ON SYNCHRONIZATION OF POISSON PROCESSES AND QUEUEING NETWORKS WITH SERVICE AND SYNCHRONIZATION NODES

Balaji Prabhakar[1], Nicholas Bambos[2] and T.S. Mountford[3]

Abstract

This paper investigates the dynamics of a synchronization node in isolation, and of networks of service and synchronization nodes. A synchronization node consists of M infinite capacity buffers, where tokens arriving on M distinct random input flows are stored (there is one buffer for each flow). Tokens are held in the buffers until one is available from each flow. When this occurs, a token is drawn from each buffer to form a group-token, which is instantaneously released as a synchronized departure.

Under independent Poisson inputs, the output of a synchronization node is shown to converge weakly (and in certain cases strongly) to a Poisson process with rate equal to the minimum rate of the input flows. Hence synchronization preserves the Poisson property, as do superposition, Bernoulli sampling and M/M/1 queueing operations.

We then consider networks of synchronization and exponential server nodes with Bernoulli routing and exogenous Poisson arrivals, extending the standard Jackson Network model to include synchronization nodes. It is shown that if the synchronization skeleton of the network is acyclic (i.e. no token visits any synchronization node twice although it may visit a service node repeatedly), then the distribution of the joint queue-length process of only the *service* nodes is *product form* (under standard stability conditions) and easily computable. Moreover, the network output flows converge weakly to Poisson processes.

Finally, certain results for networks with finite capacity buffers are presented, and the limiting behavior of such networks as the buffer capacities become large is studied.

Keywords: synchronization, Poisson processes, queueing networks

1 Introduction

In this paper we analyze a basic synchronization operation on point processes, implemented by a system called the *synchronization node* (or synchronization queue), and study networks of service and synchronization nodes. We suppose that tokens arrive at a synchronization node according to $M \in \mathbb{Z}^+$ point processes

$$\mathbf{X}_i(t) = \sum_{n=-\infty}^{\infty} \mathbb{1}_{\{t_n^i = t\}}, \quad i \in \{1, 2, 3, ..., M\}, \tag{1}$$

[1]Departments of Electrical Engineering and Computer Science, Stanford University.

[2]Department of Engineering-Economic Systems & Operations Research and (by courtesy) Dept. of Electrical Engineering, Stanford University, Stanford CA 94305. Research supported in part by NSF grants NCR-9116268, DMI-9216034, and a National Young Investigator Award NCR-9258507.

[3]Department of Mathematics, University of California at Los Angeles, Los Angeles CA 90024. Research supported by NSF grant DMS 9157461, FAPESP, and a grant from the Sloan Foundation.

where t_n^i is the arrival time of the n^{th} token of the i^{th} flow. The synchronization node has M infinite capacity buffers, one for each flow, where arriving tokens are queued up (Figure 1). The synchronization operation consists of holding tokens in the buffers until one is available from each flow. As soon as this happens, exactly one token is taken from each buffer to form a group which is instantaneously released as a synchronized departure. The point processes $\mathbf{X}_i(\cdot)$ are assumed to be jointly stationary and ergodic, and the n-index numbering is such that $... < t_{-1}^i < 0 \leq t_0^i < t_1^i < ... < t_n^i < t_{n+1}^i < ...$ pathwise for $i \in \{1, 2, 3, ..., M\}$.

The synchronization operation is assumed to begin at time 0, all buffers being empty before that time. Thus, the departure time of the n^{th} $(n \geq 0)$ token group is

$$t_n^s = \max\{t_n^i, \ i \in \{1, 2, 3, ..., M\}\}$$

and the synchronized process is

$$\mathbf{S}(t) = \sum_{t_n^s} 1\!\!1_{\{t_n^s = t\}}. \tag{2}$$

The process $\mathbf{S}(t)$ is obviously not stationary with respect to time shifts, since the extra condition that the buffers are empty at time 0 is imposed on the system.

One of our primary goals is to study the departure flow of the synchronization node

$$\mathbf{S}_r(t) = \sum_{t_n^s \geq r} 1\!\!1_{\{t_n^s - r = t\}},$$

which is just $\mathbf{S}(t)$ viewed from time r onwards, as $r \to \infty$. Although some simple results are shown for general arrival processes, the most interesting ones are for the case where the $\mathbf{X}_i(t)$ are independent Poisson processes. In this case, it is shown that $\mathbf{S}_r(t)$ converges weakly to a Poisson process as $r \to \infty$. That is, *synchronization preserves the Poisson property* of flows, as do other basic operations like superposition, Bernoulli splitting and M/M/1 queueing [15, 18]. This key property can be exploited for analyzing whole networks of service and synchronization nodes, as discussed later.

Our original motivation for analyzing the synchronization queue comes from the following canonical model of *parallel processing*. Consider M processors working in parallel on an ongoing global computation, consisting of an infinite sequence of consecutive tasks. Each processor produces a sequence of *local events* corresponding to completions of subtasks assigned to it for execution. A *global event* corresponds to the completion of a task, and each task is considered completed only when all its corresponding subtasks have been executed. Assuming that the processing times for subtasks are independent and exponentially distributed with rate λ, the processes representing local events at each processor are independent Poisson processes of equal rate λ (load-balanced processors). A key question about this basic parallel processing paradigm is the determination of the statistics of the global event flow, which tracks the completion of tasks across all processors. It turns out that this is also Poisson. The model is basically a synchronization node with independent Poisson inputs of equal rate. Besides parallel processing, synchronization nodes (and networks of service and synchronization ones) have applications in various other areas, including database concurrency control, flexible manufacturing systems, communication protocols, etc.

Jackson Networks (and their extensions) [6, 7, 13, 15, 18] are very popular models for manufacturing and computer networks, and have been extensively used in analyzing the performance of these systems. This is mainly due to the fact that the queue-size process of a Jackson

Network is a Markov chain with a stationary distribution that is of the so-called "product-form" type. That is, one can compute the joint distribution of all the queue lengths in the network by treating each queue in isolation and simply taking the product of the individual distributions. Moreover, Jackson Networks have the additional property that all output processes are Poisson. There are three basic operations performed on traffic flows in the standard Jackson Network model: superposition, Bernoulli routing, and $\cdot$/M/1 queueing. Given that synchronization preserves the Poisson property, an interesting question that arises is whether it is possible to include synchronization as a basic operation in the model, along with the three others named above (see Figure 2), and yet retain the product-form stationary distributions for queue lengths. It turns out that this is possible, allowing one to generalize the familiar Jackson Network model (and other quasi-reversible networks [15, 18]) to include both service and synchronization nodes.

Understanding the dynamics of synchronization operations is important for the design of modern communication, computer and manufacturing systems. A powerful modelling framework for studying the logical and algorithmic aspects of concurrency and synchronization is provided by Petri nets [3, 9, 19, 20]. Unfortunately, their performance analysis and evaluation is very difficult in a general setup, due to the emerging complexities [11, 19, 20]. Recent interesting approaches for analyzing the dynamics of queueing networks with synchronization operations [3, 4, 5] have provided structural results (stability, existence of stationary states, stochastic bounds) under general stationary input flows. The synchronization node exhibits an essentially pathological behavior, in the sense that it is inherently unstable. This has been the subject of interesting studies (see [1, 12]) in the past. However, due to this inherent instability, the nature of the departure flow has remained elusive and no joint treatment of service and synchronization nodes in a Markovian setup has been possible.

This paper contributes to the understanding of the dynamics of synchronization nodes in isolation, as well as in interaction with other synchronization and service nodes. The issue concerning the nature of the node departure flows is resolved by introducing a technique which uses "ghost tokens" to account for deficits of real ones. We briefly describe it at the end of this section and use it extensively in the proofs where its effectiveness is demonstrated. The paper also extends the Jackson Network paradigm to a generalized network model including synchronization nodes (Figure 2). This retains the basic "product-form" property of stationary distributions, while it has more extensive modelling power than the standard model.

We proceed by introducing some additional notation needed to describe the dynamics of the synchronization node. For any $i \in \{1, 2, 3, ...M\}$, let $N_i(t) = \sum_{n=0}^{\infty} \mathbb{1}_{\{0<t_n^i \leq t\}}$ be the number of arrivals of $\mathbf{X}_i(t)$ to the synchronization node in $(0, t]$ (we assume $N_i(t) = 0$ a.s. for $t < 0$), and let $\lambda_i = E[N_i(1)]$ be the *arrival rate* of $\mathbf{X}_i(t)$. Further, let $N_s(t) = \sum_{n=0}^{\infty} \mathbb{1}_{\{0<t_n^s \leq t\}}$ be the number of synchronized departures in $(0, t]$. Then the number of tokens in the i^{th} buffer at time t, $Q_i(t) = N_i(t) - N_s(t)$, is an almost surely right-continuous process with left-hand limits.

Observe that, by definition of the synchronization operation, at least one of the M buffers must be empty at any given instant. Specifically, in the *2-input case*, this implies $\min\{Q_1(t), Q_2(t)\} = 0$ for all $t > 0$. Thus

$$Q(t) = Q_1(t) - Q_2(t) \tag{3}$$

completely specifies the status of the buffers at time t. Note also that $Q(t) = N_1(t) - N_2(t)$. An alternative definition for $\mathbf{S}(t)$ in terms of $Q(t)$ is given by the equation

$$\mathbf{S}(t) = \mathbf{X}_1(t)\mathbb{1}_{\{Q(t^-)<0\}} + \mathbf{X}_2(t)\mathbb{1}_{\{Q(t^-)>0\}}. \tag{4}$$

Intuitively, this means that an event occurs for $\mathbf{S}(\cdot)$ at some time t, iff just prior to t (i.e. at t^-) one of the buffers is nonempty and there is an arrival to the other buffer. Due to the right continuity of the paths of $Q(t)$, it is necessary in to use $Q(t^-)$ in (4) rather than $Q(t)$.

The focus of this paper is on the canonical case where the input flows are independent Poisson processes. Considerable emphasis is placed on the special case of *equal arrival rates*, because it naturally captures *load balancing* considerations that are essential in many practical situations. The rest of the paper is organized as follows. In Section 2 the synchronization node with independent Poisson inputs is studied in isolation. We establish the following results. If there is a unique input of minimum rate, then the synchronization process converges to that input process strongly (in total variation); otherwise, it converges weakly (and provably not strongly) to a Poisson process of rate equal to the minimum rate of the inputs. Similar results are shown to be true for finite buffers in the limit as their sizes tend to infinity. The analysis is based on exploiting the transience or null-recurrence of the Markovian queue-length process and its effect on the synchronization process.

In Section 3 we focus on networks of synchronization and exponential service nodes, with Bernoulli routing and independent Poisson exogenous arrivals, generalizing the standard Jackson Network model. It is shown that if the *synchronization skeleton* of the network is *acyclic* (a token visits a synchronization node only once, but may visit a service node any number of times), then the stationary distribution of the joint queue-size process of the service nodes alone is product-form under standard stability conditions. Moreover, all network output flows converge weakly to Poisson processes. The results extend to networks with quasi-reversible service nodes. To prove the key results a special technique is employed, wherein "ghost tokens" are injected into non-Poisson flows substituting for real tokens and turning them into Poisson ones. By showing that the ghost token flow eventually dwindles to zero, we are able to prove that the real token flows asymptotically converge to Poisson processes. This technique may have applications in other problems dealing with the convergence of stochastic flows.

2 The Synchronization Node

Figure 1 illustrates the synchronization queue (node) in the M-input and 2-input cases. In Theorem 1 we show that the synchronization of two independent Poisson processes of *equal rate* converges weakly to a Poisson process. Theorem 2 generalizes this result to the M-flow case. In Theorem 3 we prove that strong convergence is *impossible* when synchronizing independent Poisson processes of equal rate. This group of results provides a precise characterization of the convergence mode of the synchronization process.

In a few sentences we describe the synchronization of arbitrarily distributed point processes, when they are of *unequal rates*. Suppose that we have two stationary and ergodic processes $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ of rates λ_1 and λ_2 respectively. Further suppose that $\lambda_1 < \lambda_2$ and that synchronization begins at time 0 (both buffers being empty before that time). Since $\mathbf{X}_2(t)$ has a higher rate than $\mathbf{X}_1(t)$, there is a finite random time τ such that $N_2(t) > N_1(t)$ for all $t > \tau$. This causes buffer 2 to never empty after τ, and so $\mathbf{S}(t) = \mathbf{X}_1(t)$ for all $t > \tau$. The following lemma summarizes this simple observation.

Lemma 1 *Let the processes $\mathbf{X}_i(t)$, $i = 1, 2, ..., M$, be jointly stationary and ergodic with rates*

λ_i such that $\lambda_1 < \lambda_2 \leq ... \leq \lambda_M$ (i.e. there is a unique process of minimum rate). Then there is an almost surely finite random time τ, such that $\mathbf{S}(t) = \mathbf{X}_1(t)$ for all $t > \tau$. Hence, $\mathbf{S}_r(t)$ converges in total variation [2] to the input process with the slowest rate as $r \to \infty$.

2.1 Synchronizing Two Independent Poisson Processes of Equal Rates

We now consider the more interesting problem of synchronizing independent Poisson processes of *equal rates*. We begin with the simpler 2-input case. Let $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ be two independent Poisson processes of rate λ. Define $Q(t) = Q_1(t) - Q_2(t)$, as in Equation (3). Then $Q(t)$ is a continuous-time, null-recurrent, birth-death chain on $\mathbb{Z}$, with birth and death rates equal to λ. The null-recurrence implies that $\lim_{t\to\infty} P(Q(t) = 0) = 0$.

We will show that $\mathbf{S}_r(t)$ converges weakly to a Poisson process, by showing that the distributional difference between $\mathbf{S}_r(t)$ and the process

$$\mathbf{P}(t) = \mathbf{X}_1(t)\mathbb{1}_{\{Q(t^-)<0\}} + \mathbf{X}_2(t)\mathbb{1}_{\{Q(t^-)>0\}} + \mathbf{Y}(t)\mathbb{1}_{\{Q(t^-)=0\}}, \tag{5}$$

goes to zero as $r \to \infty$. The process $\mathbf{Y}(\cdot)$ in (5) is a rate λ Poisson process, independent of $\mathbf{X}_1(\cdot)$ and $\mathbf{X}_2(\cdot)$. Lemma 2 shows that $\mathbf{P}(t)$ is a rate λ Poisson process, thus establishing the claim. For a discussion of weak convergence in the context of point processes see [10].

Recall (see [7], for example) that λ_t is said to be the $\mathcal{G}_t$-intensity of the stochastic point process $\mathbf{Z}(t)$ adapted to some history $\mathcal{G}_t$, iff λ_t satisfies the following conditions: it is a non-negative $\mathcal{G}_t$-progressive process, such that $\int_0^t \lambda_s ds < \infty$ almost surely for all $t \geq 0$, and

$$E\left[\int_0^\infty C_s d\mathbf{Z}(s)\right] = E\left[\int_0^\infty C_s \lambda_s ds\right]$$

for all non-negative $\mathcal{G}_t$-predictable processes C_t. We then have the following two facts (see [7]):

<u>*Fact 1:*</u> If λ_t is the $\mathcal{F}_t$-intensity of the point process $\mathbf{Z}(t)$ and λ_t is $\mathcal{G}_t$-progressive for some history $\mathcal{G}_t$ such that $\mathcal{F}_t^Z \subseteq \mathcal{G}_t \subseteq \mathcal{F}_t$ ($\mathcal{F}_t^Z$ being the internal history of $\mathbf{Z}(t)$), then λ_t is also the $\mathcal{G}_t$-intensity of $\mathbf{Z}(t)$.

<u>*Fact 2:*</u> Let $\mathbf{Z}(t)$ have $\mathcal{F}_t$-intensity λ_t and let $\mathcal{G}_t$ be some history such that $\mathcal{G}_\infty$ is independent of $\mathcal{F}_t$ for all $t \geq 0$. Then λ_t is also the $\mathcal{F}_t \vee \mathcal{G}_t$-intensity of $\mathbf{Z}(t)$, where $\mathcal{F}_t \vee \mathcal{G}_t$ is the smallest σ-algebra containing both $\mathcal{F}_t$ and $\mathcal{G}_t$.

Note that a process may have a constant intensity λ with respect to a filtration, and this makes it a Poisson process *with respect to that particular filtration.* In order for it to be a rate λ Poisson process in the standard sense, its intensity with respect to its own history (i.e. the minimal filtration to which it is adapted) must be λ. This is what Lemma 2 establishes.

Lemma 2 *(2-Input Case) If $\mathbf{X}_1(t)$, $\mathbf{X}_2(t)$ and $\mathbf{Y}(t)$ are independent Poisson processes with intensity λ and $Q(t)$ is the system size process of Equation (3), then*

$$\mathbf{P}(t) = \mathbf{X}_1(t)\mathbb{1}_{\{Q(t^-)<0\}} + \mathbf{X}_2(t)\mathbb{1}_{\{Q(t^-)>0\}} + \mathbf{Y}(t)\mathbb{1}_{\{Q(t^-)=0\}}$$

is a Poisson process with intensity λ.

Proof Observe that $\mathbf{P}(t)$ is adapted to $\mathcal{F}_t = \sigma\{\mathbf{X}_1(s), \mathbf{X}_2(s), \mathbf{Y}(s);\ 0 \leq s \leq t\}$ and that, due to Fact 2 and the mutual independence of the processes $\mathbf{X}_1(t)$, $\mathbf{X}_2(t)$ and $\mathbf{Y}(t)$, all three of them have $\mathcal{F}_t$-intensity λ. Further, the functions $1\!\!1_{\{Q(t^-)<0\}}$, $1\!\!1_{\{Q(t^-)>0\}}$ and $1\!\!1_{\{Q(t^-)=0\}}$ are $\mathcal{F}_t$-predictable (being left-continuous and $\mathcal{F}_t$-adapted). It readily follows that

$$E\left[\int_0^\infty C_s\ d\mathbf{P}(s)\ \right] = \lambda\ E\left[\int_0^\infty C_s\ ds\ \right],$$

for all $\mathcal{F}_t$-predictable processes C_t. Therefore, by Watanabe's Characterization Theorem (see [7]), $\mathbf{P}(t)$ is a Poisson process with respect to $\mathcal{F}_t$. Since the internal history of $\mathbf{P}(t)$ is contained in $\mathcal{F}_t$, Fact 1 implies that $\mathbf{P}(t)$ is Poisson with respect to its own history. ∎

We next show that the sequence of point processes $\{\mathbf{S}_r(t), r \in \mathbb{R}_+\}$ converges weakly [10] to a Poisson process as $r \to \infty$. As a consequence of Lemma 2 it suffices to show that for every bounded continuous function f with compact support, the random variable $\int_{\mathbb{R}^+} f(s)\ d\mathbf{S}_r(s)$ converges in distribution to $\int_{\mathbb{R}^+} f(s)\ d\mathbf{P}(s)$. Intuitively, a comparison of equations (4) and (5) shows that $\mathbf{S}(s)$ equals $\mathbf{P}(s)$ when $Q(s^-) \neq 0$. In other words, so long as one of the two buffers is non-empty the synchronized process equals the process arriving to the other (the empty) buffer, which is Poisson. Thus, the synchronized process is a random mixture of Poisson processes (depending on which buffer is nonempty) and the identically zero random process (if both buffers are empty). We exploit the *null recurrence* of $Q(\cdot)$ to establish that eventually the chance that both buffers are empty is arbitrarily small. Thus the "nuisance" term $\mathbf{Y}(s) 1\!\!1_{\{Q(s^-)=0\}}$ in (5) goes to zero in distribution and we get the desired weak convergence.

Theorem 1 *(2-Input Case) The synchronization process viewed from time r onwards, $\mathbf{S}_r(t)$, converges weakly to a rate λ Poisson process as $r \to \infty$.*

Proof Let f be a bounded continuous function with support in $[0, N]$. Set

$$x_r = \int_0^N f(s) d\left(\mathbf{Y}(r+s) 1\!\!1_{\{Q((r+s)^-)=0\}}\right), \quad y_r = \int_0^N f(s) d\mathbf{S}_r(s), \quad \text{and } z_r = \int_0^N f(s) d\mathbf{P}(r+s).$$

Note that $x_r + y_r = z_r$ and that z_r equals z_0 in distribution. We want to show that y_r converges in distribution to z_0. Now

$$\begin{aligned} E|x_r| &\leq E\left[\int_0^N |f(s)| d\left(\mathbf{Y}(r+s) 1\!\!1_{\{Q((r+s)^-)=0)\}}\right)\right] \\ &\leq |f|_{max}\ E\left[\int_0^N d\left(\mathbf{Y}(r+s) 1\!\!1_{\{Q((r+s)^-)=0)\}}\right)\right] \\ &= |f|_{max}\ \lambda\ E\left[\int_0^N 1\!\!1_{\{Q((r+s)^-)=0)\}} ds\right] \\ &= |f|_{max}\ \lambda \int_0^N P(Q((r+s)^-) = 0) ds, \end{aligned}$$

the last equality following from Fubini's theorem. Since $P(Q((r+s)) = 0)$ goes to zero as $r \to \infty$, by dominated convergence the right-most term goes to zero as $r \to \infty$. Therefore, $x_r \to 0$ in distribution. Now z_r equals z_0 in distribution and we conclude (see [8]; Theorem 4.4.6) that $y_r = z_r - x_r$ converges to z_0 in distribution, as was required. ∎

2.2 Extension to the M-input case

We extend the results to the M-input synchronization node, briefly discussing the additional arguments needed. Consider $M \in \mathbb{Z}^+$ independent Poisson processes $\mathbf{X}_i(t)$, $i \in \{1, 2, 3, ..., M\}$, all having rates equal to λ. Let $\mathbf{S}(t)$ be their synchronization and let $Q_i(t) = N_i(t) - N_s(t)$ be the queue length process in the i^{th} buffer at time t. For $i \in \{1, 2, ..., M\}$, define

$$\begin{aligned} f_i^M(t) &= \begin{cases} 1 & \text{if} \quad Q_i(t^-) = 0 \quad \text{and} \quad Q_j(t^-) > 0 \quad \text{for every} \quad j \neq i \\ 0 & \text{otherwise} \end{cases} \\ f_0^M(t) &= 1 - \sum_{i=1}^{M} f_i^M(t) \end{aligned} \tag{6}$$

The $f_i^M(t), i > 1$ indicate that at t^- only the i^{th} buffer is empty, and $f_0^M(t)$ indicates that at t^- more than one buffer is empty. Thus,

$$\{\ f_0^M(t) = 1\ \} = \bigcup_{i \neq j} \{\ Q_i(t^-) = Q_j(t^-) = 0\ \}. \tag{7}$$

Analogous to (4) we define $\mathbf{S}^M(t)$ as

$$\mathbf{S}^M(t) = \sum_{i=1}^{M} f_i^M(t)\, \mathbf{X}_i(t). \tag{8}$$

and show that as $r \to \infty$, $\mathbf{S}_r^M(t) = \mathbf{S}^M(r+t) \mathbb{1}_{\{r+t>0\}}$ converges weakly to the process

$$\mathbf{P}^M(t) = \sum_{i=1}^{M} f_i^M(t)\, \mathbf{X}_i(t) \ + \ f_0^M(t)\, \mathbf{Y}(t), \tag{9}$$

where $\mathbf{Y}(t)$ is a Poisson process of rate λ independent of $\mathbf{X}_i(t)$, $i \in \{1, 2, 3, ..., M\}$ as $r \to \infty$.

Lemma 3 *(M-Input Case) If $\mathbf{X}_i(t)$, $i \in \{1, 2, ..., M\}$ and $\mathbf{Y}(t)$ are independent Poisson processes of rate λ, then $\mathbf{P}^M(t) = \sum_{i=1}^{M} f_i^M(t)\, \mathbf{X}_i(t) \ + \ f_0^M(t)\, \mathbf{Y}(t)$ is a Poisson process of rate λ, where $f_i^M(t)$ $i \in \{1, 2, ..., M\}$ and $f_0^M(t)$ are defined in Equation (6).*

Proof Similar to the proof of Lemma 2. ∎

In the next lemma we show that the chance that more than one buffer is empty goes to zero eventually. This is used in Theorem 2 in the same way as null-recurrence of $Q(t)$ was used in Theorem 1.

Lemma 4 *With $f_0^M(t)$ defined as in (6),* $\lim_{t \to \infty} E[f_0^M(t)] = 0$.

Proof Since $E[f_0^M(t)] = P(f_0^M(t) = 1)$, from Equation (7) we see that it suffices to show that $\sum_{i \neq j} P(\{\ Q_i(t^-) = Q_j(t^-) = 0\ \}) \to 0$, as $t \to \infty$. Letting $B_{ij}^t = \{\ Q_i(t^-) = Q_j(t^-) = 0\}$, it suffices to show that $P[B_{ij}^t] \to 0$ for every i, j. Define $Q_{ij}(t) = Q_i(t) - Q_j(t) = N_i(t) - N_s(t) - (N_j(t) - N_s(t)) = N_i(t) - N_j(t)$. Then, $Q_{ij}(t)$ is a null-recurrent birth-death chain. Since $B_{ij}^t \subset \{Q_{ij}(t^-) = 0\}$, $P[Q_{ij}(t^-) = 0] \to 0$ implies that $P[B_{ij}^t] \to 0$ for each i, j. ∎

Theorem 2 *(M-Input Case) The synchronization of M independent Poisson processes viewed from time r onwards,*

$$\mathbf{S}_r^M(t) = \sum_{i=1}^{M} f_i^M(r+t)\,\mathbf{X}_i(r+t)\,\mathbb{1}_{\{r+t>0\}},$$

converges weakly to a Poisson process as $r \to \infty$.

Proof As in Theorem 1, let the function f be continuous with support in $[0, N]$. This time set $x_r = \int_0^N f(t)\ d\left(f_0^M(r+t)\,\mathbf{Y}(r+t)\right), y_r = \int_0^N f(t)\ d\mathbf{S}_r^M(t), z_r = \int_0^N f(t)\ d\mathbf{P}^M(r+t), z = \int_0^N f(t)\ d\mathbf{P}^M(t)$ and use the fact (from Lemma 4) that $E[f_0^M(r+t)] \to 0$ as $r \to \infty$ to conclude that x_r goes to zero in distribution. The rest follows from the method of Theorem 1. ∎

2.3 Convergence and Coupling

A point process $\mathbf{X}(t)$ is said to *couple in finite time* with another point process $\mathbf{Y}(t)$, if there is a random time $\tau < \infty$, such that $\mathbf{X}(s) = \mathbf{Y}(s)$ for all $s > \tau$ almost surely. This is exactly what happens under the conditions of Lemma 1. A natural question that arises is whether the synchronization, $\mathbf{S}^M(t)$, of independent Poisson processes of equal rate couples in finite time with some other Poisson process $\mathbf{Z}(t)$. This will strengthen the results of Theorems 1 and 2. Unfortunately, as Theorem 3 shows, such a coupling is not possible. This negative result further characterizes the convergence mode of the synchronization process.

The basic idea of the proof is this: If the synchronization process $\mathbf{S}^M(\cdot)$ couples with a Poisson process, say $\mathbf{Z}(\cdot)$, then the epochs of the two processes coincide after a finite random time. But the inter-epoch times of $\mathbf{Z}(\cdot)$ are exponentially distributed, and are always *strictly* stochastically dominated by the inter-epoch times of $\mathbf{S}^M(\cdot)$, and hence a coupling between $\mathbf{Z}(\cdot)$ and $\mathbf{S}^M(\cdot)$ cannot occur. We argue as follows. It is clear that the inter-epoch times of $\mathbf{S}^M(\cdot)$ are atleast exponentially long because we are always waiting for one component token to arrive at the synchronization buffers. And almost surely, given any T, there is a time $t > T$ at which more than one synchronization buffer is empty. When this happens the time to the next synchronization is the sum of *more than one* exponential time. This causes the strict stochastic dominance and the resulting lack of coupling.

Theorem 3 *When synchronizing M independent Poisson processes of rate λ, coupling the synchronization process $\mathbf{S}^M(t)$ in finite time with a Poisson process of rate λ is impossible.*

Proof Arguing by contradiction suppose that there is a rate λ Poisson process $\mathbf{Z}(t)$ and a finite random time $\tau < \infty$, such that $\mathbf{S}^M(s) = \mathbf{Z}(s)$ for all $s > \tau$ almost surely. Define $N_z(t) = \sum_{n=0}^{\infty} \mathbb{1}_{\{0 < t_n^z \le t\}}$ to be the number of points of $\mathbf{Z}(t)$ in $(0, t]$, where the t_n^z are the epoch times for $\mathbf{Z}(t)$. By the Central Limit Theorem

$$\frac{N_z(t) - \lambda t}{\sqrt{\lambda t}} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1), \tag{10}$$

where $\mathcal{N}(0,1)$ is the Gaussian distribution with zero mean and unit variance. Now, $\tau < \infty$ implies that $N_z(\tau) < \infty$ almost surely, which further implies that $N_z(\tau)/\sqrt{\lambda t} \to 0$ almost surely, as $t \to \infty$. Thus,

$$\frac{N_z(t) - N_z(\tau) - \lambda t}{\sqrt{\lambda t}} = \frac{N_z(t) - \lambda t}{\sqrt{\lambda t}} + \frac{N_z(\tau)}{\sqrt{\lambda t}} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1),$$

and for any $\epsilon > 0$ we have

$$\lim_{t\to\infty} P\left(N_z(t) - N_z(\tau) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) = \lim_{t\to\infty} P\left(\frac{N_z(t) - N_z(\tau) - \lambda t}{\sqrt{\lambda t}} \geq \epsilon\right) = \Phi(\epsilon), \tag{11}$$

where $\Phi(\epsilon) = \frac{1}{\sqrt{2\pi}} \int_\epsilon^\infty e^{-x^2/2}\, dx$. We now observe that given a $K \in \mathbb{R}^+$, in order to have at least K synchronizations in $[0, t]$, we must have at least K arrivals to each of the M buffers in $[0, t]$, so

$$\{\ N_s(t) \geq K\ \} \subset \bigcap_{i=1}^{M} \{\ N_i(t) \geq K\ \}.$$

Taking $K = \lambda t + \epsilon\sqrt{\lambda t}$, we get

$$\begin{aligned} P\left(N_s(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) \leq P\left(\bigcap_{i=1}^{M}\{\ N_i(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\ \}\right) &= \prod_{i=1}^{M} P\left(N_i(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) \\ &= \left(P\left(N_1(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right)\right)^M, \end{aligned}$$

since the Poisson processes $N_i(t)$ are independent with equal rates. Applying the Central Limit Theorem to $N_1(t)$, we get

$$\lim_{t\to\infty} P\left(N_1(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) = \Phi(\epsilon).$$

Thus,

$$\lim_{t\to\infty} P\left(N_s(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) \leq \lim_{t\to\infty}\left(P\left(N_1(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right)\right)^M = (\Phi(\epsilon))^M < \Phi(\epsilon), \tag{12}$$

since $\Phi(\epsilon) < 1$. On the other hand, since we have supposed that the coupling exists, we have that $\mathbf{S}^M(s) = \mathbf{Z}(s)$ for all $s > \tau$ almost surely, so $N_s(t) - N_s(\tau) = N_z(t) - N_z(\tau)$ for every $t > \tau$. Moreover, since $N_s(\tau) < \infty$, we have $N_s(\tau)/\sqrt{\lambda t} \to 0$ almost surely. Therefore, as $t \to \infty$

$$\frac{N_s(t) - \lambda t}{\sqrt{\lambda t}} - \frac{N_s(\tau)}{\sqrt{\lambda t}} = \frac{N_s(t) - N_s(\tau) - \lambda t}{\sqrt{\lambda t}} = \frac{N_z(t) - N_z(\tau) - \lambda t}{\sqrt{\lambda t}} \xrightarrow{\mathcal{D}} \mathcal{N}(0,1).$$

Thus,

$$\lim_{t\to\infty} P\left(N_s(t) \geq \lambda t + \epsilon\sqrt{\lambda t}\right) = \Phi(\epsilon),$$

which contradicts (12). Therefore, no such coupling exists. ■

2.4 The Case of Finite Buffers

We next consider the synchronization of Poisson processes which arrive at buffers of finite capacity. The queue-length process $\vec{Q}(t) = (Q_1(t), Q_2(t), ..., Q_M(t))$ is a finite state Markov chain and therefore has a stationary distribution. Consider the 2-input case first. Let tokens arriving according to independent Poisson processes of rates λ_1 and λ_2 be queued into buffers of sizes N_1

and N_2 respectively. Tokens arriving at full buffers are blocked and rejected. The state space, S, of the Markov chain is the set $\{(0,N_2),...,(0,1),(0,0),(1,0),...,(N_1,0)\}$. The corresponding equilibrium distribution, $\{\pi(i,j),\ (i,j)\in S\}$, is given by $\pi(i,j) = ca^ib^j$, where $a = \lambda_1/\lambda_2$, $b = \lambda_2/\lambda_1$ and c is a normalizing constant. Note that the stationary distribution is *product-form* in the 2-input case.

Now consider the M-queue case ($M \geq 3$). Let $\mathbf{X}_i(t),\ i \in \{1,2,..,M\}$ be independent rate λ Poisson processes, arriving at buffers of size N_i. Since at least one queue is empty at any given time, we see that the rate matrix R of $\vec{Q}(t)$ has just two types of entries. The first type is

$$R((y_1,..,y_j,..,y_k,..,y_M),(y_1,..,y_j,..,y_k+1,..,y_M)) \;=\; \lambda_k$$

for $y_j = 0$ and $y_k = 0,..,N_k-1,\ k \neq j$, corresponding to the case that the j-th queue is empty and there is an arrival in the k-th one ($k \neq j$). The second type is

$$R((y_1,..,y_j,..,y_k,..,y_M),(y_1-1,..,y_j,..,y_k-1,...,y_M-1)) \;=\; \lambda_j$$

for $y_j = 0$ and $y_k > 0,\ k \neq j$, corresponding to the case that only the j-th queue is empty and there is an arrival to that queue, triggering a synchronized departure from all queues. In light of the result for the 2-input case one wonders whether the stationary distribution is product-form. Unfortunately, as simple examples show, this is *not* true.

A consequence of the null-recurrence or transience of the joint queue-size process of infinite-buffer synchronization is this: *When synchronizing independent Poisson processes of equal rate at* finite buffers*, the synchronized process can be made to be as close to a Poisson process in distribution as desired, by making the buffers suitably large.* The details are as follows.

Suppose that we are synchronizing M independent Poisson processes $\mathbf{X}_i(t), i \in \{1,2,...,M\}$ of rate equal to λ. All the synchronization buffers are assumed to be of capacity k. Therefore the joint queue-size process $\vec{Q}^k(\cdot) = (Q_1^k(\cdot), Q_2^k(\cdot),...,Q_M^k(\cdot))$ is a positive recurrent Markov chain. Define the functions $g_i^k(t),\ i = 0,1,2,...,M$ as follows

$$\begin{aligned}
g_i^k(t) &= \begin{cases} 1 & \text{if} \quad Q_i^k(t^-) = 0 \quad \text{and} \quad Q_j^k(t^-) > 0 \quad \text{for every} \quad j \neq i \\ 0 & \text{otherwise} \end{cases} \\
g_0^k(t) &= 1 - \sum_{i=1}^{M} g_i^k(t).
\end{aligned}$$

The synchronized process, $\mathbf{S}^{M,k}(t)$, is given by

$$\mathbf{S}^{M,k}(t) = \sum_{i=1}^{M} g_i^k(t)\,\mathbf{X}_i(t).$$

We are interested in the asymptotic distribution of the process $\mathbf{S}^{M,k}(t)$ as $k \to \infty$.

For each finite k there is an equilibrium distribution for $\vec{Q}^k(t)$. Arguing as in Lemma 4, we get that $\lim_{k\to\infty} E(g_0^k(t)) = \lim_{k\to\infty} P(g_0^k(t) = 1) = 0$. And, an argument similar to Lemma 3 shows that the process

$$\mathbf{P}^{M,k}(t) = \sum_{i=1}^{M} g_i^k(t)\,\mathbf{X}_i(t) \;+\; g_0^k(t)\,\mathbf{Y}(t)$$

is Poisson, where $\mathbf{Y}(t)$ is a rate λ Poisson process independent of all the other processes. Given that $\lim_{k\to\infty} E(g_0^k(t)) = 0$, it follows in a manner similar to the proof of Theorem 2 that $\mathbf{S}^{M,k}(t)$ converges weakly to a Poisson process as $k \to \infty$.

3 Generalized Jackson Networks of Service & Synchronization Nodes

In this section we show how synchronization operations may be included in queueing networks, extending the classical Jackson Network model and its associated properties.

For the sake of completeness, we recall a few basic facts about Jackson Networks. In their standard form, such networks [6, 13, 15, 18] consist of independent exponential service nodes, where arriving tokens (jobs) are served and then Bernoulli-routed to other nodes (or back to the same one). The net input process at a node is the superposition of all token flows that arrive (either from outside or after being routed) at that node. All exogenous arrivals are independent Poisson processes. The overall queue-length process is a Markov chain. When the total average arrival rate is less than the service rate at every node of the network, the queue-length process admits a product-form stationary distribution. Moreover, in stationarity, the flows of tokens departing from the network are independent Poisson processes.

An example of a *Generalized Jackson Network*, including synchronization operations, is shown in Figure 2. After being served or synchronized, tokens are Bernoulli-routed to various nodes. The net input process at each node is the superposition of all arriving processes, synchronized or otherwise. In Figure 2, for example, the net input into node 5 is a synchronization of outputs from nodes 1 and 2 superposed with an output flow from node 4. Exogenous arrivals are Poisson, and the buffers of all nodes have infinite capacity.

It is easy to see that the Generalized Jackson Network model retains its Markovian nature with respect to a system state consisting of the queue-length processes of all nodes (service and synchronization). However, the synchronization queues are transient or null-recurrent and, hence, induce the same behavior on the overall Markov chain. Under stability conditions on the service nodes, we are interested in determining whether the distribution of the queue-lengths at these nodes (exclusively) is asymptotically product-form. We find this to be true when the synchronization skeleton is acyclic.

3.1 Generalized Jackson Networks with Acyclic Synchronization Skeleton

Throughout this section we assume that a token *cannot visit a synchronization node twice*, almost surely, although it may visit a service node any number of times. Consider a queueing network $\mathbf{N}$ consisting of M exponential server nodes and N synchronization nodes. Given is a set of *routing probabilities* $\{p_{ij}\}_{1 \le i,j \le M+N}$ where p_{ij} is the probability that a token joins node j immediately after leaving node i (i and j can be either of the service or synchronization types). Nodes i and j are said to be *connected by a route* if there is a strictly positive probability of a token arriving at j after leaving i, either directly or through a series of intermediate nodes (note that the route is directed from i to j). Otherwise, nodes i and j are said to be *disconnected.* We impose the following key condition: Every synchronization node is disconnected from itself, and hence the synchronization skeleton is acyclic.

We say that a service node j is an *offspring* of a synchronization node i if there is a route from i to j. The acyclicity of the synchronization skeleton partitions the service nodes into stages $\mathbf{S}_k$, $k = 0, 1, 2, ...$, where $\mathbf{S}_k$ consists of service nodes that are the offspring of precisely k synchronization nodes. Thus the network $\mathbf{N}$ can be decomposed into stages of smaller Jackson subnetworks of service nodes only (possibly with feedback), which interface through the synchronization nodes. A *K-stage* network is one in which at least one service node is the offspring of K synchronization nodes. The example of Figure 7 is a 3-stage network. It is composed of the subnetworks $\mathbf{N}_1, \mathbf{N}_2, \mathbf{N}_3, \mathbf{N}_4, \mathbf{N}_5$ and $\mathbf{N}_6$ each consisting of service nodes only. Each $\mathbf{N}_i$ is a standard Jackson network in itself. For this example, $\mathbf{S}_0 = \mathbf{N}_1 \cup \mathbf{N}_2 \cup \mathbf{N}_3$, $\mathbf{S}_1 = \mathbf{N}_4 \cup \mathbf{N}_5$ and $\mathbf{S}_2 = \mathbf{N}_6$. Note that all the synchronization nodes have been arranged on the boundaries between the service subnetworks.

We discuss the *stability of a service node* before stating the main theorem of this section. Service nodes that are not the offspring of any synchronization node are stable if the net input flow (exogenous arrivals and tokens routed from this or other nodes) has an average rate that is strictly smaller than the service rate of the node. In considering the stability of service nodes whose input may contain a synchronized process, the acyclicity of the synchronization skeleton allows us to treat the synchronized process as an exogenous input. Further, Equations (8) and (9) imply that a synchronized flow process is stochastically dominated by a Poisson process of asymptotically equal rate. Hence, we may replace the synchronized process by its associated Poisson process as far as stability is concerned. Accordingly, a service node which is the offspring of a synchronization node is said to be stable if the net input process (consisting of exogenous arrivals, synchronized or otherwise, and routed tokens) has an average rate strictly smaller than the service rate of the node.

The following theorem, which is the main result of this section concerns the behavior of the network $\mathbf{N}$ under independent Poisson inputs. For ease of exposition, we state and prove the theorem assuming that the network $\mathbf{N}$ consists of *precisely two stages.* The acyclicity of the synchronization skeleton allows one to generalize the results to an arbitrary number of stages by induction.

Theorem 4 *Consider a 2-stage Generalized Jackson Network* $\mathbf{N}$ *consisting of* M *service nodes and* N *synchronization nodes subject to independent Poisson inputs at time 0, having been empty prior to that time. Let*

$$\vec{Q}(t) = (Q_1(t), Q_2(t), ..., Q_M(t))$$

be the joint queue-size process at the service nodes and let

$$\vec{\mathbf{D}}(t) = (\mathbf{D}_1(t), \mathbf{D}_2(t), ..., \mathbf{D}_L(t))$$

be the vector of departure processes from $\mathbf{N}$ *at time* $t > 0$. *If all the service nodes of the network* $\mathbf{N}$ *are stable, then as* $t \to \infty$, *(1)* $\vec{Q}(t)$ *converges to a product-form distribution with geometrically distributed marginals and (2)* $\vec{\mathbf{D}}(r + \cdot)$, *the vector of network departure processes viewed from time* r *onwards, converges in distribution to a vector of independent Poisson processes as* $r \to \infty$.

Proof Decompose the service nodes of $\mathbf{N}$ into the two stages $\mathbf{S}_0$ and $\mathbf{S}_1$ depending on whether they are the offspring of a synchronization node or not. Since $\mathbf{S}_0$ forms a classical Jackson Network that is assumed to be stable, both parts of the theorem follow easily for nodes in $\mathbf{S}_0$. In particular, the departure processes of $\mathbf{S}_0$ couple with independent Poisson processes after a finite random τ [18]. Now, some of these departure processes are input to $\mathbf{S}_1$, either directly or through synchronization nodes. One can now see that in order to complete the proof of the theorem we first need to establish two separate results: (1) The queue-size process at an exponential server queue with synchronized inputs converges to a geometric distribution and the corresponding departure process converges to a Poisson process, and (2) the joint queue-size process at service nodes in $\mathbf{S}_0$ and $\mathbf{S}_1$ is asymptotically product-form.

These two results, which are interesting in their own right, are proved in the next two subsections (Theorems 5 and 6). The proof of Theorem 4 is completed at the end of Section 3.3.

3.2 The SM/M/1 Queue: Synchronized Poisson Inputs with Exponential Service

Consider the canonical model shown in the left-hand-side of Figure 3. $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ are independent Poisson processes of rates λ_1 and λ_2 respectively and $\mathbf{S}(t)$ is their synchronization, as defined in Section 2. $\mathbf{S}(t)$ forms the input to an exponential server queue of service rate μ. $Q_1(t)$ and $Q_2(t)$ are the number of tokens in the two queues of the synchronization node, while $Q(t)$ is the number of tokens in the service queue. This the simplest possible network, consisting of a synchronization node and an exponential service node in tandem. We call it the *SM/M/1 Queue*. It can also be viewed as a resequencing queue (see [4]) with non-integrable delay sequence. We are interested in the asymptotic distribution of the queue-size at the service node and in the statistics of the departure process $\mathbf{D}(t)$. The following result easily extends to the M-input case; for the sake of simplicity and clarity of exposition we state and prove it for the 2-input case of Figure 3.

Theorem 5 *Consider the queue-size process, $Q(t)$, of an SM/M/1 queue with service rate μ. If $\rho = \lambda/\mu < 1$, where $\lambda = \min\{\lambda_1,\ \lambda_2\}$, then*

$$\lim_{t\to\infty} P(Q(t) = k) = (1-\rho)\,\rho^k,$$

and the departure process $\mathbf{D}(t)$ from the service node converges weakly (when $\lambda_1 = \lambda_2$, and strongly otherwise), to a Poisson process as $t\to\infty$.

Proof If $\lambda_1 > \lambda_2$, then $\mathbf{S}(t) = \mathbf{X}_2(t)$ a.s. after some finite random time. Theorem 5 follows trivially since the service queue becomes a standard M/M/1 queue, and $Q(t)$ couples in finite time almost surely with the queue-length process of a standard M/M/1 queue fed by the Poisson process $\mathbf{X}_2(t)$ alone. Thus, the output process from the service node converges strongly to a Poisson process.

The more interesting case of $\lambda = \lambda_1 = \lambda_2$ is easier studied by the introduction of the modification shown in the right-hand-side of Figure 3. The essential difference is that the net input process to the exponential server queue, $\mathbf{A}(t)$, is now a superposition of the output process $\mathbf{S}(t)$ of the synchronization node and a spurious process $\mathbf{Y}(t)\,1\!\!1_{\{Q_1(t^-)=0,Q_2(t^-)=0\}}$, where $\mathbf{Y}(t)$ is a rate λ Poisson process, independent of $\mathbf{X}_1(\cdot)$ and $\mathbf{X}_2(\cdot)$ and of the service process. The analysis of the modified system provides a solution to the original problem as follows.

We colour the tokens of $\mathbf{A}(t)$ blue or yellow depending on whether they originate from $\mathbf{S}(t)$ or $\mathbf{Y}(t)$. From Lemma 2, $\mathbf{A}(t)$ is a rate λ Poisson process, independent of the server. Thus, the service node is a standard M/M/1 queue, and its queue-length process $Q'(t)$, consisting of both blue and yellow tokens, converges in total variation to a geometric distribution with parameter $\rho = \lambda/\mu$ [16, 18]. Since we are interested in showing that $Q(t)$, the number of blue tokens, converges to a geometric distribution with parameter ρ, and $Q(t) = Q'(t) - Q_Y(t)$, it suffices to show that $\lim_{t\to\infty} P(Q_Y(t) \geq 1) = 0$.

Due to the memoryless property of the exponential server, we may stipulate that blue tokens take absolute priority over yellow tokens for service. Specifically, if an arriving blue token finds a yellow token in service, then it gets the remainder of the service owed to the yellow token (which is again exponential, independent of past service times). After the blue token departs the yellow token is served from the beginning, so long as there are no other blue tokens in the queue. Therefore, the yellow tokens are invisible to the blue tokens in the sense that they do not alter the dynamics of the blue tokens. We will occasionally refer to the yellow tokens as *ghost-tokens*. In Figure 3 yellow (ghost) tokens are represented by circles, and blue tokens are represented by discs.

We show that $\lim_{t\to\infty} P(Q_Y(t) \geq 1) = 0$ in three steps. Steps 1 and 2 ensure that yellow tokens do not remain in the system indefinitely, and Step 3 shows that the number of fresh yellow arrivals is dwindling. Together, all three imply the desired result. Thus, at large times the ghost-tokens get "exorcized" and the queue is left with the blue (real) tokens.

Step 1: *Given $\epsilon > 0$ we can find a large enough $M \in \mathbb{Z}^+$ and $t_1 \in \mathbb{R}$ such that $P(Q'(t) \leq M) > 1 - \epsilon$ for all $t > t_1$.*
Proof Since $\lim_{t\to\infty} P(Q'(t) \leq M) = \sum_{k=1}^{M} (1-\rho)\,\rho^k$, we may first choose M and then t_1 with the desired properties.

Step 2: *For ϵ and t_1 as in Step 1, we have that for every $t > t_1$ there exists a large enough $N \in \mathbb{Z}^+$, such that $P(Q'(s) = 0 \ for \ some \ s \in [t, t+N]) > 1 - 2\epsilon$.*
Proof Let the random variable T_M denote the time needed for the service queue to become empty, given that it starts at time t with M or fewer tokens in its buffer. By the stationarity of the arrival and service processes, the distribution of T_M is independent of the starting time t. Since the queue is stable, $T_M < \infty$ almost surely [16, 18]. Thus, there exists a large enough N such that $P(T_M < N) > 1 - \epsilon$. Moreover, from Step 1 we have $P(Q'(t) \leq M) > 1 - \epsilon$ for every $t > t_1$. Since $\{Q'(t) \leq M; T_M < N\} \subset \{Q'(s) = 0$ for some $s \in [t, t+N]\}$, it follows that $P(Q'(s) = 0 \ for \ some \ s \in [t, t+N]) > 1 - 2\epsilon$.

Step 3: *Given $\epsilon > 0$, there exists a t_2 such that $P(N_Y[t, t+N] \geq 1) < \epsilon$ for all $t > t_2$, where $N_Y[t, t+N]$ is the number of yellow (ghost) arrivals to the service queue in $[t, t+N]$.*
Proof By Chebyshev's inequality and the definition of stochastic intensity, we get

$$P(N_Y[t,t+N] \geq 1) \ \leq \ E(N_Y[t,t+N]) = E\left[\int_t^{t+N} d\Big(\mathbf{Y}(s)\, \mathbb{1}_{\{Q_1(s^-)=0=Q_2(s^-)\}}\Big)\right] =$$

$$= \lambda E\left[\int_t^{t+N} \mathbb{1}_{\{Q_1(s^-)=0=Q_2(s^-)\}}\ ds\right] = \lambda\left[\int_t^{t+N} P\left(Q_1(s^-) = 0 = Q_2(s^-)\right)\ ds\right].$$

Using the dominated convergence theorem and the null-recurrence of $(Q_1(\cdot), Q_2(\cdot))$, we see that the rightmost expression goes to zero as $t\to\infty$, implying the desired result.

The fact that $\lim_{t\to\infty} P(Q_Y(t) \geq 1) = 0$ now follows easily from the previous three steps. Indeed, note that $\{N_Y[t,t+N] = 0\} \cap \{Q'(s) = 0$ for some $s \in [t, t+N]\} \subseteq \{Q_Y(t+N) = 0\}$. Therefore, Steps 1, 2 and 3 imply that, for any time $t > \max(t_1, t_2)$, we have that $P(Q_Y(t+N) \geq 1) < 3\epsilon$. Thus, as $t\to\infty$, the chance that there are yellow (ghost) tokens in the service queue becomes arbitrarily very small, fading away to zwro. This concludes the proof of part 1 of the theorem.

To show part 2, let $\mathbf{D}'(t) = \mathbf{D}(t) + \mathbf{D}_Y(t)$ be the overall departure process of the service node, consisting of both blue and yellow (ghost) tokens. In order to prove that $\mathbf{D}(t)$ converges weakly to a Poisson process, it is enough to show that for any bounded, continuous function f with compact support

$$\int f(s)d(\mathbf{D}_Y(t+s)) \to 0 \tag{13}$$

in distribution as $t \to \infty$ ([8], Section 4.4), since we already know that $\mathbf{D}'(t)$ converges in total variation (and hence weakly) to a Poisson process. But this follows from the fact that $E[|\int f(s)d\,(\mathbf{D}_Y(t+s)|] \leq |f|_{max} E[\int d\,(\mathbf{D}_Y(t+s)] = |f|_{max}\, \mu E[\int \mathbb{1}_{\{Q_Y((t+s)^-)>0; Q((t+s)^-)=0\}} ds]$, the equality following from the fact that the stochastic intensity of $\mathbf{D}_Y(t)$ with respect to the history of the service queue (including that of blue and yellow tokens) is $\mu \mathbb{1}_{\{Q_Y(t^-)>0; Q(t^-)=0\}}$ (see [7]). Using the dominated convergence theorem and the fact that $P(Q_Y(t) > 0)$ goes to zero (from the first part of this proof), we see that the last term in the previous expression goes to zero, implying (13). This completes the proof of the theorem. ∎

It is clear that the same proof also works when three or more independent Poisson arrival processes are synchronized before arriving at an exponential server queue.

3.3 Interfacing Service Stages Through Synchronization Queues

The next step is to interface stages of service nodes through synchronization queues. The study is centered on a representative simple network shown in Figure 4, but teh arguments extend naturally to teh general case. In Figure 4, $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$ are independent Poisson processes of rates λ_1 and λ_2 arriving at independent exponential service queues 1 and 2 with rates μ_1 and μ_2 respectively, such that both M/M/1 queues are stable ($\lambda_1 < \mu_1$, $\lambda_2 < \mu_2$). The input $\mathbf{X}_3(t)$ to service queue 3 is the synchronization of the departures from nodes 1 and 2.

Let $Q_i(t)$ be the queue-size process at service node $i \in \{1, 2, 3\}$. Suppose all queues are initially empty and μ_3 is larger than $\min\{\lambda_1, \lambda_2\}$ (so that service queue 3 is also stable). Let $Q_A(t)$ and $Q_B(t)$ be the queue-lengths of the synchronization buffers holding departures from nodes 1 and 2 respectively. We want to show that the joint process $(Q_1(t), Q_2(t), Q_3(t))$ tends to a limit that is product-form with appropriate marginals as $t \to \infty$, and that the output process $\mathbf{D}_3(t)$ converges to a Poisson process. The convergence of each of the marginals of $(Q_1(t), Q_2(t), Q_3(t))$ to the appropriate geometric distribution is quite obvious, given the previous discussion of the SM/M/1 queue; what we need to prove is that the *joint* distribution of $Q_1(t), Q_2(t), Q_3(t)$ is asymptotically product-form.

In classical Jackson Networks the main reason leading to a product-form stationary distribution is *quasi-reversibility* ([15, 18]), which is an equilibrium property. In networks with service and synchronization nodes no global equilibrium can be reached, due to the null-recurrence or transience of the queue-length processes of the synchronization queues. However, the basic idea of introducing "ghost tokens" (as in the analysis of the SM/M/1 queue) to make up for a deficit of real tokens proves useful again in showing the following result.

Theorem 6 *For the network in Figure 4, we have that*

$$\lim_{t \to \infty} P(Q_1(t) = i, Q_2(t) = j, Q_3(t) = k) = (1-\rho_1)(1-\rho_2)(1-\rho_3)\rho_1^i \rho_2^j \rho_3^k \tag{14}$$

for $i, j, k \in \mathbb{Z}_+$, where $\rho_1 = \frac{\lambda_1}{\mu_1}$, $\rho_2 = \frac{\lambda_2}{\mu_2}$, $\rho_3 = \frac{\min\{\lambda_1,\lambda_2\}}{\mu_3}$, and the departure process from each node converges weakly (and in special cases strongly) to a Poisson process as $t \to \infty$.

Proof If $\lambda_1 > \lambda_2$, then $\mathbf{X}_3(t) = \mathbf{D}_2(t)$ after some finite time almost surely. Hence, $\mathbf{X}_3(t)$ couples in finite time with a rate λ_2 Poisson process. Using standard quasi-reversibility arguments [15, 18], it is easy to see that the distribution of $(Q_2(t),\ Q_3(t))$ tends to a product-form limit as $t \to \infty$, while the service node outputs converge in total variation to Poisson processes. Since arrival times at node 3 coincide with departure times from node 2 after a finite random time, node 1 is decoupled from the system and the result follows.

Next suppose that $\lambda_1 = \lambda_2 = \lambda$. It is clear that the limiting distribution of $(Q_1(t),\ Q_2(t))$ is product-form with appropriate marginals. And, from Theorem 5, it follows that $Q_3(t)$ tends to a geometric distribution. However, what is not immediate is that $(Q_1(t),\ Q_2(t),\ Q_3(t))$ *jointly* converges to a product-form distribution. We argue this as follows.

Let us modify the actual system by introducing the process $\mathbf{Y}_1(t) = \mathbf{Y}(t)\mathbb{1}_{\{Q_A(t^-)=0=Q_B(t^-)\}}$ of ghost token arrivals, as in Figure 4. $\mathbf{Y}(t)$ is a Poisson process of rate λ and is independent of all other arrival and service processes. The following steps applied to the modified system implies the result.

<u>*Step 1*</u>: The only arrivals to the network are the tokens of the exogenous arrival processes $\mathbf{X}_1(t)$ and $\mathbf{X}_2(t)$, and the ghost arrival process $\mathbf{Y}(t)$; all these processes are mutually independent. Suppose that these arrivals have been coming into the network since time $-\infty$, but there was no synchronization operation being performed before time 0. That is, the departures on $\mathbf{D}_1(t)$ and $\mathbf{D}_2(t)$ were simply allowed to exit the network, without being synchronized and going through node 3. In this case, it is immediate that at time 0 all three service queues are in stationarity and their joint distribution is product form. Moreover, at time 0 all three departure processes are Poisson (rate λ) and $Q_3(0)$ consists only of yellow (ghost) tokens.

<u>*Step 2*</u>: Suppose that the synchronization operation begins at time 0; that is, departures from service nodes 1 and 2 are synchronized and driven through node 3. We can then write

$$\mathbf{X}_3(t) = \mathbf{Y}(t)\mathbb{1}_{\{t\le 0\}} + \Big\{\mathbf{Y}(t)\mathbb{1}_{\{Q_A(t^-)=0=Q_B(t^-)\}} + \mathbf{D}_1(t)\,\mathbb{1}_{\{Q_B(t^-)>0\}} + \mathbf{D}_2(t)\,\mathbb{1}_{\{Q_A(t^-)>0\}}\Big\}\,\mathbb{1}_{\{t>0\}}.$$

Since the processes $\mathbf{D}_1(t)$ and $\mathbf{D}_2(t)$ are Poisson (and independent) for $t > 0$, an application of Lemma 2 shows that $\mathbf{X}_3(t)$ is Poisson of rate λ. Nodes 1 and 2 are in equilibrium at any positive time. Because of the quasi-reversibility of exponential server nodes [15, 18], it follows that $(Q_1(t),\ Q_2(t))$ is independent of $\{\mathbf{D}_1(s),\ \mathbf{D}_2(s)\}_{s<t}$ (in equilibrium, the present state is independent of past departures). This implies $(Q_1(t),\ Q_2(t))$ is independent of $\{\mathbf{X}_3(s)\}_{s<t}$ (by definition of $\mathbf{X}_3(t)$). Hence it is independent of $Q_3(t)$. We therefore conclude that under the previous scenario of equilibrium at nodes 1 and 2, the quantities $Q_1(t),\ Q_2(t),\ Q_3(t)$ are mutually independent for any fixed $t > 0$, leading to a product-form distribution.

<u>*Step 3*</u>: Finally, we need to show that $Q_3(t)$ will eventually consist of only non-yellow (real) tokens with arbitrarily high probability, and the yellow (ghost) tokens will dwindle, restoring the operation of the actual system. This can directly be proved by arguing as in Theorem 5.

The previous arguments show that the service node queue-lengths $(Q_1(t),\ Q_2(t),\ Q_3(t))$ have a product-form stationary distribution. Moreover, if $\lambda_1 = \lambda_2$ the departure flows from all nodes converge weakly to Poisson processes as $t \to \infty$. On the other hand, if $\lambda_1 \neq \lambda_2$ the flows converge strongly (in total variation) to Poisson processes and hence also weakly. This concludes the proof of the theorem. ■

Completion of the proof of Theorem 4: Given the preceding results, the two parts of Theorem 4 follow immediately. First, use ghost tokens as in the proof of Theorem 6 to argue that $\vec{Q}(t)$ converges to a product-form distribution with appropriate marginals. Then, observe that $\vec{\mathbf{D}}(t)$ consists of processes that are either departures from stages $\mathbf{S}_0$ or $\mathbf{S}_1$, or from the synchronization nodes. Departures from $\mathbf{S}_0$ are clearly mutually independent Poisson processes. The use of ghost tokens to make the synchronized departures be Poisson processes implies that departures from $\mathbf{S}_1$ converge to Poisson processes that are mutually independent, and independent of departures from $\mathbf{S}_0$. Since all routing is assumed to be Bernoulli, if the departures from a synchronization node are split into two or more processes, then these processes will converge to mutually independent Poisson processes. (To make this point clear in the context of Figure 7, the two flows that arise by a splitting of the synchronized departure process from the 2-input synchronization node will asymptotically be independent Poisson processes. Thus, all flows that arrive at the synchronization node feeding network $\mathbf{N}_6$ are asymptotically mutually independent Poisson processes.) In conclusion, $\vec{\mathbf{D}}(t)$ converges to a vector of mutually independent Poisson processes as $t \to \infty$. ■

4 Final Remarks and Conclusion

The previous results on networks assumed that the synchronization skeleton was acyclic. For networks in which the synchronization skeleton is not acyclic, the situation is not clear. Indeed, even for the simplest of cases, there is a problem of "ill-posedness" which is related to the problem of deadlocks in Petri nets [9]. Consider the network shown in Figure 5. The output from the service queue is $\mathbf{E}(t)$, and a fraction ($p \in (0,1)$) of it is Bernoulli-routed back to a buffer of the synchronization queue as $\mathbf{F}(t)$, while the remainder forms the output flow $\mathbf{D}(t)$. When trying to synchronize $\mathbf{F}(t)$ and $\mathbf{X}(t)$ we run into the following problem. If the server is currently idle ($Q(t) = 0$), then it is *impossible* to have any new arrivals into the service node, as there will not be any tokens getting routed back to the synchronization node. On the other hand, for the example in Figure 6, no such problems occur and one may attempt to compute the limiting queue-length and departure process distributions. We are exploring this line of research since the arguments presented in this paper do not naturally extend to the case where the synchronization skeleton is non-acyclic.

Towards the end of Section 2 we saw that with large enough finite synchronization buffers, one can still have the synchronization process be as close to a Poisson process in distribution as desired. This key point may be applied to the networks of Section 3 to obtain a result of the following sort: In a Generalized Jackson Network with an acyclic synchronization skeleton and finite synchronization buffers, the joint distribution of the equilibrium queue-length process (note that equilibrium now exists) at all the exponential server nodes is asymptotically product-form and the equilibrium departure processes converge weakly to independent Poisson processes, as the size of the synchronization buffers goes to infinity. We assume that tokens that arrive at full buffers are blocked and discarded.

In conclusion, the dynamics of the synchronization node have been analyzed, both in isolation and in networks of service and synchronization nodes with an acyclic synchronization skeleton. The main vehicle of analysis was the introduction of "ghost tokens" at synchronization nodes. It has been shown that synchronization preserves the Poissonian nature of flows. Furthermore, the results obtained on networks generalize those about product-form queue-size distributions and the Poisson nature of departures in classical open Jackson Networks. Networks with a non-acyclic synchronization skeleton have not been successfully investigated.

References

[1] V. Anantharam, Modelling the Flow of Coalescing Data Streams Through a Processor, J. Applied Probability, v. 25, pp. 184-193, 1988.

[2] S. Asmussen, Applied Probability and Queues, Wiley, New York, 1987.

[3] F. Baccelli, G. Cohen, G. J. Oldser, J-P Quadrat, Synchronization and Linearity - An Algebra for Discrete Event Systems, Willey Series on Probability and Mathematical Statistics, New York, 1992.

[4] F. Baccelli and A. M. Makowski, Queueing Models for Systems with Synchronization Constraints, Proceedings of the IEEE, v. 77, no 1, pp. 138-161, 1989.

[5] F. Baccelli, W. A. Massey and D. Towsley, Acyclic Fork-Join Queueing Networks, Journal of the Assocation of Computing Machinery, v. 36, no 3, pp. 615-642, 1989.

[6] F. Baskett, K. Chandy, R. R. Muntz, F. G. Palacios, Open, closed and mixed networks of queues with different classes of customers, Journal of the Assocation of Computing Machinery, v. 22, pp. 248-260, 1975.

[7] P. Brémaud, Point Processes and Queues: Martingale Dynamics, Springer Series in Statistics, Springer Verlag, New York, 1981.

[8] K. L. Chung, A Course in Probability Theory, Second Edition. Academic Press, 1974.

[9] F. Commoner, Deadlocks in Petri Nets, Applied Data Research, CA-7606-2311, 1972.

[10] D. J. Daley and D. Vere-Jones, An Intorduction to the Theory of Point Processes, Springer Series in Statistics, Springer Verlag, New York, 1988.

[11] G. Florin and S. Natkin, On Open Synchronized Queueuing Networks, International Workshop on Timed Petri Nets, Torino, Italy, July 1985.

[12] J. M. Harrison, Assembly-like Queues, J. Applied Probability, v. 10, pp. 354-367, 1973.

[13] J. R. Jackson, Networks of waiting lines, Operations Research, v. 15, pp. 254-265, 1957.

[14] J. Jacod and A. N. Shiryaev, Limit Theorems for Stochastic Processes, A Series of Comprehensive Studies in Mathematics, Springer Verlag, Heidelberg, 1987.

[15] F. Kelly, Reversibility and Stochastic Networks. Wiley, London, 1979.

[16] R. M. Loynes, The stability of a queue with non-independent inter-arrival and service times, Proceedings of the Cambridge Philosophical Society, v. 58, pp. 497-520, 1962.

[17] K. Matthes, J. Kerstan and J. Mecke, Infinitely Divisible Point Processes, Akademie Verlag, Berlin, 1978.

[18] J. Walrand, An Introduction to Queueing Networks, Prentice Hall, Englewood Cliffs, 1988.

[19] Petri Nets and Performance Models, Proceedings of the 4th International Workshop, Melbourne, Australia, December 2-5, 1991, IEEE Computer Society Press, Los Alamitos, California.

[20] Petri Nets and Performance Models, Proceedings of the 3th International Workshop, Kyoto, Japan, December 11-13, 1989, IEEE Computer Society Press, Los Alamitos, California.

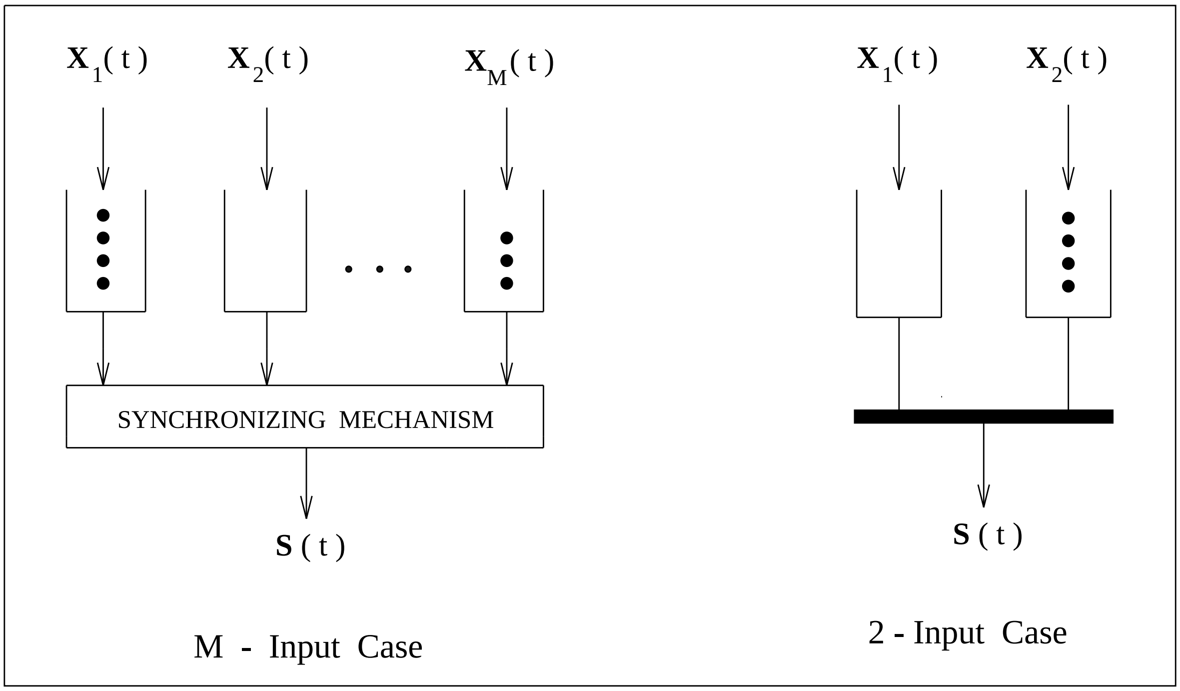

Figure 1: Synchronization of point processes. At least one queue is always empty.

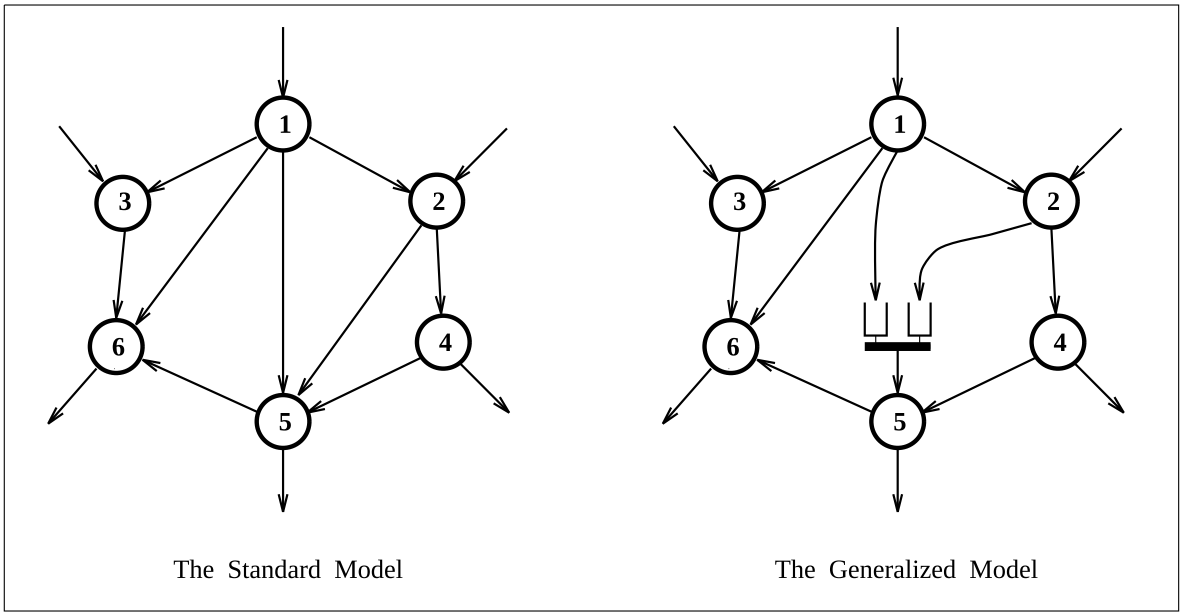

Figure 2: Inclusion of synchronization nodes in networks of service nodes. Circles represent service nodes, bars synchronization ones, and arrows arrivals and routing of tokens.

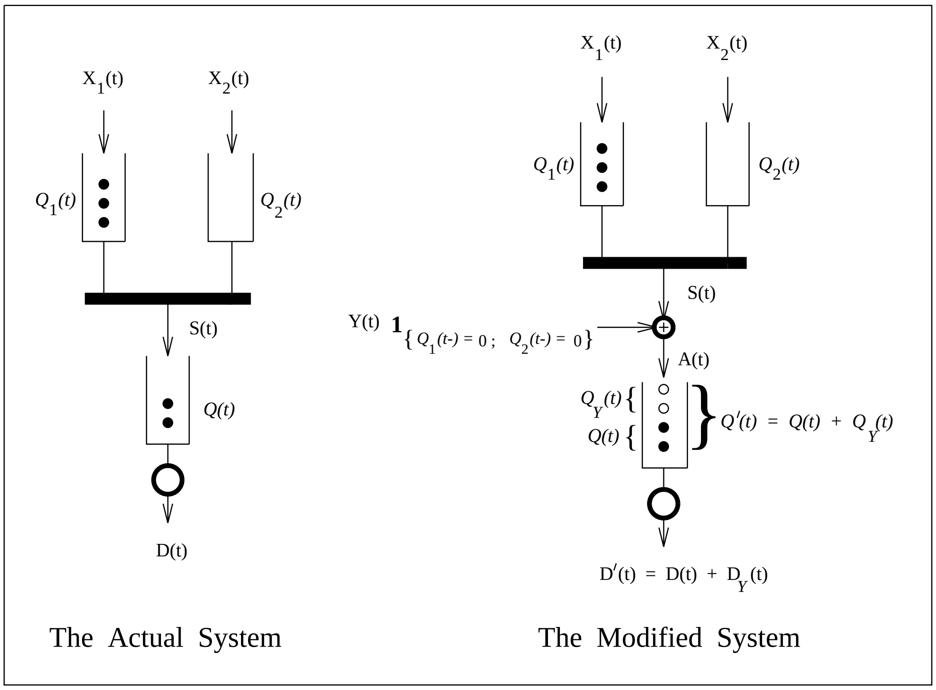

Figure 3: The SM/M/1 queue (on the left) with two synchronized independent Poisson inputs and an exponential server. On the right, the modified system used to analyze the SM/M/1 queue is shown, in which "ghost tokens" (represented by empty circles) are injected to the service queue.

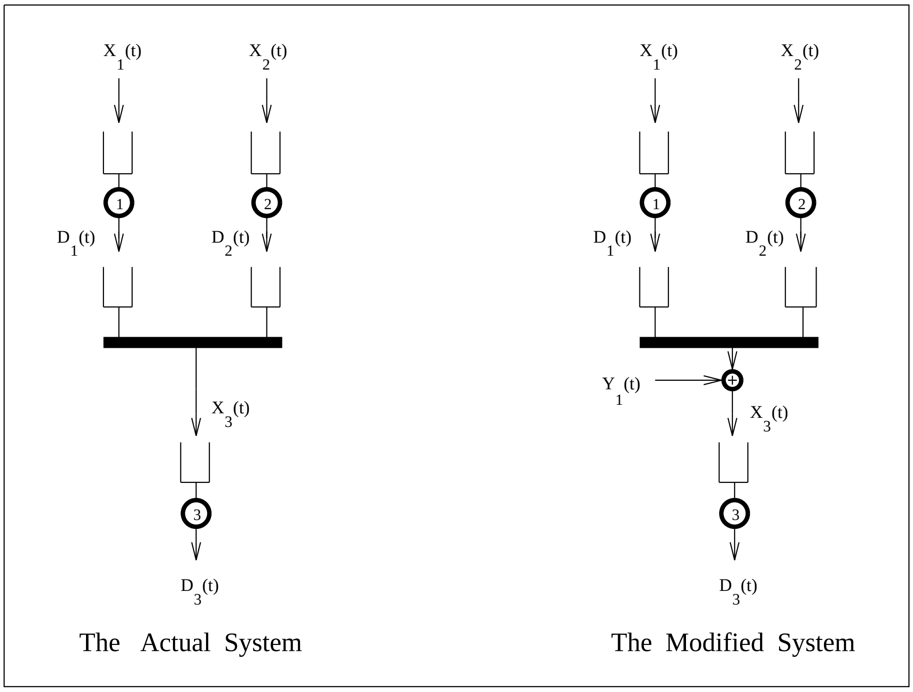

Figure 4: Interfacing service queues through synchronization ones. In the modified system the controlled flow $\mathbf{Y}_1(t)$ injects "ghost tokens" into the flow $\mathbf{X}_3$, making the latter Poisson by accounting for the deficit of actual tokens.

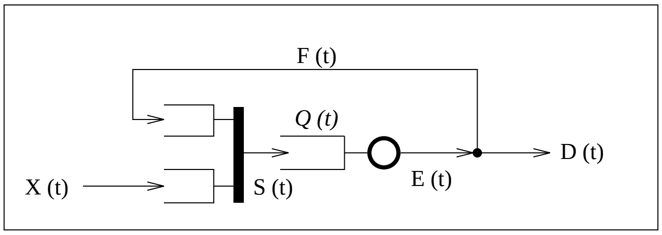

Figure 5: A simple network with feedback of tokens to the synchronization node. The node will eventually stall.

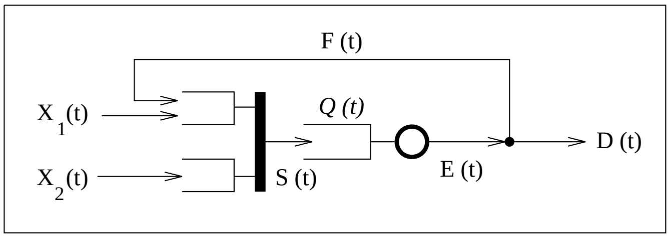

Figure 6: Feeding all synchronization buffers by non-feedback flows prevents stalling of the synchronization node.

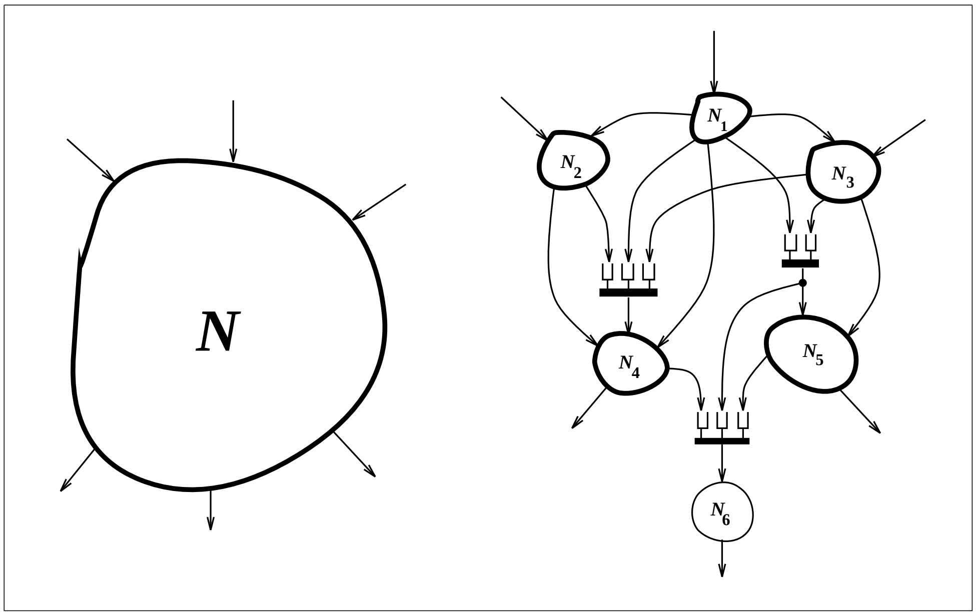

Figure 7: A network with an acyclic synchronization skeleton. Local networks of only service nodes may be non-acyclic.