

THE TROUBLE WITH DIVERSITY: FORK-JOIN NETWORKS WITH HETEROGENEOUS CUSTOMER POPULATION

BY VIÊN NGUYEN

Sloan School of Management, M.I.T.

Consider a feedforward network of single-server stations populated by multiple job types. Each job requires the completion of a number of tasks whose order of execution is determined by a set of deterministic precedence constraints. The precedence requirements allow some tasks to be done in parallel (in which case tasks would “fork”) and require that others be processed sequentially (where tasks may “join”). Jobs of a given type share the same precedence constraints, interarrival time distributions and service time distributions, but these characteristics may vary across different job types. We show that the heavy traffic limit of certain processes associated with heterogeneous fork-join networks can be expressed as a semimartingale reflected Brownian motion with polyhedral state space. The polyhedral region typically has many more faces than its dimension, and the description of the state space becomes quite complicated in this setting. One can interpret the proliferation of additional faces in heterogeneous fork-join networks as (i) articulations of the fork and join constraints and (ii) consequences of the *disordering effects* that occur when jobs fork and join in their sojourns through the network.

1. Introduction and summary. We consider in this paper the class of feedforward fork-join networks with heterogeneous customer populations. The network, which consists of d single-server stations, is populated by multiple job types. Each job requires the completion of a number of tasks whose order of execution is determined by a set of deterministic precedence constraints. The precedence requirements allow some tasks to be performed in parallel, whereupon tasks would *fork*, or *split*, and require that others be processed sequentially, in which case tasks may *join*, or *match*. (Incidentally, fork-join networks are also known as “split and match” networks.) Jobs of a given type share the same precedence constraints, interarrival time distributions and service time distributions, but these characteristics may vary across different job types. We restrict attention to the case where the network is *feedforward*, that is, stations are numbered in such a way that jobs always flow from lower numbered stations to higher numbered ones.

We present a heavy traffic analysis for processing networks of the type just described. In a previous paper [17], Nguyen showed that when the customer population is *homogeneous*—that is, when all customers share the same

Received November 1992; revised July 1993.

AMS 1991 subject classifications. 60K25, 60J65, 60F17, 90B30.

Key words and phrases. Fork-join networks, heterogeneous customer populations, reflected Brownian motion, nonsimple polyhedral state space, diffusion approximations, heavy traffic analysis.

precedence requirements, interarrival time distributions and service time distributions—the heavy traffic limit of the network is described by a d -dimensional semimartingale reflected Brownian motion (SRBM) whose state space is a nonsimple convex polyhedral cone in the nonnegative orthant. Unlike the corresponding results for conventional queueing networks (that is, networks with strictly sequential processing), the number of faces in the polyhedral region is greater than its dimension, which throughout this paper we will denote by d . One can interpret the presence of additional faces as articulations of synchronization constraints embodied in the fork and join constructs.

In this paper, we show that the heavy traffic limit of certain processes associated with *heterogeneous* fork–join networks can also be expressed as d -dimensional SRBM's with polyhedral state space. We will also show that the polyhedral region typically has many more faces than its homogeneous counterpart and that the description of the state space becomes vastly more complicated in this setting. This result is surprising when compared to those associated with conventional queueing networks, where the form of the limiting process does not change in the presence of multiple customer types. (This is an example of the “state-space collapse” phenomenon described in [18] and [20].) The proliferation of additional faces in heterogeneous fork–join networks arises from the *disordering effects* that occur when jobs fork and join in their sojourns through the network.

Processing systems that are characterized by parallel as well as sequential processing exist in many industrial settings. Readers may refer to Baccelli and Makowski [5], Mandelbaum and Avi-Itzhak [15] and Nguyen [17] for a survey of several interesting applications. From a practical point of view, the generalization of [17] to allow multiple job types constitutes an important extension. Most current studies of fork–join networks assume that all customers are statistically similar [5]. Baccelli and Liu [4] consider a fork–join network in which a job may send *batches* of tasks (which may include more than one task) to processing stations, and jobs of different types send batches of different sizes. Still, Baccelli and Liu require that all jobs share the same feedforward deterministic routing structure. Such a model can represent, for example, systems in which some processing stations are capable of performing more than one kind of task; Baccelli and Liu are motivated by multiprocessor systems running parallel programs.

The recent works by Adler, Mandelbaum, Nguyen and Schwerer [1–3] propose a processing network model for studying new product development. The model they describe, which they simply call a processing network model, is more encompassing than the class of fork–join networks studied here. The key restriction in this paper, which is not assumed in [1–3], is that jobs must visit workstations in a *feedforward* manner. The possibility of feedback in the network is an important generalization to be considered in future work. Moreover, as the work by Adler, Mandelbaum, Nguyen and Schwerer demonstrates, heterogeneity in the customer population is an essential characteristic that must be captured.

The paper is organized as follows. We give a formal description of the model in Section 2 and define the processes of interest Section 3. In order to state the heavy traffic limit results for these processes, one must refer to a “sequence of systems.” Section 4 defines such a sequence. Before stating the heavy traffic limit theorems, we introduce some additional notation and preliminary results in Section 5. The main theorems are then summarized in Section 6, where we also illustrate the heavy traffic limit theorem for several special cases. The proofs of these theorems are then given in Section 7. Finally, some concluding remarks are given in Section 8.

We end this section with some technical preliminaries. The space $\mathbf{D}^r[0, \infty)$ is the r -dimensional product space of functions $f: [0, \infty) \rightarrow \mathfrak{R}^r$ that are right continuous on $[0, \infty)$ and have left limits on $(0, \infty)$. The space $\mathbf{D}^r[0, \infty)$ is endowed with the Skorohod topology [6]. For X^n a sequence of processes in $\mathbf{D}^r[0, \infty)$ and $X \in \mathbf{D}^r[0, \infty)$, we write $X^n \Rightarrow X$ to mean X^n converges to X in distribution.

For $f: [0, \infty) \rightarrow \mathfrak{R}$, set $\|f\|_t \equiv \sup_{0 \leq s \leq t} |f(s)|$, and for a vector-valued function $f = (f_1, f_2, \dots, f_r)': [0, \infty) \rightarrow \mathfrak{R}^r$, we let

$$\|f\|_t \equiv (\|f_1\|_t, \dots, \|f_r\|_t)'.$$

A sequence of functions $\{f^n\}$ converges to a function f uniformly on compact sets (u.o.c.) if, for each $t \geq 0$, $\|f^n - f\|_t \rightarrow 0$ as $n \rightarrow \infty$. For a sequence of stochastic processes $\{X^n\}$ on $\mathbf{D}^r[0, \infty)$ and X a process in $\mathbf{D}^r[0, \infty)$, we write $X^n \rightarrow X$ u.o.c. if, almost surely, X^n converges to X uniformly on compact sets.

In our heavy traffic limit theorems, the weak limit obtained is a semimartingale reflected Brownian motion whose state space is a polyhedral cone in the nonnegative orthant. A Brownian motion process having drift vector θ and covariance matrix Γ will be denoted as $(\theta, \Gamma)\text{BM}$; likewise, a semimartingale reflected Brownian motion with these drift and covariance parameters, reflection matrix \mathbf{R} and state space \mathcal{S} is denoted as $(\mathcal{S}, \theta, \Gamma, \mathbf{R})\text{SRBM}$.

2. Model description. The network under consideration consists of d single-server stations and hosts p types of jobs. Jobs of type q arrive to the network at rate α_q . Each type q job requires completion of a number of activities. Hereafter, we refer to a job-type/activity pair as a *task* or a *class*, interchangeably. Task k receives service from station $j = s(k)$, and we denote by τ_k the mean service time for task k . Letting \mathcal{A}_q denote the set of tasks (or classes) in job type q , set

$$\mathcal{A} \equiv \mathcal{A}_1 \cup \dots \cup \mathcal{A}_p = \{1, \dots, K\}.$$

Our convention will be to index workstations by $i, j = 1, \dots, d$, job types by $q, r = 1, \dots, p$ and tasks by $k, l = 1, \dots, K$. For notational convenience, we define $\lambda_k = \alpha_q$ for all tasks $k \in \mathcal{A}_q$ and we write $q(k)$ to mean the job type q for which $k \in \mathcal{A}_q$.

The order in which tasks are executed is determined by a set of deterministic precedence constraints, which are articulated by way of a $K \times K$ prece-

precedence matrix $\mathbf{P} = (P_{kl})$ defined as follows:

$$(2.1) \quad P_{kl} \equiv \begin{cases} 1, & \text{if task } k \text{ is an immediate predecessor for task } l, \\ 0, & \text{otherwise.} \end{cases}$$

(Because all elements of the precedence matrix \mathbf{P} are 0's and 1's, routing of tasks is clearly deterministic within each job type.) We assume that there exists a column and row permutation of \mathbf{P} such that the resulting matrix is strictly upper triangular. In terms of the model, this means that each task is performed exactly once and is never repeated. From the precedence matrix \mathbf{P} , we can now define the set of immediate predecessors as

$$(2.2) \quad \mathcal{P}(l) \equiv \{k \in \mathcal{A}: P_{kl} = 1\}.$$

From the modelling point of view, $\mathcal{P}(l)$ is the set of tasks that must be completed immediately prior to commencement of task l .

For the two job types depicted in Figure 1, $\mathcal{A}_1 = \{1, 2, 3, 4\}$, $\mathcal{A}_2 = \{5, 6, 7\}$, $\mathcal{P}(4) = \{1, 3\}$ and the precedence matrix \mathbf{P} is given by

$$\mathbf{P} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We allow tasks to map to stations in a many-to-one fashion, implying that while each station may be capable of performing several types of tasks, each task is performed at exactly one station. We define the *constituency* of station $i = 1, \dots, d$ as the set of tasks that are served at station i :

$$(2.3) \quad \mathcal{E}(i) = \{k \in \mathcal{A}: s(k) = i\}.$$

We write $c(i)$ to mean the cardinality of the constituency set $\mathcal{E}(i)$. Next, define the predecessor and successor station mappings for station $j = 1 \dots, J$ as follows:

$$(2.4) \quad \pi(j) \equiv \{i = s(k): k \in \mathcal{P}(l), l \in \mathcal{E}(j)\},$$

$$(2.5) \quad \sigma(j) \equiv \{i = s(k): l \in \mathcal{P}(k), l \in \mathcal{E}(j)\}.$$




FIG. 1. Tasks and precedence constraints of two job types.

That is, $\pi(j)$ denotes the set of stations whose output feeds directly into station j ; analogously, $\sigma(j)$ is the set of stations that receive input from station j . In conjunction with the predecessor and successor station mappings, we now define a $d \times d$ “routing” matrix $\mathbb{P} = \mathbb{P}_{ij}$ whose elements are given by

$$(2.6) \quad \mathbb{P}_{ij} \equiv \begin{cases} 1, & \text{if } i \in \pi(j), \\ 0, & \text{otherwise.} \end{cases}$$

We assume that there is a column and row permutation of \mathbb{P} such that the resulting matrix is strictly upper triangular; in terms of the model, this means that we assume jobs traverse the network in such a way that jobs never return to a station it previously visited. In addition, this condition implies that there are no precedence constraints among the tasks at each station. This constitutes the *feedforward* routing assumption stated in Section 1.

It will be useful to think of new arrivals to the network as originating from a “dummy” station 0. With that interpretation in mind let us define, for each job type q ,

$$(2.7) \quad \mathcal{A}_q^0 = \{k \in \mathcal{A}_q : P_{lk} = 0 \text{ for all } l \in \mathcal{A}_q\}$$

$$(2.8) \quad \mathcal{A}_q^e = \{k \in \mathcal{A}_q : P_{kl} = 0 \text{ for all } l \in \mathcal{A}_q\}.$$

Clearly, processing of a type q job begins with those tasks $k \in \mathcal{A}_q^0$ and ends with the tasks $k \in \mathcal{A}_q^e$. Next, set $\mathcal{P}(k) = \{0\}$ if $\mathcal{P}(k) = \emptyset$ (or, equivalently, if $k \in \mathcal{A}_{q(k)}^0$), let $s(0) = 0$ and redefine (2.4) and (2.5) accordingly. Finally, we define

$$\sigma(0) = \{i : \pi(i) = \{0\}\}.$$

Thus, the stations in $\sigma(0)$ receive only external arrivals and the feedforward assumption guarantees that $\sigma(0) \neq \emptyset$. Figure 2 shows how the two job types depicted in Figure 1 are processed in a network consisting of four workstations. For this example, we have $\pi(1) = \pi(2) = \{0\}$, $\pi(3) = \{0, 2\}$, $\pi(4) = \{1, 3\}$ and $\sigma(0) = \{1, 2\}$.




FIG. 2. A sample fork-join network.

A node j is said to be a fork node if it contains a task $k \in \mathcal{E}(j)$ such that $k \in \mathcal{P}(l)$ for more than one task l . Similarly, station j is a join node if $\mathcal{P}(k)$ contains more than one element for some constituent task $k \in \mathcal{E}(j)$. At a join node, a type k task is said to be complete, or a unit, if all of its predecessors $l \in \mathcal{P}(k)$ have completed their service.

We assume that tasks are served at each station in a first-in–first-out (FIFO) manner. At nodes that do not involve a joining of tasks, this simply means that the tasks are served in the order of their arrival to the station. At join nodes, the arrival time of a (complete) task is defined to be the time at which its *last* predecessor completes service. Observe that this service discipline considers only local, or station-level, information. For the special case of fork–join networks with homogeneous customers, we argued in [17] that such a policy is equivalent to the global scheme of serving tasks in the order of the arrival of their associated *jobs*. In this setting, where multiple customer types traverse possibly different routes through the network, one observes a very different phenomenon in which jobs may overtake each other. As an example, consider the scenario depicted in Figure 3, which shows the status of five jobs in their intermediate stages of processing. The network contains three type 1 jobs and two type 2 jobs, which arrived at the network in the order 1a, 1b, 2a, 1c, 2b. Each job type follows the routing requirements described in Figure 2. Let us focus our discussion on the buffer between stations 3 and 4, hereafter referred to as buffer (3, 4). Job 2a—the third job to arrive at the network—has overtaken the first two arrivals, 1a and 1b, and is the first job to reach buffer (3, 4). Moreover, if server 1 completes his next task before server 3 can finish his, the first job to be processed by server 4 will be job 2a. In particular, the first job processed at station 4 may *not* be 1a, the first job to arrive at the network. The policy of serving tasks in the order of their arrival at a station may therefore result in serving jobs out of order. As such, a downstream station may incur delays due to the need for resequencing tasks (e.g., if the station is a join node) that have been overtaken by

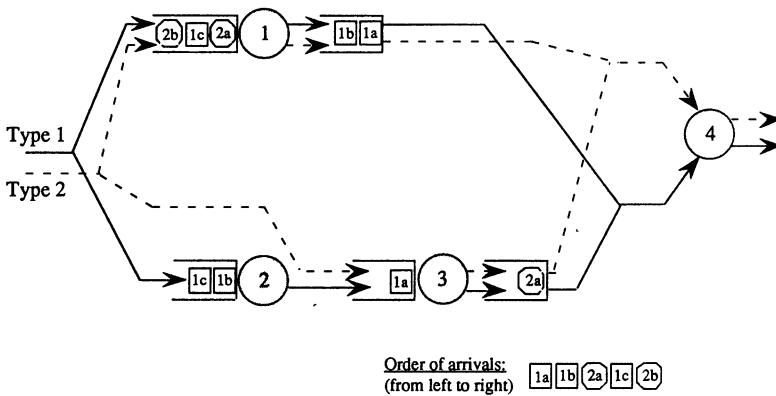


FIG. 3. A fork–join network with jobs in intermediate stages of processing.

other tasks at upstream stations (e.g., due to forks). In Figure 3, for instance, server 4 must remain idle even though each of his incident buffers is nonempty. One goal of this paper is to investigate how such a disordering is manifested in the heavy traffic limit.

3. Representations for processes of interest. To construct the basic stochastic processes associated with the fork-join network, let us assume a probability space (Ω, \mathcal{F}, P) on which are defined sequences of unitized random variables $\{u_q(i), i \geq 1\}$ and $\{v_k(i), i \geq 1\}$, $q = 1, \dots, p$, $k = 1, \dots, K$, where $u_q(i)$ and $v_k(i)$ are strictly positive with unit mean. As will be explained in the next section, we require very weak assumptions regarding the joint distribution of these sequences of unitized variables. However, readers may find it helpful to think in terms of the concrete case where each is a sequence of i.i.d. random variables and the sequences are mutually independent. From these sequences, the interarrival times and service times for the network are constructed by setting the interarrival time for the i th job of type q to be $\alpha_q^{-1}u_q(i)$, and the service time for task k of this job to be $\tau_k v_k(i)$. Recall that α_q is the average arrival rate for new type q jobs and τ_k is the mean service time for task k . Also recall that $\lambda_k = \alpha_q$ for each $k \in \mathcal{A}_q$.

To construct the external arrival process for type q jobs, set $u_q(0) \equiv 0$ and define

$$N_q(t) \equiv \max \left\{ m : \sum_{i=0}^m \alpha_q^{-1} u_q(i) \leq t \right\}.$$

For $k = 1, \dots, K$, let $V_k(t)$ be the partial sums process associated with the service times for tasks k ,

$$V_k(t) \equiv \sum_{i=1}^{\lfloor t \rfloor} \tau_k v_k(i).$$

Next, define, for each $k \in \mathcal{E}(j)$ and $j = 1, \dots, d$,

$$(3.1) \quad L_{jk}(t) \equiv V_k(N_{q(k)}(t)) = \tau_k v_k(1) + \dots + \tau_k v_k(N_{q(k)}(t)).$$

The process $L_{jk}(t)$ is the *class k total workload input process* for station j ; it is the sum of all task k service times associated with those jobs that enter the network during $[0, t]$. Note that $L_{jk}(t)$ includes service times corresponding to tasks that may not arrive at station j until after time t . Set

$$(3.2) \quad \xi_j(t) = \left(\sum_{k \in \mathcal{E}(j)} L_{jk}(t) \right) - t;$$

because t is the *potential* amount of work that can be processed in t units of time, $\xi_j(t)$ is the difference between the workload input and the potential workload output, and for this reason it is called the *total workload netflow process* at station j .

Let us choose an “external” station $j \in \sigma(0)$, fixing j until further notice. For each $k \in \mathcal{E}(j)$, set $A_{jk}(t) \equiv N_{q(k)}(t)$. Next let $M_{jk}(t) \equiv L_{jk}(t)$ and $X_j(t) \equiv$

$\xi_j(t)$. Because station j hosts only external arrivals, $A_{jk}(t)$ is the number of tasks k that have actually arrived at station j by time t . Similarly, $M_{jk}(t)$ is the amount of task k work that has arrived at station j in $[0, t]$ and $X_j(t)$ is the corresponding *immediate* workload netflow process at this station.

From Section 2.2 of [10], we can verify that the processes W_j and I_j are uniquely defined by the following three statements:

$$(3.3) \quad W_j(t) = X_j(t) + I_j(t) \geq 0 \quad \text{for all } t \geq 0;$$

$$(3.4) \quad I_j(\cdot) \text{ is continuous and nondecreasing with } I_j(0) = 0;$$

$$(3.5) \quad I_j(\cdot) \text{ increases only at times } t \text{ when } W_j(t) = 0;$$

moreover, I_j is given by the *continuous* mapping

$$(3.6) \quad I_j(t) = \sup_{0 \leq s \leq t} \{X_j(s)\}^-.$$

One interprets I_j as the *cumulative idleness process* for server j and W_j as the *immediate workload process* at station j . That is, $W_j(t)$ corresponds to the sum of the impending service times for (complete) tasks waiting at station j at time t , plus the remaining service time of any task that may be in service. If we define $Z_j(t)$ to be the *total* amount of work for server j that is present *anywhere* in the system at time t , then

$$(3.7) \quad Z_j(t) = \xi_j(t) + I_j(t),$$

and $Z_j(t) = W_j(t)$ as a consequence of $j \in \sigma(0)$. Hereafter we refer to $Z_j(t)$ as the *total workload process* for server j .

Next, let $\eta_j(t)$ be the arrival time at station j of the customer in service at station j at time t if $W_j(t) > 0$, and set $\eta_j(t) = t$ otherwise. Letting $Y_{jk}(t)$ be the amount of time server j has spent serving tasks k in $[0, t]$, it follows from the FIFO service discipline at each station that

$$(3.8) \quad Y_{jk}(t) = M_{jk}(\eta_j(t)) + \varepsilon_{1k}(t),$$

where $\varepsilon_{1k}(t)$ is the amount of service the current task has received if that task is of class k and $\varepsilon_{1k}(t) = 0$ otherwise. As a matter of definition, $\eta_j(t)$ is bounded by the immediate workload process as follows:

$$(3.9) \quad W_j(\eta_j(t)) \leq t - \eta_j(t) \leq W_j(\eta_j(t)) + \varepsilon_{2j}(t),$$

where $\varepsilon_{2j}(t) = 0$ if $W_j(\eta_j(t)) = 0$ and otherwise $\varepsilon_{2j}(t)$ is the service time of the customer currently occupying server j . Letting $S_k = \{S_k(t), t \geq 0\}$ be the counting process associated with the task k service times $\{\tau_k v_k(1), \tau_k v_k(2), \dots\}$, the number of tasks k to have departed from station j by time t , denoted as $D_{jk}(t)$, is then given by

$$(3.10) \quad D_{jk}(t) = S_k(Y_{jk}(t)).$$

Finally, defining $U_{jk}(t)$ to be the total amount of *partial* work associated with tasks k that is present *anywhere* in the system at time t , it follows from the previous definitions that

$$(3.11) \quad U_{jk}(t) = L_{jk}(t) - Y_{jk}(t).$$

Moreover, it is a trivial consequence that $Z_j(t) = \sum_{k \in \mathcal{E}(j)} U_{jk}(t)$.

In an inductive manner, these definitions can be extended to all stations in the network. Consider a station j such that all immediate predecessor stations have been “treated,” that is, if $i \in \pi(j)$, then for each $l \in \mathcal{E}(i)$ the processes $X_i(t)$, $W_i(t)$, $I_i(t)$, $Z_i(t)$, $D_{il}(t)$ and $U_{il}(t)$ have been defined. For each task $k \in \mathcal{E}(j)$ one defines its arrival process to be

$$(3.12) \quad A_{jk}(t) \equiv \begin{cases} N_{q(k)}(t), & \text{if } k \in \mathcal{A}_q^0, \\ \min_{l \in \mathcal{P}(k)} D_{s(l)l}(t), & \text{otherwise.} \end{cases}$$

One can interpret $A_{jk}(t)$ as the number of *complete* tasks k that have arrived at station j by time t . (We take the convention that work is associated with complete tasks, so incomplete tasks present zero work to the server.) The immediate workload input process and immediate netflow process for station j are defined, respectively, via

$$(3.13) \quad M_{jk}(t) \equiv V_k(A_{jk}(t)) = \tau_k [v_k(1) + \cdots + v_k(A_{jk}(t))]$$

and

$$(3.14) \quad X_j(t) \equiv \left(\sum_{k \in \mathcal{E}(j)} M_{jk}(t) \right) - t.$$

The workload process W_j , the idleness process I_j , the total workload process Z_j , the departure process D_{jk} and the partial workload process U_{jk} are then defined exactly as in (3.3)–(3.11). The vector processes N , V , L , A , M , X , W , I , Z , D and U are then defined in the obvious manner.

The throughput time of a job is the length of time between the job’s arrival and its subsequent departure from the system. Let $T_q(t)$ be the throughput time of the next type q job to enter the network after time t . The intermediate process $T_{qk}(t)$, $k \in \mathcal{A}_q$, is defined to be the “throughput time through task k ,” which is the amount of elapsed time until task k is completed. As a matter of definition, we have the relationship

$$T_q(t) \equiv \max_{k \in \mathcal{A}_q^c} \{T_{qk}(t)\}.$$

To define the intermediate processes $T_{qk}(t)$, we first define for each job type q the process

$$\Phi_q(t) \equiv \alpha_q^{-1} u(1) + \cdots + \alpha_q^{-1} u(N_q(t)) + \alpha_q^{-1} u(N_q(t) + 1),$$

interpreted as the arrival epoch of the next type q job to enter the network after time t . For each task $k \in \mathcal{A}_q^0$, let

$$\Phi_{qk}(t) \equiv \Phi_q(t) \quad \text{and} \quad T_{qk}(t) \equiv W_{s(k)}(\Phi_{qk}(t)).$$

Because a type q job begins immediately with tasks $k \in \mathcal{A}_q^0$, $\Phi_{qk}(t)$ is the arrival time of this task at station $s(k)$. Furthermore, because tasks are served in a first-in–first-out manner, the amount of time this task must spend at station $s(k)$ is precisely the amount of work found at station $s(k)$ immediately after its arrival (which includes the service time associated with the new arrival). Thus $T_{qk}(t)$ is the total sojourn time of the job through task k .

For other stations in the network, the random processes $\Phi_{qk}(t)$ and $T_{qk}(t)$ are inductively defined as follows. Suppose that k is a task such that $T_{ql}(t)$ has been defined for each $l \in \mathcal{P}(k)$, and set

$$(3.15) \quad \Phi_{qk}(t) \equiv \Phi_q(t) + \max_{l \in \mathcal{P}(k)} T_{ql}(t)$$

$$(3.16) \quad T_{qk}(t) \equiv \max_{l \in \mathcal{P}(k)} T_{ql}(t) + W_{s(k)}(\Phi_{qk}(t)).$$

Recall that the arrival time of a task is the time at which its last predecessor task is completed. (If task k requires a join, there may be a delay between completion times of its multiple predecessor tasks.) Thus, $\max_{l \in \mathcal{P}(k)} T_{ql}(t)$ is the amount of time that elapses until task k “arrives” at station $s(k)$, $\Phi_{qk}(t)$ is precisely its time of arrival and $T_{qk}(t)$ is the throughput time through task k .

4. A sequence of systems in heavy traffic. The limit theorems stated here apply to systems that satisfy conditions of “heavy traffic.” For $k \in \mathcal{E}(j)$, let $\rho_{jk} = \lambda_k \tau_k$ be the workload factor at station j associated with tasks k , and define the total traffic intensity at station j to be

$$(4.1) \quad \rho_j \equiv \sum_{k \in \mathcal{E}(j)} \rho_{jk} = \sum_{k \in \mathcal{E}(j)} \lambda_k \tau_k.$$

The system is said to be *stable* if $\rho_j < 1$ for $j = 1, \dots, J$, and it is said to be in *heavy traffic* if ρ_j is “approximately” 1 for each j . The precise formulation of our heavy traffic limit theorem requires the construction of a “sequence of systems,” indexed by n , whose corresponding traffic intensities $\rho_j^{(n)}$ converge to 1 for all j .

Recall that the interarrival times and service times for the network are defined in terms of the basic sequences of unitized random variables $\{u_q(i): i \geq 1\}$, $\{v_k(i): i \geq 1\}$, $q = 1, \dots, p$, $k = 1, \dots, K$. To construct a sequence of fork–join networks we further require sequences of positive constants $\{\alpha_q^{(n)}, n \geq 1\}$, $\{\tau_k^{(n)}, n \geq 1\}$, $q = 1, \dots, p$, $k = 1, \dots, K$. In the n th system of the sequence, the interarrival times and service times are taken to be $u_q^{(n)}(i) \equiv u_q(i)/\alpha_q^{(n)}$ and $v_k^{(n)}(i) \equiv \tau_k^{(n)}v_k(i)$, respectively. For the n th system, $\alpha_q^{(n)}$ is the arrival rate of type q jobs and $\tau_k^{(n)}$ is the mean service time for task k . Setting $\lambda_k^{(n)} = \alpha_q^{(n)}$ for $k \in \mathcal{A}_q$, define the traffic intensities $\rho_j^{(n)}$ as in (4.1) using $\lambda_k^{(n)}$ and $\tau_k^{(n)}$ in place of λ_k and τ_k .

The convention here is to denote a parameter or a process associated with the n th system by the superscript “(n)”: For example, $N_q^{(n)}$ refers to the

external arrival process for type q jobs in the n th system. Define the centered processes

$$\begin{aligned}\hat{N}_q^{(n)}(t) &\equiv N_q^{(n)}(t) - \alpha_q^{(n)}t, & \hat{V}_k^{(n)}(t) &\equiv V_k^{(n)}(t) - \tau_k^{(n)}t, \\ \hat{A}_{jk}^{(n)}(t) &\equiv A_{jk}^{(n)}(t) - \lambda_k^{(n)}t, & \hat{L}_{jk}^{(n)}(t) &\equiv L_{jk}^{(n)}(t) - \rho_{jk}^{(n)}t, \\ \hat{S}_k^{(n)}(t) &\equiv S_k^{(n)}(t) - (\tau_k^{(n)})^{-1}t, & \hat{M}_{jk}^{(n)}(t) &\equiv M_{jk}^{(n)}(t) - \rho_{jk}^{(n)}t.\end{aligned}$$

The results in this paper apply to processes that have been “scaled.” Let $X^{(n)}$ denote a “generic” process associated with the n th system. The scaled version of the process $X^{(n)}$, denoted as X^n , is defined via

$$X^n(t) \equiv n^{-1/2}X^{(n)}(nt).$$

Hereafter, when we say a “scaled” process and write the process with a superscript n , we mean a process whose space and time dimensions have been scaled in the manner specified previously.

It is assumed that the following conditions hold for the input processes of the network. First, the arrival rates and mean service times converge to finite constants, $\lambda_k^{(n)} \rightarrow \lambda_k$ and $\tau_k^{(n)} \rightarrow \tau_k$, $k = 1, \dots, K$. This implies that $\rho_j^{(n)} \rightarrow \rho_j = \sum_{k \in \mathcal{E}(j)} \lambda_k \tau_k$. Furthermore, it is assumed that there exists a d -vector $\theta = (\theta_1, \dots, \theta_d)$ such that, for each $j = 1, \dots, d$, $-\infty < \theta_j < \infty$ and

$$(4.2) \quad n^{1/2}(\rho_j^{(n)} - 1) \rightarrow \theta_j \quad \text{as } n \rightarrow \infty.$$

Condition (4.2) is called the *heavy traffic condition*. It requires not only that $\rho_j = 1$ at each station, but also that the rate of convergence is “sufficiently fast” and is uniform for all stations. Finally, it is assumed that there is a $d \times d$ covariance matrix Ω such that the following functional central limit theorem holds as $n \rightarrow \infty$:

$$(4.3) \quad (\hat{N}^n, \hat{V}^n, \hat{L}^n) \Rightarrow (N^*, V^*, L^*), \text{ where } L^* \text{ is a } (0, \Omega) \text{ Brownian motion and } N^* \text{ and } V^* \text{ are also Brownian motions with zero drift.}$$

To explore the implications and restrictions of assumption (4.3), write the centered and scaled netflow process (3.2) as

$$\xi_j^n(t) = \sum_{k \in \mathcal{E}(j)} \hat{L}_{jk}^n(t) + n^{1/2}(\rho_j^{(n)} - 1)t.$$

It follows from (3.1) and (3.2) and assumptions (4.2) and (4.3) that

$$(4.4) \quad L_{jk}^*(t) = V_k^*(\lambda_k t) + \tau_k N_{q(k)}^*(t)$$

and

$$(4.5) \quad \xi_j^*(t) = \sum_{k \in \mathcal{E}(j)} L_{jk}^*(t) + \theta_j t.$$

Defining the $d \times c$ constituency matrix C with elements

$$(4.6) \quad C_{ik} \equiv \begin{cases} 1, & \text{if } k \in \mathcal{E}(i), \\ 0, & \text{otherwise,} \end{cases}$$

and setting

$$(4.7) \quad \Gamma = C\Omega C',$$

one can conclude that

$$(4.8) \quad (\hat{N}^n, \hat{V}^n, \hat{L}^n, \xi^n) \Rightarrow (N^*, V^*, L^*, \xi^*) \quad \text{where } \xi^* \text{ is } (\theta, \Gamma)\text{BM.}$$

Next, recall that S_k is the counting process associated with the partial sums process V_k . From [23], (4.3) implies that

$$(4.9) \quad \hat{S}_k^n \Rightarrow -\tau_k^{-3/2} V_k^*.$$

Finally, consider the special case in which $\{u_q(i), i \geq 1\}$ and $\{v_k(i), i \geq 1\}$ are mutually independent sequences of i.i.d. random variables such that $u_q(i)$ and $v_k(i)$ have squared coefficients of variation c_{aq}^2 and c_{sk}^2 , respectively (the squared coefficient of variation of a random variable is defined to be its variance divided by the square of its mean). Then $N_q^{(n)}$ is a renewal process with rate $\lambda_q^{(n)}$, and a simple application of the functional central limit theorem for renewal processes [6] proves that $\hat{N}_q^n \Rightarrow N_q^*$, where N_q^* is $(0, \lambda c_{aq}^2)$ BM. Because $L_{jk}^{(n)}$ is a compound renewal process, \hat{L}^n converges to $(0, \Omega)$ BM by Theorem 2.1 of [22]. In particular, the covariance matrix is of the form

$$(4.10) \quad \Omega_{kl} = \begin{cases} \lambda_k \tau_k^2 (c_{sk}^2 + c_{aq(k)}^2), & k = l, \\ \lambda_k \tau_k \tau_l c_{aq(k)}^2, & k \neq l, q(k) = q(l), \\ 0, & \text{otherwise.} \end{cases}$$

It is not necessary to assume i.i.d. sequences for the convergence in (4.3) to hold. See Glynn [9] for examples regarding sequences with dependencies and nonstationarity.

5. Additional notation and preliminaries. We devote this section to developing some additional notation and preliminary results that will be useful for future reference. For a station $j \in \sigma(0)$ we define its “depth” to be $d(j) = 0$. The depths of all other stations are then defined recursively as

$$(5.1) \quad d(j) = 1 + \max\{d(i) : i > 0, i \in \pi(j)\}.$$

It follows from the feedforward structure that such a notion of depth is well defined and that $d(j) \leq d$ for all stations $j = 1, \dots, d$.

In Section 2 we defined $\pi(i)$ to be the set of predecessor *stations* to station i . It will also be useful to define the set of *tasks* preceding station i . Recall that $\mathcal{E}(i)$ is the constituency of station i and $c(i)$ is the cardinality of this set. Rather than labeling tasks in $\mathcal{E}(i)$ by $k = 1, \dots, K$, we now enumerate these tasks by $a_1^i, \dots, a_{c(i)}^i$. Set $\mathcal{T}(0) \equiv \{0\}$ and, for each station $i = 1, \dots, d$, define

$$(5.2) \quad \mathcal{T}(i) = \{x = (x_1, \dots, x_{c(i)}) : x_l \in \mathcal{P}(a_l^i)\}.$$

Each element of $\mathcal{T}(i)$ is a vector of $c(i)$ components, and each component corresponds to a predecessor task of a constituent task in station i .

Set

$$\mathcal{F}^1(i) \equiv \mathcal{F}(i),$$

$$\mathcal{F}^2(i) \equiv \left\{ (x^1, x^2) : x^1 \in \mathcal{F}(i); x^2 = (x_1^2, \dots, x_{c(i)}^2), x_l^2 \in \mathcal{F}(s(x_l^1)) \right\},$$

and define $\mathcal{F}^h(i)$ recursively as follows:

$$(5.3) \quad \mathcal{F}^h(i) \equiv \left\{ x = (x^1, \dots, x^h) : x^1 \in \mathcal{F}(i), x_{l_1}^2 \in \mathcal{F}(s(x_{l_1}^1)), \dots, \right. \\ \left. x_{l_1 l_2, \dots, l_{h-1}}^h \in \mathcal{F}(s(x_{l_1 l_2, \dots, l_{h-1}}^{h-1})) \right\},$$

where $l_1 = 1, \dots, c(i)$, $l_2 = 1, \dots, c(s(x_{l_1}^1))$, \dots , $l_{h-1} = 1, \dots, c(s(x_{l_1 l_2, \dots, l_{h-2}}^{h-2}))$ and we take the convention that $c(0) = 0$. For a station i with depth greater than or equal to h and $x \in \mathcal{F}^h(i)$, we define the set of indices

$$(5.4) \quad \mathcal{L}^h(x) \equiv \left\{ l = (l_1, \dots, l_h) : l_1 = 1, \dots, c(i); l_2 = 1, \dots, c(s(x_{l_1}^1)); \right. \\ \left. \dots; l_h = 1, \dots, c(s(x_{l_1 l_2, \dots, l_{h-1}}^{h-1})) \right\}$$

In addition, let

$$(5.5) \quad \mathcal{L}_j^h(x) = \left\{ l = (l_1, \dots, l_h) \in \mathcal{L}^h(x) : s(x_{l_1, \dots, l_h}^h) = j \right\}.$$

Taking the network in Figure 2 as an example, we have $d(4) = 2$ and $\mathcal{F}^2(4) = (w, x, y, z)$, where

$$(5.6) \quad \begin{aligned} w^1 &= (1, 5), & w_1^2 &= (0, 0), & w_2^2 &= (0, 0), \\ x^1 &= (1, 6), & x_1^2 &= (2, 0), & x_2^2 &= (0, 0), \\ y^1 &= (3, 5), & y_1^2 &= (0, 0), & y_2^2 &= (2, 0), \\ z^1 &= (3, 6), & z_1^2 &= (2, 0), & z_2^2 &= (2, 0). \end{aligned}$$

For the moment, consider $x \in \mathcal{F}^2(4)$, for which we have

$$\begin{aligned} \mathcal{L}^2(x) &= \{(1, 1), (1, 2), (2, 1), (2, 2)\}, \\ \mathcal{L}_2^2(x) &= \{(1, 1)\}. \end{aligned}$$

One may think of each element $x \in \mathcal{F}^{d(i)}(i)$ as describing a “path” of tasks traversed by the various job types on their sojourns toward station i . Moreover, $\mathcal{L}_j^h(x)$ identifies the particular “branch(es)” in the path x that would include a visit to station j . Using the notation established above, the following lemma can be verified directly.

LEMMA 5.1. *For each station $i = 1, \dots, d$ and for each sequence of numbers b_{kl} associated with tasks k and l ,*

$$\sum_{k \in \mathcal{E}(i)} \left(\max_{l \in \mathcal{P}(k)} b_{kl} \right) = \max_{x \in \mathcal{F}(i)} \left(\sum_{l \in \mathcal{L}(x)} b_{a_l x_l} \right).$$

6. The main results.

THEOREM 6.1. *Suppose that assumptions (4.2) and (4.3) hold. Then*

$$(\xi^n, U^n, Z^n, W^n, I^n) \Rightarrow (\xi^*, U^*, Z^*, W^*, I^*),$$

where, for each $i = 1, \dots, d$ and $k \in \mathcal{E}(i)$, the following hold: $U_{00}^* \equiv 0$;

$$(6.1) \quad \xi^* \text{ is a } (\theta, \Gamma) \text{ Brownian motion};$$

$$(6.2) \quad Z_i^* = \xi_i^* + I_i^*;$$

$$(6.3) \quad X_i^* = \xi_i^* - \sum_{k \in \mathcal{E}(i)} \rho_{ik} \left(\max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*}{\rho_{s(l)l}} \right);$$

$$(6.4) \quad W_i^* = X_i^* + I_i^*;$$

$$(6.5) \quad U_{ik}^* = \rho_{ik} \left[Z_i^* + \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*}{\rho_{s(l)l}} - \sum_{m \in \mathcal{E}(i)} \rho_{im} \left(\max_{n \in \mathcal{P}(m)} \frac{U_{s(n)n}^*}{\rho_{s(n)n}} \right) \right];$$

$$(6.6) \quad I_i^* \text{ is continuous and nondecreasing with } I_i^*(0) = 0;$$

$$(6.7) \quad I_i^* \text{ increases only at times } t \text{ with } W_i^*(t) = 0.$$

Set

$$(6.8) \quad \gamma_{kl} = \begin{cases} 0, & \text{if } k = 0 \text{ or } l = 0 \text{ and } k \neq 0, \\ 1 - \rho_{s(k)k}, & \text{if } k = l \neq 0 \text{ and } l \neq 0, \\ -\rho_{s(l)l}, & \text{if } k \neq l \neq 0 \text{ and } l \neq 0. \end{cases}$$

For notational convenience, we henceforth write s^m to mean $s(x_{i_1}^m, \dots, i_m)$. Denoting by h the depth of a station i , define for each $x \in \mathcal{S}^h(i)$ and $j = 1, \dots, d$ the following factors:

$$(6.9) \quad \begin{aligned} \beta_{ij}(x) = & \sum_{l \in \mathcal{L}_j^1(x)} \rho_{ia_l^j} + \sum_{l=(l_1, l_2) \in \mathcal{L}_j^2(x)} \rho_{ia_{l_1}^j} \gamma_{x_{l_1}^j} a_{l_2}^{j-1} \\ & + \dots + \sum_{l=(l_1, \dots, l_h) \in \mathcal{L}_j^h(x)} \rho_{ia_{l_1}^j} \gamma_{x_{l_1}^j} a_{l_2}^{j-1} \dots \gamma_{x_{l_{h-1}}^j} a_{l_h}^{j-h}. \end{aligned}$$

We then define the convex polyhedral cone \mathcal{S} to be

$$(6.10) \quad \mathcal{S} = \bigcap_{i=1}^d \bigcap_{x \in \mathcal{S}^{d(i)}(i)} \left\{ z = (z_1, \dots, z_d) : z_i - \sum_{j=1}^d \beta_{ij}(x) z_j \geq 0 \right\}.$$

It is easily verified that $z \geq 0$ if $z \in \mathcal{S}$. For each $i = 1, \dots, d$, we also define the boundary set

$$(6.11) \quad F^i = \bigcup_{x \in \mathcal{S}^{d(i)}(i)} \left\{ z \in \mathcal{S} : z_i - \sum_{j=1}^d \beta_{ij}(x) z_j = 0 \right\}.$$

THEOREM 6.2. *For each $i = 1, \dots, d$, the following hold:*

$$(6.12) \quad \xi^* \text{ is a } (\theta, \Gamma) \text{ Brownian motion;}$$

$$(6.13) \quad Z_i^* = \xi_i^* + I_i^* \in \mathcal{S};$$

$$(6.14) \quad W_i^* = Z_i^* - \max_{x \in \mathcal{F}^{d(i)}(i)} \left(\sum_{j=1}^d \beta_{ij}(x) Z_j^* \right);$$

$$(6.15) \quad I_i^* \text{ is continuous and nondecreasing with } I_i^*(0) = 0;$$

$$(6.16) \quad I_i^* \text{ increases only at times } t \text{ with } Z^*(t) \in F^i.$$

That is, the process Z^ as defined in Theorem 6.1 is a d -dimensional SRBM whose state space is the convex polyhedral cone \mathcal{S} . The SRBM Z^* has drift θ , covariance matrix Γ and reflection matrix $R = I$, where I is the d -dimensional identity matrix.*

REMARK. That such an SRBM is well defined in the strong sense follows from the feedforward structure of the network. Readers may refer to Nguyen [16] for details of the proof.

To interpret Theorem 6.2, note that statements (6.13) and (6.15) are reiterations of the characterizations given in (3.7) and (3.4), respectively. Moreover, the approximation (6.12) of the netflow process by a Brownian motion was justified in Section 4 under assumptions (4.2) and (4.3). Next, recall that each element $x \in \mathcal{F}^{d(i)}(i)$ describes a “path” of tasks traversed by the various constituent jobs on their way to station i . Equation (6.14) states that the immediate workload at station i is the minimum amount of work found among all the “paths” leading up to station i . In other words, (6.14) articulates the constraint that a task at station i cannot be processed until all of its predecessor tasks have been completed (this is the definition of a join node). With this interpretation in mind, statement (6.16) is then equivalent to (3.5). Thus, each idleness process is associated with potentially multiple boundaries on the state space \mathcal{S} . In [17], Nguyen showed that the additional faces correspond to the fork and join constraints in the network. As we will demonstrate in an example, the polyhedral state space associated with heterogeneous fork-join networks typically has many more faces than its homogeneous counterpart. These additional faces may be interpreted as results of the disordering effects that occur when jobs fork and join in their sojourns through the network.

EXAMPLE 1 (The sample fork-join network). The heavy traffic limit of the network pictured in Figure 2 is given by the following:

$$(6.17) \quad \xi^* \text{ is a } (\theta, \Gamma) \text{ Brownian motion;}$$

$$(6.18) \quad Z_i^* = \xi_i^* + I_i^*;$$

$$(6.19) \quad I_i^* \text{ is continuous and nondecreasing with } I_i^*(0) = 0;$$

$$(6.20) \quad I_i^* \text{ increases only at times } t \text{ with } Z_i^*(t) = 0, i = 1, 2;$$

$$(6.21) \quad I_3^* \text{ increases only at times } t \text{ with } Z_3^*(t) - \rho_{33} Z_2^*(t) = 0;$$

I_4^* increases only at times t such that one of the following four conditions hold:

$$(6.22) \quad Z_4^*(t) - Z_1^*(t) = 0,$$

$$(6.23) \quad Z_4^*(t) - (\rho_{44}\rho_{36} - \rho_{47}\rho_{33})Z_2^*(t) - Z_3^*(t) = 0,$$

$$(6.24) \quad Z_4^*(t) - \rho_{44}Z_1^*(t) + \rho_{47}\rho_{33}Z_2^*(t) - \rho_{47}Z_3^*(t) = 0,$$

$$(6.25) \quad Z_4^*(t) - \rho_{47}Z_1^*(t) - \rho_{44}\rho_{36}Z_2^*(t) - \rho_{44}Z_3^*(t) = 0.$$

Hence, setting

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -\rho_{33} & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & -(\rho_{44}\rho_{36} - \rho_{47}\rho_{33}) & -1 & 1 \\ -\rho_{44} & \rho_{47}\rho_{33} & -\rho_{47} & 1 \\ -\rho_{47} & -\rho_{44}\rho_{36} & -\rho_{44} & 1 \end{bmatrix}$$

and letting $\mathcal{S} = \{x: Ax \geq 0\}$, Z^* is a four-dimensional SRBM with drift θ , covariance matrix Γ , reflection matrix $R = I$ and whose state space is the convex polyhedral cone \mathcal{S} with seven faces. We now turn to the interpretation of conditions (6.20)–(6.25). Because the immediate workload is identical to the total workload process at stations 1 and 2, condition (6.20) states that the idleness processes at these stations increase only when the immediate workload at the respective stations is zero. Emulating the arguments in Harrison and Nguyen [12, 13], one may think of $\lambda_2 Z_2^*(t)$ as the number of tasks 2 occupying station 2 at time t . Consequently, $\tau_3 \lambda_1 Z_2^*(t) = \rho_{33} Z_2^*(t)$ is the amount of work associated with tasks 3 destined for station 3 that still reside at station 2 at time t . Hence $W_3^*(t) = Z_3^*(t) - \rho_{33} Z_2^*(t)$ and condition (6.21) specifies that the idleness process at station 3 increases only when there is no immediate work at that station.

Conditions (6.22)–(6.25) describe the four scenarios under which server 4 is forced to remain idle. It can be verified that conditions (6.22), (6.23), (6.24) and (6.25) correspond to the paths described by w , z , y and x of equation (5.6), respectively. The first path, w , contains the predecessor tasks 1 and 5, hence condition (6.22) states that the total amount of work destined for station 4 is contained in those tasks currently waiting for service at station 1 either in the form of task 1 or task 5. That is, the immediate work content in buffer (1, 4) (the buffer joining stations 1 and 4), given by $Z_4^*(t) - Z_1^*(t)$, is zero. Path z , on the other hand, contains the predecessor tasks 3 and 6, and one can apply a similar argument to arrive at the conclusion that condition (6.23) corresponds to the scenario in which buffer (3, 4) is empty. That is, the immediate work content in buffer (3, 4) is given by $Z_4^*(t) - (\rho_{44}\rho_{36} - \rho_{47}\rho_{33})Z_2^*(t) - Z_3^*(t)$. Because station 4 is a join node for both types of customers, it seems clear that both buffers (2, 4) (the buffer joining stations 2

and 4) and (3, 4) (the buffer joining stations 3 and 4) must be *nonempty* if server 4 is to remain busy. However, these are not the only times at which server 4 may be idle. Condition (6.24), which corresponds to path x , states that the total work destined for station 4 is completely contained in tasks 3 at station 2 and tasks 5 at station 1. That is, buffer (1, 4) contains no tasks corresponding to type 2 jobs and buffer (3, 4) does not have any type 1 work. Because there are no complete tasks for station 4, the server is forced to remain idle although both of its incident buffers may be nonempty. Similarly, condition (6.25) describes the converse situation in which buffer (1, 4) has only type 2 tasks and buffer (3, 4) contains only type 1 tasks.

EXAMPLE 2 (Multiple customer types with common routing structure). Consider a fork-join network in which all customers share the same routing constraints, but the interarrival times and service times may be different across job types. This class of networks corresponds to a particularly simple case of the networks considered in this paper. For stations $j \in \sigma(0)$, it follows from (6.5) that $U_{ik}^* = \rho_{ik} Z_i^*$. By induction on the depth of stations, one can verify that equation (6.5) becomes

$$(6.26) \quad U_{ik}^* = \rho_{ik} Z_i^* + \rho_{ik} \max_{l \in \mathcal{P}(k)} Z_{s(l)}^* - \rho_{ik} \sum_{m \in \mathcal{E}(i)} \rho_{im} \max_{n \in \mathcal{P}(m)} Z_{s(n)}^*.$$

Because all customer types share the same the same routing structure, the sets $\{s(n) : n \in \mathcal{P}(m)\}$ are identical for each task $m \in \mathcal{E}(i)$. Hence the last term on the right-hand side of (6.26) is equal to

$$-\rho_{ik} \left(\sum_{m \in \mathcal{E}(i)} \rho_{im} \right) \max_{j \in \pi(i)} Z_j^* = -\rho_{ik} \max_{j \in \pi(i)} Z_j^*$$

and equation (6.26) reduces to

$$U_{ik}^* = \rho_{ik} Z_i^*.$$

Hence, for $i = 1, \dots, d$ and $k \in \mathcal{E}(i)$, the following hold:

$$\begin{aligned} \xi^* & \text{ is a } (\theta, \Gamma) \text{ Brownian motion;} \\ Z_i^* & = \xi_i^* + I_i^*; \\ W_i^* & = Z_i^* - \max_{j \in \pi(i)} Z_j^*; \\ U_{ik}^* & = \rho_{ik} Z_i^*; \\ I_i^* & \text{ is continuous and nondecreasing with } I_i^*(0) = 0; \\ I_i^* & \text{ increases only at times } t \text{ with } W_i^*(t) = 0. \end{aligned}$$

Fork-join networks with one customer type (i.e., homogeneous fork-join networks) are clearly a special subset of the networks discussed in this section. It is straightforward to verify that the results above agree with those given in Nguyen [17].

EXAMPLE 3 (Feedforward multiclass queueing networks). Consider now the feedforward multiclass queueing network studied by Peterson [18]. The

networks described in Peterson [18] are essentially similar to those considered here with one important exception: The networks in [18] require that $\mathcal{P}(k)$ contains at most one element for each task k , that is, there are no join nodes. (Peterson's work does not explicitly consider the case in which tasks may fork, but the inclusion of the forking structure would not pose much hardship to his analysis.) Recall that $Z_i^*(t) = \sum_{k \in \mathcal{E}(i)} U_{ik}^*(t)$. Hence (6.2)–(6.5) imply

$$(6.27) \quad U_{ik}^* = \rho_{ik} \left[W_i^* + \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*}{\rho_{s(l)l}} \right],$$

$$(6.28) \quad Z_i^* = W_i^*(t) + \sum_{k \in \mathcal{E}(i)} \max_{l \in \mathcal{P}(k)} \rho_{ik} \frac{U_{s(l)l}^*}{\rho_{s(l)l}}.$$

We denote by $p(k)$ the one predecessor task of task k , and define $p(0) = 0$. Setting $p^1(k) \equiv p(k)$, we recursively define $p^h(k) \equiv p(p^{h-1}(k))$. Letting $U_{00}^* = W_0^* = 0$, equations (6.27) and (6.28) thus reduce to

$$U_{ik}^* = \rho_{ik} \left[W_i^* + \frac{U_{s(p^1(k))p^1(k)}^*}{\rho_{s(p^1(k))p^1(k)}} \right],$$

$$Z_i^* = W_i^*(t) + \sum_{k \in \mathcal{E}(i)} \rho_{ik} \frac{U_{s(p^1(k))p^1(k)}^*}{\rho_{s(p^1(k))p^1(k)}}.$$

Similarly to the previous example, we can use induction to show that, for a station i of depth h ,

$$U_{ik}^* = \rho_{ik} \left[W_i^* + \sum_{g=1}^h W_{s(p^g(k))}^* \right],$$

$$Z_i^* = W_i^*(t) + \sum_{g=1}^h \sum_{k \in \mathcal{E}(i)} \rho_{ik} W_{s(p^g(k))}^*.$$

Readers can verify that this agrees with the result obtained by Peterson [18].

THEOREM 6.3. *Under assumptions (4.2) and (4.3), $(T_1^n, \dots, T_p^n) \Rightarrow (T_1^*, \dots, T_p^*)$, where*

$$T_q^*(t) = \max_{k \in \mathcal{A}_q^c} T_{qk}^*(t),$$

$$T_{qk}^*(t) = \max_{l \in \mathcal{P}(k)} T_{ql}^*(t) + W_{s(k)}^*, \quad T_{q0}^*(t) \equiv 0.$$

If we denote by l_q the PERT/CPM “longest path operator” for type q jobs, Theorem 6.3 implies the representation

$$(6.29) \quad T_q^*(t) = l_q(W_{s(k)}^*, k \in \mathcal{A}_q).$$

As discussed in Nguyen [17], expression (6.29) is an example of Reiman's “snapshot” principle [19]. That is, in the heavy traffic scaling, the fluctuation

in workload levels is insignificantly small relative to the length of time that a job spends in the system, hence a “snapshot” of the system at the time of a job’s arrival remains representative throughout the job’s sojourn in the network. Equation (6.29) expresses the remarkable result that sojourn time analysis of a fork-join network may be phrased in terms of the familiar longest path analysis of PERT/CPM methods, where traditional task times are now replaced by waiting times at stations corresponding to the tasks.

7. Proofs. By Skorohod’s representation theorem and the continuity of Brownian motions, we can and will assume that the convergence in (4.3) is almost surely uniform on compact sets; that is, we henceforth assume

$$(7.1) \quad (\hat{N}^n, \hat{V}^n, \hat{L}^n) \rightarrow (N^*, V^*, L^*) \quad \text{u.o.c.}$$

We begin the proof of Theorem 6.1 with a few preliminary results. The first lemma is an immediate consequence of assumptions (4.2) and (7.1).

LEMMA 7.1. $\xi^n \rightarrow \xi^*$ u.o.c., where ξ^* is a (θ, Γ) Brownian motion with $\Gamma = C\Omega C'$.

LEMMA 7.2. For $k = 1, \dots, K$ and $j = 1, \dots, d$, let $\varepsilon_{1k}^n(t) = n^{-1/2}\varepsilon_{1k}^{(n)}(nt)$ and $\varepsilon_{2j}^n(t) = n^{-1/2}\varepsilon_{2j}^{(n)}(nt)$. Then $\varepsilon_{1k}^n \rightarrow 0$ and $\varepsilon_{2j}^n \rightarrow 0$ u.o.c.

PROOF. Note that

$$0 \leq \varepsilon_{1k}^{(n)}(t) \leq \max_{1 \leq i \leq N_{q(k)}^{(n)}(t)} \tau_k^{(n)} v_k(i)$$

and

$$0 \leq \varepsilon_{2j}^{(n)}(t) \leq \max_{k \in \mathcal{E}(j)} \max_{1 \leq i \leq N_{q(k)}^{(n)}(t)} \tau_k^{(n)} v_k(i).$$

The lemma follows directly from assumption (7.1) and Lemma 3.3 of Iglehart and Whitt [14]. \square

Let $\bar{\eta}_j^n(t) = n^{-1}\eta_j^{(n)}(nt)$ and $\hat{\eta}_j^n(t) = n^{-1/2}(nt - \eta_j^{(n)}(nt))$.

LEMMA 7.3. If $W_j^n \rightarrow W_j^*$ u.o.c., then $\bar{\eta}_j^n \rightarrow e$ u.o.c., where $e(t) = t$.

PROOF. Because $\bar{\eta}_j^n(t) \leq t$, it follows from equation (3.9) that, for each $t \geq 0$,

$$\begin{aligned} \|e(\cdot) - \bar{\eta}_j^n(\cdot)\|_t &\leq n^{-1/2} \|W_j^n(\bar{\eta}_j^n(\cdot))\|_t + n^{-1/2} \|\varepsilon_2^n(\cdot)\|_t \\ &\leq n^{-1/2} \|W_j^n(\cdot)\|_t + n^{-1/2} \|\varepsilon_2^n(\cdot)\|_t. \end{aligned}$$

From Lemma 7.2 and the assumption that $W_j^n \rightarrow W_j^*$ u.o.c., it follows that $\|e(\cdot) - \bar{\eta}_j^n(\cdot)\|_t \rightarrow 0$. \square

LEMMA 7.4. *If $W_j^n \rightarrow W_j^*$ u.o.c., then $\hat{\eta}_j^n \rightarrow W_j^*$ u.o.c.*

PROOF. It follows from equation (3.9) that, for each $t \geq 0$,

$$(7.2) \quad \begin{aligned} \|\hat{\eta}_j^n(\cdot) - W_j^*(\cdot)\|_t &\leq \|W_j^*(\cdot) - W_j^n(\cdot)\|_t \\ &\quad + \|W_j^n(\bar{\eta}_j^n(\cdot)) - W_j^n(\cdot)\|_t + \|\varepsilon_2^n(\cdot)\|_t. \end{aligned}$$

As a result of Lemma 7.3 and the continuity of W_j^* , the first two terms on the right-hand side of (7.2) converge to zero. Invoking Lemma 7.2, one concludes that $\hat{\eta}_j^n \rightarrow W_j^*$ u.o.c. \square

The proof of Theorem 6.1 proceeds by induction on the depth of stations. We begin with the following result for stations of depth 0.

LEMMA 7.5. *Theorem 6.1 holds for all stations of depth 0, namely, stations $j \in \sigma(0)$.*

PROOF. Note that $\mathcal{P}(k) = \emptyset$ for each task $k \in \mathcal{E}(j)$ when $j \in \sigma(0)$. Because $X_j^n = \xi_j^n$ for $j \in \sigma(0)$, it follows from Lemma 7.1, equations (3.3), (3.6) and (3.7) and from the continuous mapping theorem that $(X_j^n, W_j^n, I_j^n, Z_j^n) \rightarrow (X_j^*, W_j^*, I_j^*, Z_j^*)$ u.o.c., where $X_j^* = \xi_j^*$, $I_j^*(t) = \sup_{0 \leq s \leq t} \{X_j^*(s)\}^-$, $W_j^*(t) = X_j^*(t) + I_j^*(t)$ and $Z_j^*(t) = \xi_j^*(t) + I_j^*(t)$. Because

$$U_{jk}^n(t) = L_{jk}^n(t) - L_{jk}^n(\bar{\eta}_j^n(t)) + \rho_{jk}^{(n)} \hat{\eta}_j^n(t) - \varepsilon_{1k}^n(t),$$

it follows from Lemma 7.2, Lemma 7.4, and the continuity of L_{jk}^* that $U_{jk}^n \rightarrow U_{jk}^*$ u.o.c., where $U_{jk}^*(t) = \rho_{jk} W_j^*(t) = \rho_{jk} Z_j^*(t)$. Joint convergence of the processes of interest is a natural consequence of their continuity. \square

PROOF OF THEOREM 6.1. Define

$$\hat{Y}_{jk}^n(t) = n^{-1/2} (Y_{jk}^{(n)}(nt) - \rho_{jk}^{(n)} t) \quad \text{and} \quad \bar{Y}_{jk}^n(t) = n^{-1} Y_{jk}^{(n)}(nt),$$

and note that as a consequence of (3.11),

$$(7.3) \quad \hat{Y}_{jk}^n(t) = \hat{L}_{jk}^n(t) - U_{jk}^n(t).$$

With Lemma 7.5 we may assume inductively that the convergence in Theorem 6.1 has been established for all stations with depth h or less. Consider a station j with $d(j) = h + 1$. For each task $k \in \mathcal{E}(j)$, it follows from (3.12) that

$$\hat{A}_{jk}^n(t) = \begin{cases} \hat{N}_{q(k)}^n(t), & \text{if } \mathcal{P}(k) = \emptyset, \\ \min_{l \in \mathcal{P}(k)} \left\{ \hat{S}_{s(l)l}^n(\bar{Y}_{s(l)l}^n(t)) + \frac{1}{\tau_l^{(n)}} \hat{Y}_{s(l)l}^n(t) \right\}, & \text{otherwise.} \end{cases}$$

[Note that $\lambda_k = \lambda_l$ for all $l \in \mathcal{P}(k)$.] Setting $\bar{A}_{jk}^n(t) \equiv n^{-1}A_{jk}^{(n)}(nt)$, equation (3.13) gives

$$\hat{M}_{jk}^n(t) = \hat{V}_k^n(\bar{A}_{jk}^n(t)) + \tau_k^{(n)}\hat{A}_{jk}^n(t).$$

Because $d(s(l)) \leq h$ for all tasks $l \in \mathcal{P}(k)$, it follows from (7.3) and the induction hypothesis that $\hat{Y}_{s(l)l}^n \rightarrow Y_{s(l)l}^*$ u.o.c., where $Y_{s(l)l}^*(t) = L_{s(l)l}^*(t) - U_{s(l)l}^*(t)$, and $\bar{Y}_{s(l)l}^n \rightarrow \rho_{s(l)l} \cdot e$ u.o.c., where $e(t) = t$. Consequently, if $\mathcal{P}(k) \neq \emptyset$, $\hat{M}_{jk}^n \rightarrow M_{jk}^*$ u.o.c. with

$$\begin{aligned} M_{jk}^*(t) &= V_k^*(\lambda_k t) + \tau_k \min_{l \in \mathcal{P}(k)} \left\{ -\tau_l^{-3/2} V_l^*(\rho_{s(l)l} t) \right. \\ &\quad \left. + \tau_l^{-1} L_{s(l)l}^*(t) - \tau_l^{-1} U_{s(l)l}^*(t) \right\} \\ (7.4) \quad &= V_k^*(\lambda_k t) + \tau_k \min_{l \in \mathcal{P}(k)} \left\{ -\tau_l^{-1} V_l^*(\lambda_l t) \right. \\ &\quad \left. + \tau_l^{-1} (V_l^*(\lambda_l t) + \tau_l N_{q(l)}^*(t)) - \tau_l^{-1} U_{s(l)l}^*(t) \right\} \\ &= V_k^*(\lambda_k t) + \tau_k N_{q(k)}^*(t) - \tau_k \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*(t)}{\tau_l}, \end{aligned}$$

and the last equality follows because $q(k) = q(l)$ for $l \in \mathcal{P}(k)$. On the other hand, if $\mathcal{P}(k) = \emptyset$, we have $\hat{M}_{jk}^n \rightarrow M_{jk}^*$ u.o.c. with

$$M_{jk}^* = V_k^*(\lambda_k t) + \tau_k N_{q(k)}^*(t).$$

Applying the continuous mapping theorem to equation (3.14), $X_j^n \rightarrow X_j^*$ u.o.c., where

$$\begin{aligned} X_j^*(t) &= \sum_{k \in \mathcal{E}(j)} (V_k^*(\lambda_k t) + \tau_k N_{q(k)}^*(t)) - \sum_{k \in \mathcal{E}(j)} \max_{l \in \mathcal{P}(k)} \frac{\tau_k}{\tau_l} U_{s(l)l}^*(t) + \theta_j t \\ &= \xi_j^*(t) - \sum_{k \in \mathcal{E}(j)} \rho_{jk} \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*(t)}{\rho_{s(l)l}}, \end{aligned}$$

and we use the convention that $\max \emptyset = 0$. That $(W_j^n, I_j^n, Z_j^n) \rightarrow (W_j^*, I_j^*, Z_j^*)$ u.o.c. is again a consequence of the continuous mapping theorem by virtue of equations (3.3), (3.6) and (3.7). All that remains is to prove the convergence of U_j^n . From Lemmas 7.3 and 7.4, we have $\bar{\eta}_j^n \rightarrow e$ u.o.c. and $\hat{\eta}_j^n \rightarrow W_j^*$ u.o.c. Consequently, it follows from (3.8) that

$$Y_j^n(t) = \hat{M}_{jk}^n(\bar{\eta}_j^n(t)) - \rho_{jk}^{(n)} \hat{\eta}_j^n(t) + \varepsilon_{1k}^n(t),$$

from which we can conclude $Y_j^n \rightarrow Y_j^*$ u.o.c., where

$$(7.5) \quad Y_j^*(t) = M_{jk}^*(t) - \rho_{jk} W_j^*(t).$$

Because $U_{jk}^n(t) = \hat{L}_{jk}^n(t) - \hat{Y}_{jk}^n(t)$, it follows from (7.4) and (7.5) that $U_{jk}^n \rightarrow U_{jk}^*$ u.o.c. and

$$\begin{aligned} U_{jk}^*(t) &= L_{jk}^*(t) - T_{jk}^*(t) \\ &= L_{jk}^*(t) - (V_k^*(\lambda_k t) + \tau_k N_{q(k)}^*(t)) + \tau_k \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*(t)}{\tau_l} + \rho_{jk} W_j^*(t) \\ &= \rho_{jk} Z_j^*(t) + \rho_{jk} \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*(t)}{\rho_{s(l)l}} - \rho_{jk} \sum_{m \in \mathcal{E}(j)} \rho_{jm} \max_{n \in \mathcal{P}(m)} \frac{U_{s(n)n}^*(t)}{\rho_{s(n)n}}. \quad \square \end{aligned}$$

PROOF OF THEOREM 6.2. Define γ_{kl} as in (6.8), and rewrite equation (6.5) as

$$(7.6) \quad U_{ik}^* = \rho_{ik} \left[Z_i^* + \sum_{m \in \mathcal{E}(i)} \gamma_{km} \left(\max_{n \in \mathcal{P}(m)} \frac{U_{s(n)n}^*}{\rho_{s(n)n}} \right) \right].$$

Recall that a_l^i , $l = 1, \dots, c(i)$, enumerate the elements of $\mathcal{E}(i)$. Applying Lemma 5.1 to (7.6), we obtain

$$(7.7) \quad U_{ik}^* = \rho_{ik} \left[Z_i^* + \max_{x \in \mathcal{F}(i)} \left(\sum_{l=1}^{c(i)} \gamma_{ka_l^i} \frac{U_{s(x_l)x_l}^*}{\rho_{s(x_l)x_l}} \right) \right].$$

Similarly, equation (6.3) is equivalent to the following expression as a result of Lemma 5.1:

$$(7.8) \quad X_i^* = \xi_i^* - \max_{x \in \mathcal{F}(i)} \left(\sum_{l=1}^{c(i)} \rho_{ia_l^i} \frac{U_{s(x_l)x_l}^*}{\rho_{s(x_l)x_l}} \right).$$

Substituting (6.2) and (6.3) in (6.4) and applying (7.8), we have

$$W_i^* = Z_i^* - \max_{x \in \mathcal{F}(i)} \left(\sum_{l=1}^{c(i)} \rho_{ia_l^i} \frac{U_{s(x_l)x_l}^*}{\rho_{s(x_l)x_l}} \right).$$

Substituting (7.7) in the above expression, we obtain

$$\begin{aligned} (7.9) \quad W_i^* &= Z_i^* - \max_{x^1 \in \mathcal{F}(i)} \left(\sum_{l_1=1}^{c(i)} \rho_{ia_{l_1}^i} \left[Z_{s(x_{l_1}^1)}^* + \max_{x_{l_1}^2 \in \mathcal{F}(s(x_{l_1}^1))} \sum_{l_2=1}^{c(s(x_{l_1}^1))} \gamma_{x_{l_1}^1 a_{l_2}^{s(x_{l_1}^1)}} \frac{U_{s(x_{l_1}^2) x_{l_1}^2}^*}{\rho_{s(x_{l_1}^2) x_{l_1}^2}^*} \right] \right) \\ &= Z_i^* - \max_{x = (x^1, x^2) \in \mathcal{F}^2(i)} \left(\sum_{l_1=1}^{c(i)} \rho_{ia_{l_1}^i} Z_{s(x_{l_1}^1)}^* \right. \\ &\quad \left. + \sum_{l_1=1}^{c(i)} \sum_{l_2=1}^{c(s(x_{l_1}^1))} \rho_{ia_{l_1}^i} \gamma_{x_{l_1}^1 a_{l_2}^{s(x_{l_1}^1)}} \frac{U_{s(x_{l_1}^2) x_{l_1}^2}^*}{\rho_{s(x_{l_1}^2) x_{l_1}^2}^*} \right). \end{aligned}$$

For notational convenience, we henceforth write s^m to mean $s(x_{l_1}^m, \dots, x_{l_m}^m)$. Substituting (7.7) in (7.9) recursively, one can verify that, for a station of

depth h ,

$$(7.10) \quad W_i^* = Z_i^* - \max_{x \in \mathcal{I}^h(i)} \left(\sum_{l_1=1}^{c(i)} \rho_{ia_{l_1}^i} Z_{s^1}^* + \sum_{l_1=1}^{c(i)} \sum_{l_2=1}^{c(s^1)} \rho_{ia_{l_1}^i} \gamma_{x_{l_1}^1 a_{l_2}^1} Z_{s^2}^* \right. \\ \left. + \cdots + \sum_{l_1=1}^{c(i)} \cdots \sum_{l_h=1}^{c(s^{h-1})} \rho_{ia_{l_1}^i} \gamma_{x_{l_1}^1 a_{l_2}^1} \cdots \gamma_{x_{l_1, \dots, l_{h-1}}^{h-1} a_{l_h}^{h-1}} Z_{s^h}^* \right).$$

It is straightforward to verify that (7.10) is equivalent to equation (6.14) and the theorem is thus proved. \square

REMARK. Substituting (6.2)–(6.4) in equation (6.5), we obtain

$$U_{ik}^* = \rho_{ik} \left[W_i^* + \max_{l \in \mathcal{P}(k)} \frac{U_{s(l)l}^*}{\rho_{s(l)l}} \right].$$

Because $Z_i^*(t) = \sum_{k \in \mathcal{E}(i)} U_{ik}^*(t)$, it follows that

$$(7.11) \quad Z_i^* = W_i^* + \sum_{k \in \mathcal{E}(i)} \max_{l \in \mathcal{P}(k)} \rho_{ik} \frac{U_{s(l)l}^*}{\rho_{s(l)l}}.$$

Proceeding in the same manner as in the proof of Theorem 6.2, we can show that for a station i of depth h , (7.11) can be written as

$$(7.12) \quad Z_i^* = W_i^* + \max_{x \in \mathcal{I}^h(i)} \left\{ \sum_{l \in \mathcal{I}^1(x)} \rho_{ia_{l_1}^i} W_{s^1}^* + \cdots + \sum_{l \in \mathcal{I}^h(x)} \rho_{ia_{l_1}^i} W_{s^h}^* \right\}.$$

Readers may recognize that equation (7.12) is an “inverse” formulation of the relationship described in (6.14). It states that the amount of total work in the system for station i is the maximum of the amount of immediate work destined for station i found along each path to that station.

PROOF OF THEOREM 6.3. Define $\bar{\Phi}_q^n(t) \equiv n^{-1} \Phi_q^n(nt)$ and $\bar{\Phi}_{qk}^n(t) \equiv n^{-1} \Phi_{qk}^n(nt)$. In addition, note that

$$t \leq \Phi_q^n(t) \leq t + \max_{1 \leq i \leq N_q^n(t)+1} u_k(i) / \lambda_k^{(n)};$$

hence by Lemma 3.3 of Iglehart and Whitt [14],

$$(7.13) \quad \bar{\Phi}_q^n \rightarrow e \quad \text{u.o.c.}$$

We begin with tasks $k \in \mathcal{A}_q^0$, for which

$$\bar{\Phi}_{qk}^n(t) \equiv \bar{\Phi}_q^n(t) \quad \text{and} \quad T_{qk}^n(t) \equiv W_{s(k)}^n(\bar{\Phi}_{qk}^n(t)).$$

It follows from (7.13) and Theorem 6.1 that $\bar{\Phi}_{qk}^n \rightarrow e$ u.o.c. and $T_{qk}^n \rightarrow T_{qk}^*$, where $T_{qk}^* = W_{s(k)}^*(t)$. The theorem is then proved by applying induction on (3.15) and (3.16). \square

8. Concluding remarks. We presented in this paper a heavy traffic analysis of feedforward fork-join networks with heterogeneous customers. We made several assumptions to simplify the exposition, but the results proved here apply for more general networks as well. For example, we assumed that each station is staffed by a single server. Using the machinery developed by Chen and Shanthikumar [7], one can extend these results to fork-join networks of multiserver queues. Second, whereas we assumed that all servers are reliable, it is possible to analyze networks in which stations may experience server breakdown [13, 8]. Lastly, batch arrivals can be accommodated within the framework presented here [19]. (The model discussed by Baccelli and Liu [4] is an example of such networks.) In this case, the issue reduces to calculating Ω , the covariance matrix of the total workload input process [19].

Acknowledgments. I would like to thank Professors J. Michael Harrison and Avi Mandelbaum for the many helpful conversations throughout the course of this research.

REFERENCES

- [1] ADLER, P., MANDELBAUM, A., NGUYEN, V. and SCHWERER, E. (1993). From project to process management in engineering: An empirically-based framework for analysis of product development. *Management Sci.* To appear.
- [2] ADLER, P., MANDELBAUM, A., NGUYEN, V. and SCHWERER, E. (1992). From project to process management in engineering: Strategies for improving development cycle time. Unpublished manuscript.
- [3] ADLER, P., MANDELBAUM, A., NGUYEN, V. and SCHWERER, E. (1993). Managerial and methodological challenges in engineering process management: Lessons from a case study. Paper presented at UCLA Design Conference, Los Angeles. To appear.
- [4] BACCELLI, F. and LIU, Z. (1990). On the execution of parallel programs on multiprocessor systems—A queueing theory approach. *J. Assoc. Comput. Mach.* **37** 373–414.
- [5] BACCELLI, F. and MAKOWSKI, A. M. (1989). Queueing model for systems with synchronization constraints. *Proc. IEEE* **77** 138–161.
- [6] BILLINGSLEY, P. (1968). *Convergence of Probability Measures*. Wiley, New York.
- [7] CHEN, H. and SHANTHIKUMAR, J. G. (1992). Fluid limits and diffusion approximations for networks of multi-server queues in heavy traffic. Unpublished manuscript.
- [8] CHEN, H. and WHITT, W. (1992). Diffusion approximations for open queueing networks with server interruptions. Unpublished manuscript.
- [9] GLYNN, P. W. (1990). Diffusion approximations. In *Stochastic Models* (D. P. Heyman and M. J. Sobel, eds.) 145–198. North-Holland, Amsterdam.
- [10] HARRISON, J. M. (1985). *Brownian Motion and Stochastic Flow Systems*. Wiley, New York.
- [11] HARRISON, J. M. (1988). Brownian models of queueing networks with heterogeneous customer populations. In *Proceedings of the IMA Workshop on Stochastic Differential Systems*. Springer, New York.
- [12] HARRISON, J. M. and NGUYEN, V. (1990). The QNET method for two-moment analysis of open queueing networks. *Queueing Systems Theory Appl.* **6** 1–32.
- [13] HARRISON, J. M. and NGUYEN, V. (1993). Brownian models of multiclass queueing networks: Current status and open problems. *Queueing Systems Theory Appl.* **13** 5–40.
- [14] IGLEHART, D. L. and WHITT, W. (1970). Multiple channel queues in heavy traffic. I and II. *Adv. in Appl. Probab.* **2** 150–177; 355–364.

- [15] MANDELBAUM, M. and AVI-ITZHAK, B. (1968). Introduction to queueing with splitting and matching. *Israel J. Tech.* **6** 376–382.
- [16] NGUYEN, V. (1990). Heavy traffic analysis of processing networks with parallel and sequential tasks. Ph.D. dissertation, Dept. Operations Research, Stanford Univ.
- [17] NGUYEN, V. (1993). Processing networks with parallel and sequential tasks: Heavy traffic analysis and Brownian limits. *Ann. Appl. Probab.* **3** 28–55.
- [18] PETERSON, W. P. (1991). Diffusion approximations for networks of queues with multiple customer types. *Math. Oper. Res.* **16** 90–118.
- [19] REIMAN, M. I. (1984). Open queueing networks in heavy traffic. *Math. Oper. Res.* **9** 441–458.
- [20] REIMAN, M. I. (1988). A multiclass feedback queue in heavy traffic. *Adv. in Appl. Probab.* **20** 179–207.
- [21] VARMA, S. (1990). Heavy and light traffic approximations for queues with synchronization constraints. Ph.D. dissertation, Dept. Electrical Engineering, Univ. Maryland.
- [22] WHITT, W. (1971). Weak convergence theorems for priority queues: Preemptive-resume discipline. *J. Appl. Probab.* **8** 74–94.
- [23] WHITT, W. (1980). Some useful functions for functional limit theorems. *Math. Oper. Res.* **5** 67–85.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
SLOAN SCHOOL OF MANAGEMENT
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139