Bounds for the Coupling Time in Queueing Networks Perfect Simulation

Jantien Dopper, Bruno Gaujal, Jean-Marc Vincent

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Bounds for the Coupling Time in Queueing Networks Perfect Simulation

Jantien G. Dopper — Bruno Gaujal — Jean-Marc Vincent

N° 5828

Février 2006

Thème NUM

Rapport de recherche

ISSN 0249-6399

Bounds for the Coupling Time in Queueing Networks Perfect Simulation

Jantien G. Dopper [*] [†] , Bruno Gaujal [‡] [†] , Jean-Marc Vincent [§] [†]

Thème NUM — Systèmes numériques
Projet MESCAL

Rapport de recherche n° 5828 — Février 2006 — 22 pages

Abstract: In this paper, the duration of perfect simulations for Markovian finite capacity queuing networks is studied. This corresponds to hitting time (or coupling time) problems in a Markov chain over the Cartesian product of the state space of each queue. We establish an analytical formula for the expected simulation time in the one queue case and we provide simple bounds for acyclic networks of queues with losses. These bounds correspond to sums on the coupling time for each queue and are either almost linear in the queue capacities under light or heavy traffic assumptions or quadratic, when service and arrival rates are similar.

Key-words: Perfect simulation, Markov chain, Hitting time

[*] Jantien Dopper, Mathematical institute, Leiden University, NL (jgdopper@math.leidenuniv.nl) (this author was partially supported by a grant from INRIA)

[†] This work was partially supported by the French ACI SurePath project and SMS ANR

[‡] INRIA, Laboratoire ID-IMAG, Mescal Project INRIA-UJF-CNRS-INPG, 51, avenue Jean Kuntzmann, F-38330 Montbonnot, France (Bruno.Gaujal@imag.fr)

[§] UJF, Laboratoire ID-IMAG, Mescal Project INRIA-UJF-CNRS-INPG, 51, avenue Jean Kuntzmann, F-38330 Montbonnot, France (Jean-Marc.Vincent@imag.fr)

Bornes sur le temps de couplage de simulations parfaites de réseaux de files d'attente

Résumé : Dans cet article, nous étudions la durée de simulations parfaites de réseaux de files d'attente Markoviens, de capacités finies. Nous montrons que cette durée est liée à des problèmes de temps d'atteinte, ou de temps de couplage sur une chaîne de Markov définie sur l'espace d'état produit de chaque file d'attente. Nous établissons une formule analytique pour le temps moyen de simulation pour une seule file à partir de laquelle, nous construisons des bornes simples pour des réseaux acycliques de files d'attentes avec pertes. Ces bornes correspondent à la somme des temps de couplage pour chaque file considérée en isolation et sont quasi-linéaires en les capacités des files dans des conditions de fort ou de faible trafic et sont quadratiques en les capacités quand les taux d'arrivées et les taux de services sont proches.

Mots-clés : Simulation Parfaire, Chaîne de Markov, Temps d'Atteinte

1 Introduction

Markov chains are an important tool in modelling systems. Amongst others, Markov chains are being used in the theory of queueing systems, which itself is used in a variety of applications as performance evaluation of computer systems and communication networks. In modelling any queueing system, one of the main points of interest is the long run behavior of the system. For an irreducible, ergodic (i.e. aperiodic and positive-recurrent) Markov chain with probability matrix P, this long run behavior is described by the unique vector π which satisfies the linear system

$$\pi = \pi P.$$

We shall refer to the vector π as the stationary distribution. In most of the applications, the state space S of the Markov chain is finite and the chain is irreducible and aperiodic. Because in a finite, irreducible and aperiodic Markov chain all states are positive recurrent, the chain is ergodic. So for a finite Markov chain, irreducibility and aperiodicity are sufficient conditions for the existence of a unique stationary distribution or steady-state. However, it may be hard to compute this stationary distribution, especially when the finite state space is huge which is frequent in queuing models. In that case, steady-state simulation [6] can be used.

The classical method for simulation has been Monte Carlo simulation for many years. This method is based on the fact that an irreducible aperiodic finite Markov chain with transition matrix P and initial distribution $\mu^{(0)}$, the distribution $\mu^{(n)}$ of the chain at time n converges to π as n gets very large. That is:

$$\lim_{n\to\infty} \mu^{(n)} = \lim_{n\to\infty} \mu^{(0)} P^n = \pi.$$

So after running the Markov chain long enough, the states of the chain will not depend anymore on the initial state. However, the question is how long is long enough? That is, when is n sufficiently large so that $|\mu^{(n)} - \pi| \leqslant \epsilon$ for a certain $\epsilon > 0$? Moreover, the samples generated by this method will always be biased.

There are two approaches to estimate the steady-state from simulations. The method based on one long-run uses the ergodic theorem of Markov chains and estimates the stationary probability of a state s by the proportion of visits in s on one trajectory of the process. The drawback of this method is the auto-correlation of the sample. Moreover, a warm-up period is necessary in order to begin estimation when the process is near the steady-state. This initial transient problem remains open in many situations, and computation of confidence intervals needs elaborated techniques [3] depending highly on the model structure and parameter values.

The replication method consists in running independent finite trajectories. The advantage is to obtain independent samples of the steady state and classical convergence theorems could be used to compute confidence intervals. The drawback is the importance of the transient part in each of the replications. The simulation time could then be prohibitive. Discussion on one long-run versus replications could be found in [12]. Extensions to regenerative methods [5] decompose the long-run trajectories into “independent” batches. The simulation strategy replication, one long run with batches is discussed in [1].

In 1996, Propp an Wilson[7] solved these problems of Markov chain simulation by proposing an algorithm which returns exact samples of the stationary distribution. The striking difference between Monte Carlo simulation and this new algorithm is that Propp and Wilson do not simulate into the future, but go backwards in time. The main idea is, while going backwards in time, to run several simulations, starting with all $s \in S$ until the state at $t = 0$ is the same for all of them. If the output is the same for all runs, then the chain has coupled. Because of this coupling element and going backwards, this algorithm has been called Coupling From The Past (from now on: CFTP). A more detailed description of this algorithm will be presented in section 2.

When the coupling from the past technique is applicable, we get in a finite time one state with steady-state distribution. Then either we use a one long-run simulation from this state avoiding the estimation of the initial transient problem or we replicate independently the CFTP algorithm to get a sample of independent steady-state distributed variables. The analysis of the choice could be done exactly as in [1]. The replication technique has been applied successfully in finite capacity queueing networks with blocking and rejection (very large state-space) [10]. The efficiency of the simulation allows also the estimation of rare events (blocking probability, rejection rate) is done in [9].

The aim of this paper is to study the simulation time needed to generate one state, steady-state distributed, in the context of queueing networks with finite capacities. We will apply CFTP to networks of queues and study the coupling time τ of CFTP (*i.e.* the smallest time t for which the chain couples). Our main interest is setting bounds on the expected coupling time. We obtain exact analytical values for the expected simulation time for one M/M/1/C queue. As for networks of queues, we show how upper bounds on the mean simulation time can be obtained as sums of coupling times for each queue. This is used to provide explicit bounds which are linear in the queues capacities for acyclic networks with losses, under light or heavy traffic. However, when in input rate and the service rate are close, the bounds become quadratic in the capacities.

The paper is organized as follows. We first introduce the coupling from the past algorithm in Section 2. Then we show general properties of the coupling time for open Markovian queueing networks in Section 3. We will investigate the M/M/1/c queue in Section 4 providing exact computation for the expected coupling time and the case of acyclic networks in Section 5 where bounds are derived, together with several experimental tests assessing their quality.

2 Coupling from the Past

Let $\{X_n\}_{n\in\mathbb{N}}$ be an irreducible and aperiodic discrete time Markov chain with a finite state space $\mathcal{S}$ and a transition matrix $P = (p_{i,j})$. Let

$$\phi : \mathcal{S} \times \mathcal{E} \to \mathcal{S},$$

encode the chain, which means that it verifies the property $\mathbb{P}(\phi(i,e) = j) = p_{i,j}$ for every pair of states $(i,j) \in \mathcal{S}$ and for any e, a random variable distributed on $\mathcal{E}$. The function ϕ could be considered as a construction algorithm and e the innovation for the chain. In the context of discrete event systems, e is an *event* and ϕ is the *transition function*. Now, the evolution of the Markov chain is described as a stochastic recursive sequence

$$X_{n+1} = \phi(X_n, e_{n+1}), \tag{1}$$

with X_n the state of the chain at time n and $\{e_n\}_{n\in\mathbb{N}}$ an independent and identically distributed sequence of random variables.

Let $\phi^{(n)} : S \times \mathcal{E}^n \to \mathcal{S}$ denote the function whose output is the state of the chain after n iterations and starting in state $s \in S$. That is,

$$\phi^{(n)}(s, e_{1\to n}) = \phi(\ldots\phi(\phi(s,e_1), e_2), \ldots, e_n). \tag{2}$$

This notation can be extended to set of states. So for a set of states $A \subset \mathcal{S}$ we note

$$\phi^{(n)}(A, e_{1\to n}) = \left\{\phi^{(n)}(s, e_{1\to n}), s \in A\right\}.$$

theorem 1. *Let ϕ be a transition function on $\mathcal{S}\times\mathcal{E}$. There exists an integer l^* such that*

$$\lim_{n\to+\infty}\left|\phi^{(n)}(\mathcal{S}, e_{1\to n})\right| = l^* \text{ almost surely.}$$

This result is based on the following lemma and the fact that $\mathcal{S}$ is finite.

Lemma 2. *The sequence of integers $\{a_n\}_{n\in\mathbb{N}}$ defined by $a_n = \left|\phi^{(n)}(\mathcal{S}, e_{1\to n})\right|$, is non-increasing.*

This is clear because $\phi^{(n)}(\mathcal{S}, e_{1\to n}) = \phi(\phi^{(n-1)}(\mathcal{S}, e_{1\to n-1}), e_n)$, and the cardinal a_n of the image of $\phi^{(n-1)}(\mathcal{S}, e_{1\to n-1})$ by $\phi(., e_n)$ is less or equal than the cardinal a_{n-1} of $\phi^{(n-1)}(\mathcal{S}, e_{1\to n-1})$.

To complete the proof of the theorem, consider an arbitrary sequence of events $\{e_n\}_{n\in\mathbb{N}}$. Lemma 2 implies that the sequence $\{a_n\}_{n\in\mathbb{N}}$ converges to a limit l. Because the sizes of these sets belong to the finite set $\{1, \cdots, |\mathcal{S}|\}$, there exists $n_0 \in \mathbb{N}$ such that

$$a_{n_0} = \left|\phi^{(n_0)}\left(\mathcal{S}, e_{1\to n_0}\right)\right| = l.$$

Consider now l^* the minimal value of l among all possible sequences of events. Then there exists a sequence of events $\{e_n^*\}_{n\in\mathbb{N}}$ and an integer n_0^* such that

$$\left|\phi^{(n_0^*)}\left(\mathcal{S}, e^*_{1\to n_0^*}\right)\right| = l^*.$$

As a consequence of the Borel-Cantelli's Lemma, almost all sequences of events $\{e_n\}_{n\in\mathbb{N}}$ include the pattern $e^*_{1\to n_0^*}$. Consequently, the limit of the cardinality of $\phi^{(n)}\left(\mathcal{S}, e_{1\to n}\right)$ is less than l^*. The minimality of l^* finishes the proof.

Definition 3. *The system couples if*

$$\lim_{n\to+\infty}\left|\phi^{(n)}\left(\mathcal{S}, e_{1\to n}\right)\right| = 1 \text{ with probability } 1.$$

Then the *forward coupling time* τ^f defined by

$$\tau^f = \min\{n \in \mathbb{N};\ \text{such that } \left|\phi^{(n)}\left(\mathcal{S}, e_{1\to n}\right)\right| = 1\},$$

is almost surely finite. The coupling property of a system ϕ depends only on the structure of ϕ. The probability measure on $\mathcal{E}$ does not affect the coupling property, provided that all events in $\mathcal{E}$ has a positive probability. Moreover, the existence of some pattern $e^*_{1\to n_0^*}$ that ensures coupling, guarantees that τ^f is stochastically upper bounded by a geometric distribution

$$\mathbb{P}(\tau^f \geqslant k.n_0^*) \leqslant \left(1 - p(e_1^*).p(e_2^*)\dots p(e^*_{n_0})\right)^k; \tag{3}$$

where $p(e) > 0$ is the probability of event e.

At time τ^f, all trajectories issued from all initial states at time 0 have collapsed in only one trajectory. Unfortunately, the distribution of X_{τ^f} is not stationary. In [4] an example is given that illustrates why it is not possible to consider that this process has the stationary regime.

In fact, this iteration scheme could be reversed in time as it is usually done in the analysis of stochastic point processes. For that, one needs to extend the sequence of events to negative indexes and build the reversed scheme on sets by

$$A_n = \phi^{(n)}\left(\mathcal{S}, e_{-n+1\to 0}\right).$$

It is clear that the sequence of sets A_n is non-decreasing ($A_{n+1} \subset A_n$). Consequently, the system couples if the sequence A_n converges almost surely to a set with only one element. Almost surely, there exists a finite time τ^b, the *backward coupling time*, defined by

$$\tau^b = \min\{n \in \mathbb{N};\ \text{such that } \left|\phi^{(n)}\left(\mathcal{S}, e_{-n+1\to 0}\right)\right| = 1\}.$$

Proposition 4. *The backward coupling time τ^b and the forward coupling time τ^f have the same probability distribution.*

For a detailed proof of this proposition, we refer to [11]. Here is the main idea of the proof. Compute the probability

$$\mathbb{P}(\tau^f > n) = \mathbb{P}(\left|\phi^{(n)}\left(\mathcal{S}, e_{1\to n}\right)\right| > 1).$$

Since the process $\{e_n\}_{n\in\mathbb{Z}}$ is stationary, shifting the process to the left leads to

$$\mathbb{P}(\left|\phi^{(n)}\left(\mathcal{S}, e_{1\to n}\right)\right| > 1) = \mathbb{P}(\left|\phi^{(n)}\left(\mathcal{S}, e_{-n+1\to 0}\right)\right| > 1) = \mathbb{P}(\tau^b > n).$$

Hence, if we want to make any statement about the probability distribution of the coupling time τ^b of CFTP, we can use the conceptually easier coupling time τ^f.

The main result of the backward scheme is the following theorem [7].

theorem 5. *Provided that the system couples, the state when coupling occurs for the backward scheme, is steady state distributed.*

From this fact, a general algorithm (1) sampling the steady state can be constructed.

Algorithm 1 Backward-coupling simulation (general version)

```
for all s ∈ S do
  y(s) ← s {choice of the initial value of the vector y, n = 0}
end for
repeat
  e ← Random_event; {generation of e_{−n+1}}
  for all s ∈ S do
    y(s) ← y(φ(s, e));
    {y(s) state at time 0 of the trajectory issued from s at time −n + 1}
  end for
until All y(x) are equal
return y(x)
```

The complexity c_ϕ of this algorithm is $c_\phi = O(\tau^b.|\mathcal{S}|)$. The coupling time τ^b is of fundamental importance for the efficiency of the sampling algorithm. To improve its complexity, we could reduce the factor $|\mathcal{S}|$ and reduce the coupling time. When the state space is partially ordered by a partial order $\prec$ and the transition function is monotone for each event e, it is sufficient to simulate trajectories starting from the maximal and minimal states [7]. Denote by M and m the set of maximal, respectively minimal elements of $\mathcal{S}$ for the partial order $\prec$. The monotone version of algorithm (1) is given by algorithm (2). In this case, we need to store the sequence of events in order to preserve the coherence between the trajectories driven from $M \cup m$.

Algorithm 2 Backward-coupling simulation (monotone version)

```
n=1;
R[n]=Random_event;{array will R stores the sequence of events }
repeat
  n=2.n;
  for all s ∈ M ∪ m do
    y(s) ← s {Initialize all trajectories at time −n}
  end for
  for i=n downto n/2+1 do
    R[i]=Random_event; {generates all events from time −n + 1 to n/2 + 1}
  end for
  for i=n downto 1 do
    for all s ∈ M ∪ m do
      y(s) ← Φ(y(s), R[i])
    end for
  end for
until All y(s) are equal
return y(s)
```

The doubling scheme (first step in the loop) leads to a complexity

$$c_\phi \leqslant 2K.\tau^b.(|M| + |m|), \tag{4}$$

where K is a constant.

3 Open Markovian queueing networks

Consider an open network Q consisting of K queues $Q_1, \ldots, Q_K$. Each queue Q_i has a finite capacity, denoted by C_i, $i = 1, \ldots K$. Thus the state space of a single queue Q_i is $\mathcal{S}_i = \{0, \ldots C_i\}$. Hence, the state space $\mathcal{S}$ of the network is $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_K$. The state of the system is described by a vector $s = (s_1, \ldots, s_K)$ with s_i the number of customers in queue Q_i. The state space is partially ordered by the component-wise ordering and there are a maximal state M when all queues are full and a minimal state when all queues are empty.

The network evolves in time due to exogenous customer arrivals from outside of the network and to service completions of customers. After finishing his service at a server, a customer is either directed to another queue by a certain routing policy or leaves the network. A routing policy determines to which queue a customer will go, taking into account the global state of the system. Moreover, the routing policy also decides what happens with a customer if he is directed to a queue which buffer is filled with C_i customers.

An event in this network is characterized by the movements of some clients between queues modeling the routing strategy and the Poisson process defining the occurrence rate of the event. For example consider the acyclic queueing network (figure 1) is characterized by 4 queues and 6 events.

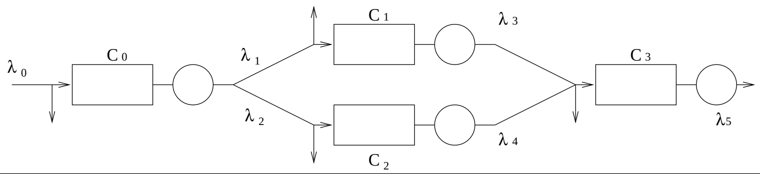

	rate	origin	destination	enabling condition	routing policy
e_0	λ_0	Q_{-1}	Q_0	none	rejection if Q_0 is full
e_1	λ_1	Q_0	Q_1	$s_0 > 0$	rejection if Q_1 is full
e_2	λ_2	Q_0	Q_2	$s_0 > 0$	rejection if Q_2 is full
e_3	λ_3	Q_1	Q_3	$s_1 > 0$	rejection if Q_3 is full
e_4	λ_4	Q_2	Q_3	$s_2 > 0$	rejection if Q_3 is full
e_5	λ_5	Q_3	Q_{-1}	$s_3 > 0$	none

Figure 1: Network with rejection

Since the network is open, clients are able to enter and leave the network. We assume that customers who enter from outside the network to a given queue arrive according to a Poisson process. Furthermore, suppose that the service times at server i are independent and exponentially distributed with parameter μ_i.

Definition 6. *An event e is an application defined on S that associates to each state $x \in S$ a new state denoted by $\phi(x, e)$. The function ϕ is the transition function of the network.*

For example, to event e_1 (fig 1) we get

$$\phi(., e_1) : (s_0, s_1, s_2, s_3) \longmapsto \begin{cases} (s_0 - 1, s_1 + 1, s_2, s_3) & \text{if } (s_0 \geqslant 1) \text{ and } (s_1 < C_1); \\ (s_0 - 1, s_1 + 1, s_2, s_3) & \text{if } (s_0 \geqslant 1) \text{ and } (s_1 = C_1)(Q_1 \text{ full}); \\ (s_0, s_1, s_2, s_3) & \text{if } (s_0 \geqslant 0)(Q_0 \text{ empty}). \end{cases}$$

Definition 7. *An event e is monotone if $\phi(x, e) \leqslant \phi(y, e)$ for every $x, y \in \mathcal{S}$ with $x \leqslant y$.*

It is clear that the previous event e_1 is monotone. Moreover usual events such as routing with overflow and rejection, routing with blocking and restart, routing with a index policy rule (eg Join the shortest queue) are monotone events [2, 10].

Denote by $\mathcal{E} = \{e_1, \ldots, e_M\}$ the finite collection of events of the network. With each event e_i is associated a Poisson process with parameter λ_i. If an event occurs which does not satisfy the enabling condition the state of the system is unchanged.

To complete the construction of the discrete-time Markov chain, the system is uniformized by a Poisson process with rate $\Lambda = \sum_i^M \lambda_i$. Hence, one can see this Poisson process as a clock which determines when an event transition takes place. To choose which specific transition actually takes place, the collection $\mathcal{E}$ of events of the network is randomly sampled with

$$p_i = \mathbb{P}\left(\text{event } e_i \text{ occurs}\right) = \frac{\lambda_i}{\Lambda}.$$

By construction, the following proposition should be clear.

Proposition 8. *The uniformized Markov chain has the same stationary distribution as the queueing network, and so does the embedded discrete time Markov chain.*

Provided that events are monotone, the CFTP algorithm can be applied on queueing networks to build steady-state sampling of the network.

In our example of Figure 1 we ran the CFTP algorithm and produced samples of coupling time. The parameters used for the simulation are the following. Queues capacity : $\forall i = 1..4, C_i = 10$. Event rates: $\lambda_1 = 1.4$, $\lambda_2 = 0.6$, $\lambda_3 = 0.8$, $\lambda_4 = 0.5$ and $\lambda_5 = 0.4$. The global input rate λ_0 is varying. The number of samples used to estimate the mean coupling time is 10000. The result is displayed in Figure 2.

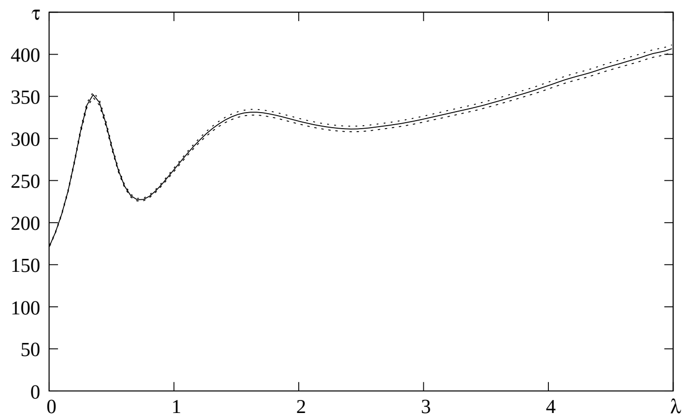

Figure 2: Mean coupling time for the acyclic network of Figure 1 when the input rate varies from 0 to 4, with 95% confidence intervals.

This type of curve is of fundamental importance because the coupling time corresponds to the simulation duration and is involved in the simulation strategy (long run versus replication). These first results can be surprising because they exhibit a strong dependence on parameters values. The aim of this paper is now to understand more deeply what are the critical values for the network and to build bounds on the coupling time that are non-trivial.

Let N_i be the function from $\mathcal{S}$ to $\mathcal{S}_i$ with $N_i\left(s_1, \ldots, s_K\right) = s_i$. So N_i is the number of customers in queue Q_i. As in section 2, τ^b refers to the *bacvkward coupling time* of the chain, which is in case the coupling time from the past of the queueing network.

Definition 9. *Let τ_i^b denote the backward coupling time on coordinate i of the state space. Thus τ_i^b is the smallest n for which*

$$\left|\left\{N_i\left(\phi^{(n)}\left(s, e_{-n+1}, \ldots, e_0\right)\right), s \in \mathcal{S}\right\}\right| = 1.$$

Because coordinate s_i refers to queue Q_i, the random variable τ_i^b represents the coupling time from the past of queue Q_i. Once all queues in the network have coupled, the CFTP algorithm returns one value and hence the chain has coupled. Thus

$$\tau^b = \max_{1 \leqslant i \leqslant K} \{\tau_i^b\} \leqslant_{st} \sum_{i=1}^{K} \tau_i^b. \tag{5}$$

By taking expectation and interchanging sum and expectation we get:

$$\mathbb{E}\left[\tau^b\right] = \mathbb{E}\left[\max_{1 \leqslant i \leqslant K} \{\tau_i^b\}\right] \leqslant \mathbb{E}\left[\sum_{i=1}^{K} \tau_i^b\right] = \sum_{i=1}^{K} \mathbb{E}\left[\tau_i^b\right] \tag{6}$$

It follows from Proposition 4 that τ^b and τ^f have the same distribution. The same holds for τ_i^f and τ_i^b. Hence $\mathbb{E}\left[\tau_i^b\right] = \mathbb{E}\left[\tau_i^f\right]$ and

$$\mathbb{E}\tau^b \leqslant \sum_{i=1}^{K} \mathbb{E}\left[\tau_i^f\right]. \tag{7}$$

The bound given in Equation 7 is interesting because $\mathbb{E}\left[\tau_i^f\right]$ is sometimes amenable to explicit computations, as shown in following sections. In order to derive those bounds, one may provide yet other bounds, by making the coupling state explicit.

Definition 10. *The hitting time $h_{j\to k}$ in a Markov chain X_n is defined as*

$$h_{j\to k} = \inf_{\mathbb{N}}\{n \text{ s.t. } X_n = k | X_0 = j\} \text{ with } j, k \in \mathcal{S}.$$

The hitting time $h_{j\to k}$ with $j, k \in \mathcal{S}_i$ is the hitting of a single queue Q_i of the network. Now $h_{0\to C_i}$ represents the number of steps it takes the queue Q_i to go from state 0 to state C_i. Now we take queue Q_i out of the network and examine it independently. Suppose that $h_{0\to C_i} = n$ for the sequence of events $e_1, \dots e_n$. Because of monotonicity of ϕ we have

$$\phi^{(n)}(0, e_1, \dots, e_n) \leqslant \phi^{(n)}(s, e_1, \dots, e_n) \leqslant \phi^{(n)}(C_i, e_1, \dots, e_n) = 0,$$

with $s \in \mathcal{S}_i$. Hence, coupling has occurred. So $h_{0\to C_i}$ is an upper bound on the forward coupling of queue Q_i. The same argumentation holds for $h_{C_i\to 0}$. Thus

$$\mathbb{E}\left[\tau_i^f\right] \leqslant \mathbb{E}\left[\min\{h_{0\to C_i}, h_{C_i\to 0}\}\right]. \tag{8}$$

Hence,

$$\mathbb{E}\tau^b \leqslant \sum_{i=1}^{K} \mathbb{E}\left[\tau_i^f\right] \leqslant \sum_{i=1}^{K} \mathbb{E}\left[\min\{h_{0\to C_i}, h_{C_i\to 0}\}\right] \leqslant \sum_{i=1}^{K} \min(\mathbb{E}h_{0\to C_i}, \mathbb{E}h_{C_i\to 0}), \tag{9}$$

by Jensen's Inequality.

4 Coupling time in a M/M/1/C queue

In a M/M/1/C model, we have a single queue with one server. Customers arrive at the queue according to a Poisson process with rate λ and the service time is distributed according to an exponential distribution with parameter μ. In the queue there is only place for C customers. So the state space $\mathcal{S} = \{0, \dots, C\}$. If a customer arrives when there are already C customers in the queue, he immediately leaves without entering the queue. After uniformization, we get a discrete time Markov chain which is governed by the events e_a with probability $p = \frac{\lambda}{\lambda+\mu}$ and e_d with probability $q = 1 - p$. Event e_a represents an arrival and event e_d represents an end of service with departure of the customer.

In order to estimate the expectation of the coupling time from the past $\mathbb{E}[\tau^b]$ we use inequality 9. Since there is only one queue, the first two inequalities in 9 become equalities. Indeed, when applying forward simulation, the chain only can couple in state 0 or state C. This follows since for $r, s \in \mathcal{S}$ with $0 < r < s < C$ we have $\phi(r, e_a) = r+1 < s+1 = \phi(s, e_a)$ and $\phi(r, e_d) = r-1 < s-1 = \phi(s, e_d)$ So the chain cannot couple in a state s with $0 < s < C$. Furthermore we have $\phi(C, e_a) = C = \phi(C-1, e_a)$ and $\phi(0, e_d) = 0 = \phi(1, e_d)$. Hence, forward coupling can only occur in 0 or C:

$$\mathbb{E}\left[\tau^b\right] = \mathbb{E}\left[\min\{h_{0\to C}, h_{C\to 0}\}\right]. \tag{10}$$

4.1 Explicit calculation of $\mathbb{E}\left[\tau^b\right]$

In order to compute $\min\{h_{0\to C}, h_{C\to 0}\}$ we have to run two copies of the Markov chain for a M/M/1/C queue simultaneously. One copy starts in state 0 and the other one starts in state C. We stop when either the chain starting in 0 reaches state C or when the copy starting in state C reaches state 0.

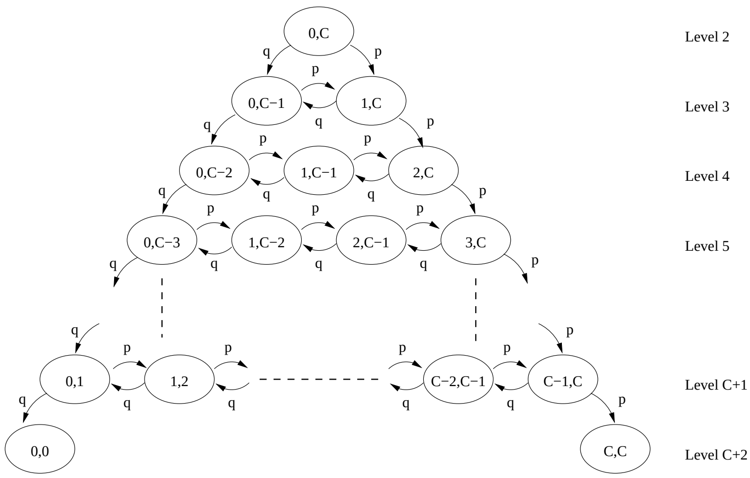

Figure 3: Markov chain $X(q)$ corresponding to $H_{i,j}$

Therefore, let us rather consider a product Markov chain $X(q)$ with state space $\mathcal{S}\times\mathcal{S} = \{(i,j), i = 0,\dots,C, j = 0,\dots,C\}$. The Markov chain $X(q)$ is also governed by the two events e_a and e_d and the function ϕ is:

$$\begin{aligned}\psi\left((x,y), e_a\right) &= \left((x+1)\wedge C, (y+1)\wedge C\right)\\ \psi\left((x,y), e_d\right) &= \left((x-1)\vee 0, (y-1)\vee 0\right).\end{aligned}$$

Without any loss of generality, we may assume that $i \leqslant j$. This system corresponds with the *pyramid Markov chain* $X(q)$ displayed in Figure 3.

Since we can only couple in 0 or C, this coupling occurs as soon as the chain $X(q)$ reaches states $(0,0)$ or (C,C). Define

$$H_{i,j} := \text{number of steps to reach state } (0,0) \text{ or } (C,C) \text{ starting from state } (i,j)$$

with $(i,j) \in \mathcal{S}\times\mathcal{S}$. By definition, $\min\{h_{0\to C}, h_{C\to 0}\} = H_{0,C}$. Now $H_{i,j}$ represents the hitting time of the coupling states $(0,0)$ and (C,C) (also called absorption time) in a product Markov chain. Using a one step analysis, we get the following system of equations for $\mathbb{E}[H_{i,j}]$:

$$\left\{\begin{array}{lcll}\mathbb{E}[H_{i,j}] &=& 1 + p\mathbb{E}[H_{(i+1)\wedge C,(j+1)\wedge C}] + q\mathbb{E}[H_{(i-1)\vee 0,(j-1)\vee 0}], & i \neq j,\\ \mathbb{E}[H_{i,j}] &=& 0, & i = j\end{array}\right. \tag{11}$$

Two states (i,j) and (i',j') are said to be at the same level if $|j-i| = |j'-i'|$. In Figure 3 we can distinguish $C+1$ levels. Because of monotonicity of ϕ, $|j-i|$ cannot increase. Hence, starting at a level with $|j-i|$, the chain will gradually pass all intermediate levels to reach finally the level with $|j-i| = 0$ in state $(0,0)$ or (C,C). Thus, starting in state $(0,C)$, the chain will run trough all levels to end up at the level with $|j-i| = 0$. So, $H_{0,C} = \min\{h_{0\to C}, h_{C\to 0}\}$. To determine $\mathbb{E}[H_{0,C}]$ we determine the mean time spent on each level and sum up over all levels.

A state (i,j) belongs to level m if $|j-i| = C+2-m$. Then state $(0,C)$ belongs to level 2 and the states $(0,0)$ and $C,C)$ belong to level $C+2$. To get from $(0,C)$ into either $(0,0)$ or (C,C), the chain $X(q)$ needs to cross all levels between the levels 2 and $C+2$. Let T_m denote time it takes to reach level $m+1$, starting in level m. Then

$$H_{0,C} = \sum_{m=2}^{C+1} T_m. \tag{12}$$

In order to determine T_m, we associate to each level m a random walk R_m on $0,\ldots,m$ with absorbing barriers at state 0 and state C. In the random walk, the probability of going up is p and of going down is $q = 1-p$. We have the following correspondence between the states of the random walk R_m and the states of $X(q)$ (see Figure 4).

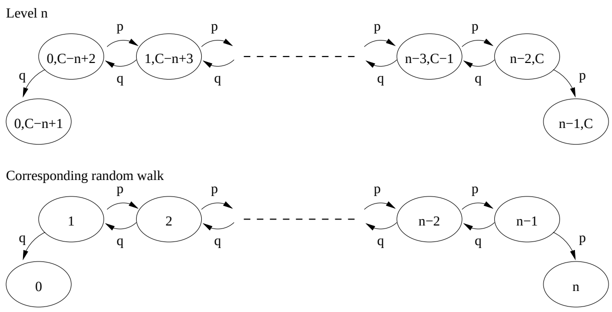

Figure 4: Relationship between level m and random walk R_m.

State 0 of R_m corresponds with state $(0, C-m+1)$ of $X(q)$,
State i of R_m corresponds with state $(i-1, C-m+1+i)$ of the $X(q)$,
$1 \leqslant i \leqslant m-1$,
State m of R_m corresponds with state $(m-1, C)$ of $X(q)$.

Now the time spent on level m in $X(q)$ is the same as the time spent in a random walk R_m before absorption. Therefore, on can use the two following results on random walks in determining T_m, which are known as ruin problems (see for example [8]).

Let $\alpha^m_{i\to 0}$ denote the probability of absorption in state 0 of the random walk R_m starting in i. Then:

$$\alpha^m_{i\to 0} = \begin{cases} \frac{a^m - a^i}{a^m - 1}, & p \neq \frac{1}{2}, \\ \\ \frac{m-i}{m}, & p = \frac{1}{2}, \end{cases} \tag{13}$$

where $a = q/p$.

Now, absorption occurs in R_m once the state 0 or C has been achieved.

Lemma 11. *Let $\widetilde{T}^m_i$ denote the mean absorption time of a random walk R_m starting in i. Then:*

$$\mathbb{E}[\widetilde{T}^m_i] = \begin{cases} \frac{i - m(1-\alpha^m_{i\to 0})}{q-p}, & p \neq \frac{1}{2}, \\ \\ i(m-i), & p = \frac{1}{2}. \end{cases} \tag{14}$$

Now, let β_i^m denote the probability that absorption occurs in $i = 0, m, $. Then

$$\beta_0^m = \sum_{i=0}^{m} \alpha_{i\to 0}^m \mathbb{P}\left(R_m \text{ starts in state } i\right), \tag{15}$$

and $\beta_m^m = 1 - \beta_0^m$. From the structure of the Markov chain $X(q)$ and the correspondence between $X(q)$ and the random walks, we have that (see Figure 4):

$$\mathbb{P}\left(\text{enter level } m+1 \text{ at } (0, C-m+1)\right) = \mathbb{P}\left(\text{absorption in 0 in } R_m\ \right) = \beta_0^m.$$

Now one has:

$$\begin{aligned}
\mathbb{E}\left[T_m\right] &= \mathbb{E}[\widetilde{T}_1^m]\beta_0^{m-1} + \mathbb{E}[\widetilde{T}_{m-1}^m]\beta_{m-1}^{m-1} \\
&= \mathbb{E}[\widetilde{T}_{m-1}^m] + \left(\mathbb{E}[\widetilde{T}_1^m] - \mathbb{E}[\widetilde{T}_{m-1}^m]\right)\beta_0^{m-1}.
\end{aligned} \tag{16}$$

4.1.1 Case q = p = 1/2

$\mathbb{E}[T_m]$ can be calculated explicitly for $p = \frac{1}{2}$. Since the random walk is symmetric, we have $\beta_0^m = \beta_n^m = \frac{1}{2}$. So:

$$\mathbb{E}\left[T_m\right] = \mathbb{E}[\widetilde{T}_1^m]\beta_0^{m-1} + \mathbb{E}[\widetilde{T}_{m-1}^m]\beta_{m-1}^{m-1} = m-1. \tag{17}$$

Hence,

$$\mathbb{E}\left[H_{0,C}\right] = \sum_{m=2}^{C+1} \mathbb{E}\left[T_m\right] = \sum_{m=2}^{C+1} m - 1 = \frac{C^2+C}{2}.$$

Lemma 12. *For a M/M/1/C with $\lambda = \mu$, $\mathbb{E}\tau^b = \frac{C^2+C}{2}$.*

4.1.2 Case $p \neq \frac{1}{2}$

Since the random walks are not symmetric, we cannot apply the same reasoning as for the case $p = \frac{1}{2}$ to compute β_0^m. Entering the random walk R_m corresponds to entering level m in $X(q)$. Since we can only enter level m in the state $(0, C-m+2)$ and $(m-2, C)$ this means we can only start the random walk in state 1 or $m-1$. Therefore (15) becomes:

$$\begin{aligned}
\beta_0^m &= \sum_{i=0}^{m} \alpha_{i\to 0}^m \mathbb{P}\left(R_m \text{ starts in state } i\right) \\
&= \alpha_{1\to 0}^m \mathbb{P}\left(R_m \text{ starts in } 1\right) + \alpha_{m-1\to 0}^m \mathbb{P}\left(R_m \text{ starts in } m-1\right) \\
&= \alpha_{m-1\to 0}^m + \left(\alpha_{1\to 0}^m - \alpha_{m-1\to 0}^m\right)\beta_0^{m-1} \\
&= \frac{a^m - a^{m-1}}{a^m - 1} + \frac{a^{m-1} - a}{a^m - 1}\beta_0^{m-1}.
\end{aligned}$$

This gives the recurrence:

$$\begin{cases}
\beta_0^m = \frac{a^m - a^{m-1}}{a^m - 1} + \frac{a^{m-1} - a}{a^m - 1}\beta_0^{m-1} & m > 2; \\
\beta_0^2 = 2.
\end{cases} \tag{18}$$

Thus we obtain,

Proposition 13. *For a M/M/1/C queue, using the foregoing notations,*

$$\mathbb{E}\tau^b = \mathbb{E}\left[H_{0,C}\right] = \sum_{m=2}^{C+1} \mathbb{E}[\widetilde{T}_{m-1}^m] + \left(\mathbb{E}[\widetilde{T}_1^m] - \mathbb{E}[\widetilde{T}_{m-1}^m]\right)\beta_0^{m-1}, \tag{19}$$

with β_0^m defined by (18) and $\mathbb{E}[\widetilde{T}_{m-1}^m]$ and $\mathbb{E}[\widetilde{T}_1^m]$ defined by (14).

4.1.3 Comparison between the cases $p = 1/2$ and $p \neq 1/2$

Proposition 14. *The coupling time in a M/M/1/C queue is maximal when the input rate λ and the service rate μ are equal.*

Proof. By definition, $\lambda = \mu$ corresponds to $p = q = 1/2$. The proof holds by induction on C. The result is obviously true when $C = 0$, because whatever q, $\mathbb{E}\left[H_{0,C}\right] = 0$.

For $C+1$, let q be an arbitrary probability with $q > 1/2$ (the case $q < 1/2$ is symmetric). We will compare the expected time for absorption of three Markov chains. The first one is the Markov chain $X := X(1/2)$ displayed in Figure 3, with $q = p = 1/2$. The second one is the Markov chain $X' = X(q)$ displayed in Figure 3 and the last one X'' is a mixture between the two previous chains: The first C levels are the same as in X while the last level $(C+1)$ is the same as in X'.

The expected absorption time for the first C levels is the same for X and for X'': $\sum_{m=2}^{C} \mathbb{E}T_m = \sum_{m=2}^{C} \mathbb{E}T'_m$. By induction, this is larger than for X': we have $\sum_{m=2}^{C} \mathbb{E}T_m = \sum_{m=2}^{C} \mathbb{E}T''_m \geqslant \sum_{m=2}^{C} \mathbb{E}T'_m$. Therefore, we just need to compare the exit times out of the last level, namely $\mathbb{E}T_{C+1}, \mathbb{E}T'_{C+1}$ and $\mathbb{E}T''_{C+1}$, to finish the proof.

We fist compare $\mathbb{E}T_{C+1}$ and $\mathbb{E}T''_{C+1}$. In both cases, the Markov chain enters level $C+1$ in state $(0,1)$ with probability $1/2$.

Equation (17) says that $\mathbb{E}T_{C+1} = C$ and Equation (14) gives after straightforward computations, $\mathbb{E}T''_{C+1} = 1/2 \frac{1-C(1-\alpha^C_{1\to 0})}{q-p} + 1/2 \frac{C-1-C(1-\alpha^C_{C-1\to 0})}{q-p} = \frac{C}{2q}\frac{a^C - a}{a^C - 1} \leqslant C/(2q) < C = \mathbb{E}T_{C+1}$.

In order to compare $\mathbb{E}T'_{C+1}$ and $\mathbb{E}T''_{C+1}$, let us first show that β_0^m is larger than $1/2$, for all $m \geqslant 2$. This is done by an immediate induction on Equation (18). If $\beta_0^{m-1} \geqslant 1/2$, then $\beta_o^m \geqslant \frac{2a^m - a^{m-1} - a}{2(a^m-1)}$ Now, $\frac{2a^m - a^{m-1} - a}{2(a^m-1)} \geqslant 1/2$ if $2a^m - a^{m-1} - a \geqslant a^m - 1$, *i.e.* after recombining the terms, $(a-1)(a^{m-1}-1) \geqslant 0$. This is true as soon as $q \geqslant 1/2$.

To end the proof, it is enough to notice that for the chain X', time to absorption starting in 1, $\mathbb{E}\tilde{T}_1^{m'}$ is smaller that time to absorption starting in $m-1$, $\mathbb{E}\tilde{T}_{m-1}^{m'}$ for all m. The difference $\mathbb{E}\tilde{T}_{m-1}^{m'} - \mathbb{E}\tilde{T}_1^{m'}$ is

$$\frac{ma^m - ma^{m-1} + ma - m - 2a^m + 2}{p\left(a^m - 1\right)\left(a-1\right)} = \frac{m(a-1)\left(\frac{a^{m-1}+1}{2} - \frac{1+a+\cdots+a^{m-1}}{m}\right)}{p\left(a^m-1\right)\left(a-1\right)} \geqslant 0,$$

by convexity of $x \mapsto a^x$. Finally,

$$\begin{aligned}\mathbb{E}T'_{C+1} &= \beta_0^{C+1}\mathbb{E}\tilde{T}_1^{C+1'} + (1-\beta_0^{C+1})\mathbb{E}\tilde{T}_C^{C+1'} \\ &\leqslant \frac{1}{2}\mathbb{E}\tilde{T}_1^{C+1'} + \frac{1}{2}\mathbb{E}\tilde{T}_C^{C+1'} = \mathbb{E}T''_{C+1}.\end{aligned}$$

□

4.2 Explicit Bounds

Equation (19) provides a quick way to compute $\mathbb{E}\left[H_{0,C}\right]$ using recurrence equation (18). However, it may also be interesting to get a simple closed form for an upper bound for $\mathbb{E}\left[H_{0,C}\right]$. This can be done using the last inequality in Equation (9) that gives an upper bound for $\mathbb{E}\left[H_{0,C}\right]$ amenable to direct computations.

$$\mathbb{E}\left[H_{0,C}\right] = \mathbb{E}\left[\min\{h_{0\to C}, h_{C\to 0}\}\right] \leqslant \min\{\mathbb{E}\left[h_{0\to C}\right], \mathbb{E}\left[h_{C\to 0}\right]\}. \quad (20)$$

The exact calculation of $\mathbb{E}\left[h_{0\to C}\right]$ can be done using a one step analysis. Let F_i be the time to go from state i to 0. Then, $h_{0\to C} = F_C$ and for all $i > 0$,

$$\mathbb{E}[F_i] = 1 + p\mathbb{E}[F_{(i+1)\wedge C}] + q\mathbb{E}[F_{i-1}]. \quad (21)$$

With an approach derived from [8] one can condition on the next event. Let T_i denote the time to go from state i to $i+1$. Then

$$\mathbb{E}\left[h_{0\to C}\right] = \sum_{i=0}^{C-1} \mathbb{E}\left[T_i\right]. \tag{22}$$

To get an expression for T_i, with $0 < i \leqslant C$, we condition on the first event. Therefore let $\mathbb{E}\left[T_i|e\right]$ denote the conditional expectation of T_i knowing that the next event is e. Since $\mathbb{E}[T_i \mid e_a] = 1$ and $\mathbb{E}[T_i \mid e_d] = 1 + \mathbb{E}\left[T_{i-1}\right] + \mathbb{E}\left[T_i\right]$, conditioning delivers the following recurrent expression for the $\mathbb{E}\left[T_i\right]$:

$$\mathbb{E}\left[T_i\right] = \begin{cases} \frac{1}{p} + \frac{q}{p}\mathbb{E}\left[T_{i-1}\right], & 0 < i < C \\ \frac{1}{p}, & i = 0. \end{cases} \tag{23}$$

By induction one can show that $\mathbb{E}\left[T_i\right] = \frac{1}{p}\sum_{k=0}^{i}\left(\frac{q}{p}\right)^k$. Hence, $\mathbb{E}\left[T_i\right] = \frac{1-\left(\frac{q}{p}\right)^{i+1}}{p-q}$ and from (22) it follows that

$$\mathbb{E}\left[h_{0\to C}\right] = \sum_{i=0}^{C-1} \frac{1-\left(\frac{q}{p}\right)^{i+1}}{p-q} = \frac{C}{p-q} - \frac{q(1-\left(\frac{q}{p}\right)^C)}{\left(p-q\right)^2}. \tag{24}$$

By reasons of symmetry, we have

$$\mathbb{E}\left[h_{C\to 0}\right] = \frac{C}{q-p} - \frac{p(1-\left(\frac{p}{q}\right)^C)}{\left(q-p\right)^2} \tag{25}$$

The curves of $\mathbb{E}\left[h_{0\to C}\right]$ and $\mathbb{E}\left[h_{C\to 0}\right]$ intersect in $C^2 + C$ when $p = q$. If $p > q$ then $\mathbb{E}\left[h_{0\to C}\right] < \mathbb{E}\left[h_{C\to 0}\right]$ and because of symmetry, if $p < q$ then $\mathbb{E}\left[h_{0\to C}\right] > \mathbb{E}\left[h_{C\to 0}\right]$. Since also $\frac{C^2+C}{2}$ is an upper bound corresponding to the critical case $p = q$ on the mean coupling time $\mathbb{E}\tau^b$, as shown in Proposition 14, one can state:

Proposition 15. *The mean coupling time $\mathbb{E}\tau^b$ of a M/M/1/C queue with arrival rate λ and service rate μ is bounded using $p = \lambda/(\lambda+u)$ and $q = 1-p$.*

Critical bound:

$$\forall p \in [0,1], \quad \mathbb{E}\tau^b \leqslant \frac{C^2+C}{2}.$$

Heavy traffic Bound:

$$\textit{if } p > \frac{1}{2}, \quad \mathbb{E}\tau^b \leqslant \frac{C}{p-q} - \frac{q(1-\left(\frac{q}{p}\right)^C)}{\left(p-q\right)^2}.$$

Light traffic bound:

$$\textit{if } p < \frac{1}{2}, \quad \mathbb{E}\tau^b \leqslant \frac{C}{q-p} - \frac{p(1-\left(\frac{p}{q}\right)^C)}{\left(q-p\right)^2}.$$

Figure 5 displays both the exact expected coupling time for a queue with capacity 10 as given by Equation (19) as well as the three explicit bounds given in Proposition 15. Note that the bounds are very accurate under light or heavy traffic ($q \leqslant 0.4$ and $q > 0.6$). In any case, the ratio is never larger than 1.2.

5 Coupling in acyclic queueing networks

This section is dedicated to the effective computation of a bound of the coupling time in acyclic networks.

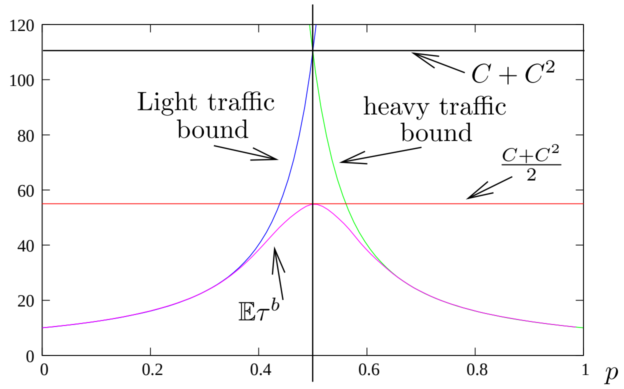

Figure 5: Expected coupling time in an M/M/1/10 queue when q varies from 0 to 1 and the three explicit bounds given in Proposition 15

If one gives a close look to the coupling time for the acyclic network given in Figure 1, one may see in Figure 2 that the coupling time has a peak when $\lambda_0 = 0.4$. This corresponds to the case when the input rate and service rate in Queue 3 are equal. This should not be surprising regarding the result for a single queue, which says that the coupling time is maximal when the rates are equal. Then a second peak occurs around $\lambda_0 = 1.4$ when coupling in Queue 0 is maximal. The rest of the curve shows a linear increase of the coupling time which may suggest an asymptotic linear dependence in λ_0. In this part, an explicit bound on the coupling time which exhibits these two features will be derived.

The first result concerns an extension of inequality (9) to distributions. The second part shows how the results for a single M/M/1/C queue can be used to get an effective computation of bounds for acyclic networks on queues.

5.1 The distribution of the coupling time in acyclic networks

In the following, the queues $Q_0, \dots Q_K$ are numbered according to the topological order of the network. Thus, no event occurring in queue Q_i has any influence on the state of queue Q_j as long as $i > j$.

Proposition 16. *The coupling time for an acyclic network is bounded in the stochastic sense by the sum of the forward coupling time of all queues:*

$$\tau^b \leqslant_{st} \tau_K^f + \cdots + \tau_0^f.$$

Proof. The proof is based on the following idea: construct a trajectory of a backward simulation over which the comparison holds. This will imply the stochastic comparison using Strassen's Theorem.

Consider a backward simulation of the network starting at time 0 until coupling occurs for the last queue, at time $-\tau_K^b$. From time $-\tau_K^b$, run a backward simulation until Queue Q_{K-1} couples. From time $-\tau_l^b - \tau_{K-1}^b$, run the backward simulation again until Queue Q_{K-2} couples. Continue this construction until the first queue has coupled at time $-\tau_K^b + \cdots + \tau_0^b$. Now, on this trajectory, the state in queue Q_0 has coupled between times $-(\tau_K^b + \cdots + \tau_0^b)$ and $-(\tau_K^b + \cdots + \tau_2^b)$. From this time on, Q_0 will remain coupled since no event in other queues may alter its state. The same property holds for queue Q_i between times $-(\tau_K^b + \cdots + \tau_i^b)$ and $-(\tau_K^b + \cdots + \tau_{i+1}^b)$, and at time 0, all queues have coupled. Finally, note that the intervals of this simulation are independent of each other so that $\sum_i \tau_i^b = \sum_i \tau_i^f$ in distribution and one gets $\tau^b \leqslant_{st} \tau_K^f + \cdots + \tau_0^f$. $\square$

This result calls for several comments. First, note that acyclicity is essential in the proof above. For networks with cycles, one would need some kind of association properties of the states of the

queues to assess something about the comparison of the distribution of τ^b and the τ_i^f's.
Second, the technique used in Section 4 to get explicit bounds for the expectation of the forward coupling time can also be used to derive an explicit form for the generating function of bounds of the distribution of the forward coupling time for each queue. Using the proposition above, this gives a generating function of a bound of the coupling time for an acyclic network of queues, up to a convolution formula.

5.2 Computation of an upper bound on the coupling time

Here, an acyclic network of ./M/1/C queues with an arbitrary topology and Bernoulli routings is considered. The events here are of only two types: exogenous arrivals (Poisson with rate γ_i in queue i) and routing of one customer from queue i to queue j after service completion in queue i (with rate μ_{ij}). Queue $K+1$ is a dummy queue representing exits: routing a customer to queue $K+1$ means that the customer exits the network forever. In case of overflow, the new customer trying to enter the full queue is lost. The service rate at queue i is also denoted $\mu_i = \sum_{i=0}^{K+1} \mu_{ij}$.

Let us introduce new random variables. $\tau^b(s_j = x)$ is the backward coupling time of the network, over the set of all intial states with the j-th coordinate equal to x. Namely,

$$\tau^b(s_j = x) = \min\left\{n \text{ s.t. } \left|\phi^{(n)}\left(\mathcal{S} \cap \{s_j = x\}, e_{-n+1}, \dots, e_0\right)\right| = 1\right\}.$$

$\tau_i^b(s_j = x)$ is the backward coupling time on coordinate i given $s_j = x$:

$$\tau_i^b(s_j = x) = \min\left\{n \text{ s.t. } \left|N_i\left(\phi^{(n)}\left(\mathcal{S} \cap \{s_j = x\}, e_{-n+1}, \dots, e_0\right)\right)\right| = 1\right\}.$$

It should be obvious that $\tau^b(s_j = x) \leqslant_{st} \tau^b$ and for all i, $\tau_i^b(s_j = x) \leqslant_{st} \tau_i^b$.
We also have the same notions for forward coupling times:

$$\tau^f(s_j = x) = \min\left\{n \in \mathbb{N};\ \text{s.t. } \left|\phi^{(n)}\left(\mathcal{S} \cap \{s_j = x\}, e_{1\to n}\right)\right| = 1\right\},$$

$\tau_i^f(s_j = x)$ being defined in the same manner, and for hitting times:

$$h_{C_i \to 0}(s_j = x) = \min\{n \in \mathbb{N};\ \text{s. t. } \phi^{(n)}\left(\mathcal{S} \cap \{s_i = C_i, s_j = x\}, e_{1\to n}\right) \in \mathcal{S} \cap \{s_i = 0\}\}.$$

Now, sweeping the list of queues in the topological order, one can construct a sequence of backward simulations in the following way.

First simulate the queueing system from the past up to coupling of queue 0. The number of steps is by definition τ_0^b. Queue Q_0 has coupled in a random state X_0. Then, run a second backward simulation up to coupling of Queue Q_1 given $s_0 = X_0^0$. This simulation takes $\tau_i^b(s_0 = X_0^0)$ steps and the state at time $t = 0$ is X_1^1 for Q_1 and X_0^1 for Q_0.

This construction goes on up to the backward simulation up to coupling of Queue Q_K given $s_0 = X_0^{K-1}, s_1 = X_1^{K-1}, \dots, s_{K-1} = X_{K-1}^{K-1}$. The last simulation takes $\tau_i^b(s_0 = X_0^{K-1}, s_1 = X_1^{K-1}, \dots, s_{K-1} = X_{K-1}^{K-1})$ steps and the coupling state of Q_K is X_K^K.

Lemma 17. *Using the previous construction,*

$$\tau^b \leqslant_{st} \sum_{i=0}^{K} \tau_i^b(s_0 = X_0^{i-1}, \dots, s_{i-1} = X_{i-1}^{i-1}),$$

and for all i, $(X_0^i, \dots, X_i^i)$ is steady state distributed for $Q_0, \dots, Q_i$. Furthermore, for all i,

$$\tau^b \quad \leqslant_{st} \quad \sum_{i=0}^{K} h_{C_i \to 0}(s_0 = X_0^{i-1}, \dots, s_{i-1} = X_{i-1}^{i-1}),$$

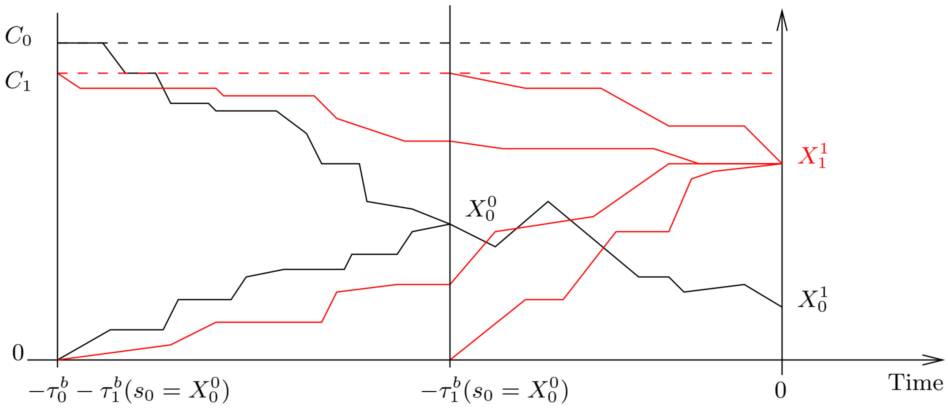

Figure 6: The trajectories of the state in Q_0 are in black while the the trajectories for Q_1 are in the lighter color. Starting at time $-\tau_0^b - \tau_1^b(s_0 = X_0^0)$, the state of Q_0 has coupled in X_0^0 at time $-\tau_1^b(s_0 = X_0^0)$. From then on, Q_0 stays coupled and Q_1 couples at time some time before 0.

Proof. From the previous sequence of backward simulations one can construct a single one by appending them in the reverse order: the backward simulation for Queue Q_K preceded by the simulation of Q_{K-1}, and so forth up to the simulation of Q_0. This is a backward simulation of the system (the last state is $(X_0^K, \ldots, X_i^K)$). This construction is illustrated in the case of two queues in tandem in Figure **??**.

A straightforward consequence, using acyclicity, is that $(X_0^i, \ldots, X_i^i)$ is steady state distributed for $Q_0, \ldots, Q_i$ for all i.

Furthermore, one gets in distribution

$$\begin{aligned}
\tau^b \quad &\leqslant_{st} \quad \sum_{i=0}^{K} \tau_i^b(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1}) \\
&= \quad \sum_{i=0}^{K} \tau_i^f(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1}) \\
&\leqslant_{st} \quad \sum_{i=0}^{K} h_{C_i \to 0}(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1}),
\end{aligned}$$

by independence of the variables given the initial states X^{i-1}.

□

Let us now consider a new circuit with one difference from the original one: all queues are replaced by infinite queues, except for queue Q_i which stays the same. In the following, all the notations related to this new network will be expressed by appending the ∞ symbol to all variables corresponding to this new circuit.

The new circuit up to Queue i is product form and using Burke's Theorem, the input stream in Queue i is Poisson. The rate of the input stream in queue i is given by ℓ_i, the solution of the flow equations:

$$\ell_i = \sum_{j<i} \ell_j \frac{u_{ji}}{\mu_j} + \gamma_i.$$

The network is said to be *stable* for Queue i as soon as $\ell_i < \mu_i$. We assume stability for all i in the following.

One can construct a sequence of backward simulations for the new network in the same way as for the original network. This provides the quantities ${}^\infty X_j^{i-1}$, ${}^\infty \tau_i^b(s_0 = {}^\infty X_0^{i-1}, \ldots, s_{i-1} = {}^\infty X_{i-1}^{i-1})$, ${}^\infty \tau_i^f(s_0 = {}^\infty X_0^{i-1}, \ldots, s_{i-1} = {}^\infty X_{i-1}^{i-1})$, and ${}^\infty h_{C_i \to 0}(s_0 = {}^\infty X_0^{i-1}, \ldots, s_{i-1} = {}^\infty X_{i-1}^{i-1})$.

The monotony property given above implies that $X_i^j \leqslant_{st} {}^{\infty}X_i^j$ and

$$h_{C_i\to 0}(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1}) \leqslant_{st} {}^{\infty}h_{C_i\to 0}(s_0 = {}^{\infty}X_0^{i-1}, \ldots, s_{i-1} = {}^{\infty}X_{i-1}^{i-1}).$$

The next step is to build yet another model. This third model is made of a single $M/M/1/C_i$ queue with three types of events, arrivals of customers with rate ℓ_i (provided that the number of customers is smaller than C_i), departures with rate μ_i (provided that the number of customers is positive) and null events (with no effect on the queue) with rate $\Lambda - \ell_i - \mu_i$.

For this isolated model, let us introduce the uniformizing probabilities $p = \ell_i/\Lambda$, $q = 1-p$ and $d = (\Lambda - \ell_i - \mu_i)/\Lambda$. Let F_k be the time to go from state k to state 0 in the isolated system. A one step analysis gives

$$\begin{aligned}\mathbb{E}[F_k] &= 1 + d\mathbb{E}[F_k] + \frac{\ell_i}{\Lambda}\mathbb{E}[F_{(k+1)\wedge C_i}] + \frac{\mu_i}{\Lambda}\mathbb{E}[F_{(k-1)}] \\ &= \frac{1}{1-d} + p\mathbb{E}[F_{(k+1)\wedge C_i}] + q\mathbb{E}[F_{(k-1)\vee 0}].\end{aligned}$$

We get the same equation as (23) except for the additional constant which is now $\frac{1}{1-d}$ instead of 1, so that the solution is the same as before up to a multiplicative factor of $\frac{1}{1-d} = \frac{\Lambda}{\ell_i+\mu_i}$. Using Equation (25), one gets

$$\mathbb{E}[F_{C_i}] = \frac{\Lambda}{\ell_i + \mu_i}\left(\frac{C_i}{q-p} - \frac{p(1-\left(\frac{p}{q}\right)^{C_i})}{(q-p)^2}\right). \tag{26}$$

Lemma 18. *Under the foregoing notations and assumptions,*

$${}^{\infty}h_{C_i\to 0}(s_0 = {}^{\infty}X_0^{i-1}, \ldots, s_{i-1} = {}^{\infty}X_{i-1}^{i-1}) = F_{C_i},$$

in distribution.

Proof. First, using Lemma 17 for the new network with infinite queues (except for Q_i), the state $({}^{\infty}X_0^{i-1}, \ldots, {}^{\infty}X_{i-1}^{i-1})$ is steady state distributed. Using Burke's Theorem, this implies that the input stream in queue Q_i is Poisson with rate ℓ_i, when one runs a simulation starting in any state in $\mathcal{S} \cap \{s_i = C_i, s_j = {}^{\infty}X_j^{i-1}, j < i\}$.

Now, during this simulation, one can couple the addition, subtraction et null events for queue Q_i in isolation and for Q_i in the complete network of infinite queues, all of them having the same laws. This implies that the state of queue Q_i in both systems is the same under that coupling. Hence, they reach 0 at the same time: ${}^{\infty}h_{C_i\to 0}(s_0 = {}^{\infty}X_0^{i-1}, \ldots, s_{i-1} = {}^{\infty}X_{i-1}^{i-1}) = F_{C_i}$ in distribution. □

We are ready to put everything together in expectation.

$$\begin{aligned}\mathbb{E}\tau^b &\leqslant_{st} \sum_i \mathbb{E}[h_{C_i\to 0}(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1})] &(27)\\ &\leqslant \sum_i \mathbb{E}[{}^{\infty}h_{C_i\to 0}(s_0 = {}^{\infty}X_0^{i-1}, \ldots, s_{i-1} = {}^{\infty}X_{i-1}^{i-1})] &(28)\\ &\leqslant \sum_i \mathbb{E}[F_{C_i}]. &(29)\end{aligned}$$

The sequence of inequalities may not hold in distribution because the variables X^i and thus $h_{C_i\to 0}(s_0 = X_0^{i-1}, \ldots, s_{i-1} = X_{i-1}^{i-1})$ are not independent.

Using (26),

$$\mathbb{E}\tau^b \leqslant \sum_i \frac{\Lambda}{\ell_i + \mu_i}\left(\frac{C_i}{q-p} - \frac{p(1-\left(\frac{p}{q}\right)^{C_i})}{(q-p)^2}\right).$$

The result of this part is summarized in the following theorem.

Theorem 19. *In an acyclic stable network of $K+1$ $./M/1/C_i$ queues with Bernoulli routing and losses in case of overflow, the coupling time from the past satisfies in expectaction,*

$$\mathbb{E}[\tau^b] \leqslant \sum_{i=0}^{K} \frac{\Lambda}{\ell_i+\mu_i}\left(\frac{C_i}{q-p} - \frac{p(1-\left(\frac{p}{q}\right)^{C_i})}{(q-p)^2}\right) \leqslant \sum_{i=0}^{K} \frac{\Lambda}{\ell_i+\mu_i}(C_i+C_i^2). \tag{30}$$

Note that this bound on the expectation is ultimately linear in the rate of any event in the system. This behavior is also noticeable for $\mathbb{E}[\tau^b]$ itself.

5.3 Some numerical experiments

In the construction of the bound given in Theorem 19, several factors may be responsible for the inaccuracy of the bound.

1. The first factor is the replacement of the max by the sum. We believe that it may be a hard task to get rid of this first approximation because of the intricate dependencies between the queues. Furthermore, experiments reported below show that this may not even be possible in many cases (see Figure 8).

2. Another factor which may increase the inaccuracy of our bounds is the fact that most events change the states of several queues at the same time, while the bound given here disregards this. In the network studied here, this may add a factor 2 between the true coupling time and the bound given in Theorem 19.

3. The most important factor which jeopardizes the quality of the bound is the stability issue. If one queues is unstable, the bound provided by Equation (31), also called the light traffic bound in Proposition 15 is very bad (as seen in Figure 5). So far we have not been able to come up with a better bound for unstable queues. However, when all queues are stable (and even more so when the load is smaller than 2/3), the bound tends to be more accurate. This is further verified in the experiments reported below.

Computations for the network displayed in Figure 1 are reported in Figures 7, 8 and 9. We have used the following parameters. The input rate is $\lambda_0 = 0.4$. the rates of the other events are $\lambda_1 = 1.4$, $\lambda_2 = 0.6$, $\lambda_3 = 0.8$, $\lambda_4 = 0.5$. The number of simulation runs is 10000. The capacity C is the same in all queues, and we let it vary from 1 to 20. The service rate in the last queue λ_5 takes three values, respectively $0.2, 0.6$ and 0.4.

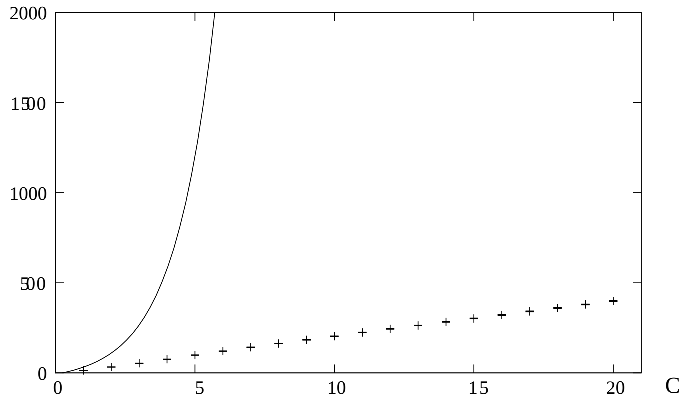

Figure 7: This figure displays the coupling time (dots) with 95% confidence intervals, and the bound given by Equation (31) when Queue Q_3 is unstable ($\lambda_5 = 2/10$), while the capacity C varies from 1 to 20.

In the first case (Fig. 7) , $\lambda_5 = 0.2$ so that queue Q_3 is unstable. Figure 7 displays the bound given by formula (31) as well as the mean coupling time computed over 10000 simulation runs. As hinted before, the bound is indeed very bad for the unstable system. A ratio larger than 10 w.r.t the true coupling time is reached when $C = 5$. It should also be noticed that the bound is convex in C while the coupling time is not.

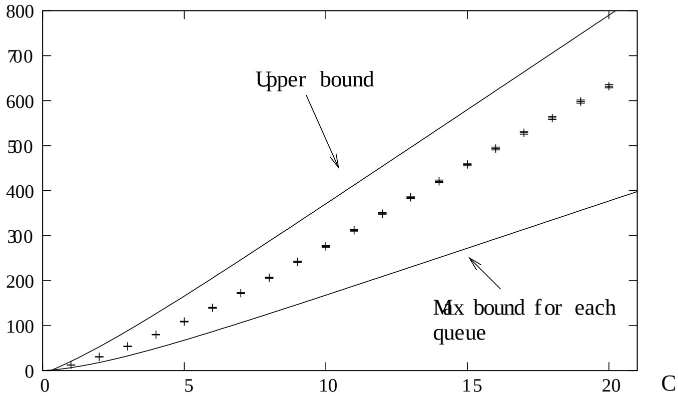

Figure 8: Here are the bound given by Equation (31), the mean coupling time (dots) with 95% confidence intervals and the maximum over Equations (31) for all queues, when Queue Q_3 is stable ($\lambda_5 = 6/10$), while the capacity C varies from 1 to 20.

In the second case (Fig. 8) , $\lambda_5 = 0.6$, and all queues are stable with a load smaller that 2/3. Figure 8 shows the bound provided by (31) and the true coupling time computed by simulation runs. Both curves appear to be almost linear in C (this is true for the bound: when q/p is small, $\mathbb{E}H_{C_i,0}$ is almost linear in C_i) and the ratio is smaller than 1.3. In that case, the curve $\max_{i\in\{0,\ldots K\}} \mathbb{E}H_{C_i,0}$ is also displayed and is below the actual coupling time. This is to be related with the first item in the comments above.

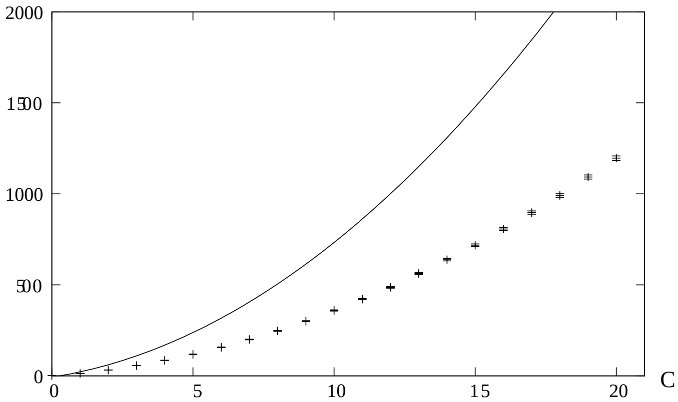

Figure 9: Display of the coupling time (dots) with 95% confidence interval and the bound given by Equation (31) when Queue Q_3 is barely unstable ($\lambda_5 = 4/10$) while the capacity C varies from 1 to 20.

The last case (Fig. 9) is for $\lambda_0 = 0.4$, so that Q_3 is barely unstable. This would correspond to the maximal coupling time for Q_3 if it were alone. Figure 9 displays the backward coupling time and the bound provided by Equation (31). For queue Q_3, we use a bound in $C_3 + C_3^2$ which is a bad approximation because of the loss of the factor 2 when compared with the bound for isolated queues. Note that the total gap has a ratio which is almost 2. In that case bothe the coupling time and the bound exhibit a convex behavior w.r.t. C.

A ratio smaller than 2 is indeed interesting because efficient perfect simulation algorithm use a doubling window technique to reduce the complexity and their running time (see Equation (4)) so that our bound gives a good estimation of the mean running time of the algorithms.

One should also note that, on a practical point of view, most actual networks which require stationary performance evaluations are indeed stable.

5.3.1 Extension to more general networks

Actually, extensive simulation runs over many examples show that the bound given in Theorem 19 is robust and also holds for more general networks with blocking and with circuits. While we have only been able to show that the light traffic bound holds for each queue, we conjecture that the heavy traffic bound and the critical bound should also hold. This would yield an overall quadratic bound: $\mathbb{E}[\tau^b] \leqslant \sum_{i=0}^{K} \frac{\Lambda}{\ell_i+\mu_i} O(C_i^2)$, for any monotone Markovian network of queues with a finite state space. Furthermore under light or heavy traffic in all queues, the bound should rather be linear: $\mathbb{E}[\tau^b] \leqslant \sum_{i=0}^{K} \frac{\Lambda}{\ell_i+\mu_i} O(C_i)$.

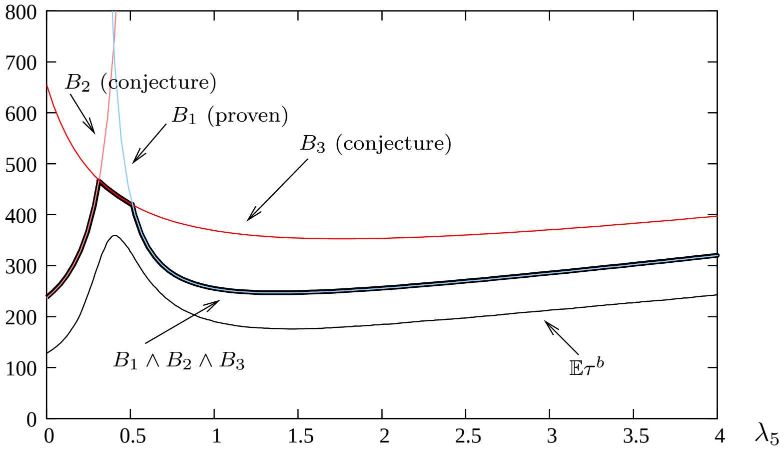

Figure 10: This figure displays the actual coupling time $\mathbb{E}\tau^b$ together with the proven light traffic bound B_1, the conjectured heavy traffic bound B_2, the conjectured critical bound B_3 and the minimum of the three bounds.

To illustrate this conjecture, we have run simulations for the network displayed in Figure 1 with the following parameters. The rates are $\lambda_0 = 0.4, \lambda_1 = 1.4$, $\lambda_2 = 0.6$, $\lambda_3 = 0.8$, $\lambda_4 = 0.5$. The capacity is fixed to 10 in all queues and we let λ_5 (the service rate in Q_3) vary from 0 to 4. As long as $\lambda_5 < 0.4$, Q_3 is unstable and our proven bound (B_1) is poor. As soon as λ_5 is large enough our bound becomes acceptable. In Figure 10, note that both the bound and the coupling time τ have a linear asymptotic growth in λ_5. The Figure also displays the heavy traffic bound B_2 and the critical bound B_3. Should these two bounds hold, the minimum of B_1, B_2, B_3 (in bold in the figure) would provide a remarkable bound on the coupling time, up to an additional constant. This issue is the subject of our current investigations.

Acknowledgments

The authors would like to thank Jérôme Vienne who partially designed the `psi2` perfect simulation software which was used to run all the simulations presented here.

References

[1] C. Alexopoulos and D. Goldsman, *To batch or not to batch?*, ACM Trans. Model. Comput. Simul., 14 (2004), pp. 76–114.

[2] P. Glasserman and D.D. Yao, *Monotone Structure in Discrete-Event Systems*, Wiley Inter-Science, Series in Probability and Mathematical Statistics, 1994.

[3] P. Glynn, *Initial transient problem for steady-state output analysis*, in Winter Simulation Conference, 2005.

[4] O. Häggström, *Finite Markov Chains and Algorithmic Applications*, Cambridge University Press, 2002.

[5] S.G. Henderson and P. Glynn, *Regenerative steady-state simulation of discrete-event systems*, ACM Trans. Model. Comput. Simul., 11 (2001), pp. 313–345.

[6] K Pawlikowski, *Steady-state simulation of queueing processes: survey of problems and solutions*, ACM Comput. Surv., 22 (1990), pp. 123–170.

[7] J. Propp and D. Wilson, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Random Structures and Algorithms, 9 (1996), pp. 223–252.

[8] S. M. Ross, *Probability models*, Academic Press, 2003.

[9] J.-M. Vincent, *Perfect simulation of monotone systems for rare event probability estimation*, in Winter Simulation Conference, Orlando, dec 2005.

[10] ——, *Perfect simulation of queueing networks with blocking and rejection*, in Saint IEEE conference, Trento, 2005, pp. 268–271.

[11] J.-M. Vincent and C. Marchand, *On the exact simulation of functionals of stationary markov chains*, Linear Algebra and its Applications, 386 (2004), pp. 285–310.

[12] W. Whitt, *The efficiency of one long run versus independent replication in steady-state simulation*, Manage. Sci., 37 (1991), pp. 645–666.

Unité de recherche INRIA Rhône-Alpes
655, avenue de l’Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)
Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)
Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399