

STATIONARITY AND STABILITY OF FORK-JOIN NETWORKS

PANAGIOTIS KONSTANTOPOULOS AND
 JEAN WALRAND,* *University of California, Berkeley*

Abstract

We consider a *fork-join* network with arrival and service times forming a stationary and ergodic process. The usual stability condition, namely that the input rate is strictly less than all the service rates, is proved to be valid in this general case. Finally we extend the result to the case where there is random routing.

STATIONARY POINT PROCESSES; PALM TRANSFORMATION; QUEUEING NETWORKS

1. Introduction

Fork-join networks are queueing networks without feedback. When a job arrives at the input, it splits into several parts which are serviced in different subnetworks and then joined together again to form the output. Each of the parts may be served in one or more queues in tandem or split again into more subparts, and so on. The network consists therefore of three elements:

1. *Queues*: They are assumed to be FCFS with one server and infinite waiting space.
2. *Forks*: Upon arrival at a fork, a job is split into as many parts as there are outgoing links and this is an instantaneous process.
3. *Joins*: The operation of a join is to merge subparts of the same part. Subparts wait at the join until the last one arrives and then they merge and leave instantaneously. Thus the departure time from a join is the maximum of the arrival times of the parts.

From the description of the network we see that the number of joins is equal to the number of forks and in fact each join has a corresponding fork. It is assumed that the network has a unique arrival and a unique departure stream. For another description of the fork-join networks see Baccelli and Massey (1986).

For $n \in \mathbb{Z}$, let T_n denote the arrival time of the n th job into the network ($T_0 \leq 0 < T_1$) and $\sigma_n = (\sigma'_n, j \in J)$ the vector of service times, where J is the set of queues of the network. We assume that $\{T_n, \sigma_n : n \in \mathbb{Z}\}$ are random variables defined on a common probability space (Ω, \mathcal{F}, P) . The random variable T_n is considered as the n th point and σ_n as the n th mark of a stationary and ergodic marked point process. Stationarity and ergodicity are with respect to some semigroup $\{\theta_t, t \in \mathbb{R}\}$ of the space Ω (see

Received 29 April 1987; revision received 4 October 1988.

* Postal address for both authors: Department of Electrical Engineering and Computer Sciences and Electronics Research Laboratory, University of California, Berkeley CA 94720, USA.

Appendix 1). If we take as Ω the space of realizations of the marked point process then θ_t is exactly the left-shift-by- t operator. (For the relevant theory see Franken et al. (1982), Neveu (1977), Baccelli and Brémaud (1987), Walrand (1988) and the brief discussion in Appendix 1.)

Let P^0 denote the Palm transformation of P with respect to the arrival point process $\{T_n\}$ and E^0 the expectation with respect to P^0 . Let $\lambda = [E^0(T_{n+1} - T_n)]^{-1}$ be the rate of the input and $\mu_j = (E^0 \sigma_n^j)^{-1}$ the service rate at queue j . We assume $0 < \lambda, \mu_j < \infty$ for all $j \in J$. As a description of the network we take the vector $W_t = (W_t^j, j \in J)$ of workloads at time t .

In Section 2 we show that, under the assumption $\lambda < \min_{j \in J} \mu_j$, there is a unique finite stationary workload process $\{\tilde{W}_t^j, t \in \mathbb{R}\}$ and that W_t converges in law to the distribution of the latter. In Section 3 we consider the possibility of random routing and show that the same results hold under a more relaxed condition (which is also almost necessary) for a more general class of networks. Finally we give some important special cases in Section 4.

The problem that we consider in this paper is not new. The first paper on stability for $G/G/1$ queues with ergodic inputs was that of Loynes (1962). See also the monograph of Baccelli and Brémaud (1987) and also Baccelli and Massey (1986) and Baccelli et al. (1987a) for special cases. The existence part of the problem of the fork-join network has been given in Baccelli and Massey (1986) and relies on the fact that there is a global description (namely the collection of the delays experienced by a job from the time of its arrival until the beginning of its service in queue j , for all $j \in J$). We show that this is not necessary and actually not feasible in certain cases as in the case of random routing. The main result that we use is that of the stability of a single queue (see Appendix 2).

2. Stability of fork-join networks

Since our proof is going to be inductive we define a sequence of stages (i.e., sets of links) as follows. As a first stage we take the set of links which are immediately accessible from the external input. Suppose now that the i th stage has been defined. Then a link of the $(i + 1)$ th stage is either a link of the previous stage or is connected to a link of it via a queue, a fork or a join (see Figure 1). Assuming that all the stages are distinct and that the network has only a finite number of elements it is clear that the network can be exhausted with a finite number of stages.

With each link l of the network we associate a point process N^l of arrivals of (parts of) jobs at this link. The objective of this section is to show that each of these processes will coincide in finite time with a stationary (under measure P) point process \tilde{N}^l with rate λ . That is, there is a finite (random) time such that, after this time the points of N^l coincide with the points of \tilde{N}^l . So the input to each queue j will eventually be stationary. This stationary input is used consequently for the construction of the stationary regime \tilde{W}^j for the workload of this queue.

Theorem 1. If $\lambda < \min_{j \in J} \mu_j$ the point process N^l associated with a link l coincides in finite time with a P -stationary and ergodic point process \tilde{N}^l with rate λ .

Proof. Suppose that the initial workload W_0 is some finite (random) non-negative vector. The statement is trivially true for the links of the first stage since their point processes are replicas of the stationary arrival stream $\{T_n\}$.

Suppose now, for the sake of induction, that each point process N^l (with points, say, T_n^l , where $T_0^l \leq 0 \leq T_1^l$), for all links l prior to stage i , coincides, in finite time, with a P -stationary process \tilde{N}^l (with points, say, \tilde{T}_n^l , where $\tilde{T}_0^l \leq 0 \leq \tilde{T}_1^l$) with rate λ . This process is defined on the original probability space (Ω, F, P) and is stationary with respect to the semigroup $\{\theta_t\}$. We shall prove that the same thing happens for the next stage. Consider the three different possibilities (see Figure 1).

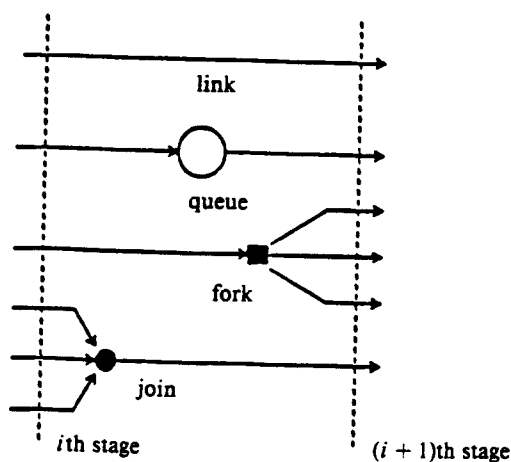


Figure 1. The definition of stages

Case 1. A certain link of stage i is joined to a link of stage $i+1$ only through a single queue j . We assumed that the point process N^j (with points T_n^j) entering queue j coincides, in finite time, with a stationary point process \tilde{N}^j (with points \tilde{T}_n^j) with rate λ . Consider the Palm transformation P^j of P with respect to the point process \tilde{N}^j . We can easily see that the marked point process with points \tilde{T}_n^j and marks σ_n^j is P^j -synchronous. (This can be seen either directly or by applying the result of Appendix 3.) Since λ (the rate of \tilde{N}^j) is strictly smaller than μ_j (which is, due to the previous remark, equal to $(E^j \sigma_n^j)^{-1}$) we can apply the results of Appendix 2 and use \tilde{N}^j to construct a finite stationary regime \tilde{W}^j for the workload of queue j . Clearly $\{\tilde{W}_t^j = \tilde{W}^j \circ \theta_t, t \in \mathbb{R}\}$ is P -stationary and ergodic. If W_n^j denotes the workload of queue j at time T_n^j we can easily see by the induction hypothesis and the results of Appendix 2 that W_n^j coincides, after a finite time, with \tilde{W}_n^j . (The latter is, by definition, the value of \tilde{W}_t^j at the time $t = \tilde{T}_n^j$.) The departure time of the n th arrival from this queue is $S_n^j = T_n^j + W_n^j + \sigma_n^j$ and, because of the previous statement, it coincides after a finite time with $\tilde{S}_n^j = \tilde{T}_n^j + \tilde{W}_n^j + \sigma_n^j$. Noting that the marked point process with points \tilde{S}_n^j and marks $(\tilde{W}_n^j, \sigma_n^j)$ is P^j -synchronous (and P -stationary) with rate λ , we can apply Campbell's formula to pass from P to P^j (see Appendix 1) and deduce that

$$\begin{aligned}
 \text{rate of } \{\tilde{S}_n^j\} &= E \sum_{n \in \mathbb{Z}} I(\tilde{S}_n^j \in [0, 1]) = E \sum_{n \in \mathbb{Z}} I(\tilde{T}_n^j + \tilde{W}_n^j + \sigma_n^j \in [0, 1]) \\
 &= \lambda E^j \int_{\mathbb{R}} I(t + \tilde{W}^j + \sigma^j \in [0, 1]) dt = \lambda.
 \end{aligned}$$

Case 2. A link l of stage i is joined to several links of stage $i + 1$ through a fork. Here the hypothesis is trivially true since the outgoing point processes are replicas of the ingoing one.

Case 3. Several links of stage i are joined to a link of stage $i + 1$ through a join a . Clearly, there is a corresponding fork s prior to a , as in Figure 2. Let $N^s = \{T_n^s\}$, $N^a = \{T_n^a\}$ be the point processes corresponding to the link prior to s and after a , respectively. By the induction hypothesis, all processes inside the subnetwork between s and a couple with P -stationary processes in finite time. In particular, $N^s = \{T_n^s\}$ couples with $\tilde{N}^s = \{\tilde{T}_n^s\}$ having rate λ . Let P^s denote the Palm transformation with respect to \tilde{N}^s . Since the delay d_n of the n th job from the time of its arrival at s until its exit time from a is a certain deterministic function of the workloads and the service times of the queues inside the subnetwork, which couple with P^s -stationary ones, it follows that $\{d_n\}$ couples with a P^s -stationary sequence $\{\tilde{d}_n\}$. Hence $T_n^a = T_n^s + d_n$, couples with $\tilde{T}_n^a = \tilde{T}_n^s + \tilde{d}_n$ in finite time. The point process $\tilde{N}^a = \{\tilde{T}_n^a\}$ is clearly P -stationary. We can also verify that it has the correct rate λ by writing a calculation similar to that of Case 1 above using Campbell's formula.

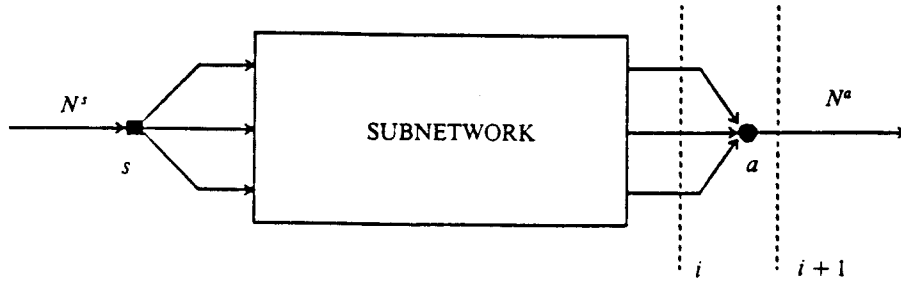


Figure 2. A join

We have therefore shown, by examining all possible cases, that the point processes associated with stage $i + 1$ couple with stationary ones with rate λ . Therefore the theorem is true.

Conclusions

1. The collection of all stationary regimes $\tilde{W} = (\tilde{W}^j, j \in J)$, which have been constructed on the same probability space (Ω, \mathcal{F}, P) gives rise to a stationary workload process $\tilde{W}_t = \tilde{W} \circ \theta_t$, which is finite since each component is by construction finite (see Appendix 2).

2. The previous theorem shows that any finite workload process $\{W_t\}$ coincides in finite time with $\{\tilde{W}_t\}$. The usual coupling argument shows that

$$P(W_t \in A) \rightarrow P(\tilde{W} \in A) \quad \text{as } t \rightarrow \infty, \quad \text{for all Borel sets } A \subset [0, \infty)$$

or, from a job point of view,

$$P^0(W_n \in A) \rightarrow P^0(\tilde{W} \in A) \quad \text{as } n \rightarrow \infty, \quad \text{for all Borel sets } A \subset [0, \infty).$$

In fact the convergence is uniform in A (i.e., we have convergence in total variation).

The following theorem shows that the process $\{\tilde{W}_t\}$ is unique.

Theorem 2. The process $\{\tilde{W}_t\}$ is the unique stationary finite workload process for the fork-join network.

Proof. Let $\{\tilde{V}_t\}$ be another finite stationary workload process on $(\Omega, F, P, \{\theta_t, t \in \mathbb{R}\})$. Then, passing on to Palm, $\tilde{W}_n - \tilde{V}_n \rightarrow 0$ as $n \rightarrow \infty$, P^0 -a.s. (since the two processes coincide in finite number of steps) where $\tilde{W}_n = \tilde{W}_{T_n}$ and $\tilde{V}_n = \tilde{V}_{T_n}$. Notice that $P^0(|\tilde{W}_n - \tilde{V}_n| > \varepsilon) = P^0(|\tilde{W}_0 - \tilde{V}_0| > \varepsilon)$ for all n , by stationarity. Therefore the latter probability is equal to zero for all $\varepsilon > 0$. This shows that $P^0(\tilde{W}_0 = \tilde{V}_0) = 1$ and hence $P^0(\tilde{W}_n = \tilde{V}_n, \text{ for all } n) = 1$. Consequently, $P(\tilde{W}_t = \tilde{V}_t, \text{ for all } t) = 1$ and the theorem is proved.

Note. The condition $\lambda < \min_{j \in J} \mu_j$ is easily seen to be almost necessary for stability in the sense that if there is a $j \in J$ such that $\lambda > \mu_j$ then W_n converges almost surely to infinity. Indeed, let j be the smallest index (in the sense of the partial ordering of the network) such that the above happens. Then, as the proof of Theorem 1 shows, the input to this queue j will eventually coincide with a stationary process with rate λ which, however, exceeds the service rate μ_j . From the instability part of the results for a single queue (see Appendix 1) we see that W_n^j converges almost surely to infinity.

3. Pseudo-fork-join networks with random routing

The networks considered in the previous section have the property that a cut made at a join a and its corresponding fork s (as in Figure 2) results in two disjoint pieces. That is, the subnetwork between s and a does not interact with the rest of the network except via s and a . We find however that we can drop this property and consider a more general network whose graph is any acyclic graph. This type of network we call a *pseudo-fork-join* network. For instance, consider the network of Figure 3. Here we can no longer separate the subnetwork between a_1 and s_1 or between a_2 and s_2 .

By the *graph* of such a network we mean a graph whose vertices are the set of queues together with one source and one destination point. Further, there is an edge between

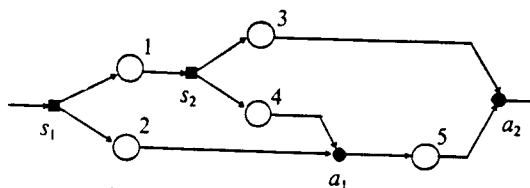


Figure 3. A pseudo-fork-join network

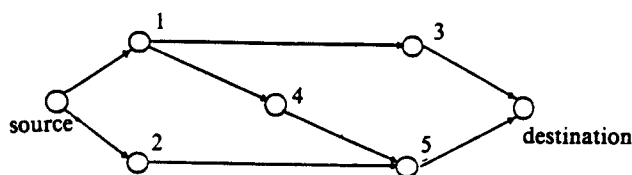


Figure 4. Graph of the network of Figure 3

two vertices if there is a direct connection between them or a connection through a fork or a join. By *source point* we mean a point without any ingoing edges and by *destination point* we mean a point without any outgoing edges. For example, the graph of the network of Figure 3 is shown in Figure 4. It is also clear how to construct the network, given an acyclic graph as in Figure 4: the source point represents the input which, in case it has more than one outgoing edge, is replaced by a fork. The destination point represents the output and is replaced by a join, if it has more than one ingoing edge. Similarly, if a vertex j has more than one outgoing edge then we replace j by a queue followed by a fork, and if a vertex j has more than one ingoing edge then we replace j by a queue preceded by a join.

A further generalization, regarding the routing of incoming jobs, can also be made. Consider a situation where, upon arrival to a fork, a job is not necessarily split into as many parts as there are outgoing links, but instead, it chooses a random subset of them to pass through. For instance, it can toss a die and choose some links independently of everything else. To make the situation as general as possible we assume that the n th external arrival carries, as a mark, besides its service requirements σ_n , the subnetwork D_n which it will follow. That is, D_n is a random variable taking values in the collection of all subnetworks of the given network. By subnetwork we mean a pseudo-fork-join network whose graph is a subgraph of the given one and has the same input and output.

Our assumption this time is that the marked point process (m.p.p.) $\Phi = \{(T_n; \sigma_n, D_n), n \in \mathbb{Z}\}$ is stationary on $(\Omega, F, P, \{\theta_t, t \in \mathbb{R}\})$. Intuitively, since the network is not fully utilized as before, we expect that $\lambda < \mu_j$, for all $j \in J$, is still a sufficient condition for stability. However, we can relax this condition and we claim that an exact condition (see Note below) for stability is:

$$(1) \quad E^0 \sigma^j I(j \in D) < E^0 \tau \quad \text{for all } j \in J$$

where $I(\cdot)$ denotes indicator function, $\sigma^j = \sigma_0^j$, $D = D_0$, $\tau = \tau_0$. (One more word about notation: by saying $j \in D$ we mean that queue j belongs to the subnetwork D ; same convention holds for other elements of the network.)

Because our model has random routing, we associate with a certain link l a marked point process Φ^l . The underlying point process associated with Φ^l is, as before, the point process of arrival times of (parts of) jobs at link l and is still denoted by N^l . The mark corresponding to a typical arrival time consists of the subnetwork, to be followed, and the service times that are required. That is, the mark contains all the necessary information for the future path.

The main difference between a pseudo-fork-join network with random routing and an ordinary fork-join network is that the order of arrivals may not be preserved throughout the network. For instance, in Figure 3, if the n th arriving job chooses to pass through queues 1 and 4 spending time d_n until it gets to join a_1 while the $(n+1)$ th arriving job chooses to pass through queue 2 spending time $d_{n+1} < d_n$, then, upon arrival at queue 5, the order of these two jobs is switched. However, if we look at the arrivals in front of queue 5 as a marked point process and show that it eventually coincides with a stationary marked point process with mean service time less than mean interarrival time, then this is all we need for queue 5 to have a stable regime.

We prove the claim by proving a theorem analogous to Theorem 1. Because of the acyclicity of a pseudo-fork-join network we can define a sequence of stages, starting from the source and ending at the destination point.

Theorem 3. If (1) holds then the m.p.p. Φ^l associated with a link l coincides in finite time with a P -stationary and ergodic m.p.p. $\tilde{\Phi}^l$ with rate $\lambda P^0(l \in D)$.

Proof. Suppose, as before, that we start the network with a finite workload vector W_0 . We follow the method of stages. For a link l of stage 1 we have immediately that $\Phi^l = \tilde{\Phi}^l$ is already stationary since it is simply a thinning of the input m.p.p. To find the rate we apply Campbell's formula (see Appendix 1) to:

$$\tilde{N}^l(B) = \sum_{n \in \mathbb{Z}} I(T_n \in B, l \in D_n) \quad \text{for all Borel subsets } B \text{ of } \mathbb{R}$$

where \tilde{N}^l is the underlying point process of $\tilde{\Phi}^l$ and we get:

$$E\tilde{N}^l(B) = \lambda E^0 \int_{\mathbb{R}} I(s \in B, l \in D) ds = \lambda |B| P^0(l \in D)$$

where $|B|$ denotes the Lebesgue measure of the set B . Hence the rate of $\tilde{\Phi}^l$ is $\lambda P^0(l \in D)$.

Now assume that, as before, all processes prior to stage i couple (in finite time) with stationary ones and have the claimed rates. We prove that the same is true for the next stage. Clearly, the case of a fork presents no extra difficulty.

Assume first that a certain queue j joins a link of stage i with a link of stage $i+1$. By the induction hypothesis, the input m.p.p. to queue j couples with a stationary one denoted by $\tilde{\Phi}^j$ which has rate $\lambda P^0(j \in D)$. If P^j denotes the Palm transformation of P with respect to \tilde{N}^j the condition for stability is (according to Appendix 2)

$$(2) \quad \lambda P^0(j \in D) < (E^j \sigma^j)^{-1}.$$

The argument is similar to that of Case 1 of Theorem 1, provided that we show that (2) is equivalent to (3):

$$(3) \quad E^0 \sigma^j I(j \in D) < E^0 \tau.$$

To this end, we use the cycle formula (see Appendix 1) to compare the measures P^0 and P^j . This gives

$$\lambda E^0 \sigma^j I(j \in D) = \lambda P^0(j \in D) E^j \sum_{n \in \mathbb{Z}} \sigma_n^j I(j \in D_n) I(\tilde{T}_0^j < T_n \leq \tilde{T}_1^j)$$

where \tilde{T}_0 and \tilde{T}_1 are two typical successive points of the point process \tilde{N}^j . Noting that the last factor of the above expression is equal to $E^j \sigma^j$ (because the indicator function $I(j \in D_n)$ is equal to 1 for exactly one n on the set $\{\tilde{T}_0 < T_n \leq \tilde{T}_1\}$) we get

$$E^j \sigma^j = \frac{E^0 \sigma^j I(j \in D)}{P^0(j \in D)}$$

which proves that (2) is equivalent to (3). Therefore, as before, the output of j will reach stationarity in finite time.

Suppose next that several links of stage i are joined to a link of stage $i + 1$ through a join a . We assume that all m.p.p.'s prior to stage i reach stationarity in finite time. We must show that the m.p.p. Φ^a at the output of a join a also reaches stationarity in finite time. This can be done as follows. Observe that a job arriving in the system at time T_n will reach the join a at time $T_n + d_n$, say, if $a \in D_n$, or will not reach a at all if $a \notin D_n$. Due to the induction hypothesis, the sequence $\{d_n\}$ couples with a P^0 -stationary sequence $\{\tilde{d}_n\}$. Define $\tilde{\Phi}^a$ using $\{\tilde{d}_n\}$. Clearly Φ^a couples with $\tilde{\Phi}^a$. We can now show that Φ^a is stationary and has rate $\lambda P^0(a \in D)$ by considering the subsystem from the source up to the join a (i.e., this subsystem contains all paths starting from the source and ending at the join a) as a 'G/G/ ∞ system with deaths' (see Appendix 3) and apply the Lemma of Appendix 3. The theorem is now proved.

Reasoning now as before, we conclude that we have constructed a stationary workload process $\{\tilde{W}_i\}$ which is finite and unique among the class of finite workload processes.

Note. We can easily see that condition (1) is also almost necessary for stability. For, if it is violated for some queue j , then the workload vector converges almost surely to infinity.

4. Final comments

(a) The origin of the method of our proof goes back to Loynes (1962) where the result is proved for a system of FCFS tandem queues which is, of course, a special fork-join network.

(b) Another important special case fitting in the framework of Section 3 is that of k parallel queues where the arriving job chooses one 'at random'. If this selection is done independently of its service requirements, the stability condition (1) becomes:

$$\lambda p_j < \mu_j$$

where p_j is the probability that queue j is chosen by the incoming job.

(c) The case of parallel queues where each time one or more of them is chosen is treated by Baccelli et al. (1987a).

(d) Finally, we note that there is no extra difficulty in applying the method of the previous section to acyclic networks constructed by joining pseudo-fork-join networks in series and/or in parallel provided that we have a single arrival stream. (For instance, we can have two copies of the network of Figure 3 in tandem.) We summarize by saying that

a stability condition (ensuring convergence towards a unique stationary regime) has been found for a general class of acyclic networks.

Appendix 1

For reference only we give some results from the theory of stationary point processes (see Neveu (1977), Franken et al. (1982), Baccelli and Brémaud (1987) and Walrand (1988)). Let (Ω, F, P) be a probability space and $\{\theta_t, t \in \mathbb{R}\}$ a P -invariant semigroup (flow) on Ω . A point process N defined on Ω is a measurable mapping from Ω into the space M of counting measures on \mathbb{R} where M is equipped with the σ -field generated by cylinder sets. If T_t denotes the left shift by t in M (i.e., $T_t \phi(B) = \phi(B + t)$, $\phi \in M$, B a Borel set in \mathbb{R} , $t \in \mathbb{R}$) then N is, by definition, P -stationary if and only if $N \circ \theta_t = T_t N$. The rate λ of N is defined as $EN(0, 1]$, i.e. the mean number of points of N in the unit interval $(0, 1]$. The Palm transformation P^0 of P with respect to N is a probability measure on Ω defined by:

$$P^0(A) = \frac{1}{\lambda |B|} E \int_B I_A \circ \theta_t N(dt), \quad A \in F.$$

Under P^0 , N is *synchronous*, namely its interarrival times $\tau_n = T_{n+1} - T_n$, $n \in \mathbb{Z}$ form a P^0 -stationary random sequence. Let now Z_n be any random element of some measurable space S such that $Z_n = Z_0 \circ \theta_{T_n}$. Then Campbell's formula says:

$$E \sum_{n \in \mathbb{Z}} f(T_n, Z_n) = \lambda E^0 \int_{\mathbb{R}} f(s, Z_0) ds$$

for any non-negative, measurable, bounded function f . Given two stationary point processes N^1, N^2 on the same $(\Omega, F, P, \{\theta_t, t \in \mathbb{R}\})$ and letting P^1, P^2 be the Palm transformations of P with respect to N^1 and N^2 , respectively, the *cycle formula* (whose validity for two jointly stationary point processes was established by Neveu (1976)) relates them as follows:

$$\lambda_2 P^2(A) = \lambda_1 E^1 \sum_{n \in \mathbb{Z}} I_A \circ \theta_{T_n^1} I(0 < T_n^1 \leq T_1^1)$$

where λ_1, λ_2 are the rates of N^1 and N^2 , respectively.

Appendix 2

The main result used in the present paper is that of the stability of a $G/G/1$ queue. We give an outline of that result. (See Loynes (1962).)

Let $\{T_n, n \in \mathbb{Z}\}$ be the arrival times and $\{\sigma_n, n \in \mathbb{Z}\}$ the service times at a $G/G/1$ queue. It is assumed that the marked point process $\{(T_n; \sigma_n), n \in \mathbb{Z}\}$ is P -stationary (or, what is equivalent, P^0 -synchronous) and ergodic. Let $\tau_n = T_{n+1} - T_n$, $\lambda = (E^0 \tau_n)^{-1}$, and $\mu = (E^0 \sigma_n)^{-1}$. It is assumed that $0 < \lambda < \mu < \infty$. Lindley's equation is a recursive equation for the workload W_n at the n th arrival instant:

$$(4) \quad W_{n+1} = (W_n + \sigma_n - \tau_n)^+, \quad n \in \mathbb{Z}.$$

A stationary regime is a random variable \tilde{W} such that

$$(5) \quad \tilde{W} \circ \theta = (\tilde{W} + \sigma - \tau)^+$$

where $\theta = \theta_{T_1}$, $\sigma = \sigma_0$, $\tau = \tau_0$. Clearly, if such a stationary regime exists, then $\{\tilde{W}_n = \tilde{W} \circ \theta_{T_n}, n \in \mathbb{Z}\}$ is a P^0 -stationary solution of (4). Denoting by M_n the workload W_0 given that $W_{-n} = 0$ we observe that M_n is non-decreasing and so we let $\tilde{W} = \lim_{n \rightarrow \infty} M_n$. An ergodic argument shows that \tilde{W} satisfies (5), is finite and is the unique solution of (5). Another ergodic argument shows that any finite solution W_n of (4) coincides in finite time with \tilde{W}_n , P^0 -a.s. and this fact is used in Theorems 1 and 3. The workload process can be extended in continuous time to form a P -stationary process (see Franken et al. (1982)).

Finally, if $\lambda > \mu$ it can be shown that any solution of (4) converges almost surely to infinity, whereas we need more information to be able to say what happens in the critical case (see Loynes (1962)).

Appendix 3

The system we describe here we call *G/G/∞ system with deaths*. Assume that particles arrive in a room at times T_n , ordered according to n , carrying marks r_n, d_n, Z_n where r_n takes only two values, 0 or 1, d_n takes real non-negative values and Z_n is fairly arbitrary. If $r_n = 0$ the particle dies immediately. If $r_n = 1$ the particle stays for a period of length d_n and exists at time $T_n + d_n$ carrying its mark Z_n which can be thought of as some characteristic of the particle. Let Φ denote the input m.p.p. with points T_n and marks r_n, d_n, Z_n and Φ^* the output m.p.p. with points $T_n + d_n$ and marks Z_n , provided that $r_n = 1$. Observe that the points $T_n + d_n$ are not ordered according to n , in general. Nevertheless, Φ^* is well-defined. Observe also that we do not need to have the variables d_n, Z_n defined if $r_n = 0$ since any such point is not contained in Φ^* . We now prove the following result.

Lemma. If Φ is a stationary random m.p.p. on some probability space $(\Omega, \mathcal{F}, P, \{\theta_t, t \in \mathbb{R}\})$ then Φ^* is also stationary with rate $\lambda P^0(r = 1)$. If P^0, P^* are the Palm transformations of P with respect to Φ, Φ^* , respectively, then

$$P^*(Z_n \in A) = P^0(Z_n \in A \mid r_n = 1).$$

Proof. Observe that the system maps the input m.p.p. deterministically to the output m.p.p. That is, knowledge of a particular realization of the input yields the corresponding realization of the output. Consider now the m.p.p. $\{(T_n + t; r_n, d_n, Z_n), n \in \mathbb{Z}\}$ as input. (This is the translation of the original input by t .) Then the corresponding output is $\{(T_n + d_n + t; Z_n), n \in \mathbb{Z}\}$ (keeping only those indices n for which $r_n = 1$). Since the distribution of the former does not depend on t it follows, due to the previous observation, that the distribution of the latter is also independent of t . This shows that Φ^* is stationary. The rate of Φ^* is:

$$\lambda^* = E \sum_n I(T_n + d_n \in [0, 1], r_n = 1)$$

which, according to Campbell's formula, is

$$\lambda^* = \lambda E^0 \int_{\mathbf{R}} I(t + d \in [0, 1]) I(r = 1) dt = \lambda P^0(r = 1).$$

To prove the last claim, we apply the formula relating P^* and P :

$$P^*(Z \in A) = \lambda^{*-1} E \sum_n I(Z_n \in A, T_n + d_n \in [0, 1], r_n = 1)$$

followed by Campbell's formula:

$$\begin{aligned} P^*(Z \in A) &= \lambda^{*-1} \lambda E^0 \int_{\mathbf{R}} I(Z \in A) I(t + d \in [0, 1]) I(r = 1) dt \\ &= \lambda^{*-1} \lambda E^0 I(Z \in A) I(r = 1) \int_{\mathbf{R}} I(0 \leq s \leq 1) ds \\ &= P^0(Z \in A \mid r = 1). \end{aligned}$$

Acknowledgement

We should like to thank the referee for a number of careful comments. The referee has also pointed to us that, when limited to the case of deterministic routing, the class of pseudo-fork-join networks seems to correspond to the class of acyclic fork-join networks considered in Baccelli et al. (1987b).

References

- BACCELLI, F. AND BRÉMAUD, P. (1987) *Palm Probabilities and Stationary Queues*. Lecture Notes in Statistics 41, Springer-Verlag, Berlin.
- BACCELLI, F. AND MASSEY, W. A. (1986) Series-parallel, fork-join queueing networks and their stochastic ordering. Rapport de Recherche No. 488, INRIA-Rocquencourt, France.
- BACCELLI, F., MAKOWSKI, A. M. AND SCHWARTZ, A. (1987a) The fork-join queue and related systems with synchronization constraints: Stochastic ordering, approximations and computable bounds. Technical Research Report, SRC, Univ. of Maryland.
- BACCELLI, F., MASSEY, W. AND TOWSLEY, D. (1987b) COINS Technical Report 87-40, University of Massachusetts.
- FRANKEN, P., KÖNIG, D., ARNDT, U. AND SCHMIDT, V. (1982) *Queues and Point Processes*. Akademie-Verlag, Berlin.
- LOYNES, R. M. (1962) The stability of a queue with non-independent inter-arrival and service times. *Proc. Camb. Phil. Soc.* 58, 497-520.
- NEVEU, J. (1976) Sur les mesures de Palm de deux processus ponctuels stationnaires. *Z. Wahrscheinlichkeitsthe.* 34, 199-203.
- NEVEU, J. (1977) Processus ponctuelles. In *Lecture Notes in Mathematics* 598, Springer-Verlag, Berlin, 249-447.
- WALRAND, J. (1988) *An Introduction to Queueing Networks*. Prentice-Hall, Englewood Cliffs, NJ.