

Delay-Optimal Server Allocation in Multiqueue Multiserver Systems With Time-Varying Connectivities

Somsak Kittipiyakul and Tara Javidi, *Member, IEEE*

Abstract—This paper considers the problem of optimal server allocation in a time-slotted system with N statistically symmetric queues and K servers when the arrivals and channels are stochastic and time-varying. In this setting, we identify two classes of “desirable” policies with potentially competing goals of maximizing instantaneous throughput versus balancing the load. Via an example, we show that these goals, in general, can be incompatible, implying an empty intersection between the two classes of policies. On the other hand, we establish the existence of a policy achieving both goals when the connectivities between each queue and each server are random and either “ON” or “OFF.” We use dynamic programming (DP) and properties of the value function to establish the delay optimality of a policy, which, at each time-slot, simultaneously maximizes the instantaneous throughput and balances the queues.

Index Terms—Communication models, intermittent connectivity, multiqueue multiserver, optimal transmission scheduling, orthogonal frequency-division multiple access (OFDMA), queuing analysis.

I. INTRODUCTION

IN the face of frequency selectivity, code-division, or spatial degrees of freedom, many wireless systems with multiple “orthogonalized subchannels” and multiple users can be viewed as multiqueue multiserver queuing systems, which enable transmission of packets in a parallel manner. Examples of such systems include orthogonalized code-division multiple-access (CDMA) systems and orthogonal frequency-division mul-

multiple access (OFDMA) systems where the total available bandwidth is divided into multiple orthogonal narrow subcarriers to be shared by users [1]. Usually data packets in such systems arrive stochastically to each user and are stored in buffers prior to transmission. In this context, there is often a limited number of orthogonal codes or subcarriers, not allowing simultaneous transmission of all queued packets. This gives rise to a scheduling problem involving the allocation of the orthogonal channels to the different data streams.

In addition, due to selective fading, for instance, the “quality” of these channels as perceived by different users varies stochastically with time and users. This introduces the notion of stochastic server quality, also known as connectivity. In the presence of reliable estimates of channel quality, i.e., channel state information (CSI), at the transmitter, the stochastic variation across users provides opportunities for selectively scheduling the transmission among users. The opportunistic scheduling among users is known to increase the overall throughput of systems with saturated queues by providing multiuser diversity gain [2]. However, delay-optimal policies must trade off between two competing goals: the desire to get maximum throughput (MT) now (which is achieved by opportunistic scheduling) and the desire to get the MT in the future (which is achieved by balancing the remaining load). The second goal, under admissible traffic regimes, accounts for future queue occupancies; the intuitive reasoning is that the queued packets should be spread over multiple queues so that the system will have a better chance of avoiding idling any servers in the future.

The focus of this paper is in finding an optimal scheduling policy with the optimal delay performance for the above system with stochastic channel state and arrival processes. In other words, we are interested in the delay performance of the system under stochastic admissible traffic. We first model the above multichannel allocation problem as a multiqueue multiserver allocation problem in Section II. We account for stochastically varying channel states, via a general notion of connectivity. In addition, we provide a detailed review of related work. In Section III, we discuss two classes of “desirable” policies: the instantaneous throughput maximizing policies and load-balancing (LB) policies. Via a simple example, we show that, in general, the intersection of these two classes of policies can be empty. This implies a complicated structure for the optimal policy in the general case. However, in Section IV, we show that when the connectivity profile has a binary form (ON–OFF connectivity), it is possible to construct a policy,

Manuscript received June 01, 2007; revised March 05, 2008. Current version published April 22, 2009. This work was supported in part by the Center for Wireless Communications, University of California at San Diego and University of California Discovery Grant Com04-10176, by the ARO-MURI Grant W911NF-04-1-0224, by the National Science Foundation CAREER Award CNS-0533035, and by the AFOSR Grant FA9550-05-01-0430. The material in this paper was presented in part at the 2004 Asilomar Conference on Signals, Systems, and Computers and at the 2004 IEEE Conference on Decision and Control.

S. Kittipiyakul was with the Electrical and Computer Engineering Department, University of California—San Diego, La Jolla, CA 92093 USA. He is now with the School of Information, Computer and Communication Technology, Sirindhorn International Institute of Technology, Thammasat University, Pathum Thani 1212, Thailand (e-mail: somsak@siit.tu.ac.th).

T. Javidi is with the Electrical and Computer Engineering Department, University of California, San Diego, La Jolla, CA 92093 USA (e-mail: tjavidi@ucsd.edu).

Communicated by E. Modiano, Associate Editor for Communication Networks.

Color versions of Figures 1–4 in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2009.2016051




Fig. 1. Multiqueue multiserver allocation problem with time-varying connectivities.

which simultaneously maximizes the instantaneous throughput and balances the load. In Section V, we show the optimality of this maximum-throughput and load-balancing (MTLB) policy when there are $N = 2$ users, using dynamic programming (DP) arguments and the properties of the DP value function. In Section VI, we show that if fractional server allocations are allowed, then the fluid version of the MTLB policy is optimal for general N . Section VII concludes the paper and discusses the direction of future studies.

II. PROBLEM FORMULATION AND ASSUMPTIONS

A. Model and Notations

We consider a multiqueue multiserver system with stochastic connectivities as shown in Fig. 1. There are N queues (users) and K servers (subcarriers). Let $U = \{1, \dots, N\}$ be the set of all queues and $V = \{1, \dots, K\}$ the set of all servers. Fixed-size packets arrive stochastically for each user and are transmitted over a set of allocated servers. Each user has an infinite buffer to store the data packets that cannot be immediately transmitted. The system is time slotted. The users have the same priority and are symmetric, i.e., they have *statistically* identical arrival and channel connectivity processes. At the beginning of each time slot, the assignment of servers to users is instantaneous and made by a centralized resource manager. The resource manager has perfect knowledge of the current queue backlogs and the connectivities, which are assumed constant during a time slot but varying independently over time slots (e.g., block fading model). We do not allow sharing of any servers and assume no error in the transmission.

The following notations are used throughout the paper. Note that we use the following conventions: lower case letters for scalar, bold-faced lower case letters for row vectors, upper case letters for matrices, and scripted upper case letters for space of matrices.

- $\mathbf{b}(t) = (b_1, \dots, b_N)$: backlogs (in units of packets) of each queue at the beginning of time slot t .
- $\mathbf{a}(t) = (a_1, \dots, a_N)$: stochastic number of fixed-length packets arrived to each queue during time slot t . The new packet arrivals at time t can be served only at time $t + 1$ or after.
- $C(t) = [c_{ij}]$: the $K \times N$ stochastic *connectivity matrix* at time slot t where $c_{ij} \in \{0, 1, \dots, c_{\max} < \infty\}$ denotes the

maximum number of packets server i can serve from queue j at time t .

- $W(t) = [w_{ij}]$: the $K \times N$ allocation matrix at the beginning of time slot t where $w_{ij} \in \{0, 1\}$ and $w_{ij} = 1$ denotes that server i is assigned to serve queue j during time t .

The dynamics of the queue length vectors under an allocation $W(t)$ is described by

$$\mathbf{b}(t+1) = [\mathbf{b}(t) - \mathbf{1}(W(t) \odot C(t))]^+ + \mathbf{a}(t), \quad t = 1, 2, \dots \quad (1)$$

where an elementwise product $W(t) \odot C(t)$ is a matrix $[w_{ij}c_{ij}]$, $\mathbf{1}$ is a K -dimensional row vector of K ones, and, for a vector $\mathbf{v} \in \mathbb{R}^N$, $[\mathbf{v}]^+ = [v_1^+, \dots, v_N^+]$ with $v_j^+ = \max\{0, v_j\}$. For the case of the ON-OFF connectivities, where $c_{\max} = 1$, the above queue dynamics reduces to

$$\mathbf{b}(t+1) = [\mathbf{b}(t) - \mathbf{1}W(t)]^+ + \mathbf{a}(t), \quad t = 1, 2, \dots \quad (2)$$

Definition 1: For a row vector $\mathbf{x} = (x_1, \dots, x_N)$ and a matrix $Y = [\mathbf{y}_1, \dots, \mathbf{y}_N]$ where \mathbf{y}_j is a column vector, a column-by-column matrix permutation Π_π corresponding to a permutation π is defined as, for any j and $k \in \{1, \dots, N\}$

$$\pi(x_j) = x_k \Leftrightarrow \Pi_\pi(\mathbf{y}_j) = \mathbf{y}_k.$$

Using the above notations and definition, we make the following symmetric assumptions on the arrival and connectivity processes. The intuition behind a symmetric system is that relabeling the queues leads to a statistically identical system.

- A1) The packet arrival processes $\{\mathbf{a}(t)\}$ are independent identically distributed (i.i.d.) across time slots and *symmetric* or *exchangeable*, i.e., the joint probability mass function (pmf) is permutation invariant. That is

$$P[\mathbf{a}(t) = \pi(\mathbf{x})] = P[\mathbf{a}(t) = \mathbf{x}] \quad (3)$$

for any t , vector \mathbf{x} , and permutation π .

- A2) The connectivity profiles $\{C(t)\}$ are i.i.d. across time slots and exchangeable across users, i.e., the joint pmf for $\{C(t)\}$ is column-by-column permutation invariant. That is

$$P[C(t) = \Pi_\pi(Y)] = P[C(t) = Y]$$

for any t , matrix Y , and column-by-column permutation matrix Π_π .

Assumption A2) is valid when the channel and mobility create a homogeneous environment for all users. Note that A1) and A2) imply independence across time but not across users, i.e., at a given time the arrivals to various queues or connectivities among users or servers need not be independent.

B. Problem Formulation

Problem (P): Considering the system described above, we wish to determine a Markov server allocation policy σ that minimizes the cost function at the finite horizon T

$$J_T^\sigma = E[\Lambda_T^\sigma | \mathcal{I}_0] \quad (4)$$

where \mathcal{I}_0 summarizes all information available at time zero and Λ_T^σ denotes the cost under the Markov policy σ over horizon T

$$\Lambda_T^\sigma = \sum_{t=0}^T \phi(\mathbf{b}(t)) \quad (5)$$

where the cost function $\phi(\mathbf{b}) = \sum_{j=1}^N g(b_j)$ and g is a convex and strictly increasing function.

We note that the restriction to Markov policies does not entail any loss of optimality because Problem (P) is a stochastic control problem with perfect observations [17]. Also, note that when g is an identity function, Problem (P) reduces to an average total backlog ($E[\sum_{t=0}^T \sum_{j=1}^N b_j(t)]$) minimization problem over horizon T . From Little's theorem [18], any optimal policy that achieves the minimum average backlog achieves the minimum average packet delay as well. Thus, our study includes the study of the average-delay minimization.

C. Related Work

Delay-optimal policies have been studied in many queuing systems with stochastically varying connectivities under different settings [3]–[12]. In this paper, we consider a statistically symmetric case of arrival and connectivity processes. In many instances, as stated in [4], “symmetry sometimes leads to rather simple optimal policies, although their optimality can be hard to establish.” The most related models to our work are those introduced in [3] and [4], while our proof technique is closely related to those developed in [6], [7], [13], [14], and [16].

In their seminal work, Tassiulas *et al.* [3] studied a single-server N -queue assignment problem where the connectivity followed an ON–OFF model, i.e., the connectivity was described by a binary vector of dimension N . They showed that a longest-connected-queue (LCQ) policy maximizes the stability region of the system (i.e., throughput optimal) and is also average-delay optimal when the arrival processes and the channel processes are statistically identical among users, i.e., the users are symmetric. Subsequently, Ganti *et al.* [4], [5] generalized the problem to a symmetric K -server, N -queue allocation problem with binary connectivity vector of dimension N .

The model studied by Ganti *et al.* is different from our model in subtle but nonetheless important aspects. The main contribution of our work is a generalization of the models and results in [3]–[5] to 1) a connectivity model of a K -by- N matrix form, 2) more general arrival processes, not restricted to only Bernoulli (and variation thereof as discussed in Section III of [4]), and 3) a generalization where any queue can be served simultaneously by multiple servers.¹ These generalizations and extensions, however, are realized at the cost of a) restricting the number of users to $N = 2$, or b) allowing for fluid (nonintegral) server allocation. Just before the final version of the paper was submitted, the authors came across an important manuscript on monotonicity in Markov reward and decisions by Koole [15]. The main objective of [15] is to study certain classes of functions (e.g., the set \mathcal{F} in Definition 7) that propagate through the DP operator. It is clear that given an assumption on the immediate

reward function belonging to such a class of functions (e.g., Fact 1), the value function at all times also belongs to that class (e.g., see the proof of Theorem 2). Hence, one can utilize the properties of the class to deduce the structure of the optimal policy partially or completely (e.g., Lemma 2).

In light of the findings of [15], we can highlight and contrast the results and the proof techniques in our work and [4]: Class of functions \mathcal{F} defined in Definition 7 propagates through the DP operator only when the number of users is two. This resembles the difficulty in extending the results in [6] beyond two dimensions as discussed in [15]. Under fluid relaxation, however, the state space \mathbb{Z}_+^N is extended to \mathbb{R}_+^N over which the class of convex functions propagates through the DP operator, forming the basis of the proof of Theorem 3. Many of the proofs in [4] rely on stochastic coupling, an alternative framework to the monotonicity approach of [15]. However, coupling arguments remain partly intuitive and hard to check, in a general setting. This can explain the extra requirements on server allocations and arrival processes in [4].

III. INSTANTANEOUS THROUGHPUT MAXIMIZING VERSUS LOAD-BALANCING POLICIES

In this section, we consider two classes of server allocation policies: class of instantaneous throughput maximizing (MT) policies and class of load-balancing (LB) policies. As discussed in the introduction, each class represents one of the competing goals: an MT policy maximizes the number of packets being served now, while an LB policy maximizes the number of nonempty queues (hence, the multiuser diversity gain and the number of packets served) in the future. In Section III-D, we demonstrate by an example that, in general, the intersection of the two classes of policies can be empty. In such cases, the optimal policies for Problem (P) remain and can be, in general, a complicated mixture of policies carefully chosen at different time from one of the above two classes of policies. To be precise, we first define the feasible allocation and non-idling feasible allocation. Then, we describe the two classes of policies mentioned above.

A. Feasible and Nonidling Allocations

Assume that at the beginning of time slot t , the state of the system is (\mathbf{b}, C) . An allocation $W = [w_{ij}]_{K \times N}$ is a *feasible allocation* for time slot t if:

- $w_{ij} \in \{0, 1\}$;
- $c_{ij} = 0 \Rightarrow w_{ij} = 0$;
- $\sum_{j=1}^N w_{ij} \leq 1, \forall i = 1, \dots, K$.

The set of all feasible allocations is denoted by $\mathcal{W}(C)$. In addition, define $\mathcal{W}(\mathbf{b}, C) \subseteq \mathcal{W}(C)$ to denote the set of all *nonidling* feasible allocation W if W also satisfies:

- $\sum_{i=1}^K w_{ij} c_{ij} \leq b_j, \forall j = 1, \dots, N$.

B. Instantaneous MT Policies

An MT allocation $W^{\text{MT}} = [w_{ij}^{\text{MT}}] \in \mathcal{W}(\mathbf{b}, C)$ is a non-idling allocation that achieves the MT at time t if for all $W = [w_{ij}] \in \mathcal{W}(\mathbf{b}, C)$

$$\sum_{j=1}^N \sum_{i=1}^K w_{ij}^{\text{MT}} c_{ij} \geq \sum_{j=1}^N \sum_{i=1}^K w_{ij} c_{ij}. \quad (6)$$

¹The constraint on the number of servers per queue in [4] is relaxed only when a relaxation of integral allocation is also allowed.

C. Load-Balancing Policies

It is reasonable to maximize the expected number of packets served in future under stochastic arrival and connectivity processes. Under assumptions A1) and A2), this is achieved via an LB policy, which distributes the *future workload* among the queues as evenly as possible to minimize the expected future server idling. This roughly ensures a larger set of feasible non-idling allocations in the future. The future workload is defined as the queue length vector after assignment, i.e., the leftover queue vector. The LCQ policy [3] and most balanced policy [5] are some examples of the LB policies. Note that the LB policies potentially sacrifice the current throughput (by giving priority to long queues) for the future throughput.

To introduce the LB policy, we need the following definition to compare queue vectors in term of their load distribution.

Definition 2: We say $\mathbf{x} \leq_{\text{LQO}} \mathbf{y}$ (\mathbf{x} is more balanced than \mathbf{y}) iff $\text{ord}(\mathbf{x}) \leq_{\text{lex}} \text{ord}(\mathbf{y})$ where vector $\text{ord}(\mathbf{v})$ has the ordered elements of \mathbf{v} in descending order, and the relation \leq_{lex} on \mathbb{R}^N is the *lexicographic ordering*.

Example: i) $(5, 1, 4, 2) \leq_{\text{LQO}} (0, 3, 5, 4)$ because $\text{ord}(5, 1, 4, 2) = (5, 4, 2, 1) \leq_{\text{lex}} (5, 4, 3, 0) = \text{ord}(0, 3, 5, 4)$.
ii) $(3, 3) \leq_{\text{LQO}} (4, 1)$, although $(3, 3)$ has more total number of packets than $(4, 1)$.

Load Balancing: An LB allocation $W^{\text{LB}} \in \mathcal{W}(\mathbf{b}, C)$ is a nonidling allocation that produces the most balanced future (leftover) queue distribution if, for all $W \in \mathcal{W}(\mathbf{b}, C)$

$$[\mathbf{b} - \mathbf{1}(W^{\text{LB}} \odot C)]^+ \leq_{\text{LQO}} [\mathbf{b} - \mathbf{1}(W \odot C)]^+. \quad (7)$$

D. Example

Here we show that the incompatibility of the MT and LB policies exists even in a single server case.

Example: If $\mathbf{b} = (6, 2)$ and $C = (1, 2)$, then the MT allocation $W^{\text{MT}} = (0, 1)$ achieves the throughput of 2 and leaves the remaining queue highly unbalanced at $\mathbf{b} - W^{\text{MT}} = (6, 0)$. The system with this unbalanced queue state is unlikely to benefit from any multiuser diversity in the next time slot. In contrast, the LB allocation $W^{\text{LB}} = (1, 0)$ sacrifices the throughput with the balancedness of the remaining queues at $(5, 2)$. In other words, here the two goals of throughput maximization and queue balancing cannot be achieved simultaneously.

IV. SPECIAL CASE: ON-OFF CHANNEL

In this section, we consider a special case of the connectivity process where c_{ij} only takes values 0 (OFF) or 1 (ON). Under this ON-OFF connectivity, we show the following: 1) a policy that simultaneously maximizes the instantaneous throughput and balances the loads always exists, and 2) for the case of two users, this maximum-throughput and load-balancing (MTLB) policy is an optimal policy for Problem (P).

A. MTLB Policy

Here we define a class of MTLB policies specific for the ON-OFF channel connectivity.

Definition 3: Given state (\mathbf{b}, C) at the beginning of time slot t , an MTLB policy chooses a nonidling feasible allocation $W^* = [w_{ij}^*] \in \mathcal{W}(\mathbf{b}, C)$ such that it satisfies the following two conditions:

C1) MT: W^* achieves the MT, i.e., for all $W = [w_{ij}] \in \mathcal{W}(\mathbf{b}, C)$

$$\sum_{j=1}^N \sum_{i=1}^K w_{ij}^* \geq \sum_{j=1}^N \sum_{i=1}^K w_{ij}. \quad (8)$$

C2) LB: W^* produces the most balanced queue configuration, i.e., for all $W = [w_{ij}] \in \mathcal{W}(\mathbf{b}, C)$

$$\mathbf{b} - \mathbf{1}W^* \leq_{\text{LQO}} \mathbf{b} - \mathbf{1}W. \quad (9)$$

Example 1: If $C = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}$ and $\mathbf{b} = [3, 3, 2, 2]$, an MTLB allocation is $W^{\text{MTLB}} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$, resulting in the leftover queues $\mathbf{b} - \mathbf{1}W^{\text{MTLB}} = [2, 3, 1, 2]$. In addition, $\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$ is another possible MTLB allocation, resulting in the leftover queues $[3, 2, 2, 1]$. For this (\mathbf{b}, C) , there are four possible MTLB allocations. Hence, the MTLB policy is not uniquely defined. However, if $\mathbf{b} = [4, 3, 3, 2]$, then there is only one MTLB allocation $\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$.

B. Existence of MTLB Policy

Due to the ON-OFF connectivity, the following result shows that there always exists an MTLB allocation satisfying conditions (C1) and (C2) at every time slot and for any (\mathbf{b}, C) .

Theorem 1: For any given (\mathbf{b}, C) , an MTLB allocation always exists. □

Proof: See Appendix I.

C. Construction of MTLB Policy

In this section, we specifically propose an algorithm to construct an MTLB assignment. We first convert the original graph of queues and servers (Fig. 1) into the following equivalent bipartite graph with proper weights on the edges.

Equivalent Bipartite Graph Construction:

- 1) Associated with each queue j , construct $m_j = \min(b_j, \sum_{i=1}^K c_{ij})$ nodes labeled as $a_{j1}, a_{j2}, \dots, a_{jm_j}$.
- 2) Let $U^{eq} = \{a_{11}, a_{12}, \dots, a_{1m_1}, a_{21}, \dots, a_{Nm_N}\}$ be the set of all such nodes.
- 3) Let $V^{eq} = \{v_i\}_{i=1}^K$ be the set of servers.
- 4) Let $E^{eq} = \{(a_{jm}, v_i) : c_{ij} = 1\}$ be the set of edges representing connectivities.
- 5) Let $\psi : E^{eq} \mapsto \mathbb{Z}_{++}$, $\psi(a_{jt}, *) = b_j - t + 1$ be the positive integer weight of all edges incident to node a_{jt} in E^{eq} .

To arrive at an MTLB allocation, we run a maximum weight matching (MWM) algorithm on the equivalent bipartite graph. In Proposition 1, we show that the resulting assignment satisfies conditions C1) and C2), hence, it is an MTLB allocation. Before we proceed with Proposition 1, we provide the following definitions.

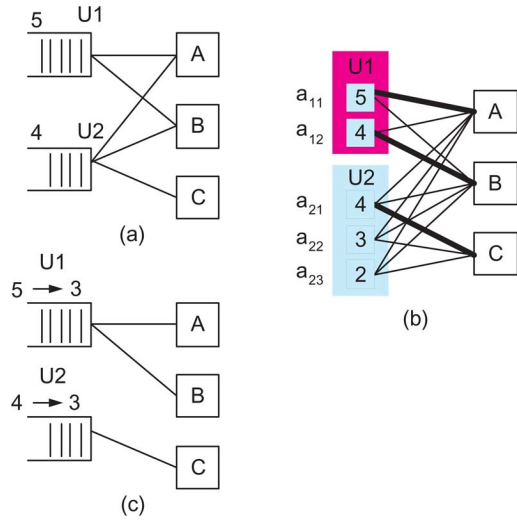


Fig. 2. Example of MTLB construction (a) queue lengths and connectivities. (b) The equivalent bipartite graph with the weights are shown at each subnode, e.g., the weights of the edges (a_{11}, A) and (a_{11}, B) are five. The thick edges indicate the MWM. (c) The edges indicate the resulted MTLB assignment. The leftover queue length after the allocation is $(3, 3)$.

Definition 4 [19]: Consider a bipartite graph (U, V, E) with two vertex sets U and V , an edge set $E \subseteq U \times V$, and a weight function $\psi : E \mapsto \mathbb{R}$. A *matching* M is a subset of E such that no two edges in M share an endpoint. The weight of a matching M is $\psi(M) = \sum_{e \in M} \psi(e)$. A matching M is an *MWM* if its weight is no less than the weight of any other matching.

Definition 5: A server allocation $W = [w_{ij}]$ and a matching M are said to be *equivalent* when: 1) M is a matching on the equivalent bipartite graph, and 2) $w_{ij} = 1$ if and only if there exists m such that (a_{jm}, v_i) is a matching edge, i.e., $(a_{jm}, v_i) \in M$.

Proposition 1: An MWM on the equivalent bipartite graph is MTLB, i.e., it satisfies conditions C1) and C2).

Proof: See Appendix I. \square

An example of the MTLB assignment based on the proposed algorithm is shown in Fig. 2. It is intuitive to see that the MWM on the equivalent bipartite graph achieves an MTLB assignment. This is because the equivalent bipartite graph, in effect, expands the individual packets that can possibly be served into nodes and basically labels each packet with the number of packets waiting behind it [see Fig. 2(b)]. The MWM selects the matching that serves the packets with the most number of packets waiting behind them. This guarantees the MT and the LB properties at the same time. Over-assignments are avoided since, in the equivalent bipartite graph, only $\min\{b_j, \sum_{i=1}^K c_{ij}\}$ packets from each node j are expanded. Again, note that since a graph can have multiple MWMs, MTLB allocation is not unique. The complexity of finding an MTLB allocation is equal to the complexity of the existing MWM algorithm applied to the equivalent bipartite graph, which is $O(K^2(N + \log K))$ [19].

V. OPTIMALITY OF MTLB POLICY FOR TWO USERS

With the existence of the MTLB policy, we proceed to establish its optimality.

Theorem 2: Consider Problem (P) with ON-OFF connectivity and $N = 2$ users. The MTLB policy is optimal for all choices of g , T , and \mathcal{I}_0 as defined in Problem (P).

Remark 1: The optimality of the MTLB policy shown in Theorem 2 implies that the maximum instantaneous throughput criterion [condition C1)] is not sufficient to guarantee the delay optimality unless it is complemented by the LB criterion [condition C2)].

To show Theorem 2, we use a similar framework as in [6], [7], [14], and [16] as follows: we first define a class of functions \mathcal{F} , which contains, for instance, any cost functions $\phi(\mathbf{b})$ of the form $\sum_{j=1}^N g(b_j)$, where g is strictly increasing and convex. We, then, show that if the cost function ϕ belongs to \mathcal{F} , the average optimal cost-to-go function (derived via a DP equation) also belongs to \mathcal{F} . The properties of \mathcal{F} are then used to show the optimality of the MTLB policy.

A. Class of Cost Functions

Here we give the definition of class- \mathcal{F} functions. This is the class of the cost functions for which the solution to Problem (P) is proved to be the MTLB policy (Theorem 2). But first we define the following for convenience.

Definition 6: $R_{ij}(\mathbf{b}) := \mathbf{b} - \mathbf{e}_i + \mathbf{e}_j$, where \mathbf{e}_m is a row vector of zeros except for the m th element which is one, is equivalent to a transfer of a packet from queue i to queue j .

Definition 7: A function $f : \mathbb{Z}_+^N \rightarrow \mathbb{R}$ belongs to the set \mathcal{F} if, for any $i \neq j \in \{1, \dots, N\}$, f satisfies the following conditions:

- B.1) monotonicity condition: $f(\mathbf{b}) \leq f(\mathbf{b} + \mathbf{e}_i)$;
- B.2) permutation invariance condition: $f(\mathbf{b}) = f(\pi(\mathbf{b}))$ for any permutation π ;
- B.3) supermodularity condition: $f(\mathbf{b} + \mathbf{e}_i) - f(\mathbf{b}) \leq f(\mathbf{b} + \mathbf{e}_i + \mathbf{e}_j) - f(\mathbf{b} + \mathbf{e}_j)$;
- B.4) coordinatewise convexity condition: $2f(\mathbf{b}) \leq f(\mathbf{b} + \mathbf{e}_i) + f(\mathbf{b} - \mathbf{e}_i)$ for $\mathbf{b} > \mathbf{0}$;
- B.5) convexity along a constant-sum line condition: $2f(\mathbf{b}) \leq f(R_{ij}(\mathbf{b})) + f(R_{ji}(\mathbf{b}))$;
- B.6) balancing advantage condition: $f(R_{ij}(\mathbf{b})) \leq f(\mathbf{b})$ if and only if $b_i \geq b_j + 1$.

The terminologies in B.3)–B.5) follow that used in [14]. Conditions B.3)–B.5) are second-order relations related to convexity over lattice spaces.² Condition B.6) establishes the optimality of the MTLB policy. The following can be easily shown.

Fact 1: Any function of the form $\phi(\mathbf{b}) = \sum_{j=1}^N g(b_j)$, where g is strictly increasing and convex, belongs to \mathcal{F} .

²Due to the symmetric assumptions [i.e., condition B.2)] in our model, conditions B.3)–B.5) are special cases of the multimodularity condition in [20].

B. Dynamic Programming Formulation

Next we use a DP approach to relate the cost function in (4) to the expected cost-to-go $V_n^\sigma(\mathbf{b}, C)$ at time $t = T - n$ (i.e., at horizon n) under a Markovian policy σ . Let allocation $W^\sigma(\mathbf{b}, C) \in \mathcal{W}(\mathbf{b}, C)$ denote the allocation at state (\mathbf{b}, C) prescribed by policy σ . Note that, in general, the action $W^\sigma(\mathbf{b}, C)$ depends on horizon n and is assumed implicitly. It is clear that the following recursion:

$$V_0^\sigma(\mathbf{b}, C) = \phi(\mathbf{b})$$

and for $n \geq 1$

$$V_n^\sigma(\mathbf{b}, C) = \phi(\mathbf{b}) + E_{\mathbf{a}, \tilde{C}}[V_{n-1}^\sigma(\mathbf{b} + \mathbf{a} - \mathbf{1}W^\sigma(\mathbf{b}, C), \tilde{C})] \quad (10)$$

is related to the cost function in (4) as

$$J_T^\sigma = E_C[V_T^\sigma(\mathbf{b}, C)] \quad (11)$$

when $\mathcal{I}_0 = \mathbf{b}$. This is due to the validity of the DP theorem for a finite horizon Markov decision process (MDP) [17]. Define

$$V_n^*(\mathbf{b}, C) := \min_{\sigma \in \mathcal{U}_n} V_n^\sigma(\mathbf{b}, C) \quad (12)$$

to be the minimum cost-to-go over the set \mathcal{U}_n of all Markovian policies at horizon n . Furthermore, we define the average optimal cost-to-go function as

$$v_n(\mathbf{b}) := E_{\mathbf{a}, C}[V_n^*(\mathbf{b} + \mathbf{a}, C)]. \quad (13)$$

In Proposition 2, we show the recursive structure of v_n .

Proposition 2: Given a horizon n , the average optimal cost-to-go at n , $v_n(\mathbf{b})$, satisfies the following recursions:

$$\begin{aligned} v_0(\mathbf{b}) &= \bar{\phi}(\mathbf{b}) := E_{\mathbf{a}}[\phi(\mathbf{a} + \mathbf{b})] \\ v_n(\mathbf{b}) &= \bar{\phi}(\mathbf{b}) + E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(C)} v_{n-1}([\mathbf{b} + \mathbf{a} - \mathbf{1}W]^+) \right]. \end{aligned} \quad (14)$$

Proof: From the recursion in (10), we have the following recursion for the optimal cost-to-go $V_n^*(\mathbf{b}, C)$:

$$\begin{aligned} V_0^*(\mathbf{b}, C) &= \phi(\mathbf{b}) \\ V_n^*(\mathbf{b}, C) &= \phi(\mathbf{b}) + \min_{W \in \mathcal{W}(\mathbf{b}, C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b} + \mathbf{a} - \mathbf{1}W, \tilde{C})] \\ &= \phi(\mathbf{b}) + E_{\mathbf{a}, \tilde{C}} \left[V_{n-1}^*(\mathbf{b} + \mathbf{a} - \mathbf{1}W^*(\mathbf{b}, C), \tilde{C}) \right] \\ &= \phi(\mathbf{b}) + v_{n-1}(\mathbf{b} - \mathbf{1}W^*(\mathbf{b}, C)) \end{aligned} \quad (15)$$

where $W^*(\mathbf{b}, C)$ denotes an optimal allocation at horizon n when the state of the queue backlogs is equal to the vector \mathbf{b} and the connectivity profile is C . In other words, $W^*(\mathbf{b}, C) \in \arg \min_{W \in \mathcal{W}(\mathbf{b}, C)} v_{n-1}(\mathbf{b} - \mathbf{1}W)$. Now taking the expectation of both sides, we have

$$\begin{aligned} v_n(\mathbf{b}) &= E_{\mathbf{a}, C}[V_n^*(\mathbf{b} + \mathbf{a}, C)] \\ &= \bar{\phi}(\mathbf{b}) + E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(\mathbf{b} + \mathbf{a}, C)} v_{n-1}(\mathbf{b} + \mathbf{a} - \mathbf{1}W) \right] \\ &= \bar{\phi}(\mathbf{b}) + E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(C)} v_{n-1}([\mathbf{b} + \mathbf{a} - \mathbf{1}W]^+) \right] \end{aligned}$$

where the last equality holds because, for any allocation $W \in \mathcal{W}(C)$, there exists $W' \in \mathcal{W}(\mathbf{b}, C)$ such that $\mathbf{b} - \mathbf{1}W' = [\mathbf{b} - \mathbf{1}W]^+$. Finally, $v_0(\mathbf{b}) = E_{\mathbf{a}, C}[V_0^*(\mathbf{b} + \mathbf{a}, C)] = \bar{\phi}(\mathbf{b})$. \square

Proof of Theorem 2

Using the above class of functions \mathcal{F} and the recursive structure of v_n in Proposition 2, we are ready to prove Theorem 2. Note that the lemmas used here are proved in Appendix II.

Proof of Theorem 2: We first show the strict monotonicity of v_n for all horizon n , using the strict monotonicity of the cost function. This is shown in the following lemma.

Lemma 1: $v_n(\mathbf{b})$ is strictly increasing on \mathbf{b} for all $n = 0, \dots, T$.

Next, we show that, for any horizon n , the MTLB policy is optimal at time $n + 1$ whenever $v_n \in \mathcal{F}$. This is shown in the following lemma.

Lemma 2: For any horizon n , if $v_n \in \mathcal{F}$, then the MTLB policy is optimal at horizon $n + 1$.

The above lemma immediately establishes Theorem 2 if we can show that $v_n \in \mathcal{F}$ for all $n = 0, \dots, T$. To show that $v_n \in \mathcal{F}$ for all $n = 0, \dots, T$, we use the following induction.

Induction basis: From Proposition 2, $v_0(\mathbf{b}) = \bar{\phi}(\mathbf{b}) = \sum_{\mathbf{a}} P_{\mathbf{a}}(\mathbf{a})\phi(\mathbf{b} + \mathbf{a})$, where $P_{\mathbf{a}}(\mathbf{a})$ is the probability of the arrival vector is \mathbf{a} . Using Fact 4, $v_0 \in \mathcal{F}$ since $\phi \in \mathcal{F}$.

Induction step: Suppose $v_n \in \mathcal{F}$. To show that $v_{n+1} \in \mathcal{F}$, we recall from Proposition 2 that

$$v_{n+1}(\mathbf{b}) = \bar{\phi}(\mathbf{b}) + E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(C)} v_n([\mathbf{b} + \mathbf{a} - \mathbf{1}W]^+) \right]. \quad (16)$$

We find that it is more convenient to work with a relaxed version of v_n , which allows the queue vector to be negative. That is, we work with \hat{v}_n where $\hat{v}_n(\mathbf{b}) = v_n([\mathbf{b}]^+)$ for $\mathbf{b} \in \mathbb{Z}^N$. This relaxation technique (used in [6], [7], [14], and [16]) removes the need for the separate treatment of various boundary cases. To facilitate this relaxation, we define an extended class of functions $\hat{\mathcal{F}}$ as follows.

Definition 8: Consider $f: \mathbb{Z}_+^N \rightarrow \mathbb{R}$. We denote $\hat{f}: \mathbb{Z}^N \rightarrow \mathbb{R}$ as an extension of f on \mathbb{Z}^N such that $\hat{f}(\mathbf{b}) = f([\mathbf{b}]^+)$. Furthermore, we define an extension $\hat{\mathcal{F}}$ of \mathcal{F} :

$$\hat{\mathcal{F}} := \left\{ \hat{f}: \mathbb{Z}^N \rightarrow \mathbb{R} : \hat{f} \text{ meets B.1)–B.6) } \right\}. \quad (17)$$

With this extension, it is clear that v_{n+1} in (16) is the restriction of \hat{v}_{n+1} to the nonnegative domains, where for $\mathbf{b} \in \mathbb{Z}^N$

$$\hat{v}_{n+1}(\mathbf{b}) = \hat{\phi}(\mathbf{b}) + E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b} + \mathbf{a} - \mathbf{1}W) \right]. \quad (18)$$

Now, we show that $v_{n+1} \in \mathcal{F}$ via the following four steps:

Assuming $v_n \in \mathcal{F}$

$$\xrightarrow{\text{Step 1}} \hat{v}_n \in \hat{\mathcal{F}} \quad (\text{Fact 2})$$

$$\xrightarrow{\text{Step 2}} E_{\mathbf{a}, C} \left[\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b} + \mathbf{a} - \mathbf{1}W) \right] \text{ satisfies B.3)–B.6) } \quad (\text{Lemmas 4 and 5})$$

$$\begin{aligned} \xrightarrow{\text{Step 3}} \hat{v}_{n+1} &\in \hat{\mathcal{F}} \quad (\text{Lemmas 1 and 3, and Facts 2 and 3}) \\ \xrightarrow{\text{Step 4}} v_{n+1} &\in \mathcal{F} \quad (\text{Fact 5}) \end{aligned}$$

where facts and lemmas in the parentheses indicate how to establish the above steps. For completeness, the statements of these facts and lemmas are listed below after the proof. All steps except Step 3 are immediate from the listed facts and lemmas. In Step 3, we first show that \hat{v}_{n+1} satisfies B.3)–B.6) (note that $\bar{\phi} \in \mathcal{F}$ and Fact 2 imply that $\bar{\phi}(\cdot) \in \hat{\mathcal{F}}$). Then, using Lemmas 1 and 3 and Fact 2, \hat{v}_{n+1} satisfies B.1) and B.2) as well. Hence, $\hat{v}_{n+1} \in \hat{\mathcal{F}}$. Note that Lemmas 4 and 5, which establish Step 2, also use the fact (Lemma 2) that the optimal allocation W at horizon $n + 1$ is MTLB if $\hat{v}_n \in \hat{\mathcal{F}}$.³ \square

Here we list the facts and the (remaining) lemmas used in the above proof. All the facts are taken from [6] and [7] and can be easily verified. The lemmas are proved in Appendix II.

Fact 2: If $f \in \mathcal{F}$, then the function $\hat{f} : \mathbb{Z}^N \rightarrow \mathbb{R}$ defined as $\hat{f}(\mathbf{b}) = f([\mathbf{b}]^+)$ is in $\hat{\mathcal{F}}$.

Fact 3: If $\hat{f}_1, \hat{f}_2, \dots$ are functions that belong to $\hat{\mathcal{F}}$, then $\hat{h}(\mathbf{b}) = \sum_l p_l \hat{f}_l(\mathbf{b})$ also belongs to $\hat{\mathcal{F}}$, where p_l are nonnegative constants.

Fact 4: If f_1, f_2, \dots are functions that belong to \mathcal{F} , then $h(\mathbf{b}) = \sum_l p_l f_l(\mathbf{b})$ also belongs to \mathcal{F} , where p_l are nonnegative constants.

Fact 5: If $\hat{f} \in \hat{\mathcal{F}}$, then the restriction of \hat{f} to nonnegative domain is in \mathcal{F} .

Lemma 3: $v_n(\mathbf{b})$ is permutation invariant on \mathbf{b} for all $n = 0, \dots, T$.

Lemma 4: Assume $N = 2$ and $\hat{v}_n \in \hat{\mathcal{F}}$. For any state \mathbf{b} , $E_{\mathbf{a}, \mathbf{C}} [\min_{W \in \mathcal{W}(\mathbf{C})} \hat{v}_n(\mathbf{b} + \mathbf{a} - \mathbf{1}W)]$ satisfies B.3)–B.5).

Lemma 5: Assume $N = 2$ and $\hat{v}_n \in \hat{\mathcal{F}}$. For any state \mathbf{b} such that $b_1 \geq b_2 + 1$, $E_{\mathbf{a}, \mathbf{C}} [\min_{W \in \mathcal{W}(\mathbf{C})} \hat{v}_n(\mathbf{b} + \mathbf{a} - \mathbf{1}W)]$ satisfies condition B.6).

Remark 2: All the above lemmas and facts, except Lemmas 4 and 5, hold for general N . Lemmas 4 and 5 are proved for $N = 2$. We detail the difficulty in extending these lemmas to general N after Theorem 3 in Section VI and also after the proof of Lemma 5 in Remark 4 in Appendix II.

Remark 3: Theorem 2, in addition, can be extended to the optimality of the MTLB policy in an expected average cost sense for an infinite horizon problem.

Corollary 1: Consider an infinite horizon version of Problem (P), where the cost is modified to be the average expected cost at each horizon. Then, the MTLB policy is optimal for any initial state $\mathcal{I}_0 = \mathbf{b}$.

Proof: Theorem 2 proves that there exists a stationary MTLB policy, which is optimal for Problem (P) for any finite horizon T . Hence, our MTLB policy achieves the minimization of the average expected cost Λ_T^π / T for any finite horizon T . Since the policy is independent of the horizon T , it is optimal

³Although Lemma 2 assumes $v_n \in \mathcal{F}$, it is easy to show that the lemma works with $\hat{v}_n \in \hat{\mathcal{F}}$ as well.

with respect to an average expected cost criterion for the infinite horizon version of the problem. \square

VI. OPTIMALITY OF MTLB POLICY UNDER FLUID RELAXATION

In the previous section, we established the optimality of the MTLB policy for a very restricted case of $N = 2$. As we will see later, the major difficulty in extending the proof to general N is due to the *integral* server allocation constraint. In this section, we study a relaxed system where we allow each server to serve a fractional (*fluid*) number of packets from multiple queues as long as the total number of packets served per server is no greater than one. In other words, we consider a real allocation $W = [w_{ij}] \in \mathbb{R}^{K \times N}$, with $w_{ij} \in [0, 1]$. We call this relaxed constraint the fluid server allocation relaxation or *fluid relaxation*. Under this fluid relaxation, we can show the optimality of the MTLB policy for general N and the ON-OFF channel model. Before we proceed, we provide a modification of $\mathcal{W}(\mathbf{b}, \mathbf{C})$ to include fluid allocations and a definition of the fluid version of the MTLB policy.

Definition 9: A class of *fluid* nonidling feasible allocations $\mathcal{W}^f(\mathbf{b}, \mathbf{C})$ is equivalent to $\mathcal{W}(\mathbf{b}, \mathbf{C})$ with the fluid server allocation condition, i.e., $W = [w_{ij}] \in \mathcal{W}^f(\mathbf{b}, \mathbf{C})$ if:

- a) $0 \leq w_{ij} \leq 1$;
- b) $c_{ij} = 0 \Rightarrow w_{ij} = 0$;
- c) $\sum_{j=1}^N w_{ij} \leq 1, \forall i = 1, \dots, K$;
- d) $\sum_{i=1}^K w_{ij} \leq b_j, \forall j = 1, \dots, N$.

Definition 10: The MTLB-F policy is a fluid version of the MTLB policy defined in Section IV-A, i.e., the MTLB-F policy chooses $W^* \in \mathcal{W}^f(\mathbf{b}, \mathbf{C})$ such that C1) and C2) are satisfied.

Example 2: It is clear that the MTLB-F policy is not uniquely defined although the leftover queue vector is. For example, for the $\mathbf{C} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}$ in Example 1, if $\mathbf{b} = [3, 3, 3, 3, 1]$, then all allocations of the form $\begin{bmatrix} 0.4 & 0.4 & x & 0.2 - x \\ 0 & 0 & y & 1 - y \end{bmatrix}$, where $x \in [0, 0.2]$ and $y = 0.7 - x$, are MTLB-F allocations, resulting in the unique leftover queues $[2.6, 2.6, 2.6, 2.6]$.

Assuming that the cost function $\phi(\mathbf{b})$ is convex on $\mathbf{b} \in \mathbb{R}_+^N$, we have the following result.

Theorem 3: For the problem (P) with the fluid server allocation relaxation, the MTLB-F policy is optimal.

Proof: See Appendix III. \square

The key element in the proof of Theorem 3 is the convexity of the cost-to-go function v_n in the fluid relaxation, for all $n = 0, \dots, T$. This convexity property directly establishes the optimality of the MTLB-F policy. However, under the integral server allocation constraint, it is hard to establish a similar convexity property over lattices. In fact, one can interpret the difficulty in establishing Lemma 4 and 5 as an indication that the properties B.3)–B.6) of the set \mathcal{F} provide a sufficient set of properties for convex functions over two-dimensional lattices, while they fail to sufficiently capture the convexity in higher dimensional lattices.

VII. CONCLUSION AND FUTURE RESEARCH

In this paper, we have considered the problem of optimal server allocation in a time-slotted system with N symmetric queues and K servers when the arrivals and channels are stochastic and time-varying. Focusing on a long-term average-delay objective, we identified the MTLB policy that achieves the instantaneous MT as well as balances the queues. Such a policy always exists when the channel connectivity follows an ON-OFF model. In such a case, we proved that the MTLB policy achieves the minimum average delay (mean response time) at any time when there are only $N = 2$ users. Although the general N case remains open, we showed that the MTLB-F policy, which is the fluid version of the MTLB policy, is optimal for the general N case under the fluid server allocation relaxation. With these results, we propose the following conjecture.

Conjecture 1: In the case of $N > 2$ statistically symmetric users with the ON-OFF channel connectivity and the integral server allocation constraint, the MTLB policy is optimal.

The conjecture is rather intuitive due to the symmetry of the users and the symmetry and convexity of the cost function. In addition, the optimality of the MTLB policy for general N in some special cases have been previously shown. For the case of single server and Bernoulli arrivals, Tassioulas and Ephremides [3] proved the optimality of the LCQ policy, which coincides with the MTLB policy for $K = 1$. For the case of multiple servers and Bernoulli arrivals but with vector connectivity (i.e., each user is either connected to all servers or none) and the one-server-per-queue constraint, the LCQ policy (a more generalized multiqueue version of the LCQ policy in [3]) is optimal [4]. This LCQ policy serves the longest connected queues—this is equivalent to the MTLB policy with the one-server-per-queue constraint. It is interesting to note the complimentary roles of the stochastic coupling and majorization techniques used in [3] and [4] and the DP technique we employed in this paper.

We would like to point out that although the ON-OFF connectivity model may seem restricted for practical wireless systems, the optimality of the MTLB policy in such ON-OFF model can lead to useful insight for the general connectivity setting. In fact, in [21], we have obtained excellent performance for the general connectivity using some MTLB-based heuristic policies.

APPENDIX I

EXISTENCE PROOF OF MTLB: PROOF OF THEOREM 1

In this section, we prove Theorem 1 and Proposition 1 using the notions of alternating, balancing, and throughput-increasing paths, which are the concepts taken from graph literature [22]. We note that some of the results here are useful in the next appendixes. Note that the discussions in this section (both the existence proof as well as the construction of the MTLB allocation) are valid for general N .

A. Alternating, Balancing, and Throughput-Increasing Paths

For convenience and simplicity of the proofs, we adopt the language of graph literature. Let $G = (V, U, C)$ be a bipartite graph with U the set of queues, V the set of servers, and edge

set $C \subseteq V \times U$ representing the set of connectivities between queues and servers. Each allocation matrix $W = [w_{i,j}]$ can be thought of an edge set where an edge $(v, u) \in W$ if $w_{v,u} = 1$. Hence, an edge set W is a (feasible) allocation (i.e., $W \in \mathcal{W}(C)$) if each vertex $v \in V$ is incident with exactly one edge in C . Furthermore, W is nonidling (i.e., $W \in \mathcal{W}(\mathbf{b}, C) \subseteq \mathcal{W}(C)$) if W is feasible and $\sum_{i=1}^K w_{iu} \leq b_u$ for each $u \in U$.

Definition 11: A vertex in V that is incident to any edge in the allocation is called *matched*, and *unmatched* otherwise. A queue or vertex u in U with $b_u - \sum_{i=1}^K w_{iu} > 0$ is called *nonempty*, and *empty* otherwise.

Definition 12: Consider a given allocation $W \subseteq \mathcal{W}(\mathbf{b}, C)$ in G and a sequence of distinct vertices

$$T = u_0, v_1, u_1, v_2, \dots, v_k, u_k$$

from a queue $u_0 \in U$ to a queue $u_k \in U$, through matched servers, with $v_i \in V$, $u_i \in U$, $(v_i, u_{i-1}) \in C \setminus W$, and $(v_i, u_i) \in W$ for each $i = 1, \dots, k$. An *alternating path* $S(W, T)$ with respect to W is a sequence of edges with vertices in T , i.e.,

$$S(W, T) := \{(v_1, u_0), (v_1, u_1), (v_2, u_1), \dots, (v_k, u_k)\}. \quad (19)$$

For brevity, we omit the nodes $v_1, u_1, v_2, \dots, v_k$ in our notation and write $S(W, T)$ as $S(W, u_0, u_k)$ to emphasize the end nodes.

Definition 13: An alternating path $S(W, u_0, u_k)$ is called a *balancing path* if it satisfies $b_{u_0} - \sum_{i=1}^K w_{i,u_0} \geq b_{u_k} - \sum_{i=1}^K w_{i,u_k} + 2$.

For convenience, we treat paths as a sequence of vertices. For example, we write

$$S(W, u_0, u_k) = (u_0, v_1, u_1, v_2, \dots, v_k, u_k)$$

to show $S(W, u_0, u_k)$ as a sequence of vertices alternatively taken from U and V starting from u_0 and ending at u_k .

Definition 14: A *throughput-increasing path* relative to W , from an unmatched server $v_0 \in V$ to a nonempty queue $u_k \in U$, is a sequence of distinct vertices (or equivalently, a sequence of edges)

$$I(W, v_0, u_k) := (v_0, u_1, v_1, u_2, \dots, v_{k-1}, u_k)$$

with $v_i \in V$, $u_i \in U$, $(v_{i-1}, u_i) \in C \setminus W$, and $(v_i, u_i) \in W$ for each i .

Definition 15: For the alternating path $S = S(W, u_0, u_k)$ given in (19), $W^a(S)$ is the *alternating allocation* of the allocation W along an alternating path S if server v_l is reassigned to serve queue u_{l-1} , $\forall l = 1, \dots, k$. If, in addition, S is a balancing path, then $W^a(S)$ is specifically called the *balancing allocation* and denoted by $W^b(S)$. In a similar fashion, $W^t(I)$ is called the *throughput-increasing allocation* if I is a throughput-increasing path and $W^t(I)$ assigns server v_l to serve queue u_{l+1} , $\forall l = 0, \dots, k-1$. Equivalently, we can write $W^a(S) = W \oplus S$, $W^b(S) = W \oplus S$, and $W^t(I) = W \oplus I$, where $A \oplus B := (A \setminus B) \cup (B \setminus A)$ for any sets A, B .




Fig. 3. Example of an alternating path and the alternating allocation from queue u_0 to queue u_2 . (a) Alternating path $S = (u_0, v_1, u_1, v_2, u_2)$. The dotted and solid lines show the connectivities while the solid lines show the allocation W . (b) The solid lines now show the alternating allocation $W^a(S) = W \oplus S$.

An example of some alternating path and alternating allocation is shown in Fig. 3. It is easy to see that $W^a(S(W, u_0, u_k)) \in \mathcal{W}(\mathbf{b}, C)$ if u_0 is nonempty under W . Obviously, $W^b(S) \in \mathcal{W}(\mathbf{b}, C)$. Note that the alternating allocations (when u_0 is nonempty) and the balancing allocations leave the cardinality of the allocation (i.e., throughput) unchanged, while throughput-increasing allocations increase the current throughput by one. In addition, when u_0 is nonempty under W , the allocations W and $W^a(S = S(W, u_0, u_k))$ result in the leftover queues that are identical except for queues u_0 and u_k . In other words, if we denote the leftover queues under W and $W^a(S)$ as $\mathbf{l} = \mathbf{b} - \mathbf{1}W$ and $\mathbf{l}^a = \mathbf{b} - \mathbf{1}W^a(S)$, respectively, then $l_{u_0}^a = l_{u_0} - 1$, $l_{u_k}^a = l_{u_k} + 1$, and $l_u^a = l_u$, for all $u \in U \setminus \{u_0, u_k\}$.

Notice that the above notion of throughput-increasing path is conceptually related to the notion of the alternating path in the graph matching literature [23]. Likewise, our notion of balancing path is related to the notion of the cost-reducing path in [22] where cost is the “unbalancedness” of the queues.

B. Proof of Existence of MTLB Policy

Proposition 3 is used to find the necessary and sufficient condition for policies to satisfy C1) and show the existence of the MTLB policy (Theorem 1).

Proposition 3: An allocation achieves the MT C1) if and only if it has no throughput-increasing paths.

Proof: Obviously, if there is a throughput-increasing path for a given allocation W , then W does not achieve the MT. To show that not having any throughput-increasing paths is a sufficient condition for achieving the MT, the proof follows the standard graph technique used in [23], which turns the problem into a maximum network flow problem. Since the proof is long and standard, we refer interested readers to [23] for a detailed proof.

Theorem 1: For any given (\mathbf{b}, C) , an MTLB allocation always exists.

Proof: From Proposition 3 and C2), it follows that every LB allocation satisfies C1). Hence the existence of an MTLB allocation is trivial. This is because the \leq_{LQO} is an order and the set $\mathcal{W}^{LB}(\mathbf{b}, C)$ is nonempty and finite. Hence, there exists a minimal element in $\mathcal{W}^{LB}(\mathbf{b}, C)$. \square




Fig. 4. Illustration for the proof of Proposition 4.

C. Necessary and Sufficient Condition for MTLB Policy

Proposition 4 gives a necessary and sufficient condition for the MTLB policy. This result will be useful in the proof of the optimality of the MTLB policy in Appendix II.

Proposition 4: Any allocation satisfying the maximum-throughput condition C1) also satisfies the LB condition C2) if and only if it has no balancing path.

Proof: Without loss of generality, consider (\mathbf{b}, C) such that $\mathcal{W}(\mathbf{b}, C) \neq \emptyset$. The only if part is obvious, i.e., if there is a balancing path S relative to an allocation $W \in \mathcal{W}(\mathbf{b}, C)$, then $\mathbf{b} - \mathbf{1}W^b(S) \leq_{LQO} \mathbf{b} - \mathbf{1}W$ but $\mathbf{b} - \mathbf{1}W \not\leq_{LQO} \mathbf{b} - \mathbf{1}W^b(S)$, i.e., W does not satisfy C2).

What remains is the *if* part: if a maximum-throughput allocation does not satisfy C2), then it has at least one balancing path. Let $W \in \mathcal{W}(\mathbf{b}, C)$ be a maximum-throughput allocation [satisfying C1)] but is not the most balanced [not satisfying C2)]. We show that a balancing path relative to W must exist. Since at least one MTLB allocation exists (by Theorem 1), we let $W^* \in \mathcal{W}(\mathbf{b}, C)$ be an MTLB allocation. If more than one MTLB allocations exist, we pick W^* such that the number of edges in the symmetric difference $W^* \oplus W = (W^* \setminus W) \cup (W \setminus W^*)$ is minimized among all MTLB allocations. That is, W^* is the “closest” MTLB allocation to W , i.e., among all MTLB allocations, W^* requires the minimum number of servers to be reassigned to get to W . Now, let G_d be the subgraph of the bipartite graph $G = (V, U, C)$ induced by the edges of $W^* \oplus W$. Color the edges of $W^* \setminus W$ green and the edges of $W \setminus W^*$ red. Direct the green edges from V to U and the red edges from U to V (see an example in Fig. 4). Let the leftover queue vectors under W and W^* be $\mathbf{l} = \mathbf{b} - \mathbf{1}W$ and $\mathbf{l}^* = \mathbf{b} - \mathbf{1}W^*$, respectively.

We claim that for every directed path P in G_d from $u_1 \in U$ to $u_2 \in U$, we have

$$l_{u_1}^* \leq l_{u_2}^*. \quad (20)$$

To see this, let $P = (u_1, \dots, u_2)$ be a directed path in G_d . By the choice of the directions for the edges, P must be alternating

between red and green edges. If $l_{u_2}^* < l_{u_1}^* - 1$, then P is a balancing path for W^* , and $\mathbf{b} - \mathbf{1}W^b(P) \leq_{\text{LQO}} \mathbf{b} - \mathbf{1}W^*$ but $\mathbf{b} - \mathbf{1}W^* \not\leq_{\text{LQO}} \mathbf{b} - \mathbf{1}W^b(P)$, contradicting to the assumption that W^* satisfy C2). Similarly, if $l_{u_2}^* = l_{u_1}^* - 1$, then by alternating the assignment of the V -vertices along P , we can get another MTLB allocation W^{**} such that the number of edges in $W^{**} \oplus W$ is strictly less than that in $W^* \oplus W$, in contradiction to the choice of W^* . Hence, we must have that $l_{u_2}^* \geq l_{u_1}^*$. Using a similar argument, we can also show that G_d is acyclic [this fact can also be observed from Fig. 4(d)].

Since both W^* and W achieve the MT, we have that $\sum_{i=1}^N l_i = \sum_{i=1}^N l_i^*$. But since W does not satisfy the LB condition C2), there must exist $u_1 \in U$ such that

$$l_{u_1} < l_{u_1}^*. \quad (21)$$

Obviously, there is a red edge directed out of u_1 . Starting from u_1 , we build an alternating red–green path P' in G_d as follows.

- 1) From an arbitrary vertex $u \in U$ (including u_1), if there is a red edge directed out of u and $l_u \leq l_{u_1} + 1$, we build P' by arbitrarily selecting one of the red edges directed out of u .
- 2) From an arbitrary $v \in V$, we build P' by following the single green edge directed out of v . Such a green edge always exist.⁴
- 3) Otherwise, stop.

Using the fact that G_d is acyclic, P' is well defined and finite. Let $u_2 \in U$ be the final vertex on the path. There are two possible cases.

Case 1) $l_{u_2} > l_{u_1} + 1$. In this case, we reverse the order of nodes in P' to arrive at a balancing path relative to W .

Case 2) There is no red edge directed out of u_2 . This means that u_2 is served at least one more packet under W^* , relative to W . Hence, $l_{u_2} > l_{u_2}^*$. This together with (20) and (21) give $l_{u_2} > l_{u_2}^* \geq l_{u_1}^* > l_{u_1}$, which means $l_{u_2} \geq l_{u_1} + 2$. Reversing the order of nodes in P' gives a balancing path relative to W .

Since there exists a balancing path in both cases, we have the assertion of the proposition. \square

The above proposition states that all MTLB allocations are such that they have no balancing path. The results in the proof that G_d is acyclic and has finite edges immediately implies the following result.

Corollary 2: The minimum number of balancing allocations required to turn any MT allocation satisfying C1) to an MTLB allocation is finite.

D. Proof of Proposition 1

The equivalence of the MWM matching on the equivalent bipartite graph and the MTLB allocation (Proposition 1 in Section IV-C) can be proved as follows.

Proof of Proposition 1: Since all weights are strictly positive, the MWM matching on the equivalent bipartite graph necessarily matches all possible servers and hence the equivalent allocation (defined in Definition 5) achieves the

⁴Otherwise, combining the red edge coming into v with W^* would have yielded a nonidling feasible allocation with additional one packet throughput, a contradiction to C1).

maximum-throughput condition C1). We prove the LB condition C2) by contradiction. Suppose the MWM M results in the allocation W that achieves the MT but does not produce the most balanced queues. From Proposition 4, we know that there must exist a balancing path $S(W, j, i)$ from some queue j to queue i such that $b_j - w_j \geq b_i - w_i + 2$, where $w_s = \sum_{m=1}^K w_{m,s}$ for $s = i, j$. Let us denote the balancing allocation of W along $S(W, j, i)$ as W^b . Let M^b be the equivalent matching to W^b . According to M , node $a_{i w_i}$ is matched and $a_{j(w_j+1)}$ is not, while the reverse is true for M^b . In other words, $\psi(M^b) - \psi(M) = b_j - w_j - (b_i - w_i + 1) \geq 1$. But this is a contradiction to the assumption that M is the MWM on the equivalent bipartite graph. \square

APPENDIX II

SUPPORTING LEMMAS FOR THEOREM 2

In this appendix, we establish the proofs for Lemmas 1–5 stated in Section V-C. The first lemma establishes the strict monotonicity of v_n for all $n = 0, \dots, T$. Hence, v_n satisfies B.1) for all n as well.

Lemma 1: $v_n(\mathbf{b})$ is strictly increasing on \mathbf{b} for all $n = 0, \dots, T$, i.e., $\mathbf{b}' > \mathbf{b} \Rightarrow v_n(\mathbf{b}') > v_n(\mathbf{b})$.

Proof: Since v_n is related to V_n^* by (13), it suffices to show the strict monotonicity of $V_n^*(\mathbf{b}, C)$ for any C . We show this by induction.

Induction basis: $V_0^*(\mathbf{b}, C) = \phi(\mathbf{b}) = \sum_{j=1}^N g(b_j)$ is strictly increasing by the assumption of g .

Induction step: Assume $V_{n-1}^*(\mathbf{b}', C) > V_{n-1}^*(\mathbf{b}, C)$ for any $\mathbf{b}' > \mathbf{b}$, then

$$\begin{aligned} V_n^*(\mathbf{b}', C) &= \phi(\mathbf{b}') + \min_{W' \in \mathcal{W}(\mathbf{b}', C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b}' - \mathbf{1}W' + \mathbf{a}, \tilde{C})] \\ &\geq \phi(\mathbf{b}') + \min_{W' \in \mathcal{W}(\mathbf{b}', C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b} - \mathbf{1}W(W') + \mathbf{a}, \tilde{C})] \\ &\geq \phi(\mathbf{b}') + \min_{W \in \mathcal{W}(\mathbf{b}, C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b} - \mathbf{1}W + \mathbf{a}, \tilde{C})] \\ &> \phi(\mathbf{b}) + \min_{W \in \mathcal{W}(\mathbf{b}, C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b} - \mathbf{1}W + \mathbf{a}, \tilde{C})] \\ &= V_n^*(\mathbf{b}, C) \end{aligned}$$

where, for each allocation $W' \in \mathcal{W}(\mathbf{b}', C)$, we define $W(W') \in \mathcal{W}(\mathbf{b}, C)$ as the allocation that assigns to each queue j the same number of servers (the same servers) as W' does unless the queue is empty, in which case it assigns only b_j . In other words, $\mathbf{1}W(W') = \mathbf{b} - [\mathbf{b} - \mathbf{1}W']^+$. In light of this, the first inequality holds by the induction hypothesis and noticing that $\mathbf{b} - \mathbf{1}W(W') = [\mathbf{b} - \mathbf{1}W']^+ \leq [\mathbf{b}' - \mathbf{1}W']^+ = \mathbf{b}' - \mathbf{1}W'$, where the $[\cdot]^+$ is removed because we know that $\mathbf{b}' \geq \mathbf{1}W'$ from $W' \in \mathcal{W}(\mathbf{b}', C)$. The second inequality holds because $W(W') \in \mathcal{W}(\mathbf{b}, C)$. The third inequality is a result of the strict monotonicity of ϕ . \square

The following Lemma 2 shows that the MTLB policy is optimal at horizon $n + 1$ if we know that $v_n \in \mathcal{F}$.

Lemma 2: If $v_n \in \mathcal{F}$, then the MTLB policy is optimal at horizon $n + 1$.

Proof: We need to show that $v_n(\mathbf{b} - \mathbf{1}W^*) = \min_{W \in \mathcal{W}(\mathbf{b}, C)} v_n(\mathbf{b} - \mathbf{1}W)$ when $W^* = [w_{ij}^*] \in \mathcal{W}(\mathbf{b}, C)$ is an MTLB allocation.

We first show that W^* must satisfy the maximum-throughput condition C1). Assume W^* does not satisfy C1), i.e., $\sum_{i,j} w_{ij}^* < L$, where L is the maximum achievable throughput. By Proposition 3, there exists at least one throughput-increasing path I . The throughput-increasing allocation $W^t(I)$ results in one more throughput and smaller leftover queues, i.e., $\mathbf{b} - 1W^t(I) < \mathbf{b} - 1W^*$. Hence, by Lemma 1, $v_n(\mathbf{b} - 1W^t(I)) < v_n(\mathbf{b} - 1W^*)$, a contradiction with the optimality of W^* .

Next, we show that W^* must also satisfy C2). Assume W^* satisfies C1) but not C2). Hence, by Proposition 4, there must exist a balancing path $S = S(W^*, i, k)$ for some queues i, k . Let W' be the corresponding balancing allocation. Let $\mathbf{I}^* = \mathbf{b} - 1W^*$ and $\mathbf{I}' = \mathbf{b} - 1W'$. Since S is a balancing path and W' is the balancing allocation, we know that $l_i^* \geq l_k^* + 2$ and $\mathbf{I}' = R_{ik}(\mathbf{I}^*)$. Using this fact and the assumption that $v_n \in \mathcal{F}$ [hence, v_n satisfies B.6)], we have $v_n(\mathbf{I}') = v_n(R_{ik}(\mathbf{I}^*)) \leq v_n(\mathbf{I}^*)$. Hence, W' is also optimal. Since any maximum-throughput allocations can be made to some MTLB allocation via some finite sequence of balancing allocations (Corollary 2), we have that any MTLB allocation is also optimal. \square

The rest of the appendix provides Lemmas 3–5, necessary to establish that $v_{n+1} \in \mathcal{F}$ if $v_n \in \mathcal{F}$, as discussed in the proof of Theorem 2. The next lemma shows that v_n satisfies B.2) for all $n = 0, \dots, T$.

Lemma 3: $v_n(\mathbf{b})$, $n = 0, \dots, T$, is permutation invariant on \mathbf{b} , i.e., $v_n(\pi(\mathbf{b})) = v_n(\mathbf{b})$ for any permutation function π .

Proof: From (13) and assumption A1), it suffices to show the permutation invariance property of $V_n^*(\mathbf{b}, C)$ for any (\mathbf{b}, C) .

Induction basis: $V_0^*(\mathbf{b}, C) = \phi(\mathbf{b}) = \sum_{j=1}^N g(b_j)$ is clearly permutation invariant.

Induction step: Assume $V_{n-1}^*(\pi(\mathbf{b}), \Pi_\pi(C)) = V_{n-1}^*(\mathbf{b}, C)$, then

$$\begin{aligned} & V_n^*(\pi(\mathbf{b}), \Pi_\pi(C)) \\ &= \min_{W \in \mathcal{W}(\pi(\mathbf{b}), \Pi_\pi(C))} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\pi(\mathbf{b}) - 1W + \mathbf{a}, \tilde{C})] \\ &\quad + \phi(\pi(\mathbf{b})) \\ &= \min_{W \in \mathcal{W}(\pi(\mathbf{b}), \Pi_\pi(C))} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\pi(\mathbf{b}) - 1W + \pi(\mathbf{a}), \Pi_\pi(\tilde{C}))] \\ &\quad + \phi(\mathbf{b}) \end{aligned}$$

where the last equality is a direct result of assumptions A1) and A2). Now, using the fact that

$$W \in \mathcal{W}(\mathbf{b}, C) \Leftrightarrow \Pi_\pi(W) \in \mathcal{W}(\pi(\mathbf{b}), \Pi_\pi(C))$$

and the induction hypotheses, we have

$$\begin{aligned} & V_n^*(\pi(\mathbf{b}), \Pi_\pi(C)) \\ &= \min_{W \in \mathcal{W}(\mathbf{b}, C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\pi(\mathbf{b}) - 1\Pi_\pi(W) + \pi(\mathbf{a}), \Pi_\pi(\tilde{C}))] \\ &\quad + \phi(\mathbf{b}) \\ &= \min_{W \in \mathcal{W}(\mathbf{b}, C)} E_{\mathbf{a}, \tilde{C}}[V_{n-1}^*(\mathbf{b} - 1W + \mathbf{a}, \tilde{C})] + \phi(\mathbf{b}) \\ &= V_n^*(\mathbf{b}, C). \end{aligned} \quad \square$$

Next, we establish Lemmas 4 and 5. Specifically, given that $\hat{v}_n \in \hat{\mathcal{F}}$, we show that $E_{\mathbf{a}, C}[\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b} - 1W + \mathbf{a})]$

satisfies B.3)–B.5) in Lemma 4 and satisfies B.6) in Lemma 5. Without loss of generality, we consider $i = 1$ and $j = 2$ in B.3)–B.6). Note that from now on we are working with the relaxed problem where overallocation is allowed, i.e., we are considering allocations in $\mathcal{W}(C)$, instead of $\mathcal{W}(\mathbf{b}, C)$. Before we proceed, for notational convenience, we write

$$\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b}) := \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b} - 1W + \mathbf{a}) \quad (22)$$

and define the set of optimal allocations as follows.

Definition 16: Define $\mathcal{X}^*(\mathbf{b}, C)$ to be the set of all optimal (not necessarily nonidling) allocations when the state of the system is (\mathbf{b}, C) . In other words, at horizon $n + 1$

$$\begin{aligned} \mathcal{X}^*(\mathbf{b}, C) &:= \left\{ W^* \in \mathcal{W}(C) : v_n([\mathbf{b} - 1W^*]^+) \right. \\ &\quad \left. = \min_{W \in \mathcal{W}(C)} v_n([\mathbf{b} - 1W]^+) \right\}. \end{aligned} \quad (23)$$

We are now ready to show the following important lemma.

Lemma 4: Assume $N = 2$ and $\hat{v}_n \in \hat{\mathcal{F}}$. For any state \mathbf{b} , $E_{\mathbf{a}, C}[\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})]$ satisfies B.3)–B.5).

Proof: By using Fact 3, it suffices to show that $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ satisfies conditions B.3)–B.5) for any realization (\mathbf{a}, C) of the arrival and connectivity processes. From B.3)–B.5) and the definition of $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ in (22), it is equivalent to show the nonnegativity of the following quantities, respectively

$$\begin{aligned} \text{[i]} \quad & \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2 - 1W) \\ & + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - 1W) \\ & - \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 - 1W) \\ & - \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_2 - 1W) \end{aligned} \quad (24)$$

$$\begin{aligned} \text{[ii]} \quad & \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 - 1W) \\ & - 2 \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - 1W) \\ & + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{e}_1 - 1W) \end{aligned} \quad (25)$$

and

$$\begin{aligned} \text{[iii]} \quad & \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{12}(\mathbf{b}') - 1W) \\ & - 2 \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - 1W) \\ & + \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{21}(\mathbf{b}') - 1W) \end{aligned} \quad (26)$$

where we let $\mathbf{b}' := \mathbf{a} + \mathbf{b}$ for convenience. The nonnegativity of (24)–(26) is shown by using the assumption that $\hat{v}_n \in \hat{\mathcal{F}}$ [hence, \hat{v}_n satisfies conditions B.1)–B.6)] and the fact that the MTLB policy is optimal at horizon $n + 1$ (Lemma 2).

Since the MTLB policy is optimal at horizon $n + 1$, we first make the following important observation.⁵

⁵As we will discuss later, this observation which is essential in proving the lemma does not hold in the case of general $N \geq 3$.

Observation 1: For $N = 2$ users, there exists an MTLB allocation $W^* \in \mathcal{X}^*(\mathbf{b}', C)$ at horizon $n + 1$ such that $W^* \in \mathcal{X}^*(\mathbf{b}', C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1, C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2, C)$.

This is because 1) adding one packet to each queue does not create any balancing paths, i.e., $W^* \in \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2, C)$; and 2) W^* can always be chosen such that it gives priority to serving queue 1, hence, adding one packet to queue 1 does not create any balancing paths, i.e., $W^* \in \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1, C)$.

Now, with this choice of $W^* \in \mathcal{X}^*(\mathbf{b}', C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1, C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2, C)$ and $\mathbf{d} := \mathbf{b}' - \mathbf{1}W^*$, we rewrite some terms in (24)–(26)

$$\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{1}W + \mathbf{e}_1 + \mathbf{e}_2) = \hat{v}_n(\mathbf{d} + \mathbf{e}_1 + \mathbf{e}_2) \quad (27)$$

$$\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{1}W) = \hat{v}_n(\mathbf{d}) \quad (28)$$

$$\min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 - \mathbf{1}W) = \hat{v}_n(\mathbf{d} + \mathbf{e}_1). \quad (29)$$

Now, we are ready to show that $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ satisfies B.3)–B.5), respectively.

i) $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ satisfies B.3)

$$\begin{aligned} & \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2 - \mathbf{1}W) \\ & + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{1}W) \\ & - \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_1 - \mathbf{1}W) \\ & - \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_2 - \mathbf{1}W) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1 + \mathbf{e}_2) + \hat{v}_n(\mathbf{d}) - \hat{v}_n(\mathbf{d} + \mathbf{e}_1) \\ & \quad - \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' + \mathbf{e}_2 - \mathbf{1}W) \\ & \geq \hat{v}_n(\mathbf{d} + \mathbf{e}_1 + \mathbf{e}_2) \\ & \quad + \hat{v}_n(\mathbf{d}) - \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - \hat{v}_n(\mathbf{d} + \mathbf{e}_2) \\ & \geq 0 \end{aligned}$$

where the equality is due to (27)–(29), the first inequality is due to the observation that $W^* \in \mathcal{X}^*(\mathbf{b}', C) \subseteq \mathcal{W}(C)$ [but not necessarily in $\mathcal{X}^*(\mathbf{b}' + \mathbf{e}_2, C)$], and the last inequality holds because $\hat{v}_n \in \hat{\mathcal{F}}$ and hence satisfying condition B.3).

ii) $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ satisfies B.4). Using (28) and (29), this is equivalent to showing

$$\hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{e}_1 - \mathbf{1}W) \geq 0.$$

To show this, we consider the following two cases, depending on whether $W^* \in \mathcal{X}^*(\mathbf{b}' - \mathbf{e}_1, C)$.

Case 1: If $W^* \in \mathcal{X}^*(\mathbf{b}' - \mathbf{e}_1, C)$, then

$$\begin{aligned} & \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{e}_1 - \mathbf{1}W) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \hat{v}_n(\mathbf{d} - \mathbf{e}_1) \geq 0 \end{aligned}$$

since \hat{v}_n satisfies B.4).

Case 2: $W^* \notin \mathcal{X}^*(\mathbf{b}' - \mathbf{e}_1, C)$. Thus, there exists a balancing path $S(W^*, 2, 1)$ from queue 2 to queue 1 and $d_2 = d_1 + 1$.

Note that since there are only two queues, this balancing path is simply a balancing path $(2, v, 1)$ for some server $v \in V$, which is connected to both queues but is being assigned to queue 1 under W^* . Hence, the balancing allocation $(W^*)^b(S)$ is in $\mathcal{X}^*(\mathbf{b}' - \mathbf{e}_1, C)$ and we have $\mathbf{1}(W^*)^b(S) = R_{12}(\mathbf{1}W^*) = \mathbf{1}W^* - \mathbf{e}_1 + \mathbf{e}_2$. In other words

$$\begin{aligned} & \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \min_{W \in \mathcal{W}(C)} \hat{v}_n(\mathbf{b}' - \mathbf{e}_1 - \mathbf{1}W) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \hat{v}_n(\mathbf{b}' - \mathbf{e}_1 - \mathbf{1}(W^*)^b(S)) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - 2\hat{v}_n(\mathbf{d}) + \hat{v}_n(\mathbf{d} - \mathbf{e}_2) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - \hat{v}_n(\mathbf{d}) - \hat{v}_n(\pi_{12}(\mathbf{d})) + \hat{v}_n(\mathbf{d} - \mathbf{e}_2) \\ & = \hat{v}_n(\mathbf{d} + \mathbf{e}_1) - \hat{v}_n(\mathbf{d}) - \hat{v}_n(\mathbf{d} + \mathbf{e}_1 - \mathbf{e}_2) + \hat{v}_n(\mathbf{d} - \mathbf{e}_2) \\ & \geq 0 \end{aligned}$$

where the second equality holds because $\mathbf{b}' - \mathbf{e}_1 - \mathbf{1}(W^*)^b(S) = \mathbf{d} - \mathbf{e}_2$, the third equality holds because \hat{v}_n satisfies B.2), the fourth equality follows from $d_2 = d_1 + 1$, and the last inequality because \hat{v}_n satisfies B.3).

iii) $\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})$ satisfies B.5). This is equivalent to showing

$$\min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{12}(\mathbf{b}') - \mathbf{1}W) - 2\hat{v}_n(\mathbf{d}) + \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{21}(\mathbf{b}') - \mathbf{1}W) \geq 0. \quad (30)$$

To show this, we consider the following three cases.

Case 1: If $W^* \in \mathcal{X}^*(R_{21}(\mathbf{b}'), C) \cap \mathcal{X}^*(R_{12}(\mathbf{b}'), C)$, then we are done since \hat{v}_n satisfies B.5).

Case 2: If $W^* \notin \mathcal{X}^*(R_{21}(\mathbf{b}'), C)$, then there exists a balancing path $S(W^*, 1, 2)$. Since $W^* \in \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1, C)$ (i.e., W^* was chosen to give priority to serving queue 1), we must have $d_1 = d_2$. Hence, $(W^*)^b(S) \in \mathcal{X}^*(R_{21}(\mathbf{b}'), C)$. Furthermore, $\mathbf{1}(W^*)^b(S) = R_{21}(\mathbf{1}W^*)$. Thus, for this case, the left-hand side of (30) becomes

$$\begin{aligned} & \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{12}(\mathbf{b}') - \mathbf{1}W) - 2\hat{v}_n(\mathbf{d}) \\ & + \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{21}(\mathbf{b}') - \mathbf{1}W) \\ & = \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{12}(\mathbf{b}') - \mathbf{1}W) - 2\hat{v}_n(\mathbf{d}) + \hat{v}_n(\mathbf{d}) \\ & = \min_{W \in \mathcal{W}(C)} \hat{v}_n(R_{12}(\mathbf{b}') - \mathbf{1}W) - \hat{v}_n(\mathbf{d}). \quad (31) \end{aligned}$$

Next we consider the two following subcases.

— *Case 2.1:* If $W^* \in \mathcal{X}^*(R_{12}(\mathbf{b}'), C)$, then (31) is equal to $\hat{v}_n(R_{12}(\mathbf{d})) - \hat{v}_n(\mathbf{d}) \geq 0$, because $d_1 = d_2$ and $\hat{v}_n \in \hat{\mathcal{F}}$.

— *Case 2.2:* If $W^* \notin \mathcal{X}^*(R_{12}(\mathbf{b}'), C)$, then there exists a balancing path $S(W^*, 2, 1)$. Hence, $(W^*)^b(S) \in \mathcal{X}^*(R_{12}(\mathbf{b}'), C)$, ensuring that $\mathbf{1}(W^*)^b(S) = R_{12}(\mathbf{1}W^*)$. Thus, (31) is equal to zero.

Case 3: If $W^* \in \mathcal{X}^*(R_{21}(\mathbf{b}'), C)$ but $W^* \notin \mathcal{X}^*(R_{12}(\mathbf{b}'), C)$, then, by the permutation invariance property of \hat{v}_n , this reduces to Case 2.1 above.

Lemma 5: Assume $N = 2$ and $\hat{v}_n \in \hat{\mathcal{F}}$. For any state \mathbf{b} such that $b_1 \geq b_2 + 1$, $E_{\mathbf{a}, C}[\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})]$ satisfies condition B.6).

Proof: To show that $E_{\mathbf{a}, C}[\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b})]$ meets condition B.6) is equivalent to show the nonnegativity of $E_{\mathbf{a}, C}[\mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b}) - \mathcal{T}_n^{\mathbf{a}, C}(R_{12}(\mathbf{b}))]$. For convenience, let us define

$$Z_n^{\mathbf{a}, C}(\mathbf{b}) := \mathcal{T}_n^{\mathbf{a}, C}(\mathbf{b}) - \mathcal{T}_n^{\mathbf{a}, C}(R_{12}(\mathbf{b})). \quad (32)$$

Using the permutation invariance property of the arrival and connectivity processes [assumptions A1) and A2)], we can rewrite $E_{\mathbf{a},C}[Z^{\mathbf{a},C}(\mathbf{b})]$ as

$$E_{\mathbf{a},C}[Z^{\mathbf{a},C}(\mathbf{b})] = \frac{1}{2}E_{\mathbf{a},C}\left[Z^{\mathbf{a},C}(\mathbf{b}) + Z^{\pi_{12}(\mathbf{a}),\Pi_{\pi_{12}}(C)}(\mathbf{b})\right].$$

Thus, it suffices to show that, for any (\mathbf{a}, C) and $b_1 \geq b_2 + 1$

$$Z^{\mathbf{a},C}(\mathbf{b}) + Z^{\pi_{12}(\mathbf{a}),\Pi_{\pi_{12}}(C)}(\mathbf{b}) \geq 0.$$

We show this by noticing that

$$\begin{aligned} & Z^{\mathbf{a},C}(\mathbf{b}) + Z^{\pi_{12}(\mathbf{a}),\Pi_{\pi_{12}}(C)}(\mathbf{b}) \\ &= Z^{\mathbf{a},C}(\mathbf{b}) + Z^{\mathbf{a},C}(\pi_{12}(\mathbf{b})) \\ &= Z^{\mathbf{a},C}(\mathbf{b}) + \mathcal{T}_n^{\mathbf{a},C}(\pi_{12}(\mathbf{b})) - \mathcal{T}_n^{\mathbf{a},C}(\pi_{12}(R_{12}(\mathbf{b}))) \\ &= \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^0) - \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^1) \\ &\quad + \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^M) - \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^{M-1}) \end{aligned} \quad (33)$$

where $M := b_1 - b_2 (\geq 1)$ and $\mathbf{b}^m := \mathbf{b} - m\mathbf{e}_1 + m\mathbf{e}_2$, for $m = 0, \dots, M$. The first equality follows from the permutation invariance property, while the second and third equalities follow from (32). Note that $\pi_{12}(\mathbf{b}) = b_2\mathbf{e}_1 + b_1\mathbf{e}_2 = \mathbf{b} - (b_1 - b_2)\mathbf{e}_1 + (b_1 - b_2)\mathbf{e}_2 = \mathbf{b}^M$, $\pi_{12}(R_{12}(\mathbf{b})) = \mathbf{b}^{M-1}$, $\mathbf{b}^{m+1} = R_{12}(\mathbf{b}^m)$, and $\mathbf{b}^{m-1} = R_{21}(\mathbf{b}^m)$.

Now notice that if $M = 1$, then the right-hand side of (33) is zero. If $M \geq 2$, we have

$$\begin{aligned} & \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^0) - \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^1) - \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^{M-1}) + \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^M) \\ &= \sum_{m=1}^{M-1} \{\mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^{m-1}) - 2\mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^m) + \mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^{m+1})\} \\ &= \sum_{m=1}^{M-1} \{\mathcal{T}_n^{\mathbf{a},C}(R_{21}(\mathbf{b}^m)) - 2\mathcal{T}_n^{\mathbf{a},C}(\mathbf{b}^m) + \mathcal{T}_n^{\mathbf{a},C}(R_{12}(\mathbf{b}^m))\} \\ &\geq 0 \end{aligned}$$

where the inequality holds because $\hat{v}_n \in \hat{\mathcal{F}}$ and, from Lemma 4, $\mathcal{T}_n^{\mathbf{a},C}(\cdot)$ satisfies condition B.5). \square

Remark 4: The proofs for Lemmas 4 and 5 are valid only for $N = 2$. The main difficulty in the extension to the general case of $N > 2$ is with Lemma 4, Observation 1, where we cannot claim that there exists an MTLB allocation $W^* \in \mathcal{X}^*(\mathbf{b}', C)$ such that $W^* \in \mathcal{X}^*(\mathbf{b}', C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1, C) \cap \mathcal{X}^*(\mathbf{b}' + \mathbf{e}_1 + \mathbf{e}_2, C)$. This reflects a major obstacle that adding one packet to queue 1 and/or queue 2 may generate a balancing path in the original optimal allocation. For general N , we will need to explore more cases and require extra convexity properties of functions in \mathcal{F} . Currently, we do not know which extra conditions are needed and how to show that these conditions of v_n carried over to v_{n+1} . Therefore, the extension to the general case of $N > 2$ remains open.

APPENDIX III

SUPPORTING LEMMAS AND PROOF OF THEOREM 3

In this appendix, we prove the optimality of the MTLB-F policy for the fluid server allocation relaxation. We assume that the cost function $\phi(\mathbf{b})$ is monotonically increasing, permutation invariant, and convex on $\mathbf{b} \in \mathbb{R}_+^N$. It is easy to see that the strict monotonicity and permutation invariance of v_n in the fluid relaxation (Lemmas 1 and 3) still hold for all n . Next we show the convexity of v_n for all n .

Lemma 6: $v_n(\mathbf{b})$ is convex on \mathbf{b} for all $n = 0, 1, \dots, T$.

Proof: It suffices to show the convexity of $V_n^*(\mathbf{b}, C)$ for every C . We show this by induction: $V_0^*(\mathbf{b}, C) = \phi(\mathbf{b})$ is convex on \mathbf{b} . Assume $V_{n-1}^*(\mathbf{b}, C)$ convex on \mathbf{b} . For $i = 1, 2$, let $W^i \in \mathcal{W}^f(\mathbf{b}^i, C)$ be an optimal allocation at time n for (\mathbf{b}^i, C) , i.e.,

$$V_n^*(\mathbf{b}^i, C) = \phi(\mathbf{b}^i) + E_{\mathbf{a},\tilde{C}}\left[V_{n-1}^*(\mathbf{b}^i + \mathbf{a} - \mathbf{1}W^i, \tilde{C})\right]. \quad (34)$$

Then, for any $\beta \in [0, 1]$, let $W^\beta = \beta W^1 + (1-\beta)W^2$ and $\mathbf{b}^\beta = \beta\mathbf{b}^1 + (1-\beta)\mathbf{b}^2$. We can easily see that $W^\beta \in \mathcal{W}^f(\mathbf{b}^\beta, C)$ since it satisfies the conditions a')-d) in Definition 9. Then, we have

$$\begin{aligned} & V_n^*(\mathbf{b}^\beta, C) \\ &= \phi(\mathbf{b}^\beta) + \min_{W \in \mathcal{W}^f(\mathbf{b}^\beta, C)} E_{\mathbf{a},\tilde{C}}\left[V_{n-1}^*(\mathbf{b}^\beta + \mathbf{a} - \mathbf{1}W, \tilde{C})\right] \\ &\leq \phi(\mathbf{b}^\beta) + E_{\mathbf{a},\tilde{C}}\left[V_{n-1}^*(\mathbf{b}^\beta + \mathbf{a} - \mathbf{1}W^\beta, \tilde{C})\right] \\ &\leq \phi(\mathbf{b}^\beta) + E_{\mathbf{a},\tilde{C}}\left[\beta V_{n-1}^*(\mathbf{b}^1 + \mathbf{a} - \mathbf{1}W^1, \tilde{C})\right] \\ &\quad + E_{\mathbf{a},\tilde{C}}\left[(1-\beta)V_{n-1}^*(\mathbf{b}^2 + \mathbf{a} - \mathbf{1}W^2, \tilde{C})\right] \\ &\leq \beta \left(\phi(\mathbf{b}^1) + E_{\mathbf{a},\tilde{C}}\left[V_{n-1}^*(\mathbf{b}^1 + \mathbf{a} - \mathbf{1}W^1, \tilde{C})\right]\right) \\ &\quad + (1-\beta) \left(\phi(\mathbf{b}^2) + E_{\mathbf{a},\tilde{C}}\left[V_{n-1}^*(\mathbf{b}^2 + \mathbf{a} - \mathbf{1}W^2, \tilde{C})\right]\right) \\ &= \beta V_n^*(\mathbf{b}^1, C) + (1-\beta)V_n^*(\mathbf{b}^2, C) \end{aligned}$$

where the second inequality follows from the induction hypothesis, the last inequality from the convexity of the cost function ϕ , and the last equality from (34). \square

The notions of alternating, balancing, and throughput-increasing paths and allocations can be generalized as follows.

Definition 17: For $\epsilon > 0$, an ϵ -alternating path from queue $u_0 \in U$ to queue $u_k \in U$ with respect to $W \in \mathcal{W}^f(\mathbf{b}, C)$ is a sequence of distinct vertices

$$S(W, u_0, u_k, \epsilon) := (u_0, v_1, u_1, v_2, \dots, v_k, u_k)$$

with $v_i \in V$, $u_i \in U$, $w_{v_i, u_i} \geq \epsilon$ for each $i = 1, \dots, k$. An ϵ -alternating path $S(W, u_0, u_k, \epsilon)$ is called an ϵ -balancing path if $(b_{u_0} - \sum_{i=1}^k w_{i, u_0}) - (b_{u_k} - \sum_{i=1}^k w_{i, u_k}) \geq 2\epsilon$.

Definition 18: For $\epsilon > 0$, an ϵ -throughput-increasing path relative to $W \in \mathcal{W}^f(\mathbf{b}, C)$ is a sequence of distinct vertices

$$I(W, v_0, u_k, \epsilon) := (v_0, u_1, v_1, u_2, \dots, v_{k-1}, u_k)$$

with a) $v_i \in V$, $u_i \in U$, and $w_{v_i, u_i} \geq \epsilon$ for each i ; b) v_0 is not fully matched, i.e., $1 - \sum_{j=1}^N w_{v_0, j} \geq \epsilon$; and c) u_k is nonempty, i.e., $b_{u_k} - \sum_{i=1}^K w_{i, u_k} \geq \epsilon$.

Definition 19: Given an ϵ -balancing path $S = S(W, u_0, u_k, \epsilon)$ for $\epsilon > 0$, an ϵ -balancing allocation $W^{b, \epsilon}(S)$ balances the queue u_0 and u_k by ϵ packets by reassigning server v_l to serve ϵ packets more from queue u_{l-1} and ϵ packets less from queue u_l , $\forall l = 1, \dots, k$.

Definition 20: Given an ϵ -throughput-increasing path $I = I(W, v_0, u_k, \epsilon)$ for $\epsilon > 0$, an ϵ -throughput-increasing alloca-

tion $W^{t,\epsilon}(I) \in \mathcal{W}^f(\mathbf{b}, C)$ achieves additional throughput of ϵ packets by assigning v_0 to serve ϵ more packets from u_1 and reassigning server v_l to serve ϵ packets more from queue u_{l+1} and ϵ packets less from queue $u_l, \forall l = 1, \dots, k-1$.

Note that the throughput-increasing and balancing paths previously considered under the integral server allocation (Definitions 13 and 14) are equivalent to the fluid versions (Definitions 17 and 18), when we take $\epsilon = 1$.

We see that similar results as in Appendix I hold for the ϵ -throughput-increasing and ϵ -balancing paths as well. The existence of the MTLB-F policy could be similarly established as in Theorem 1 using the following result.⁶

Proposition 3': An allocation achieves the MT C1) if and only if it has no ϵ -throughput-increasing paths for any $\epsilon > 0$.

Proof: Similar to the proof in Proposition 3. \square

Proposition 4': Any allocation satisfying the maximum-throughput condition C1) also satisfies the LB condition C2) if and only if it has no ϵ -balancing path for any $\epsilon > 0$.

Proof: The proof is very similar to that of Proposition 4 but here we allow $\epsilon \in (0, 1]$ packets to be reallocated. It suffices to show that if a maximum-throughput allocation $W = [w_{i,j}] \in \mathcal{W}^f(\mathbf{b}, C)$ does not satisfy C2), then W has at least one ϵ -balancing path for some $\epsilon > 0$. Let $W^* = [w_{i,j}^*] \in \mathcal{W}^f(\mathbf{b}, C)$ be an MTLB-F allocation, chosen such that $\|W^* - W\|$ is minimized, where $\|X\| := \sum_{i,j} |x_{i,j}|$ for any matrix $X = [x_{i,j}]$. Now let G_d be the weighted subgraph of the bipartite graph $G = (V, U, C)$ induced by the allocation difference matrix $W^* - W$. Specifically, G_d contains an edge (v, u) for $v \in V$ and $u \in U$ if and only if $|w_{v,u}^* - w_{v,u}| > 0$. We assign the weight of the edge (v, u) as $|w_{v,u}^* - w_{v,u}|$. Color the edges (v, u) of G_d green if $w_{v,u}^* - w_{v,u} > 0$, and red if $w_{v,u}^* - w_{v,u} < 0$. Direct the green edges from V to U and the red edges from U to V . Let the leftover queue vectors under W and W^* be $\mathbf{l} = \mathbf{b} - \mathbf{1}W$ and $\mathbf{l}^* = \mathbf{b} - \mathbf{1}W^*$, respectively.

We claim that for every directed path P in G_d from $u_1 \in U$ to $u_2 \in U$, we have

$$l_{u_1}^* \leq l_{u_2}^*. \quad (35)$$

To see this, let $P = (u_1, \dots, u_2)$ be a directed path in G_d . By the choice of the directions for the edges, P must be alternating between red and green edges. Let ϵ' be the minimum of the weights of the edges along this path ($\epsilon' > 0$ by construction of G_d). If $l_{u_1}^* > l_{u_2}^*$, then P is an ϵ -balancing path for W^* with $\epsilon = \min\{\epsilon', (l_{u_1}^* - l_{u_2}^*)/2\} > 0$, contradicting to the assumption that W^* satisfy C2).

Next, we claim that G_d is acyclic. Assume G_d is cyclic, i.e., there is a red-green alternating and directed path $P = (u_1, \dots, u_1)$ in G_d , for some $u_1 \in U$, with the minimum weight $\epsilon > 0$ along this path. Then, we get another MTLB-F allocation $W^{**} = [w_{i,j}^{**}]$ by letting $w_{v,u}^{**} = w_{v,u}^* - \epsilon$ for all green edges $(v, u) \in P$, $w_{v,u}^{**} = w_{v,u}^* + \epsilon$ for all red edges $(v, u) \in P$, and $w_{v',u'}^{**} = w_{v',u'}^*$ for all edges (v', u') not in P . We see that W^{**} is closer to W than W^* is to W because $|w_{v,u}^{**} - w_{v,u}| = |w_{v,u}^* - w_{v,u}| - \epsilon \geq 0$ for each edge (v, u) in

⁶One can argue the existence by seeing that this is equivalent to minimizing $\max(\mathbf{b} - \mathbf{1}W)$ on a compact set.

P . In other words, $\|W^{**} - W\| < \|W^* - W\|$, in contradiction to the choice of W^* . Hence, we must have that G_d is acyclic.

Now, since both W^* and W achieve the MT, we have that $\sum_{i=1}^N l_i = \sum_{i=1}^N l_i^*$. But since W does not satisfy the LB condition C2), there must exist $u_1 \in U$ such that

$$l_{u_1} < l_{u_1}^*. \quad (36)$$

Obviously, there is a red edge directed out of u_1 . Starting from u_1 , we build an alternating red-green and directed path P' in G_d as follows. 1) From an arbitrary vertex $u \in U$ (including u_1), if there is a red edge directed out of u and $l_u \leq l_{u_1}$, we build P' by arbitrarily select one of the red edges directed out of u . 2) From an arbitrary $v \in V$, we build P' by arbitrarily following one of the green edges directed out of v . Such a green edge always exist.⁷ 3) Otherwise, stop.

Using the fact that G_d is acyclic, P' is well defined and finite. Let $u_2 \in U$ be the final vertex on the path and $\epsilon' > 0$ be the minimum weight along P' . There are two possible cases.

Case 1: $l_{u_2} > l_{u_1}$. In this case, we reverse the order of nodes in P' to arrive at an ϵ -balancing path relative to W , where $\epsilon = \min\{\epsilon', (l_{u_2} - l_{u_1})/2\} > 0$.

Case 2: There is no red edge directed out of u_2 . Thus, P' arrived at u_2 via a green edge (v, u_2) for some $v \in V$. This means that u_2 is served at least $w_{v,u_2}^* - w_{v,u_2} > 0$ packets more under W^* , relative to W . Hence, $l_{u_2} > l_{u_2}^*$. This together with (35) and (36) give $l_{u_2} > l_{u_2}^* \geq l_{u_1}^* > l_{u_1}$, which means $l_{u_2} > l_{u_1}$. Reversing the order of nodes in P' gives an ϵ -balancing path relative to W , where $\epsilon = \min\{\epsilon', (l_{u_2} - l_{u_1})/2\} > 0$.

Since there exists an ϵ -balancing path in both cases for some $\epsilon > 0$, we have the assertion of the proposition.

Using the above results, we can show the following.

Theorem 3: For Problem (P) with the fluid server allocation relaxation, the MTLB-F policy is optimal.

Proof: We need to show that

$$v_n(\mathbf{b} - \mathbf{1}W^*) = \min_{W \in \mathcal{W}^f(\mathbf{b}, C)} v_n(\mathbf{b} - \mathbf{1}W), \quad (37)$$

where $W^* \in \mathcal{W}^f(\mathbf{b}, C)$ is MTLB-F. Since $\mathcal{W}^f(\mathbf{b}, C)$ is convex and compact, there exists an optimal allocation W^* . Similarly as in the proof of Lemma 2, we can show that W^* must satisfy C1) using Proposition 3' and the strict monotonicity of v_n (Lemma 1).

Now assume W^* satisfies C1) but not C2). By Proposition 4', there must exist an ϵ -balancing path $S = S(W^*, i, k, \epsilon)$ relative to W^* , for some $\epsilon > 0$ and some queues $i, k \in U$. Let $W' = W^{b,\epsilon}(S)$ be the corresponding ϵ -balancing allocation. Let $\mathbf{l}^* = \mathbf{b} - \mathbf{1}W^*$ and $\mathbf{l}' = \mathbf{b} - \mathbf{1}W'$. We then have $l'_i = l_i^* - \epsilon$, $l'_k = l_k^* + \epsilon$, and $l'_u = l_u^*$ for all $u \neq i, k$. By the convexity and permutation invariance properties of v_n , we can show that $v_n(\mathbf{l}') \leq v_n(\mathbf{l}^*)$ [4] as follows. Since \mathbf{l}^* and \mathbf{l}' differ only in the i th and k th components, it suffices to consider a function v_n of two variables. We notice that $(l_i^* - \epsilon, l_k^* + \epsilon)$ lies on the interval joining (l_i^*, l_k^*) and (l_k^*, l_i^*) . Hence, for some $\gamma \in [0, 1]$, we have $(l_i^* - \epsilon, l_k^* + \epsilon) = \gamma(l_i^*, l_k^*) + (1 - \gamma)(l_k^*, l_i^*)$. Using

⁷Otherwise, combining one of the red edges coming into v with W^* would have yielded a nonidling feasible allocation with additional positive packet throughput, a contradiction to C1).

the convexity and the permutation invariance of v_n , we obtain $v_n(l_i^* - \epsilon, l_k^* + \epsilon) \leq \gamma v_n(l_i^*, l_k^*) + (1 - \gamma)v_n(l_k^*, l_i^*) = v_n(l_i^*, l_k^*)$. Hence, $v_n(\mathbf{I}') \leq v_n(\mathbf{I}^*)$ and W' is also optimal but more balanced than W^* . Thus, we can conclude that any MTLB-F allocation is optimal. \square

ACKNOWLEDGMENT

The authors would like to thank Prof. Ladner and Tamir and the reviewers for their tremendously helpful suggestions.

REFERENCES

- [1] C. Y. Wong, R. Cheng, K. Lataief, and R. Murch, "Multiuser OFDM with adaptive subcarrier, bit, and power allocation," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 10, pp. 1747–1758, Oct. 1999.
- [2] R. Knopp and P. A. Humblet, "Information capacity and power control in single-cell multiuser communications," in *Proc. IEEE Int. Conf. Commun.*, Seattle, WA, Jun. 1995, pp. 331–335.
- [3] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 466–478, Mar. 1993.
- [4] A. Ganti, E. Modiano, and J. N. Tsitsiklis, "Optimal transmission scheduling in symmetric communication models with intermittent connectivity," *IEEE Trans. Inf. Theory*, vol. 53, no. 3, pp. 998–1008, Mar. 2007.
- [5] A. Ganti, "Transmission scheduling for multi-beam satellite systems," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, 2003.
- [6] B. Hajek, "Optimal control of two interacting service stations," *IEEE Trans. Autom. Control*, vol. AC-29, no. 6, pp. 491–499, Jun. 1984.
- [7] N. Ehsan and M. Liu, "Optimal bandwidth allocation in a delay channel," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1614–1626, Aug. 2006.
- [8] C. Lott and D. Teneketzis, "On the optimality of an index rule in multichannel allocation for single-hop mobile networks with multiple service classes," *Probab. Eng. Inf. Sci.*, vol. 14, no. 3, pp. 259–297, Jul. 2000.
- [9] S. Shakkottai, R. Srikant, and A. L. Stolyar, "Pathwise optimality of the exponential scheduling rule for wireless channels," *Adv. Appl. Probab.*, vol. 36, no. 4, pp. 1021–1045, 2004.
- [10] E. M. Yeh and A. S. Cohen, "Delay optimal rate allocation in multiaccess fading communications," in *Proc. Allerton Conf. Commun. Control Comput.*, Monticello, IL, 2004.
- [11] M. Andrews, A. Stolyar, K. Kumaran, R. Vijayakumar, K. Ramanan, and P. Whiting, "Scheduling in a queuing system with asynchronously varying service rates," *Probab. Eng. Inf. Sci.*, vol. 18, pp. 191–217, 2004.
- [12] S. Musy, "A delay optimal policy for symmetric broadcast channels," in *Proc. Int. Conf. Commun. Technol.*, Nov. 2006.
- [13] T. Javidi, N. Song, and D. Teneketzis, "Expected makespan minimization on identical machines in two interconnected queues," *Probab. Eng. Inf. Sci.*, vol. 15, no. 4, pp. 409–443, 2001.
- [14] G. Koole, "Convexity in tandem queues," *Prob. Eng. Inf. Sci.*, pp. 13–31, 2004.
- [15] G. Koole, "Monotonicity in Markov reward and decision chains: Theory and applications," *Found. Trends Stochastic Syst.*, pp. 1–76, 2006.
- [16] N. Ehsan and T. Javidi, "Delay optimal transmission policy in a wireless multiaccess channel," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3745–3751, Aug. 2008.
- [17] R. Kumar and P. Varaiya, *Stochastic Control*. Englewood Cliffs, NJ: Prentice-Hall, 1986.
- [18] D. P. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [19] K. M. Naheer, *The LEDA Platform of Combinatorial and Geometric Computing*. Cambridge, U.K.: Cambridge Univ. Press, 1999.
- [20] E. Altman, B. Gaujal, and A. Hordijk, *Discrete-Event Control of Stochastic Networks: Multimodularity and Regularity*. Berlin, Germany: Springer-Verlag, 2003.
- [21] S. Kittipiyakul and T. Javidi, "Resource allocation in OFDMA with time-varying channel and bursty arrivals," *IEEE Commun. Lett.*, vol. 11, no. 9, pp. 708–710, Sep. 2007.
- [22] N. Harvey, R. Ladner, L. Lovasz, and T. Tamir, "Semi-matchings for bipartite graphs and load balancing," in *Proc. Workshop Algorithms Data Structures*, Ottawa, ON, Canada, Jul. 2003.
- [23] U. Manber, *Introduction to Algorithm: A Creative Approach*. Reading, MA: Addison-Wesley, 1989.

Somsak Kittipiyakul received the S.B. and M.Eng. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1996 and the Ph.D. degree in electrical and computer engineering from the University of California, San Diego, in 2008.

Currently, he is with the School of Information, Computer and Communication Technology, Sirindhorn International Institute of Technology, Thammasat University, Pathum Thani, Thailand. His research interests include wireless communication and networking, stochastic resource allocation and scheduling, and queuing analysis.

Tara Javidi (S'96–M'02) studied electrical engineering at Sharif University of Technology, Tehran, Iran, from 1992 to 1996. She received the M.S. degrees in electrical engineering (systems) and in applied mathematics (stochastics) and the Ph.D. degree in electrical engineering and computer science from the University of Michigan, Ann Arbor, in 1998, 1999, and 2002, respectively.

From 2002 to 2004, she was an Assistant Professor at the Electrical Engineering Department, University of Washington, Seattle. She joined University of California, San Diego, in 2005, where she is currently an Assistant Professor of Electrical and Computer Engineering. Her research interests are in communication networks, stochastic resource allocation, stochastic control theory, and wireless communications.

Dr. Javidi was a Barbour Scholar during 1999–2000 academic year and received an NSF CAREER Award in 2004.