. The second important distinction is that with a single activity there is no question of control. In contrast, our focus is on networks with multiple matching types, in which the network performance is determined by decisions concerning which activities to undertake when – that is, by the matching decisions. This question of matching control has not been fully explored in the SPN literature.

Studies of SPNs typically follow a hierarchy in which one first considers a static planning problem, and subsequently a waiting- or holding-cost minimization problem. The second step is often facilitated by use of the “standard Brownian machinery” proposed in [5]; namely, by introducing a heavy-traffic asymptotic regime in which resources are almost fully utilized and queue-lengths can be approximated by a function of a Brownian motion. An equivalent problem of lower dimension, that can be more easily solved, is constructed based on the workload process (see [8]) whose construction, in turn, relies on the dual of a static planning problem. A class of SPNs that is closely related to matching networks is that of parallel-server networks (Figure 1 serves to visualize this similarity). That class is studied using the standard Brownian machinery in, for example, [2, 7, 11], and those papers identify resource-pooling conditions under which the workload process is one-dimensional. The solution to the one-dimensional equivalent workload formulation is used to construct a control for the original SPN that is proved to be asymptotically optimal in heavy-traffic.

In spirit, our analysis of matching networks follows this same hierarchy. The construction of the imbalance process is based on the dual to an appropriate static planning problem, and we perform an asymptotic analysis to motivate a dynamic control and prove that the proposed control is asymptotically optimal. The conventional notion of heavy traffic does not, however, apply in our setting because there is no obvious notion of resource capacity. The appropriate analogue to the heavy-traffic asymptotic regime involves

balance among the various input flows, which effectively “serve” one another. The network in Figure 2 is balanced when $\lambda_1(t) = \lambda_2(t)$ for all $t \geq 0$, and that of Figure 1 is balanced when $\lambda_1(t) + 0.5(\lambda_3(t) + \lambda_4(t)) = \lambda_2(t)$ and $\lambda_3(t) = \lambda_4(t)$ for all $t \geq 0$. Instead of the resource pooling condition, a reduction in problem dimensionality is achieved here by a dedicated-item condition which is an analogue of the local traffic condition developed in [10] for a bandwidth sharing model. The control problem does not, however, collapse here to a single dimension.

The assemble-to-order system literature. Our interpretation of the imbalance process is closely connected to the interpretation of the inventory position in an assemble-to-order (ATO) system. In such a system, stochastic demand for a set of end products is met by assembling components. Different end-products use different (possibly overlapping) subsets of components and the controller must decide dynamically which end products to assemble given the backlogged demand for products and the inventory of components; see, e.g., the survey paper [15]. The components may either be arriving stochastically, if the capacitated component production is explicitly modeled, or they may be ordered, in which case they arrive after a lead time. Regardless of the component delivery method, the relevant state descriptor is the inventory position which tracks the number of components required to satisfy product demand, and is positive (negative) if the system has extra (is short of) components. Our imbalance process is similar in that it tracks which items are plentiful and which items are in short supply.

In the assemble-to-order setting with component production, as in [12], the notion of heavy-traffic equates the rate at which products are demanded to the rate at which components are produced. When those rates are in balance, the standard Brownian machinery can be used to construct an asymptotically optimal control that specifies dynamically which products to assemble. An important similarity (and a departure from traditional queuing networks) is that, in both the ATO context and in our matching setting, capacity – be it the queue of components in the ATO setting or the queues of items in our setting (that serve as “resources” for other items) – can be “banked” and is not perishable. Yet this is where the connection to the matching setting partly breaks down. In the matching setting there is no natural notion of physical capacity as items play the dual role of “products” needing components and “components” that are used in other products. This duality requires a different notion of “heavy-traffic”.

The standard Brownian machinery is applied also in the study of ATO systems where components are ordered; see [3] and [13]. These papers use functional central limit theorems to show that, as the lead times grow large

(and so does, in particular, the demand for products during the lead time) a certain lower bound stochastic program is attained. This is also the spirit of the analysis in this paper: to use functional central limit theorem methodology to show that when the arrival rates become large, the lower bound in (1.4) can be attained. In our setting, the absence of capacitated resources allows us to cover, within a single asymptotic framework, balanced and non-balanced networks.

Notation. We let \mathbb{R} denote the real numbers and \mathbb{R}_+ denote the positive real numbers. The set of integers is \mathbb{Z} and \mathbb{N} denotes the non-negative integers. For a set \mathcal{S} , $|\mathcal{S}|$ denotes its cardinality. All vectors are assumed to be column vectors. The transpose of a vector v is denoted by v^T . The notation $|v|$ denotes the Euclidean norm of v . We let e be the vector of all 1's, and e_j be the vector of all 0's except with a 1 in the j th place. All processes considered in what follows are assumed to be right continuous with left limits, and $\mathcal{D}^d[0, \infty)$ denotes the space of such functions from $[0, \infty)$ to \mathbb{R}^d . For a process $x \in \mathcal{D}^d[0, \infty)$ and a constant $u > 0$ we let $\|x\|_{s,u} = \sup_{s \leq t \leq u} |x(t)|$ (we abbreviate to $\|x\|_u$ if $s = 0$) and define $\Delta x(t) = x(t) - x(t-)$.

For asymptotic optimality we consider a sequence of systems indexed by $n \in \mathbb{R}_+$. We use the notation \Rightarrow to denote convergence in distribution as $n \rightarrow \infty$ in the space $\mathcal{D}^d[0, \infty)$. We use the same notation for weak convergence of random variables and the correct interpretation will be clear from the context. For a sequence of random variables $\{X^n\}$ and a sequence of non-negative numbers a^n we say that $X^n = o_P(a^n)$ if $|X^n|/a^n \rightarrow 0$ in probability, as $n \rightarrow \infty$.

2. The matching model. The model consists of a set \mathcal{I} of input streams, or item classes, and a set of matchings \mathcal{J} . A matching corresponds to a subset of \mathcal{I} that contains at least two item classes. We let $\mathcal{I}(j)$ be the set of item classes participating in matching $j \in \mathcal{J}$ and $\mathcal{J}(i)$ be the set of matchings which involve item class $i \in \mathcal{I}$. The *matching matrix* $M \in \{0, 1\}^{\mathbf{I} \times \mathbf{J}}$, where $\mathbf{I} = |\mathcal{I}|$ and $\mathbf{J} = |\mathcal{J}|$, has $M_{ij} = 1$ if $i \in \mathcal{I}(j)$ and 0 otherwise. We assume that for each i , there exists at least one j such that $M_{ij} = 1$; that is, each item class is connected to at least one matching. In Figure 1,

$$(2.1) \quad M = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \\ 0 & 1 \end{pmatrix},$$

and

$$\mathcal{I} = \{1, 2, 3, 4\} \text{ and } \mathcal{J} = \{A, B\}, \text{ for } A = \{1, 2\} \text{ and } B = \{2, 3, 4\},$$

$$\mathcal{I}(A) = \{1, 2\}, \mathcal{I}(B) = \{2, 3, 4\}, \mathcal{J}(1) = \{A\}, \mathcal{J}(2) = \{A, B\},$$

$$\text{and } \mathcal{J}(3) = \mathcal{J}(4) = \{B\}.$$

Class i items arrive according to a (possibly time varying) Poisson process $A_i = (A_i(t), t \geq 0)$ with instantaneous rate $\lambda_i(t)$, so that $\Lambda_i(t) = \int_0^t \lambda_i(s) ds$ is a first order approximation for the cumulative arrivals up to time t .¹ We assume that

$$0 < \lambda_{\min} \leq \lambda_i(t) \leq \lambda_{\max} < \infty, \quad t \geq 0, i \in \mathcal{I}.$$

The control is the vector of processes $D_j = (D_j(t), t \geq 0)$, $j \in \mathcal{J}$, where $D_j(t)$ tracks the cumulative number of times matching j has been performed in $[0, t]$, and has

$$(2.2) \quad D(0) = 0, \quad \Delta D(t) \in \mathbb{N}^{\mathcal{J}}.$$

Let $q_{0,i}$ be the number of items in queue i at time 0. The number of class i items waiting at time $t \geq 0$ is then

$$Q_i(t; D) = q_{0,i} + A_i(t) - \sum_{j \in \mathcal{J}(i)} D_j(t),$$

or, in vector notation

$$(2.3) \quad Q(t; D) = q_0 + A(t) - MD(t).$$

Naturally, we only consider controls under which

$$(2.4) \quad Q(t; D) \geq 0 \text{ for all } t \geq 0.$$

Also, since matching $j \in \mathcal{J}$ is only feasible at times $t \geq 0$ in which at least one item is waiting in each of the queues $i \in \mathcal{I}(j)$, we require that for all $j \in \mathcal{J}$, $i \in \mathcal{I}(j)$ and $t > 0$,

$$(2.5) \quad \Delta D_j(t) > 0 \text{ implies } Q_i(t-; D) + \Delta A_i(t) > 0 \text{ for all } i \in \mathcal{I}(j),$$

where we define $Q(0-; D) = q_0$. A control D is *admissible* if (2.2)–(2.5) hold.

Given a non-negative and convex function $C : \mathbb{R}_+^{\mathcal{I}} \rightarrow \mathbb{R}_+$ that has $C(0) = 0$ and is strictly increasing with respect to the natural partial order on $\mathbb{R}_+^{\mathcal{I}}$, we seek to solve the problem

$$(2.6) \quad \text{minimize } \int_0^u C(Q(t; D)) dt \text{ over all admissible controls } D,$$




FIG 3. (LHS) DI violated (RHS) DI satisfied.

for any given $u > 0$, where the minimization should be interpreted in a stochastic sense. In words, we wish to minimize the finite horizon costs of having items waiting (holding costs).

Our proposed solution will be pathwise (asymptotically) optimal under our assumptions. As a corollary, if one considers $q_0 = 0$ and the particular cost function

$$(2.7) \quad C(q) = \sum_{i \in \mathcal{I}} q_i \text{ for } q \in \mathbb{R}_+^{\mathcal{I}},$$

the solution minimizes $\sum_{i \in \mathcal{I}} Q_i(t)$ for each $t \geq 0$ and, consequently, maximizes the total number of items matched in $[0, u]$.

Dedicated item condition. The following condition is the analogue of the local traffic condition in [10, Assumption 5.1], defined there for a bandwidth sharing network. In words, the condition requires that each matching has at least one class that is used by that matching and that matching only. In Figure 1, class 1 is served only by matching A and classes 3 and 4 are served only by matching B. The dedicated item condition is not satisfied, for example, in Figure 3(LHS) where matching B has no such item class, but it is satisfied in Figure 3(RHS).

DEFINITION 1 (dedicated item (DI)). For each $j \in \mathcal{J}$ there exists $i \in \mathcal{I}$ such that $M_{ij} = 1$ but $M_{ik} = 0$ for all $k \neq j$.

¹Imposing a Poisson structure simplifies discussing non-stationary arrivals but is not necessary for some of our results; see Remark 3.

ASSUMPTION 1. *The network satisfies the DI condition.*

Since a matching must have at least two classes, the dedicated item condition implies that $\mathbf{I} > \mathbf{J}$. It also implies that the matrix M (after possibly re-arranging indices) has the form $M^T = (I, M_2)$ where I is the $\mathbf{J} \times \mathbf{J}$ identity matrix and M_2 is some $\mathbf{J} \times (\mathbf{I} - \mathbf{J})$ matrix. In particular, the matrix M has rank \mathbf{J} . A useful implication of the structure of the matrix M is that $Mx \geq 0$ if and only if $x \geq 0$. In the algebraic literature this property is referred to as *inverse monotonicity* and it is this implication of the dedicated item condition that is central to our proofs.

We say that a network is *fluid-balanced* if for each $t \geq 0$, there exists $z = z(t) \geq 0$ such that $Mz = \Lambda(t)$: starting with empty queues one can (in first order) match all arrivals by time t and make the “fluid” queues empty.

Conversely, we say that a network is *not balanced* if there exists t such that $Mz \neq \Lambda(t)$ for all $z \geq 0$. In this case, no control can empty the queues completely in fluid scale – there must be some positive queues some of the time.

This separation according to fluid balance is not central to the development of our proposed control in Sections 3 and 4. However, the conditions for asymptotic optimality are stronger for non-balanced networks; see Section 6.

3. The imbalance matrix and process. Our proposed control is based on the imbalance process (1.3). The imbalance process is used to identify the lower bound on the achievable cost given in (1.4). In this section, we provide one possible construction of the matrix Y in (1.3), and relate this construction to the duality of item classes as customers and servers.

To motivate the imbalance process, suppose that no action is taken until time t . At this time, it is feasible to match all items if there exists $d \geq 0$ that solves

$$(3.1) \quad Md = q_0 + A(t), \quad d \in \mathbb{R}_+^{\mathbf{J}}.$$

The dual to (3.1) is given by the linear program

$$(3.2) \quad \begin{aligned} \max_{y \in \mathbb{R}^{\mathbf{I}}} \quad & y^T(q_0 + A(t)) \\ \text{s.t.} \quad & y^T M \leq 0. \end{aligned}$$

By strong duality (3.1) has a feasible solution d^* if and only if (3.2) has a finite optimal solution y^* (any such solution must have $(y^*)^T(q_0 + A(t)) = 0$). Suppose that $q_0 + A(t) \gg 0$, i.e., that all components of the vector $q_0 + A(t)$

are strictly positive. Then, by the dedicated item condition it must be that $d^* \gg 0$ and, by complementary slackness, that $(y^*)^T M = 0$. Thus, for $q_0 + A(t) \gg 0$, the primal has a feasible solution d^* if and only if the dual has an optimal solution y^* with $(y^*)^T(q_0 + A(t)) = 0$ and $(y^*)^T M = 0$. If, on the other hand, there exists y such that $y^T M = 0$ but $y^T(q_0 + A(t)) > 0$ there is no solution $d \geq 0$ to $Md = q_0 + A(t)$ so that, regardless of our actions, some queues *must* be positive at time t .

We construct a process $S = (S(t), t \geq 0)$ having $S(t) = 0$ if there exists $d \geq 0$ such that $Md = q_0 + A(t)$ and $S(t) \neq 0$ otherwise. We fix a matrix Y whose columns span

$$(3.3) \quad \mathcal{Y} = \{y \in \mathbb{R}^{\mathbf{I}} : y^T M = 0\},$$

and define

$$(3.4) \quad S(t) = Y^T(q_0 + A(t)) \text{ for all } t \geq 0.$$

The process $S(t)$ obtains values in the subspace

$$(3.5) \quad \mathcal{M} = \{s \in \mathbb{R}^{\mathbf{I}-\mathbf{J}} : Y^T x = s, \text{ for some } x \geq 0\}.$$

The following formalizes our heuristic motivation of the imbalance process.

LEMMA 1. *Suppose that Assumption 1 holds. For each $x \geq 0$ such that $Y^T x = 0$ there exists a unique solution d to the system of equations $Md = x$, and this solution is non-negative.*

As a corollary of this lemma we observe that the network is fluid balanced if, for all $t \geq 0$, $Y^T \Lambda(t) = 0$ as this guarantees the existence of $z(t) \geq 0$ such that $Mz = \Lambda(t)$.

EXAMPLE 1. Consider the network in Figure 1 and assume that $q_0 = 0$. Then, $y^T = (-1, 1, -a, -(1-a))$ (with arbitrary $a \in (0, 1)$) satisfies $y^T M = 0$. One possible choice for Y is

$$Y^T = \begin{pmatrix} -1 & 1 & -1/2 & -1/2 \\ 0 & 0 & -1 & 1 \end{pmatrix},$$

so that when $q_0 = 0$

$$S(t) = \begin{pmatrix} -A_1(t) + A_2(t) - \frac{1}{2}(A_3(t) + A_4(t)) \\ -A_3(t) + A_4(t) \end{pmatrix}.$$

Then, $S_2(t) = 0$ only if $A_3(t) = A_4(t)$, which is consistent with the fact that—assuming no matchings are performed by time t —all class 3 and 4 items can be matched at that time only if the exact same number of each class is present. If $S_2(t) = 0$, then $S_1(t) = 0$ only if

$$\begin{aligned} A_2(t) &= A_1(t) + \frac{1}{2}(A_3(t) + A_4(t)) \\ &= A_1(t) + A_3(t), \end{aligned}$$

so that there are enough class 2 items to match all items of classes 1, 3, and 4 that are present in the queues at time t . \square

In general, there may be multiple choices for the matrix Y . The choice of Y does not affect our results so we do not make this dependence explicit in our notation. Regardless of how Y is chosen, its rank is $\mathbf{I} - \mathbf{J}$ (recall that $\text{rank}(M) = \mathbf{J} < \mathbf{I}$).

There are standard (algebraic ways) to construct such a matrix Y . The following is a construction of Y that has an intuitive physical interpretation and is rooted in the greater relative importance of some item classes. In Figure 1, for example, class 2 is such a class. It is the only class that is used in more than one matching so that, by allocating items of this class between the two distinct matchings A and B , the system manager can control item departures. In this sense, class 2 is a “resource” for class 1, 3, and 4 “items”. Formalizing this idea is conceptually useful and facilitates an explicit construction of the matrix Y .

DEFINITION 2. We say that a class $i \in \mathcal{I}$ is a *resource class* if $|\mathcal{J}(i)| > 1$ (it participates in more than one matching). We let $\mathbb{S} \subset \mathcal{I}$ be the set of resource classes and let $\mathbb{C} = \mathcal{I} \setminus \mathbb{S}$ be the remaining classes, which we refer to as the *customer classes*.

The algorithm shown below can be used to construct an appropriate matrix Y (that is, one whose rows span (3.3)). In the presentation of this algorithm, we use the notation r to refer to a resource class and c to refer to a customer class. Then, for a row vector y in the matrix Y^T (which has \mathbf{I} entries), the notation y_r (y_c) refers to the position in the vector y associated with that resource (customer) class. For easier understanding, we illustrate each step by applying it to Example 1.

(0) Start with an “empty” matrix Y^T that has $\mathbf{I} - \mathbf{J}$ rows and \mathbf{I} columns.

$$Y^T = \begin{pmatrix} x & x & x & x \\ x & x & x & x \end{pmatrix}$$

- (1) For each resource class $r \in \mathbb{S}$:
- (i) Add a row vector y that has a 1 in the entry associated with that resource class.

$$r = 2 \text{ and } y = (x, 1, x, x)$$

- (ii) For each matching $j \in \mathcal{J}(r)$:

For each customer class $c \in \mathbb{C} \cap \mathcal{I}(j)$:

$$y_c = \frac{-1}{|\mathbb{C} \cap \mathcal{I}(j)|}.$$

$$\mathcal{J}(r) = \{A, B\}$$

$$\mathbb{C} \cap \mathcal{I}(A) = \{1\} \text{ and so } y = (-1, 1, x, x)$$

$$\mathbb{C} \cap \mathcal{I}(B) = \{3, 4\}, |\mathbb{C} \cap \mathcal{I}(B)| = 2,$$

$$\text{and so } y = (-1, 1, -1/2, -1/2)$$

- (2) For each matching $j \in \mathcal{J}$ having $|\mathbb{C} \cap \mathcal{I}(j)| \geq 2$, fill one of the remaining rows in the matrix Y with the vector y constructed as follows:
- (i) Arbitrarily designate a customer class $c(j) \in \mathbb{C} \cap \mathcal{I}(j)$.
- (ii) Let $y_{c(j)} = -1$, $y_{c'} = 1$ for all $c' \in \mathbb{C} \cap \mathcal{I}(j)$, $c' \neq c(j)$, and set all remaining entries to 0.

$$j = B, c(B) = 3, c' = 4, \text{ and } y = (0, 0, -1, 1).$$

In the application of the last step to the network in Example 1, note that matching A has a single customer class so that this step applies only to matching B . Overall, we generate exactly $\mathbf{I} - \mathbf{J}$ rows, each of size \mathbf{I} . To see this, note that step (1) creates $|\mathbb{S}|$ rows. Then, step (2) creates $|\mathbb{C}| - \mathbf{J}$ rows, because each matching j yields $|\mathbb{C}(j)| - 1$ rows and there is no overlap between local inputs. Finally, the independence of the columns of Y (rows of Y^T) follows immediately by construction as does the fact that $Y^T M = 0$.

4. The proposed discrete-review control. Consider first the case that $q_0 = 0$: the initial queues are empty. A modification of the algorithm to account for general initial conditions is provided at the end of this section.

At each decision epoch $0 < t_1 \leq t_2 \leq \dots \leq t_m \leq u$, we solve for the number of matches of each type d_m and the resulting queues q_m that minimize the instantaneous holding cost; that is, we solve for (d_m, q_m) in

$$(4.1a) \quad \min \quad C(q_m)$$

$$(4.1b) \quad \text{s.t.} \quad q_m = Q(t_{m-1}) + A(t_m) - A(t_{m-1}) - M d_m,$$

$$(4.1c) \quad q_m, d_m \geq 0.$$

The proposed control D_\star and queue-length processes are

$$D_\star(t_m) = D_\star(t_{m-1}) + d_m \text{ and } Q(t_m, D_\star) = q_m.$$

From (4.1), it is straightforward to see that the constructed D_\star is piecewise constant, satisfies (2.2)–(2.5) and is thus admissible.

The optimization (4.1) is exactly (1.5) in the Introduction because $\Delta A(t_m) = 0$ for all m almost surely when arrivals are Poisson, so that

$$Q(t_{m-1}) + A(t_m) - A(t_{m-1}) = Q(t_{m-1}) + A(t_m^-) - A(t_{m-1}) = Q(t_m^-).$$

In particular, substitution shows that the constraints in (4.1) and (1.5) are identical. This is important because the formulation (1.5) does not require tracking the arrival increments $A(t_m) - A(t_{m-1})$.

To see the connection of our proposed control with the imbalance process introduced in §3, we multiply by Y^T in (4.1b). Then, recalling that $Y^T M = 0$ and the imbalance process definition in (3.4), we find

$$\begin{aligned} Y^T q_m &= Y^T (Q(t_{m-1}) + A(t_m) - A(t_{m-1})) \\ &= Y^T Q(t_{m-1}) + S(t_m) - S(t_{m-1}). \end{aligned}$$

If $Y^T Q(t_{m-1}) = S(t_{m-1})$, then $Y^T q_m = S(t_m)$. This is true at time $t_0 = 0$ and will subsequently hold for all m by the definition of the algorithm. Then, (4.1) can be equivalently written as

$$\begin{aligned} (4.2a) \quad & \min && C(q_m) \\ (4.2b) \quad & \text{s.t.} && Y^T q_m = S(t_m), \\ (4.2c) \quad & && q_m = Q(t_{m-1}) + A(t_m) - A(t_{m-1}) - M d_m, \\ (4.2d) \quad & && q_m, d_m \geq 0. \end{aligned}$$

For each m the problem (4.2) has the lower bound given in (1.4) of the introduction (set $t = t_m$):

$$\begin{aligned} (4.3a) \quad & \min && C(q) \\ (4.3b) \quad & \text{s.t.} && Y^T q = S(t) \\ (4.3c) \quad & && q \geq 0. \end{aligned}$$

The source of the potential higher cost under the proposed discrete-review matching control in (4.2) and the lower bound in (4.3) comes from the constraint (4.2c), that prevents previously matched items from being “re-matched” more advantageously.

In particular, a lower bound on what the algorithm can achieve (or, in fact, on what any algorithm can achieve) is given by

$$\begin{aligned}
 (4.4a) \quad & \min \int_0^u C(q(s))ds, \\
 (4.4b) \quad & \text{s.t. } Y^T q = S(t), \text{ for all } 0 \leq t \leq u, \\
 (4.4c) \quad & q(t) \geq 0, \text{ for all } 0 \leq t \leq u,
 \end{aligned}$$

Obviously, given a function $\mathcal{Q} : \mathcal{M} \rightarrow \mathbb{R}_+^I$ defined as

$$(4.5) \quad \mathcal{Q}(s) \in \operatorname{argmin}_{q \geq 0} \{C(q) : Y^T q = s\},$$

$\mathcal{Q}(S(t))$ is optimal for (4.3) and, given a sample path $(S(t), t \geq 0)$, the sample path $(\mathcal{Q}(S(t)), t \geq 0)$ is optimal for (4.4).

Thus, the imbalance-based process $\mathcal{Q}(S(t))$ generates a lower bound on the original finite horizon cumulative cost objective (2.6). To understand how our algorithm overcomes the gap between the original formulation and this lower bound it is useful to formalize an equivalence between (4.4) and the original cumulative finite horizon cost formulation (2.6).

Eliminate first the requirement in (2.6) that only integer numbers of items can be matched. The original problem formulation then becomes:

$$\begin{aligned}
 (4.6) \quad & \min \int_0^u C(Q(t; D)) dt \\
 & \text{s.t. } Q(t) = q_0 + A(t) - MD(t), \text{ for all } 0 \leq t \leq u, \\
 & D(0) = 0, D \text{ is increasing,} \\
 & \int_0^\infty \mathbf{1}\{Q_i(t-) + \Delta A_i(t) = 0 \text{ for any } i \in \mathcal{I}(j)\} dD_j(t) = 0, \\
 & Q(t) \geq 0, \text{ for all } 0 \leq t \leq u.
 \end{aligned}$$

The imbalance-based problem formulation is:

$$\begin{aligned}
 (4.7) \quad & \min \int_0^u C(q(s))ds, \\
 & \text{s.t. } Y^T q(t) = S(t), \text{ for all } 0 \leq t \leq u, \\
 & A(t) - A(s) \geq q(t) - q(s), \text{ for all } 0 \leq s \leq t \leq u, \\
 & q(t) \geq 0, \text{ for all } 0 \leq t \leq u.
 \end{aligned}$$

A function $(q(t), t \geq 0)$ is said to be admissible for (4.7) if it is RCLL and satisfies all the constraints. The following establishes that, under the dedicated item condition (Assumption 1), these problems are equivalent.

THEOREM 1. *Suppose that Assumption 1 holds. If (Q, D) is an admissible solution for (4.6), then, Q is admissible for (4.7). Conversely, if Q is an admissible solution for (4.7), then there exists a process D such that (Q, D) is admissible for (4.6).*

The simple proof appears at the end of this section. Theorem 1 tells us that, to study the original problem formulation, it suffices to study the imbalance formulation (4.7). Note that the formulation (4.7) is the lower-bound problem (4.4) with the added constraint $A(t) - A(s) \geq q(t) - q(s)$; in other words, that constraint is the source of any potential cost gap. For the solution $\mathcal{Q}(S(t))$ to be feasible for (4.7) we need that

$$A(t) - A(s) \geq \mathcal{Q}(S(t)) - \mathcal{Q}(S(s)), \quad 0 \leq s \leq t \leq u.$$

This is the key observation underlying our proposed control. By spacing out the decision epochs far enough we can guarantee that, with high probability,

$$A(t_{m+1}) - A(t_m) \geq \mathcal{Q}(S(t_{m+1})) - \mathcal{Q}(S(t_m)).$$

Our control will then be able to track the trajectory of queues $\mathcal{Q}(S(t))$ at decision epochs; that is, the solutions to (4.2) and (4.3) will be the same. Assuming that the decision epochs can still be spaced close enough so that the cost build-up during a review period is negligible, our proposed discrete review matching control (4.2) will achieve a near minimum cumulative finite horizon cost.

In summary, the essential element is that there are sufficiently many arrivals during sufficiently short review periods.

REMARK 1. It is important that if (Q, D) is a feasible solution to (4.6), then the feasibility of Q for (4.7) does not require the DI condition. It follows, in particular, that regardless of any assumptions the optimal value in (4.7) serves as a lower bound for (4.6) and, in turn, for (2.6). \square

REMARK 2 (when \mathcal{Q} is not unique). If, for each $x \geq 0$, the solution $\mathcal{Q}(x)$ in (4.5) is guaranteed to be unique (say, if $C(\cdot)$ is strictly convex), the explicit form of \mathcal{Q} is not required in order to use our algorithm and generate the optimality results that follow. In the absence of such uniqueness we modify the algorithm and use the explicit \mathcal{Q} .

At review epoch t_m , if there exists $d_m \geq 0$ such that $(d_m, \mathcal{Q}(S(t_m)))$ satisfies (4.1c) (and, in turn, is optimal for (4.1a)–(4.1c)) choose this solution, i.e, set $Q(t_m, D_\star) = \mathcal{Q}(S(t_m))$ and $D_\star(t_m) = D_\star(t_{m-1}) + d_m$.

In words, at review epoch t_m the algorithm chooses $\mathcal{Q}(S(t_m))$ whenever feasible. \square

PROOF OF THEOREM 1. Let (Q, D) be an admissible solution for (4.6). Then, the first equation in (4.6) together with the fact that D is an increasing process guarantee that $A(t) - A(s) \geq Q(t) - Q(s)$ for all $0 \leq s \leq t \leq u$. Also, since $Y^T M = 0$ by construction, we have that $Y^T Q(t) = Y^T(q_0 + A(t)) - Y^T M D(t) = Y^T(q_0 + A(t)) = S(t)$. Hence, Q is feasible for (4.7). Next, we will show that if Q is a solution to (4.7) then there exists a process D such that (Q, D) is a solution to (4.6). We construct the process $D(t)$ as follows: Let $x(t) = q_0 + A(t) - Q(t)$, where q_0 is defined to equal $Q(0)$. Note that since Q is a solution to (4.7) we have, in particular, that $x(t) \geq 0$ and $Y^T x(t) = 0$. Using Lemma 1 let $D(t)$ be the unique non-negative solution to $M D(t) = x(t)$. We claim that the process $D(t)$ constructed this way is, in fact, increasing. Indeed, by construction $M(D(t) - D(s)) = x(t) - x(s)$. Since $x(t) - x(s) \geq 0$ (by the second constraint in (4.7)) and since $Y^T x(t) = Y^T x(s) = 0$, we have by Lemma 1 that $D(t) - D(s)$ must be the unique (non-negative solution) to this system. Thus, $D(t)$ is increasing. Finally, since $A(t)$ is RCLL and so is, by definition $Q(t)$, they both have a finite number of discontinuity points on any finite interval. Thus, to show that the third constraint in (4.6) holds, it suffices to show that if $(s, t]$ is an interval such that $Q_i(u-) + \Delta A_i(u) = 0$ for all $u \in (s, t]$ then $D_j(t) - D_j(s) = 0$ for all $j \in \mathcal{J}(i)$. Since $Q \geq 0$ and A is an increasing pure jump process we have on this interval that $Q(s) = 0$ for all $s < t$ and $\Delta A_i(u) = 0$ for all $u \in (s, t]$. In particular $A_i(t) - A_i(s) = 0$. Suppose to reach a contradiction that there exist $j \in \mathcal{J}(i)$ with $D_j(t) - D_j(s) > 0$. In this case $(M(D(t) - D(s)))_i > 0$ and, by our construction of D , $Q_i(t) - Q_i(s) = A_i(t) - A_i(s) - (M(D(t) - D(s)))_i < 0$ which, since $Q(s) = 0$ for all $s < t$ is a contradiction to the non-negativity of Q . \square

General initial conditions. We end this section with a modification of the algorithm to accommodate general initial conditions. Such a modification is needed only if at time 0 $q_0 \neq \mathcal{Q}(Y^T q_0)$. Otherwise, if $q_0 = \mathcal{Q}(Y^T q_0)$, then the algorithm can be used as presented.

Let

$$(4.8) \quad t_0 = \min \{t > 0 : q_0 + A(t) - \mathcal{Q}(Y^T(q_0 + A(t))) > 0\}$$

At time t_0 , $x = q_0 + A(t) - \mathcal{Q}(Y^T(q_0 + A(t))) \geq 0$ and $Y^T x = 0$ so that, applying Lemma 1, there exists $d \geq 0$ that solves

$$M d = q_0 + A(t_0) - \mathcal{Q}(Y^T(q_0 + A(t_0))),$$

and we can set $D_\star(t_0) = d$ to obtain $Q(t_0) = \mathcal{Q}(Y^T(q_0 + A(t_0)))$. From here, we can proceed as in our original algorithm.

5. Asymptotic optimality for balanced networks. We consider a sequence of systems, indexed by n , in which the arrivals are accelerated: $\lambda_i^n(t) = n\lambda_i(t)$ is the instantaneous arrival rate of class i items at time t , and $\Lambda_i^n(t) = n\Lambda_i(t)$ is the mean cumulative number of class i item arrivals. We assume, without loss of generality, that $\lambda_{max} = 1$, so that n is interpreted as the maximal aggregate arrival rate over the time horizon. The review epochs are

$$(5.1) \quad t_m^n = t_0^n + m \left(\frac{1}{n}\right)^{2/3} \quad \text{for } m = 1, 2, \dots,$$

and we have at most $\lfloor un^{2/3} \rfloor$ review epochs on $[0, u]$. If $t_0^n = 0$, then $t_m^n = m(1/n)^{2/3}$. Our convention is to superscript with n any process or quantity associated with the n^{th} network. Thus, for example, q_0^n is the initial queue-length vector in the n^{th} network. It is standard to construct non-stationary Poisson processes from unit-rate Poisson processes $(A_i, i \in \mathcal{I})$ as follows

$$A_i^n(t) = A_i(n\Lambda_i(t)), \quad \text{for } i \in \mathcal{I}.$$

Given a control D^n , the queue process Q^n is constructed as in (2.2)–(2.5). With some abuse of terminology we henceforth say that a sequence $\{D^n\}$ is an admissible control if D^n is admissible for each n (i.e., D^n satisfies (2.2)–(2.5)).

From the functional central limit theorem for renewal processes and the random time change theorem it follows, when the underlying Poisson processes are independent, that

$$(5.2) \quad \hat{A}^n = \frac{A^n - \Lambda^n}{\sqrt{n}} \Rightarrow \hat{A} = B \circ \Lambda, \quad \text{as } n \rightarrow \infty,$$

where B is a standard \mathbf{I} -dimensional Brownian motion. In turn,

$$(5.3) \quad \hat{S}^n = \frac{S^n - Y^T q_0^n - Y^T \Lambda^n}{\sqrt{n}} \Rightarrow \hat{S} = Y^T \hat{A}, \quad \text{as } n \rightarrow \infty.$$

When the system is fluid balanced, $Y^T \Lambda^n = 0$. Then, for continuous \mathcal{Q} , and assuming $q_0^n/\sqrt{n} \rightarrow \hat{q}_0$, as $n \rightarrow \infty$,

$$\frac{1}{\sqrt{n}}\mathcal{Q}(S^n) = \mathcal{Q}(Y^T q_0^n/\sqrt{n} + \hat{S}^n) \Rightarrow \mathcal{Q}(Y^T \hat{q}_0 + \hat{S}).$$

In particular, the lower-bound cost is of the order of \sqrt{n} . The fluid queues would be 0 in this case, so that \sqrt{n} is the cost of stochasticity. Following the standard notion of asymptotic optimality, we say that a control is asymptotically optimal if its optimality gap is negligible relative to \sqrt{n} .




FIG 4. An example with one resource class and a one dimensional matrix Y .

DEFINITION 3. The control $\{D_\star^n\}$ is asymptotically optimal if

$$(5.4) \quad \frac{1}{\sqrt{n}} \left| \int_0^u C(Q^n(t, D_\star^n)) dt - \int_0^u C(Q(S^n(t))) dt \right| \Rightarrow 0, \text{ as } n \rightarrow \infty.$$

The informal weak-convergence argument in the previous paragraph suggests that our proof of asymptotic optimality hinges on the following continuity assumption. Below, \mathcal{M} is as in (3.5) and, recall, $|v|$ is the Euclidean norm of a vector v .

ASSUMPTION 2 (Lipschitz selection). *There exists a function $Q(\cdot)$ and a constant κ , such that, for each $s \in \mathcal{M}$, $Q(s)$ is a minimizer in (4.5) and*

$$|Q(s^1) - Q(s^2)| \leq \kappa |s^1 - s^2|, \quad s^1, s^2 \in \mathcal{M}.$$

EXAMPLE 2. The matching network in Figure 4 has the matching matrix

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

and it clearly satisfies the DI assumption. The unique (up to multiplicative constant) vector that spans \mathcal{Y} is $y = (-1, 1, -1)^T$. Consider the cost function

$$C(q) = \sum_{i \in \mathcal{I}} c_i q_i^m$$

for $m \geq 2$. The function C is convex and strictly increasing if $c_i > 0$, $i = 1, 2, 3$. Next, the solution to

$$\operatorname{argmin}_{q \geq 0} \{C(q) : -q_1 + q_2 - q_3 = s\}$$

for $s \in \mathcal{M}$, is

$$Q_i(s) = [-s]^+ \left(\frac{(2c_i)^{-\frac{1}{m-1}}}{(2c_1)^{-\frac{1}{m-1}} + (2c_3)^{-\frac{1}{m-1}}} \right), \quad i = 1, 3,$$

and

$$Q_2(s) = [s]^+,$$

so that Q is, in particular, Lipschitz continuous in s (regardless of m). \square

THEOREM 2. *Fix $u > 0$ and suppose that Assumptions 1 and 2 hold, that $C(\cdot)$ is Lipschitz continuous, and that the network is fluid-balanced. If $q_0^n/\sqrt{n} \rightarrow \hat{q}_0$, then,*

(i) *For any admissible control $\{D^n\}$ and all n ,*

$$(5.5) \quad C(Q^n(t, D^n)) \geq C(Q(S^n(t))), \text{ for all } t \geq 0 \text{ almost surely.}$$

(ii) *For the proposed control $\{D_\star^n\}$, as $n \rightarrow \infty$,*

$$(5.6) \quad \begin{aligned} \frac{1}{\sqrt{n}} \|Q^n - Q(S^n)\|_{t_0^n, u} &\Rightarrow 0, \\ \frac{1}{\sqrt{n}} \|C(Q^n(\cdot, D_\star^n)) - C(Q(S^n(\cdot)))\|_{t_0^n, u} &\Rightarrow 0, \end{aligned}$$

and $\{D_\star^n\}$ is asymptotically optimal.

REMARK 3 (General arrival processes). We restricted Theorem 2 to Poisson arrivals as this facilitates covering stationary and non-stationary arrivals in one result by having the weak convergence (5.2). However, neither the Poisson assumption nor independence between the \mathbf{I} components of A^n is required. The key requirement in the proof of Theorem 2 is that the bounds in Lemma A.2 in our appendix hold. This would be the case, for example, if A_1^n, \dots, A_I^n are renewal processes with finite 5th moment for the inter-arrival time (see the proof of Lemma A.2 and the references therein). \square

We postpone the proof of Theorem 2 to Section 7, after we have considered the case of non-balanced networks.

6. Asymptotic optimality for non-balanced networks. The fluid optimization problem is

$$(6.1) \quad \begin{aligned} \min \quad & \int_0^u C(q^n(s)) ds \\ \text{s.t.} \quad & q^n(t) = q_0^n + \Lambda^n(t) - Mz^n(t), \text{ for all } 0 \leq t \leq u, \\ & z^n(0) = 0, z^n \text{ is increasing,} \\ & q^n(t) \geq 0, \text{ for all } 0 \leq t \leq u. \end{aligned}$$

When the network is balanced there exists a solution $(z^n(t), t \geq 0)$ to $Mz^n(t) = \Lambda^n(t)$ so that if $q_0^n = o(n)$, the queues can be kept small in fluid scale. For networks that are not fluid-balanced, the optimal queue-length fluid trajectory is non zero and is, specifically, constructed using \mathcal{Q} as follows: multiplying the first constraint by Y^T on both sides and using $Y^T M = 0$, results in the equation $Y^T q^n(t) = Y^T(q_0^n + \Lambda^n(t))$. With

$$(6.2) \quad \bar{q}^n(t) = \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t))),$$

$\int_0^u C(\bar{q}^n(s))ds$ provides then a lower bound on the optimal cost for (6.1). Further, if \bar{q}^n is such that

$$(6.3) \quad \bar{q}^n(t) - \bar{q}^n(s) \leq \Lambda^n(t) - \Lambda^n(s), \quad s \leq t \leq u,$$

then $x_{s,t} = \Lambda^n(t) - \Lambda^n(s) - (\bar{q}^n(t) - \bar{q}^n(s)) \geq 0$ satisfies $Y^T x_{s,t} = 0$ so that by Lemma 1, there exists a non-negative solution $d_{s,t}$ to $Md = x_{s,t}$. Constructing z^n by setting $z^n(0) = 0$ and $z^n(t) - z^n(s) = d_{s,t} \geq 0$, we have that (\bar{q}^n, z^n) is feasible and optimal for (6.1).

What we seek to achieve with our control is to track the stochastic fluctuations of the optimal trajectory around the fluid

$$\mathcal{Q}(S^n(t)) - \mathcal{Q}(E[S^n(t)]) = \mathcal{Q}(S^n(t)) - \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t))).$$

To that end, we impose conditions that guarantee that \mathcal{Q} is well behaved in the vicinity of $Y^T(q_0^n + \Lambda^n(t))$. For simplicity of exposition, we focus for the remainder of this section on **stationary arrivals, i.e., on the case that $\lambda(t) \equiv \lambda > 0$** , so that $\Lambda^n(t) = n\lambda t$. The proofs apply to non-stationary arrivals but the conditions in that case are less transparent and we relegate them to the [appendix](#).

ASSUMPTION 3 (Contraction). *There exists a function \mathcal{Q} that satisfies Assumption 2 and a constant $\eta < 1$ such that*

$$\mathcal{Q}(Y^T \lambda) \leq \eta \lambda.$$

ASSUMPTION 4 (Homogeneous cost function). *The function $C(\cdot)$ is homogeneous; that is, there exists δ such that for all $x \in \mathbb{R}_+$ and all $\kappa > 0$ $C(\kappa x) = \kappa^\delta C(x)$.*

Notice that, if the network is fluid balanced, $Y^T \lambda = 0$ and $\mathcal{Q}(Y^T \lambda) = 0$, so that Assumption 3 is trivially satisfied. In general, whether Assumption 3 holds or not is a property of both the cost function $C(\cdot)$ and the rate λ . Verifying this condition merely requires solving the optimization problem (4.5) at the single point $s = Y^T \lambda$. Consider, for example, the network Figure 4 with linear holding costs $C(x) = h^T x$ with $h_1 = 0$ and $h_2 = h_3 > 0$. Suppose that arrivals are stationary with instantaneous rate $\lambda = (2, 1, 1)$. Then, $\mathcal{Q}(Y^T \lambda) = (2, 0, 0) \not\prec (2, 1, 1)$ so that Assumption 3 is violated. In this same example, if there exists $\epsilon \in (0, 1)$ such that $\lambda_2 = (1 + \epsilon)$ and $\lambda_1 = 2, \lambda_3 = 1 < \lambda_2$ then $\mathcal{Q}(Y^T \lambda) = (2 - \epsilon, 0, 0) < (2, 1 + \epsilon, 1)$ and the assumption is satisfied.

The intuition behind this requirement is as follows: to keep the queue of class 3 small, the ‘‘capacity’’ of resource class 2, λ_2 , must be strictly greater than the input of class 3. The reason the inequality is strict is that the controller must have enough flexibility to guard against stochastic fluctuations in the arrival process that result in no class 2 jobs being present when class 3 jobs need them.

Note that with $\bar{q}_0^n = 0$, Assumption 3 implies that

$$\begin{aligned} \bar{q}^n(t) - \bar{q}^n(s) &= \mathcal{Q}(Y^T n \lambda t) - \mathcal{Q}(Y^T n \lambda s) = n(t - s) \mathcal{Q}(Y^T \lambda) \\ &< \eta n(t - s) \lambda = \Lambda^n(t) - \Lambda^n(s). \end{aligned}$$

The fact that

$$(6.4) \quad |\bar{q}_i^n(t) - \bar{q}_i^n(s)| \leq \eta(\Lambda_i^n(t) - \Lambda_i^n(s)),$$

which, in words, means that optimal fluid queue of each item class increases slower than the arrivals, is the only consequence of Assumption 3 that we use in our proofs. If (6.4) can be verified directly by computing \mathcal{Q} , then the requirement that arrivals are time-homogeneous is not needed. **In fact, homogeneity of the cost function, as in Assumption 4, is not necessary.** More general (but less transparent) conditions are specified in the appendix; see Remark 5.

THEOREM 3. *Suppose that the network is not balanced. If, in addition to the Assumptions of Theorem 2, $\lambda(t) \equiv \lambda$ and Assumptions 3 and 4 hold, then the conclusions of Theorem 2 continue to hold.*

Non-Lipschitz cost functions. The assumption that C is Lipschitz continuous in Theorems 2 and 3 can be relaxed. But that relaxation does not

come for free. The difference between the minimum achievable cost and the lower bound cost may be larger than \sqrt{n} , and the definition of asymptotic optimality must be modified accordingly. Assuming that \mathcal{Q} is Lipschitz continuous,

$$\begin{aligned} \mathcal{Q}(S^n(t)) - \bar{q}^n(t) &= \mathcal{Q}(S^n(t)) - \mathcal{Q}(Y^T(q_0^n + A^n(t))) \\ &= \mathcal{Q}(q_0^n + \Lambda^n(t) + \sqrt{n}\widehat{S}^n(t)) - \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t))) \\ &\approx \sqrt{n}\widehat{S}^n(t) \approx \sqrt{n}. \end{aligned}$$

In turn, the lower bound cost should satisfy, in order of magnitude, that

$$C(\mathcal{Q}(S^n(t))) - C(\bar{q}^n(t)) \approx \mathcal{L}(\bar{q}^n(t)) (\mathcal{Q}(S^n(t)) - \bar{q}^n(t)) \approx \sqrt{n}\mathcal{L}(\bar{q}^n(t)),$$

where

$$\mathcal{L}(\kappa) = \sup_{q, \tilde{q}: |q|, |\tilde{q}| \leq \kappa} \frac{|C(\tilde{q}) - C(q)|}{|\tilde{q} - q|},$$

is the local Lipschitz constant of the convex function C . A control is asymptotically optimal if its distance from the optimum is negligible compared to the “cost of stochasticity” $\sqrt{n}\mathcal{L}(\bar{q}^n(t))$.

DEFINITION 4. The control $\{D_\star^n\}$ is asymptotically optimal if, given $\epsilon > 0$, there exists K such that, for all $\delta > 0$,

$$(6.5) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \left| \int_0^u C(Q^n(t, D_\star^n)) dt - \int_0^u C(\mathcal{Q}(S^n(t))) dt \right| > \delta \bar{L}_u^n(K) \right\} \leq \epsilon,$$

where

$$\bar{L}_u^n(K) = \mathcal{L}(K\sqrt{n} + \|\bar{q}^n\|_u).$$

REMARK 4 (Lipschitz C). When the cost function C is Lipschitz continuous, $\bar{L}_u^n(K) \leq \beta$ for all n and K , where β is the Lipschitz constant of C . In this case, asymptotic optimality reduces to the simpler requirement that, for all $\epsilon, \delta > 0$,

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \left| \int_0^u C(Q^n(t, D_\star^n)) dt - \int_0^u C(\mathcal{Q}(S^n(t))) dt \right| > \delta \right\} \leq \epsilon,$$

which is consistent with Definition 3. \square

THEOREM 4. Theorems 2 and 3 hold for non-Lipschitz cost functions with (ii) in Theorem 2 replaced by

(ii) For the proposed control $\{D_\star^n\}$, as $n \rightarrow \infty$,

$$(6.6) \quad \frac{1}{\sqrt{n}} \|Q^n - \mathcal{Q}(S^n)\|_{t_0^n, u} \Rightarrow 0 \text{ as } n \rightarrow \infty.$$

Furthermore, given $\epsilon > 0$, there exists K such that for all $\delta > 0$

$$(6.7) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \|C(Q^n(\cdot, D_\star^n)) - C(Q(S^n(\cdot)))\|_{t_0^n, u} > \delta \bar{L}_u^n(K) \right\} \leq \epsilon,$$

and D_\star^n is asymptotically optimal in the sense of Definition 4.

To make the connection between the bound in (6.7) in Theorem 4 and that in Theorems 2 and 3 more concrete, recall that for Lipschitz continuous C $\limsup_{n \rightarrow \infty} \bar{L}_u^n(K) < \infty$ and the optimality gap is $o_P(\sqrt{n})$. If, alternatively, one has separable costs of the form $C(q) = \sum_i c_i q_i^m$ then, $\bar{L}_u^n(0) = m(\|\bar{q}^n\|_u)^{m-1}$. Here, if $\|\bar{q}^n\|_u = 0$, we have a gap that is $o_P(\sqrt{n}^m)$, but if $\|\bar{q}^n\|_u \approx n$ (as one expects, e.g., in the non-balanced case), the optimality gap is $o_P(n^{m-1}\sqrt{n})$.

7. Proof essentials. We prove Theorem 4, which is the most general statement of Theorems 2 and 3. The key to its proof is the following lemma showing that the queue length under the proposed control tracks the path of $Q(S(t))$. Let $r^n(u)$ be the number of review epochs by time u . That is, $r^n(u) = \max\{m \in \mathbb{N} : t_m^n < u\}$. Then $r^n(u) \leq r^n = \lfloor un^{2/3} \rfloor$ and $r^n = \lfloor un^{2/3} \rfloor$ if $t_0^n = 0$.

LEMMA 2. Fix $u \geq 0$ and suppose that the conditions of Theorem 4 hold. Then, under the proposed control $\{D_\star^n\}$,

$$(7.1) \quad \liminf_{n \rightarrow \infty} \mathbb{P} \{Q^n(t_m^n) = Q(S^n(t_m^n)) \text{ for all } m = 1, \dots, r^n(u)\} = 1,$$

and for any $\epsilon > 0$ there exists $K(\epsilon)$ and $t_0(\epsilon)$ such that

$$(7.2) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \{t_0^n \geq t_0(\epsilon)/\sqrt{n}\} \leq \epsilon,$$

$$(7.3) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \{\|Q^n - Q(S^n)\|_{t_0^n} \geq K(\epsilon)\sqrt{n}\} \leq \epsilon,$$

and

$$(7.4) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \{\|Q^n\|_u \vee \|Q(S^n)\|_u \geq K(\epsilon)\sqrt{n} + \|\bar{q}^n\|_u\} \leq \epsilon.$$

The proof of Lemma 2 requires Assumptions 3 and 4, and this is the only place those assumptions are used in the proof of Theorem 4.

PROOF OF THEOREM 4. From Remark 1, the further lower bound (4.4) and its solution (4.5) it follows that for all n and all $t \geq 0$,

$$C(Q^n(t, D^n)) \geq C(Q(S^n(t))),$$

which immediately proves (5.5).

In what follows we write $Q^n(\cdot)$ instead of $Q^n(\cdot, D_\star^n)$. We first prove that (6.6) implies (6.7) and the asymptotic optimality. Fix $\epsilon > 0$ and define the event

$$\mathcal{A}^n(\epsilon) = \left\{ \|Q^n\|_u \vee \|\mathcal{Q}(S^n)\|_u \leq K(\epsilon/16)\sqrt{n} + \|\bar{q}^n\|_u, \right. \\ \left. \|Q^n - \mathcal{Q}(S^n)\|_{t_0^n} \leq K(\epsilon/16)\sqrt{n}, t_0^n \leq t_0(\epsilon/16)/\sqrt{n} \right\},$$

with $K(\cdot)$ and $t_0(\cdot)$ as in Lemma 2. Then, $\mathbb{P}\{\mathcal{A}^n(\epsilon)\} \geq 1 - 3\epsilon/16 \geq 1 - \epsilon/4$ for all sufficiently large n and, on $\mathcal{A}^n(\epsilon)$,

$$\|C(Q^n) - C(\mathcal{Q}(S^n))\|_{t_0^n, u} \leq \bar{L}_u^n(K(\epsilon/16)) \|Q^n - \mathcal{Q}(S^n)\|_{t_0^n, u},$$

so that for all $\delta > 0$

$$\mathbb{P} \left\{ \frac{1}{\sqrt{n}} \|C(Q^n) - C(\mathcal{Q}(S^n))\|_{t_0^n, u} > \frac{\delta}{2} \bar{L}_u^n(K(\epsilon/16)) \right\} \\ \leq \mathbb{P}\{(\mathcal{A}^n(\epsilon))^c\} + \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \|Q^n - \mathcal{Q}(S^n)\|_{t_0^n, u} > \frac{\delta}{2} \right\},$$

and from (6.6) it follows that

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \|C(Q^n) - C(\mathcal{Q}(S^n))\|_{t_0^n, u} > \frac{\delta}{2} \bar{L}_u^n(K(\epsilon/16)) \right\} \leq \frac{\epsilon}{2},$$

which establishes (6.7). In turn,

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \left| \int_{t_0^n}^u C(Q^n(t, D_\star^n)) dt - \int_{t_0^n}^u C(\mathcal{Q}(S^n(t))) dt \right| \right. \\ \left. > \frac{\delta}{2} u \bar{L}_u^n(K(\epsilon/16)) \right\} \leq \frac{\epsilon}{2}.$$

Since this is true for all $\delta > 0$, it is also true (replacing δu without δ) that

$$(7.5) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \left| \int_{t_0^n}^u C(Q^n(t, D_\star^n)) dt - \int_{t_0^n}^u C(\mathcal{Q}(S^n(t))) dt \right| \right. \\ \left. > \frac{\delta}{2} \bar{L}_u^n(K(\epsilon/16)) \right\} \leq \frac{\epsilon}{2},$$

for all $\delta > 0$. For asymptotic optimality it remains to prove that, for all $\delta > 0$,

$$(7.6) \quad \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \frac{1}{\sqrt{n}} \left| \int_0^{t_0^n} C(Q^n(t, D_\star^n)) dt - \int_0^{t_0^n} C(\mathcal{Q}(S^n(t))) dt \right| \right. \\ \left. > \frac{\delta}{2} \bar{L}_u^n(K(\epsilon/16)) \right\} \leq \frac{\epsilon}{2}.$$

On $\mathcal{A}^n(\epsilon)$, however,

$$\begin{aligned} \frac{1}{\sqrt{n}} \int_0^{t_0^n} |C(Q^n(t, D_\star^n)) - C(Q(S^n(t)))| dt \\ \leq \frac{1}{\sqrt{n}} t_0(\epsilon/16) K(\epsilon/16) \bar{L}_u^n(K(\epsilon/16)) \\ \leq \frac{\delta}{2} \bar{L}_u^n(K(\epsilon/16)), \end{aligned}$$

where n is taken sufficiently large so that $\frac{1}{\sqrt{n}} t_0(\epsilon/16) K(\epsilon/16) \leq \delta/2$. This proves (7.6) and, together with (7.5), establishes asymptotic optimality.

It only remains to prove (6.6). To that end, note that

$$\|Q^n - Q(S^n)\|_{t_0^n, u} = \max_{m=0,1,2,\dots,r^n(u)} \sup_{t_m^n \leq t \leq t_{m+1}^n} |Q^n(t) - Q(S^n(t))|.$$

Since matches are only made at review epochs we have that

$$Q^n(t) = Q^n(t_m^n) + A^n(t) - A^n(t_m^n),$$

for $t \in [t_m^n, t_{m+1}^n)$. Thus,

$$\begin{aligned} \frac{1}{\sqrt{n}} \|Q^n - Q(S^n)\|_{t_0^n, u} \\ = \max_{m=0,1,2,\dots,r^n(u)} \sup_{t_m^n \leq t \leq t_{m+1}^n} \frac{1}{\sqrt{n}} \left| Q^n(t_m^n) - Q(S^n(t_m^n)) + A^n(t) - A^n(t_m^n) \right. \\ \left. + Q(S^n(t_m^n)) - Q(S^n(t)) \right|. \end{aligned}$$

We note that $A^n(t) - A^n(s) = \Lambda^n(t) - \Lambda^n(s) + \sqrt{n}(\hat{A}^n(t) - \hat{A}^n(s))$. Recalling that $\lambda_{max} = 1$, we have $|\Lambda^n(t) - \Lambda^n(s)| \leq n\lambda_{max}(t - s) = n(t - s)$. Since $|t - t_m^n| \leq t_{m+1}^n - t_m^n < n^{-2/3}$ for $t \in [t_m^n, t_{m+1}^n]$ it follows that

$$\begin{aligned} \frac{1}{\sqrt{n}} |A^n(t) - A^n(t_m^n)| &\leq \left| \hat{A}^n(t_{m+1}^n) - \hat{A}^n(t_m^n) \right| + n^{1/2}(t_{m+1} - t_m^n) \\ &= \left| \hat{A}^n(t_{m+1}^n) - \hat{A}^n(t_m^n) \right| + n^{-1/6}. \end{aligned}$$

This results in the upper bound

$$(7.7) \quad \frac{1}{\sqrt{n}} \|Q^n - Q(S^n)\|_{t_0^n, u} \leq \max_{m=0,1,2,\dots,r^n(u)} \frac{1}{\sqrt{n}} |Q^n(t_m^n) - Q(S^n(t_m^n))|$$

$$(7.8) \quad + \max_{m=0,1,2,\dots,r^n(u)} \left| \widehat{A}^n(t_{m+1}^n) - \widehat{A}(t_m^n) \right| + n^{-1/6}$$

$$(7.9) \quad + \max_{m=0,1,2,\dots,r^n(u)} \sup_{t_m^n \leq t \leq t_{m+1}^n} \frac{1}{\sqrt{n}} |\mathcal{Q}(S^n(t_m^n)) - \mathcal{Q}(S^n(t))|.$$

To complete the proof, we argue that each of the terms (7.7)–(7.9) converges weakly to zero. The term (7.7) converges in probability to 0 by (7.1) in Lemma 2. The weak convergence of \widehat{A}^n to a continuous limit process in (5.2) implies that the term (7.8) converges weakly to 0. Finally, to prove that (7.9) converges weakly to 0, note that for $t \in [t_m^n, t_{m+1}^n)$, letting $\bar{y} = \max_{kl} |Y_{kl}|$

$$\begin{aligned} |S^n(t_m^n) - S^n(t)| &= |S^n(t_m^n) - S^n(t) - nY^T(\Lambda(t_m^n) - \Lambda(t))| \\ &\quad + |Y^T n(\Lambda(t) - \Lambda(t_m^n))| \\ &\leq |S^n(t_m^n) - S^n(t) - nY^T(\Lambda(t_m^n) - \Lambda(t))| + 2\bar{y}n^{1/3} \\ &\leq \sqrt{n} \left| \widehat{S}^n(t_m^n) - \widehat{S}^n(t) \right| + 2\bar{y}n^{1/3}, \end{aligned}$$

and we used again the fact $\Lambda(t) - \Lambda(t_m^n) \leq (t - t_m^n) \leq n^{-2/3}$ for all $t \in [t_m^n, t_{m+1}^n)$. The Lipschitz continuity of \mathcal{Q} (Assumption 2) then implies

$$\frac{1}{\sqrt{n}} |\mathcal{Q}(S^n(t_m^n)) - \mathcal{Q}(S^n(t))| \leq \kappa \left(\left| \widehat{S}^n(t_m^n) - \widehat{S}^n(t) \right| + 2\bar{y}n^{-1/6} \right),$$

so that

$$\begin{aligned} &\max_{m=0,1,2,\dots,r^n(u)} \sup_{t_m^n \leq t \leq t_{m+1}^n} \frac{1}{\sqrt{n}} |\mathcal{Q}(S^n(t_m^n)) - \mathcal{Q}(S^n(t))| \\ &\leq \max_{m=0,1,2,\dots,r^n(u)} \sup_{t_m^n \leq t \leq t_{m+1}^n} \kappa \left(\left| \widehat{S}^n(t_m^n) - \widehat{S}^n(t) \right| + 2\bar{y}n^{-1/6} \right) \end{aligned}$$

The right-hand side converges to 0 by the weak convergence of \widehat{S}^n to a continuous limit (5.3), and we conclude that (7.9) converges weakly to 0 and, in turn, that (6.6) holds. \square

8. Numerical experiments. For our experiments, we use the networks in Figures 1 and 3(ii). We refer to these as network *I* and network *II* respectively. The DI condition holds for both networks. We consider separable quadratic cost of the form $C(q) = \sum_i c_i q_i^2$ with the coefficients $c^T = (2, 1, 5, 7)$ for network I and $c^T = (3, 1, 5, 7, 0)$ for network II.

We fix the horizon to be $[0, 1]$. Given a sample path of the arrivals $(A(t), 0 \leq t \leq 1)$ we generate the corresponding sample paths of $(\mathcal{Q}(S(t)), 0 \leq t \leq 1)$ and of $(Q(t, D_\star), 0 \leq t \leq 1)$. In the figures, $q_i^L = (\mathcal{Q}_i(S(t), t \geq 0))$,

is the trajectory of the lower bound queue trajectory and $q_i^a = (Q_i(t, D_\star), t \geq 0)$ is the trajectory of queue i under the proposed algorithm. We also compute the associated costs $\int_0^1 C(Q(S(t)))dt$ and $\int_0^1 C(Q(t, D_\star))dt$. The fact that the gap between q_i^L and q_i^a and the gap between the costs are all small is the numerical manifestation of our mathematical result in Theorem 4. Our experiments show an impressive precision when the proposed solution is applied to a concrete network.

Figures 5 and 6 depict the sample paths for the stationary balanced and non-balanced networks. Assumption 4 holds for this example since C is a homogenous function. It is also checked (by solving numerically for $Q(Y^T \lambda)$) that Assumption 3 holds for Network I but is violated for Network II. Nevertheless, the algorithm performs well in both settings. The costs appear in a box in the top-left corner. Each figure lists also the arrival-rates that are used.

Finally, we consider a non-stationary setting with sinusoidal arrival-rate functions of the form $\lambda_i(t) = a_i + b_i \sin(c_i t)$. The arrival rates are specified by three vectors $a^T = (a_1, \dots, a_I)$, $b^T = (b_1, \dots, b_I)$ and $c^T = (c_1, \dots, c_I)$. All sample paths are initialized with empty queues at time $t = 0$. The result is displayed in Figure 7. While verifying the conditions for non-stationary non-balanced networks (as in the appendix) is complex, it is evident that the algorithm performs extremely well in the test settings.

9. Concluding remarks. We have introduced a matching queue model. The control question of interest is how to match items in order to minimize holding costs. We established that a simple myopic discrete review matching control performs well, both analytically (by proving asymptotic optimality) and numerically (through simulation), for the broad class of networks that satisfy the dedicated item condition. The central idea of this work is the identification of an imbalance process that facilitates the construction of a lower bound and the design of a policy that asymptotically achieves this bound.

The analogue in parallel-server networks to the imbalance process is the workload process. In that setting, the workload process and workload formulations have played a central role in solving a variety of queueing control problems. A natural question is, then, whether the imbalance process can play a similar role in solving a variety of matching control problems. In what follows, we hint to how it could.

Imbalance-based control problems. The purpose of the discussion here is to illustrate how the results of this paper may be leveraged towards the control of more elaborate matching models. We do not pursue the construction of




FIG 5. Balanced (TOP) Network I: $\lambda^T = (300, 600, 300, 300)$, (BOTTOM) Network II: $\lambda^T = (300, 600, 600, 300, 300)$.

a formal framework or the proof of asymptotic optimality. Our derivation is purely formal and intended to illustrate how the imbalance process can potentially be used to construct a lower dimension control problem in other matching models.




FIG 6. Non-balanced (TOP) Network I: $\lambda^T = (450, 600, 200, 250)$, (BOTTOM) Network II: $\lambda^T = (350, 750, 500, 230, 270)$.

A problem of input regulation/admission control serves this illustration purpose. Suppose that arrivals of class i follow a (possibly non-stationary) Poisson process $E_i = (E_i(t), t \geq 0)$ with instantaneous rate $\lambda_i(t)$. In contrast to our original model assume here that the controller may accept some of




FIG 7. Non-stationary (TOP) Network I: $a^T = (300, 600, 300, 300)$, $b^T = (75, 150, 50, 25)$, $c^T = (15, 10, 5, 10)$ (BOTTOM) Network II: $a^T = (300, 600, 600, 300, 300)$, $b^T = (75, 100, 75, 25, 0)$, $c^T = (12.5, 15, 10, 5, 0)$.

the arriving items and reject others upon arrival. Let $R_i = (R_i(t), t \geq 0)$ be a process that counts the number of class i items rejected by time t . This is a control process.

Suppose there is a penalty p_i for rejecting a class i item. The relevant control problem is to minimize

$$\int_0^u C(Q(t, R, D))dt + p'R(u),$$

where we now have two controls – an admission control R and a matching control D . Let

$$A_i(t) = E_i(t) - R_i(t)$$

be the cumulative class i items that enter the system by time t . The constraints (2.2)–(2.5) remain valid, and the constraint that R_i can increase only in points of increase of E_i must be added.

In the same spirit, the definition of the imbalance process remains

$$S^R(t) = Y^T(q_0 + A(t)).$$

It is then natural to consider the following analogue of the imbalance formulation (4.7)

$$\begin{aligned} \min \quad & \int_0^u C(q(s))ds + p'R(u), \\ \text{s.t.} \quad & Y^T q(t) = S^R(t), \text{ for all } 0 \leq t \leq u, \\ & A(t) - A(s) \geq q(t) - q(s), \text{ for all } 0 \leq s \leq t \leq u, \\ & E(t) - E(s) \geq R(t) - R(s) \text{ for all } 0 \leq s \leq t \leq u, \\ & q(t) \geq 0, \text{ for all } 0 \leq t \leq u, \end{aligned}$$

Using, instead of $S^R(t)$, the control invariant process $S(t) = q_0 + E(t)$, we re-write the above as

$$\begin{aligned} \min \quad & \int_0^u C(q(s))ds + p'R(u), \\ \text{s.t.} \quad & Y^T(q(t) + R(t)) = S(t), \text{ for all } 0 \leq t \leq u, \\ & A(t) - A(s) \geq q(t) - q(s), \text{ for all } 0 \leq s \leq t \leq u, \\ & E(t) - E(s) \geq R(t) - R(s), \text{ for all } 0 \leq s \leq t \leq u, \\ & q(t) \geq 0, \text{ for all } 0 \leq t \leq u, \end{aligned}$$

which leads to the lower bound

$$\begin{aligned} \min \quad & \int_0^u C(q(s))ds + p'R(u), \\ \text{s.t.} \quad & Y^T(q(t) + R(t)) = S(t), \text{ for all } 0 \leq t \leq u, \\ & q(t) \geq 0, \text{ for all } 0 \leq t \leq u, \\ & R(t) \leq E(t), \text{ for all } 0 \leq t \leq u. \end{aligned}$$

Despite the addition of the constraint on $R(t)$, a useful re-shuffling interpretation is maintained. Our lower bound, recall, had the interpretation of allowing the controller to take corrective actions by “unmatching” items that were already matched. Here, we are allowing to “un-reject” items that were previously rejected. The solution to this lower bound may not be as simple as (4.4). Nevertheless, it is a significant simplification relative to the original formulation.

From a solution point of view, it stands to reason that the following modification of the action in the review epochs will generate near optimal solutions:

$$\begin{aligned} \min \quad & C(q_m) + p' r_m \\ \text{s.t.} \quad & Y^T(q_m + \sum_{k=1}^m r_k) = S(t_m), \\ & q_m = Q(t_{m-1}) + E(t_m) - E(t_{m-1}) - r_m - Md_m, \\ & r_m \leq E(t_m) - E(t_{m-1}), \\ & d_m, q_m, r_m \geq 0. \end{aligned}$$

A formal study of this heuristic derivation and, more broadly, of a variety of control problems for matching networks seems a promising (and challenging) direction for future research.

APPENDIX

PROOF OF LEMMA 1. First note that every $x \geq 0$ that satisfies $Y^T x = 0$ (as in the statement of the lemma) is in the image of M . This follows from the fact that Y spans the orthogonal space to M and thus x (being orthogonal to Y) must be in the column space of M : there must exist d such that $Md = x$. The result is now immediate from the DI condition. As noted in the paragraph following Assumption 1, with the DI condition $Md = x \geq 0$ implies $d \geq 0$. \square

PROOF OF LEMMA 2. For $x \in \mathcal{M}_t^n = \{y \in \mathbb{R}^{\mathbf{I}-\mathbf{J}} : Y^T(q_0^n + \Lambda^n(t)) + y \in \mathcal{M}\}$, define

$$\mathcal{Q}_t^n(x) = \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t)) + x), \text{ and } \widehat{\mathcal{Q}}_t^n(x) = \mathcal{Q}_t^n(x) - \bar{q}^n(t), \text{ and } x \in \mathcal{M}_t^n.$$

Recall that $\bar{q}^n(t) = \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t)))$, so that $\widehat{\mathcal{Q}}_t^n(x)$ captures the effect of second order perturbations around the optimal fluid trajectory; see (A.11) below.

LEMMA A.1. *Suppose that there exists a function $c(t), t \geq 0$ such that $\lambda(t) = \lambda c(t)$ for some vector $\lambda > 0$ and $0 < \bar{c}_{\min} < c(t) < \bar{c}_{\max} < \infty$. Under Assumption 4, there exists $\gamma > 0$ so that*

$$(A.2) \quad |\widehat{Q}_t^n(x) - \widehat{Q}_s^n(y)| \leq \gamma|x - y| + g(s, t)(|x| + |y|),$$

for all $x \in \mathcal{M}_t^n, y \in \mathcal{M}_s^n$ and $0 < s \leq t \leq u$ where $g(\cdot, \cdot)$ is such that,

$$(A.3) \quad \bar{g} = \sup_{s \leq t \leq u} g(s, t) < \infty,$$

and for any function $h(\epsilon) > 0$ satisfying $\epsilon^{-3/4}h(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$,

$$(A.4) \quad \lim_{\epsilon \rightarrow 0} \sup_{t \in [h(\epsilon), u]} \epsilon^{-1/4}g(t - \epsilon, t) = 0.$$

Note that, by setting $c(t) = t$, Lemma A.1 applies in particular to stationary arrivals as assumed in §6.

REMARK 5 (Beyond the stated assumptions). In all that follows, the only properties of the primitives $C(\cdot)$ and $\lambda(\cdot)$ that we use are (A.2)–(A.4) and (6.4). If the network is balanced, (A.2) (with $g \equiv 0$) follows trivially from the Lipschitz continuity of \mathcal{Q} in Assumption 2 and recalling that $\bar{q}^n(t) = 0$ in that case. The latter also guarantees (6.4). Consequently, Lemma 2 applies, in particular, to the setting of Theorem 2 which is, in turn, a corollary of the proof of Theorem 4.

More broadly, If \mathcal{Q} is given in closed form and (A.2)–(A.4) and (6.4) can be verified directly, Lemma A.1 and the proof of Theorem 4 apply. \square

We also require the following auxiliary lemma.

LEMMA A.2. *For any finite constant $\alpha > 0$,*

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{m=1, \dots, r^n(u)} |A^n(t_{m+1}^n) - A^n(t_m^n) - n(\Lambda(t_{m+1}^n) - \Lambda(t_m^n))| \leq \alpha n^{1/3} \right\} = 1,$$

and

$$(A.5) \quad \liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{m=1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) \sqrt{n} \left(|\widehat{S}^n(t_{m+1}^n)| + |\widehat{S}^n(t_m^n)| \right) \leq \alpha n^{1/3} \right\} = 1.$$

The proofs of Lemmas A.1–A.2 are found at the end of this [appendix](#).

Let η be as in Assumption 3 and choose α so that

$$\alpha < \frac{\lambda_{\min}(1 - \eta)}{2 + \gamma \bar{y} \mathbf{I}},$$

where γ is the Lipschitz constant in Lemma A.1 and $\bar{y} = \max_{kl} |Y_{kl}|$. For each n , define the events

$$\mathcal{A}_1^n = \left\{ \max_{m=1, \dots, r^n(u)} |A^n(t_{m+1}^n) - A^n(t_m^n) - n(\Lambda(t_{m+1}^n) - \Lambda(t_m^n))| \leq \alpha n^{1/3} \right\}.$$

$$\mathcal{A}_2^n = \left\{ \max_{m=1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) \sqrt{n} \left(|\widehat{S}^n(t_{m+1}^n)| + |\widehat{S}^n(t_m^n)| \right) \leq \alpha n^{1/3} \right\}.$$

Let

$$\mathcal{A}^n = \mathcal{A}_1^n \cap \mathcal{A}_2^n.$$

It then follows from Lemma A.2 that, given $\epsilon > 0$,

$$\mathbb{P}(\mathcal{A}^n) \geq 1 - \epsilon.$$

for all sufficiently large n .

We claim, and will establish shortly, that on \mathcal{A}^n ,

$$(A.6) \quad A_i^n(t_{m+1}^n) - A_i^n(t_m^n) - (\mathcal{Q}_i(S^n(t_{m+1}^n)) - \mathcal{Q}_i(S^n(t_m^n))) \geq 0, \\ i \in \mathcal{I}, n \in \mathbb{N}, m = 1, \dots, r^n(u).$$

Assuming that (A.6) holds, equation (7.1) is argued now by induction on the review epochs starting with $m = 1$. Recall the optimization problem (4.2)

$$(A.7) \quad \min \quad C(q_m)$$

$$(A.8) \quad \text{s.t.} \quad Y^T q_m = S^n(t_m^n)$$

$$(A.9) \quad q_m = Q^n(t_{m-1}^n) + A^n(t_m^n) - A^n(t_{m-1}^n) - M d_m,$$

$$(A.10) \quad q_m, d_m \geq 0.$$

At time t_0^n we have by construction (recall (4.8)) that $Q^n(t_0^n) = \mathcal{Q}(S^n(t_0^n))$. At t_1^n we choose (q_1, d_1) that satisfy

$$q_1 = \mathcal{Q}(S^n(t_1^n)) \text{ and } M d_1 = A^n(t_1^n) - A^n(t_0^n) - (\mathcal{Q}(S^n(t_1^n)) - \mathcal{Q}(S^n(t_0^n))).$$

The vector q_1 is, by definition, optimal for (A.7)–(A.10) (recall (4.5)). Further, setting $x = A^n(t_1^n) - A^n(t_0^n) - (\mathcal{Q}(S^n(t_1^n)) - \mathcal{Q}(S^n(t_0^n)))$ we have by (A.6) that $x \geq 0$ and $Y^T x = 0$ so that by Lemma 1, d_1 as above exists. Thus, (q_1, d_1) is feasible and optimal for (A.7)–(A.9).

Assuming next that

$$Q^n(t_{m-1}^n) = Q(S^n(t_{m-1}^n)).$$

for all $m = 0, \dots, k-1$ (and, in particular for $m = k-1$), we choose (q_k, d_k) so that

$$\begin{aligned} q_k &= Q(S^n(t_k^n)) \\ Md_k &= A^n(t_k^n) - A^n(t_{k-1}^n) - (Q(S^n(t_k^n)) - Q(S^n(t_{k-1}^n))). \end{aligned}$$

Such a choice is feasible building again on (A.6) and Lemma 1. This concludes the induction argument.

It remains to prove (A.6). By definition

$$\begin{aligned} Q(S^n(t)) &= Q(Y^T(q_0^n + \Lambda^n(t)) + \sqrt{n}\widehat{S}^n(t)) \\ &= Q_t^n(\sqrt{n}\widehat{S}^n(t)) \\ &= \widehat{Q}_t^n(\sqrt{n}\widehat{S}^n(t)) + \bar{q}^n(t), \end{aligned}$$

for all $t \geq 0$, where Q_t^n and \widehat{Q}_t^n are as defined prior to Lemma A.1. By (A.2) we have

$$\begin{aligned} (A.11) \quad & |Q(S^n(t_{m+1}^n)) - Q(S^n(t_m^n)) - (\bar{q}^n(t_{m+1}^n) - \bar{q}^n(t_m^n))| \\ &= |\widehat{Q}_t^n(\sqrt{n}\widehat{S}^n(t_{m+1}^n)) - \widehat{Q}_t^n(\sqrt{n}\widehat{S}^n(t_m^n))| \\ &\leq \gamma\sqrt{n} \left| \widehat{S}^n(t_{m+1}^n) - \widehat{S}^n(t_m^n) \right| + g(t_m^n, t_{m+1}^n)\sqrt{n} \left(|\widehat{S}^n(t_m^n)| + |\widehat{S}^n(t_{m+1}^n)| \right). \end{aligned}$$

The second element on the right hand side is bounded by $\alpha n^{1/3}$ on \mathcal{A}^n . Furthermore, on \mathcal{A}^n ,

$$\begin{aligned} \sqrt{n} \left| \widehat{S}^n(t_{m+1}^n) - \widehat{S}^n(t_m^n) \right| &= |S^n(t_{m+1}^n) - S^n(t_m^n) - nY^T(\Lambda(t_{m+1}^n) - \Lambda(t_m^n))| \\ &\leq \bar{y} |A^n(t_{m+1}^n) - A^n(t_m^n) - n(\Lambda(t_{m+1}^n) - \Lambda(t_m^n))| \\ &\leq \mathbf{I}\alpha\bar{y}n^{1/3}, \end{aligned}$$

We conclude that, on \mathcal{A}^n ,

$$(A.12) \quad |Q(S^n(t_{m+1}^n)) - Q(S^n(t_m^n)) - (\bar{q}^n(t_{m+1}^n) - \bar{q}^n(t_m^n))| \leq \alpha n^{1/3}(1 + \gamma\bar{y}\mathbf{I}),$$

and, in turn,

$$\begin{aligned} & A_i^n(t_{m+1}^n) - A_i(t_m^n) - (Q_i(S^n(t_{m+1}^n)) - Q_i(S^n(t_m^n))) \\ & \geq n(\Lambda_i(t_{m+1}^n) - \Lambda_i(t_m^n)) - (\bar{q}^n(t_{m+1}^n) - \bar{q}^n(t_m^n)) - \alpha n^{1/3}(2 + \gamma\bar{y}\mathbf{I}) \end{aligned}$$

$$\begin{aligned} &\geq (1 - \eta)n(\Lambda_i(t_{m+1}^n) - \Lambda_i(t_m^n)) - \alpha n^{1/3}(2 + \gamma\bar{y}\mathbf{I}) \\ &\geq ((1 - \eta)\lambda_{\min} - \alpha(2 + \gamma\bar{y}\mathbf{I}))n^{1/3}. \end{aligned}$$

on \mathcal{A}^n where $\eta < 1$ in the second inequality is as in (6.4). We conclude that (A.6) holds and, in turn, so does (7.1).

We next prove (7.2). Recall the definition of t_0^n in (4.8). By subtracting and adding terms we have, for all $t \geq 0$,

$$\begin{aligned} (A.13) \quad q_0^n + A^n(t) - \mathcal{Q}(S^n(t)) &\geq \bar{q}^n(0) + \Lambda^n(t) - \bar{q}^n(t) \\ &\quad - |A^n(t) - \Lambda^n(t)| - |q_0^n - \bar{q}^n(0)| \\ &\quad - |\mathcal{Q}(S^n(t)) - \bar{q}^n(t)| \end{aligned}$$

Recalling that $S^n(t) = Y^T(q_0^n + A^n(t))$ and $\bar{q}^n(t) = \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t)))$ we have, by Assumption 2, that

$$\begin{aligned} |\mathcal{Q}(S^n(t)) - \bar{q}^n(t)| &= |\mathcal{Q}(Y^T(q_0^n + A^n(t))) - \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t)))| \\ &\leq \kappa\bar{y}|A^n(t) - \Lambda^n(t)|, \end{aligned}$$

for all $t \geq 0$. By the FCLT for the Poisson process it then holds that, given $\epsilon > 0$, there exists $K(\epsilon)$ such that

$$\mathbb{P}\{\|A^n - \Lambda^n\|_u \geq K(\epsilon)\sqrt{n}\} \leq \epsilon.$$

By the conditions of the theorem $|q_0^n - \bar{q}^n(0)| \leq c\sqrt{n}$ for some constant c and we conclude that, given $\epsilon > 0$, there exists a (re-defined) constant $K(\epsilon)$ such that

$$\begin{aligned} \mathbb{P}\{q_0^n + A^n(t) - \mathcal{Q}(S^n(t)) \geq \bar{q}^n(0) + \Lambda^n(t) - \bar{q}^n(t) - K(\epsilon)\sqrt{n}, \text{ for all } t \leq u\} \\ \geq 1 - \epsilon. \end{aligned}$$

Next, by (6.4), $|\bar{q}^n(t) - \bar{q}^n(0)| \leq \eta\Lambda^n(t)$. Thus, for all $t \geq 0$,

$$\begin{aligned} \bar{q}^n(0) + \Lambda^n(t/\sqrt{n}) - \bar{q}^n(t/\sqrt{n}) &\geq \bar{q}^n(0) + \Lambda^n(t/\sqrt{n})(1 - \eta) \\ &\geq n\lambda_{\min}t(1 - \eta)/\sqrt{n}, \end{aligned}$$

and we can choose $t_0(\epsilon)$ such that $\bar{q}^n(0) + \Lambda^n(t + 0/\sqrt{n}) - \bar{q}^n(t_0/\sqrt{n}) > K(\epsilon)\sqrt{n}$. In turn,

$$(A.14) \quad \mathbb{P}\{t_0^n \geq t_0(\epsilon)/\sqrt{n}\} \leq \epsilon$$

which proves (7.2). We turn to (7.3).

As no actions are taken on $[0, t_0^n)$, we have that $\sup_{s \leq t_0^n} |Q^n(s) - q_0^n| \leq |A^n(t_0^n)|$. In turn, for each $K > 0$,

$$\mathbb{P}\{\sup_{s \leq t_0^n} |Q^n(s) - q_0^n| > K\sqrt{n}\} \leq \mathbb{P}\{|A^n(t_0^n)| \geq K\sqrt{n}\}.$$

Using (7.2), it follows from the functional strong law for A^n that, given $\epsilon > 0$, there exists $K_1(\epsilon)$ satisfying

$$\begin{aligned} \mathbb{P}\{|A^n(t_0^n)| \geq K(\epsilon)\sqrt{n}\} &\leq \mathbb{P}\{t_0^n > t_0(\epsilon/2)/\sqrt{n}\} \\ &\quad + \mathbb{P}\left\{\|A^n\|_{t_0(\epsilon/2)/\sqrt{n}} \geq K_1(\epsilon)\sqrt{n}\right\} \leq \epsilon/2. \end{aligned}$$

Since $|q_0^n - \bar{q}^n(0)| = |q_0^n - \mathcal{Q}(Y^T q_0^n)| \leq c\sqrt{n}$ by the assumption of the theorem, we then have the existence of a (re-defined) constant $K_1(\epsilon)$ such that

$$\mathbb{P}\{\|Q^n - \bar{q}^n(0)\|_{t_0^n} > K_1(\epsilon)\sqrt{n}\} \leq \epsilon/2.$$

Recall that $\bar{q}^n(0) = \mathcal{Q}(Y^T(q_0^n + \Lambda^n(t)))$ so that by Assumption 2

$$\|\mathcal{Q}(S^n) - \bar{q}^n(0)\|_{t_0^n} \leq \|\mathcal{Q}(S^n) - \mathcal{Q}(Y^T(q_0^n + \Lambda^n))\|_u \leq \kappa\sqrt{n}\|\widehat{S}^n\|_u,$$

and it follows from the FCLT for \widehat{S}^n that for each ϵ there exists $K_2(\epsilon)$ such that

$$\mathbb{P}\{\|\mathcal{Q}(S^n) - \bar{q}^n(0)\|_{t_0^n} \geq K_2(\epsilon)\} \leq \epsilon/2.$$

Taking $K(\epsilon) = 2 \max\{K_1(\epsilon), K_2(\epsilon)\}$ we have that

$$\begin{aligned} \mathbb{P}\{\|Q^n - \mathcal{Q}(S^n)\|_{t_0^n} > K(\epsilon)\sqrt{n}\} &\leq \mathbb{P}\{\|Q^n - \bar{q}^n(0)\|_{t_0^n} > K_1(\epsilon)\sqrt{n}\} \\ &\quad + \mathbb{P}\{\|\mathcal{Q}(S^n) - \bar{q}^n(0)\|_{t_0^n} > K_2(\epsilon)\sqrt{n}\} \leq \epsilon, \end{aligned}$$

which proves (7.3).

Finally, to prove equation (7.4) recall that (6.6) in Theorem 4 does not build on (7.4) so we may use it here without having a circular argument. Note that

$$\begin{aligned} \|\mathcal{Q}(S^n)\|_u \vee \|Q^n\|_u &\leq \|\mathcal{Q}(S^n)\|_u + \|\mathcal{Q}(S^n) - Q^n\|_u \\ \text{(A.15)} \quad &\leq \|\mathcal{Q}(S^n)\|_u + \|\mathcal{Q}(S^n) - Q^n\|_{t_0^n} + \|\mathcal{Q}(S^n) - Q^n\|_{t_0^n, u} \end{aligned}$$

By (6.6), for each ϵ there exists $\bar{K}_1(\epsilon)$ such that

$$\text{(A.16)} \quad \mathbb{P}\{\|\mathcal{Q}(S^n) - Q^n\|_{t_0^n, u} > \bar{K}_1(\epsilon)\sqrt{n}\} \leq \epsilon/3.$$

By (7.3), there exists $K_2(\epsilon)$ such that

$$\text{(A.17)} \quad \mathbb{P}\{\|\mathcal{Q}(S^n) - Q^n\|_{t_0^n} > \bar{K}_2(\epsilon)\sqrt{n}\} \leq \epsilon/3.$$

Using, as before, the fact that

$$\|\mathcal{Q}(S^n) - \mathcal{Q}(Y^T(q_0^n + \Lambda^n))\|_u \leq \kappa\sqrt{n}\|\widehat{S}^n\|_u,$$

we have, by the FCLT for \widehat{S}^n , the existence of a constant $K_3(\epsilon)$ such that

$$(A.18) \quad \mathbb{P}\{\|\mathcal{Q}(S^n)\|_u \geq \bar{K}_3(\epsilon) + \|\bar{q}^n\|_u\} \leq \epsilon/3.$$

Combining (A.15)–(A.16) and letting $K(\epsilon) = 3 \max\{\bar{K}_1(\epsilon), \bar{K}_2(\epsilon), \bar{K}_3(\epsilon)\}$ we conclude that

$$\mathbb{P}\{\|\mathcal{Q}(S^n)\|_u \vee \|Q^n\|_u > K(\epsilon)\sqrt{n} + \|\bar{q}^n\|_u\} \leq \epsilon.$$

This proves (7.4) and completes the proof of the lemma. \square

PROOF OF LEMMA A.1. Assume that $q_0 = 0$, $C(\cdot)$ is homogeneous and there exists a function $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\lambda_i(t) = \lambda_i c(t)$, and $0 < \bar{c}_{\min} < c(t) < \bar{c}_{\max} < \infty$ for all $t \geq 0$. Let $\varpi(t) = \int_0^t c(s)ds$. Then, $\Lambda_i(t) = a_i \varpi(t)$.

Fix $s, t > 0$ and consider the optimization problem $\min\{C(q) : Y^T q = Y^T \Lambda(t) + x\}$, for $x \in \mathbb{R}$. Due to the assumed homogeneity of C this problem, for $t > 0$, is equivalent in terms of its optimal-solution set, to the problem $\min\{C(q) : Y^T q = Y^T \lambda + x/\varpi(t)\}$, where $\lambda = (\lambda_1, \dots, \lambda_{\mathbf{I}})$ is as in the definition of $\lambda_i(t)$. To see this, replace q with $q/\varpi(t)$ and apply the homogeneity of $C(\cdot)$. In particular, for all $t > 0$,

$$\bar{q}(t) = \mathcal{Q}(Y^T \lambda) \varpi(t), \quad \mathcal{Q}(Y^T \Lambda(t) + x) = \varpi(t) \mathcal{Q}(Y^T a + x/\varpi(t)),$$

and

$$\widehat{\mathcal{Q}}_t(x) = \mathcal{Q}(Y^T \Lambda(t) + x) - \bar{q}(t) = \varpi(t) (\mathcal{Q}(Y^T a + x/\varpi(t)) - \mathcal{Q}(Y^T \lambda)),$$

for all such t . Then, for $s, t > 0$,

$$\begin{aligned} \widehat{\mathcal{Q}}_t(x) - \widehat{\mathcal{Q}}_s(y) &= \varpi(t) (\mathcal{Q}(Y^T \lambda + x/\varpi(t)) - \mathcal{Q}(Y^T a)) \\ &\quad - \varpi(s) (\mathcal{Q}(Y^T \lambda + y/\varpi(s)) - \mathcal{Q}(Y^T \lambda)) \\ &= \varpi(s) (\mathcal{Q}(Y^T \lambda + x/\varpi(t)) - \mathcal{Q}(Y^T \lambda)) \\ &\quad - \varpi(s) (\mathcal{Q}(Y^T \lambda + y/\varpi(s)) - \mathcal{Q}(Y^T \lambda)) \\ &\quad + (\varpi(t) - \varpi(s)) (\mathcal{Q}(Y^T \lambda + x/\varpi(t)) - \mathcal{Q}(Y^T \lambda)) \end{aligned}$$

Since by Assumption 2,

$$|\mathcal{Q}(Y^T \lambda + x/\varpi(t)) - \mathcal{Q}(Y^T \lambda + y/\varpi(s))| \leq \kappa |x/\varpi(t) - y/\varpi(s)|$$

and

$$|\mathcal{Q}(Y^T \lambda + x/\varpi(t)) - \mathcal{Q}(Y^T \lambda)| \leq \kappa|x|/\varpi(t),$$

it follows that

$$\left| \widehat{\mathcal{Q}}_t(x) - \widehat{\mathcal{Q}}_s(y) \right| \leq \varpi(s)\kappa|x/\varpi(t) - y/\varpi(s)| + (\varpi(t) - \varpi(s))\kappa|x|/\varpi(t).$$

Furthermore, since

$$\begin{aligned} \varpi(s)|x/\varpi(t) - y/\varpi(s)| &= |\varpi(s)x/\varpi(t) - x + x - y| \\ &\leq \frac{|\varpi(s) - \varpi(t)|}{|\varpi(t)|}|x| + |x - y|, \end{aligned}$$

we find that

$$\left| \widehat{\mathcal{Q}}_t(x) - \widehat{\mathcal{Q}}_s(y) \right| \leq \kappa|x - y| + 2\kappa \frac{|\varpi(s) - \varpi(t)|}{|\varpi(t)|}|x|.$$

Next, $|\varpi(s) - \varpi(t)|/|\varpi(t)| \leq \bar{c}_{\max}|t - s|/\bar{c}_{\min}t$ implies

$$\left| \widehat{\mathcal{Q}}_t(x) - \widehat{\mathcal{Q}}_s(y) \right| \leq \kappa|x - y| + 2\kappa \frac{\bar{c}_{\max}}{\bar{c}_{\min}} \frac{|t - s|}{t} (|x| + |y|).$$

Finally, letting $\gamma = \kappa$ and defining for $0 < s \leq t$

$$g(s, t) = 2\kappa \frac{\bar{c}_{\max}}{\bar{c}_{\min}} \frac{|t - s|}{t},$$

we see that

$$\epsilon^{-1/4}g(t - \epsilon, t) = 2\kappa \frac{\bar{c}_{\max}}{\bar{c}_{\min}} \frac{\epsilon^{3/4}}{t} \leq \frac{\epsilon^{3/4}}{h(\epsilon)} \text{ for all } t \in [h(\epsilon), u].$$

so that the result follows recalling that, by assumption, $h(\epsilon)$ satisfies that $\epsilon^{-3/4}h(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$. \square

PROOF OF LEMMA A.2. The first part follows from [1] (see the explanation in the proof of Lemma 4.1 in [12]). We do not repeat the proof. For the second part, from strong approximation theorems follows the existence of a \mathbf{I} -dimensional Brownian motion B and a constant c (depending on u) such that

$$\limsup_{n \rightarrow \infty} \mathbb{P}\{\|A(n \cdot) - n \cdot - B(n \cdot)\|_u > cn^{1/4}(\log n)\} = 0.$$

Recalling that $\widehat{A}^n(t) = (A(n\Lambda(t)) - n\Lambda(t))/\sqrt{n}$ and that $\Lambda(t) \leq t$ ($\lambda_{max} = 1$), we have

$$\left\| \widehat{A}^n(\cdot) - \frac{B(n\Lambda(\cdot))}{\sqrt{n}} \right\|_u \leq \frac{1}{\sqrt{n}} \|A(n \cdot) - n \cdot - B(n \cdot)\|_{\lambda_{max}u},$$

so that

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \sup_{0 \leq t \leq u} \sqrt{n} |\widehat{A}^n(t) - B(n\Lambda(t))/\sqrt{n}| > cn^{1/4}(\log n) \right\} = 0;$$

see e.g. [9, Theorem 2.1]. Since $\widehat{S}^n = Y^T \widehat{A}^n$ (5.3), we have

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \sqrt{n} \|\widehat{S}^n(\cdot) - Y^T B(n\Lambda(\cdot))/\sqrt{n}\|_u > cn^{1/4}(\log n) \right\} = 0.$$

For simplicity of notation we let $B^n(t) = B(n\Lambda(t))/\sqrt{n}$. Then, for each $m = 1, \dots, r^n(u)$,

$$\begin{aligned} \sqrt{n} \left(|\widehat{S}^n(t_{m+1}^n)| + |\widehat{S}^n(t_m^n)| \right) &\leq \sqrt{n} |Y^T B^n(t_m^n)| + \sqrt{n} |Y^T B^n(t_{m+1}^n)| \\ &\quad + \sqrt{n} |Y^T B^n(t_{m+1}^n) - \widehat{S}^n(t_{m+1}^n)| \\ &\quad + \sqrt{n} |Y^T B^n(t_m^n) - \widehat{S}^n(t_m^n)| \\ &\leq \sqrt{n} \bar{y} \left(|B^n(t_{m+1}^n)| + |B^n(t_m^n)| \right) \\ &\quad + 2\sqrt{n} \|\widehat{S}^n - Y^T B^n\|_u. \end{aligned}$$

where, as before, $\bar{y} = \max_{kl} |Y_{kl}|$.

For all sufficiently large n , $cn^{1/4} \log n < (\alpha/2)n^{1/3}$ so that, since by (A.3) $\bar{g} = \sup_{s \leq t \leq u} g(s, t) < \infty$,

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{m=1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) 2\sqrt{n} \|\widehat{S}^n - Y^T B^n\|_u \leq \alpha n^{1/3} \right\} = 1,$$

and to establish (A.5) it remains to prove that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{P} \left\{ \max_{m=1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) \sqrt{n} \bar{y} \left(|B^n(t_{m+1}^n)| + |B^n(t_m^n)| \right) > \frac{\alpha}{2} n^{1/3} \right\} \\ = 0. \end{aligned}$$

Fix $0 < \eta < 1/6$ and let $m^n = \lceil n^{1/6+\eta} \rceil$. Let

$$x_m^n = g(t_m^n, t_{m+1}^n) \sqrt{n} \bar{y} \left(|B^n(t_{m+1}^n)| + |B^n(t_m^n)| \right).$$

Then,

$$(A.19) \quad \mathbb{P} \left\{ \max_{m=1, \dots, r^n(u)} x_m^n > \frac{\alpha}{2} n^{1/3} \right\} \leq \mathbb{P} \left\{ \max_{m=1, \dots, m^n} x_m^n > \frac{\alpha}{2} n^{1/3} \right\}$$

$$+ \mathbb{P} \left\{ \max_{m=m^n+1, \dots, r^n(u)} x_m^n > \frac{\alpha}{2} n^{1/3} \right\}.$$

We will argue that each of the elements on the right-hand side converges to 0 as $n \rightarrow \infty$ starting with the second element.

Define $h(\epsilon) = \epsilon^{\frac{3}{4} - \frac{3}{2}\eta}$. we observe that $\epsilon^{-3/4}h(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$ and for all $n \in \mathbb{N}$ and $m \geq m^n$, $t_m^n = \lceil n^{1/6+\eta} \rceil / n^{2/3} \geq n^{-1/2+\eta} = h(n^{-2/3})$. Since, by construction, $t_{m+1}^n - t_m^n = n^{-2/3}$, we have by (A.4) (replacing $\epsilon = n^{-2/3}$) that

$$(A.20) \quad \max_{m=m^n+1, \dots, r^n(u)} n^{1/6} g(t_m^n, t_{m+1}^n) \rightarrow 0$$

as $n \rightarrow \infty$. Also, $B^n(t) = B(n\Lambda(t))/\sqrt{n}$ has the same law as $B(\Lambda(t))$ and, in particular, since $\lambda_{max} = 1$, $\|B^n\|_u$ has the same law as $\|B\|_u$. Then,

$$\begin{aligned} & \mathbb{P} \left\{ \max_{m=m^n+1, \dots, r^n(u)} x_m^n > \frac{\alpha}{2} n^{1/3} \right\} \\ & \leq \mathbb{P} \left\{ \max_{m=m^n+1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) \sqrt{n} 2\bar{y} \|B\|_u > \frac{\alpha}{2} n^{1/3} \right\} \\ & = \mathbb{P} \left\{ \max_{m=m^n+1, \dots, r^n(u)} g(t_m^n, t_{m+1}^n) n^{1/6} 2\bar{y} \|B\|_u > \frac{\alpha}{2} \right\} \rightarrow 0 \end{aligned}$$

where the convergence follows from (A.20) and the well known fact that, $\lim_{K \rightarrow \infty} \mathbb{P}\{\|B\|_u \geq K\} = 0$ for any $u \geq 0$; see e.g. [14, Exercice II.1.23].

We turn to treat the first element on the right-hand side of (A.19). Provided that $t_0^n \leq \frac{1}{2}n^{-(1/2-\eta)}$, $t_m^n \leq 2n^{\eta-1/2}$ for all $m = 1, \dots, m^n$ and all sufficiently large n . In turn, for all sufficiently large n and on the event that $t_0^n \leq t_0/\sqrt{n}$, it holds that $|B^n(t_m^n)| + |B^n(t_{m+1}^n)| \leq 2\|B^n\|_{2n^{\eta-1/2}}$, so that, by (A.3),

$$\begin{aligned} & \mathbb{P} \left\{ \max_{m=1, \dots, m^n} g(t_m^n, t_{m+1}^n) \sqrt{n} \bar{y} (|B^n(t_{m+1}^n)| + |B^n(t_m^n)|) > \alpha n^{1/3} \right\} \\ & \leq \mathbb{P} \left\{ 2\bar{g}\bar{y}\sqrt{n}\|B\|_{2/\sqrt{n}} > \alpha n^{1/3} \right\} + \mathbb{P} \left\{ t_0^n \geq \frac{1}{2}n^{-(1/2-\eta)} \right\} \\ & = \mathbb{P} \left\{ 2\sqrt{2}n^{1/4}\bar{g}\bar{y}\|B\|_1 > \alpha n^{1/3} \right\} + \mathbb{P} \left\{ t_0^n \geq \frac{1}{2}n^{-(1/2-\eta)} \right\} \\ & = \mathbb{P} \left\{ \|B\|_1 > \frac{\alpha}{2\sqrt{2}\bar{g}\bar{y}} n^{1/12} \right\} + \mathbb{P} \left\{ t_0^n \geq \frac{1}{2}n^{-(1/2-\eta)} \right\} \rightarrow 0, \end{aligned}$$

as $n \rightarrow \infty$. In the second line of the display we use the fact that $\frac{n^{1/4}}{\sqrt{2}}B(2t/\sqrt{n})$ is equal in law to $B(t)$. The convergence to 0 of the first item on the last line

follows, as before, from the fact $\lim_{K \rightarrow \infty} \mathbb{P}\{\|B\|_u > K\} = 0$ for any $u \geq 0$. The fact that $\mathbb{P}\{t_0^n \geq \frac{1}{2}n^{1/2-\eta}\} \rightarrow 0$ follows from (7.2) in Lemma (2). Note that (7.2) does not rely on Lemma (A.2) and hence can be used here. \square

Acknowledgements. We thank Rakesh Vohra for suggestions that motivated this work as well as the SAMSI working group on patient flow for valuable discussions on the ideas in this paper. We are extremely grateful to Kibaek Kim for writing the simulation and optimization code for our numerical experiments. Finally, we are grateful to anonymous referees whose observations and suggestions led to significant improvements to the original manuscript.

REFERENCES

- [1] ATA, B. and KUMAR, S., Heavy traffic analysis of open processing networks with complete resource pooling: Asymptotic optimality of discrete review policies. *Ann. Appl. Prob.*, 2005. [MR2115046](#)
- [2] BELL, S. L. and WILLIAMS, R. J., Dynamic scheduling of a parallel server system in heavy traffic with complete resource pooling: Asymptotic optimality of a threshold policy. *Electronic Journal of Probability*, 10(33):1044–1115, 2005. [MR2164040](#)
- [3] DOGRU, M. K., REIMAN, M. I., and WANG, Q., A stochastic programming based inventory policy for assemble-to-order systems with application to the W model. *Operations Research*, 58(4):849–864, 2010. [MR2683480](#)
- [4] HARRISON, J. M., Assembly-like queues. *Journal of Appl. Prob.*, 10:354–367, 1973. [MR0356276](#)
- [5] HARRISON, J. M., The BIGSTEP approach to flow management in stochastic processing networks. In F. Kelly, S. Zachary, and I. Ziedins, editors, *Stochastic Networks: Theory and Applications*, pages 57–90. Oxford University Press, 1996.
- [6] HARRISON, J. M., Correction: Brownian models of open processing networks: Canonical representation of workload. *Ann. Appl. Prob.*, 16(3):1703–1732, 2006. [MR2260079](#)
- [7] HARRISON, J. M. and LOPEZ, M. J., Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, 33:339–368, 1999. [MR1742575](#)
- [8] HARRISON, J. M. and VAN MIEGHEM, J. A., Dynamic control of brownian networks: State space collapse and equivalent workload formulations. *Ann. Appl. Prob.*, 7:747–771, 1996. [MR1459269](#)
- [9] HORVATH, L., Strong approximation of renewal processes. *Stochastic Processes and Their Applications*, 18(1):127–138, 1984. [MR0757352](#)
- [10] KANG, W. N., KELLY, F. P., LEE, N. H., and WILLIAMS, R. J., State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. *Ann. Appl. Prob.*, 19(5):1719–1780, 2009. [MR2569806](#)
- [11] MANDELBAUM, A. and STOLYAR, S., Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. *Operations Research*, 52:836–855, 2004. [MR2104141](#)
- [12] PLAMBECK, E. L. and WARD, A. R., Optimal control of a high-volume assemble-to-order system. *Mathematics of Operations Research*, 31(3):453–477, 2006. [MR2254418](#)

- [13] REIMAN, M. I. and WANG, Q., Asymptotically optimal inventory control for assemble-to-order systems with identical lead times, 2013. Working Paper.
- [14] REVUZ, D. and YOR, M., *Continuous Martingales and Brownian Motion*, volume 293. Springer Verlag, 1999. [MR1725357](#)
- [15] SONG, J. S. and ZIPKIN, P., Supply chain operations: Assemble-to-order and configure-to-order systems. In *Handbooks in Operations Research and Management Science*, volume XXX, pages 561–593, 2003.