Exact Simulation of the Stationary Distribution of M/G/c Queues

Professor Karl Sigman
Columbia University
New York City
USA

Conference in Honor of Søren Asmussen
Monday, August 1, 2011
Sandbjerg Estate

Outline

We will present two different algorithms for simulating (exactly) from the stationary distribution of customer delay for the stable ($\rho = \lambda/\mu < c$) FIFO M/G/c queue. (c servers in parallel, Poisson arrivals, iid service times.)

Our first algorithm is for the special case when $\rho = \lambda/\mu < 1$ (super stable case). This algorithm involves the general method of dominated coupling from the past (DCFTP) and we use the single-server queue operating under the *processor sharing (PS)* discipline as a sample-path upper bound. The algorithm is shown to have finite expected termination time if and only if service times have finite second moment.

Outline

Our second algorithm is for the general case of $\rho < c$. Here we use discrete-time processes and regenerative simulation methods, in which as regeneration points, we use return visits to state 0 of a corresponding random assignment (RA) model which serves as a sample-path upper bound.

Both algorithms yield, as output, a stationary copy of the entire Kiefer-Wolfowitz workload vector.

Muti-server queue (1)

Here we consider the FIFO M/G/c queue. Poisson arrivals at rate λ, iid service times S with general distribution $G(x) = P(S \leq x)$, mean $E(S) = 1/\mu$, and the stability condition $\rho = \lambda/\mu < c$. When $c = 1$, this is the classic "M/G/1" queue and it has a stationary distribution for customer delay D, that is known via the *Pollaczek-Kintchine* formula (Laplace transform of D):

$$E^{-sD} = \frac{1-\rho}{1-\rho E(e^{-sS_e})}, \; s \geq 0,$$

where S_e is distributed as G_e, the equilibrium distribution of service, it has density $\mu P(S > x)$.

Muti-server queue (2)

This implies that D can be expressed (in distribution) as a geometric sum of iid S_e rvs:

$$D = \sum_{j=1}^{L} Y_j, \tag{1}$$

where the $\{Y_j\}$ are iid distributed as the equilibrium distribution of service, with cumulative distribution function given by $G_e(x) = \mu \int_0^x P(S > y)dy,\ x \geq 0$, and independently L has a geometric distribution, $P(L = k) = \rho^k(1 - \rho),\ k \geq 0$.

Muti-server queue (3)

It is reasonable to assume that we could simulate from both G and G_e, and of course we can simulate a geometric rv. Thus we have an exact simulation algorithm under such assumptions:

Algorithm for simulating D for the FIFO M/G/1 queue

1. Generate L geometrically distributed with parameter ρ.
2. If $L = 0$, set $D = 0$. Otherwise generate L iid copies $Y_1, \dots, Y_L$ distributed as G_e, and set

$$D = \sum_{j=1}^{L} Y_j.$$

Muti-server queue (3)

When $c \geq 2$, no such formula for D is known. At arrival epochs t_n with iid interarrival times $T_n = t_n - t_{n-1}$ ($t_0 \stackrel{\text{def}}{=} 0$), define the vector $\mathbf{W}_n$ defined recursively by

$$\mathbf{W}_n = R(\mathbf{W}_{n-1} + S_n \mathbf{e} - T_n \mathbf{f})^+, \; n \geq 1, \tag{2}$$

where $\mathbf{W}_n = (W_n(1), \ldots, W_n(c))$, $\mathbf{e} = (1, 0, \ldots 0)$, $\mathbf{f} = (1, 1, \ldots, 1)$, R places a vector in ascending order, and $^+$ takes the positive part of each coordinate. $D_n = W_n(1)$ is then customer delay in queue (line) of the n^{th} customer. This is called the *Kiefer-Wolfowitz workload vector* and when $\rho < c$ it is known to converge to a unique stationary distribution π. (It is notoriously complicated to analyze π.)

Muti-server queue (4)

In continuous time, the Kiefer-Wolfowitz workload vector is denoted by

$$\mathbf{V}(t) = (V(1,t), V(2,t), \ldots, V(c,t)), \ t \geq 0,$$

and

$$\mathbf{W}_n = \mathbf{V}(t_n-), \ n \geq 0,$$

the workload found by the n^{th} customer (not including their own service time).
When arrivals are Poisson (as we are assuming), PASTA implies that the two processes have the same stationary distribution, π. So we can, and will, work in continuous time instead of discrete time.

Muti-server queue (4)

We take a from the past stationary version,

$$\{\mathbf{V}(t) : \ t \le 0\},$$

and our objective is to simulate a copy of $\mathbf{V}(0) \sim \pi$.

We shall assume that $\rho < 1$, the system is *super stable*: the corresponding M/G/1 queue will also be stable.

Muti-server queue (5)

Lemma

Let $V_1(t)$ denote total work in system at time t for the FIFO M/G/1 queue, and let $V_c(t) = \sum_{i=1}^{c} V(i,t)$ denote total work in system at time t for the corresponding FIFO $M/G/c$ queue, where $V_1(0) = V_c(0) = 0$ and both are fed exactly the same input of Poisson arrivals and iid service times. Then

$$P(V_c(t) \leq V_1(t),\ \text{for all } t \geq 0) = 1. \tag{3}$$

(Workload is defined as the sum of all remaining service times in the system.)

Muti-server queue (6)

KEY IDEA:

1. If we were to start off both the $c-$server and the single-server models empty at time $t = -\infty$ while feeding them exactly the same input, then both would have their stationary distributions at time 0 and their workload would be ordered at all times due to the Lemma.
2. Moreover, if we walk backwards in time from the origin, and detect the first time $-\tau \leq 0$ at which the single-server model is empty, then from the Lemma, the $c-$server model would be empty as well.
3. We then could construct a sample of $\mathbf{V}(0)$ (having the stationary distribution π) by starting off empty at time $-\tau$ and using the Kiefer-Wolfowitz recursion forwards in time from time $-\tau$ to 0.

Muti-server queue (7)

We now proceed to show how to accomplish this. The main problem is how to “walk backwards in time" in stationarity for the single-server model, how do we do that?

Muti-server queue (8)

Outline of the approach/solution:

1. Workload for single-server queues is invariant under changes of disciplines: FIFO, LIFO, Processor-sharing (PS), pre-emptive LIFO, random choice, etc., all have exactly the same sample-paths for workload $\{V_1(t) : t \geq 0\}$.
2. Under PS, it is known that :

$$\{\mathbf{X}(t) : t \geq 0\} = \{(L(t), Y_1(t), \ldots, Y_{L(t)}(t)) : t \geq 0\},$$

   where $L(t)$ denotes number of customers in service at time t, and $Y_i(t)$ their remaining service times, is a Markov process with stationary distribution $(L, Y_1, \ldots, Y_L)$ exactly as from the Pollaczek-Kintchine formula; L is geometric with parameter ρ, and the Y_i are iid G_e.

Muti-server queue (9)

3. When started with its stationary distribution (which we know how to simulate from), the time-reversal of this Markov process is the Markov process representing this same PS model, except the Y_i are now the *ages* of the service times: It too has Poisson arrivals, and iid $\sim G$ service times. (This means that the departure process of the PS model (when stationary) is Poisson at rate λ.)
4. Thus we can simulate the time-reversal PS model until it empties, all the while recording the departure times and the service times attached to those departures. We then feed these service times and interarrival times back into an initially empty multi-server model *forward* in time-using the Kiefer-Wolfowitz recursion-to construct $\mathbf{V}(0)$.

Muti-server queue (10)

Algorithm for simulating $\mathbf{V}(0)$ distributed as π

1. Set $t = 0$ (time). Generate a vector $(L, Y_1, \ldots, Y_L)$ distributed as the stationary distribution and set $\mathbf{X}(0) = (L(0), Y_1(0), \ldots, Y_{L(0}(0)) = (L, Y_1, \ldots, Y_L)$.
2. If $L = 0$, then stop simulating and set $\tau = 0$. Otherwise, continue to simulate (as a discrete-event simulation with iid interarrival times $T \sim exp(\lambda)$ and iid service times $S \sim G$) the time-reversed PS model until time $\tau = \min\{t \geq 0 : L(t) = 0\}$: Each of the $L > 0$ customers' service times are to be served simultaneously at rate $r = 1/L$ until the time of the *next event*, either a new arrival or a departure; reset $t =$ this new time.

Muti-server queue (11)

3. If the next event is an arrival, then generate a service time S for this customer distributed as G (keep a record of its value and place it in service), generate the next interarrival time T distributed as $exp(\lambda)$ and reset $L = L + 1$, set $r = 1/L$.
4. If the next event is a departure, then record this as the next departure time and record the service time of the customer associated with it and reset $L = L - 1$. If $L = 0$, then stop simulating, set $\tau = t$.
5. If $\tau > 0$ after stopping the simulation, then let $t_1, \dots, t_k$ and $S_1, \dots, S_k$ denote the $k \geq 1$ recorded departure times (in order of departure), and the associated service times, that occurred up to time τ (with $t_k = \tau$ the last such departure time). Define the interdeparture times $T_i = t_i - t_{i-1}$, $0 \leq i \leq k$, with $t_0 = 0$.

Muti-server queue (12)

6. We now construct $\mathbf{V}(0)$ as follows: If $\tau = 0$, then set $\mathbf{V}(0) = \mathbf{0}$. Otherwise: Reset $(S_1, \ldots, S_k) = (S_k, \ldots, S_1)$ and $(T_1, \ldots, T_k) = (T_k, \ldots, T_1)$ (that is, place them in *reverse order*). (They have the *conditional* distribution of iid input given τ resulted in k departures, so they are no longer iid.) Using $(S_1, \ldots, S_k)$ and $(T_1, \ldots, T_k)$ as the input, construct $\mathbf{W}_k$ (initializing with $\mathbf{W}_0 = \mathbf{0}$), by using the Kiefer-Wolfowitz recursion from $n = 1$ up until $n = k$. Now set $\mathbf{V}(0) = \mathbf{W}_k$.

Final comments

1. $E(\tau) < \infty$ if and only if $E(S^2) < \infty$ because τ (given it is > 0) is the stationary excess (equilibrium) distribution of a M/G/1 busy period B. $E(\tau) = \rho E(B_e) = \rho E(B^2)/2E(B)$.
2. At time $t = 0$, we actually need both the stationary remaining service times and their ages. This is because the customers in service at time $t = 0$ have (via the inspection paradox) service times (age plus excess) distributed as the *spread distribution*. If G has a density $g(x)$, then the spread has density $h(x) = \mu x g(x)$.
3. Our method of using the PS queue also works for exactly simulating the stationary distribution of general networks with iid routes, Poisson arrivals and c FIFO single-server stations; but we must have the harsh condition that $\rho < 1$. If $((i_1, S(1)), (i_2, S(2)), \ldots, (i_K, S(K))$ denotes a route of random length $K \geq 1$ and we define $S = \sum_{i=1}^{K} S(i)$, then total work brought by a customer is $E(S)$ and we define $\rho = \lambda E(S)$.

Muti-server queue-B (1)

Now we present our second algorithm and we allow for any $\rho < c$.

$$\mathbf{W}_{n+1} = R(\mathbf{W}_n + S_n\mathbf{e} - A_n\mathbf{f})^+, \; n \geq 0, \tag{4}$$

Simulating the stationary distr. of a reg. proc. (1)

Suppose that $\mathbf{X} = \{X(t) : t \geq 0\}$ is a positive recurrent non-delayed regenerative process, with iid cycle lengths generically denoted by T distributed as $F(x) = P(T \leq x),\ x \geq 0$ with finite and non-zero mean $E(T) = 1/\lambda$. A generic length T cycle is thus $C = \{X(t) : 0 \leq t < T\}$. From regenerative process theory, the (marginal) stationary distribution π is given by (expected value over a cycle divided by the expected cycle length)

$$\pi(\cdot) = \lambda E \int_0^T I\{X(t) \in \cdot\} dt. \tag{5}$$

Simulating the stationary distr. of a reg. proc. (2)

Due to Assmusen, Glynn and Thorisson (1992):

Proposition

1. *Suppose we can and do sequentially simulate iid copies of* $C = \{X(t) : 0 \le t < T\}$ *(the first cycle), denoted by* $C_n = \{X_n(t) : 0 \le t < T_n\},\ n \ge 1$*, having iid cycle lengths* $\{T_n\}$ *distributed as* F*.*
2. *Suppose further that we can and do simulate (independently) one copy* T^e *distributed as the equilibrium distribution having density function* $f_e(t) = \lambda P(T > t) = \lambda \bar{F}(t),\ t \ge 0$*.*
3. *Let* $\tau = \min\{n \ge 1 : T_n \ge T^e\}$*.*
4. *Use cycle* C_τ *to construct* $X^* = X_\tau(T^e)$

Then the simulated random element X^* *is distributed as* π*.*

Simulating the stationary distr. of a reg. proc. (3)

Proof.

Conditional on $T^e = t$, it holds that $\tau = \min\{n \geq 1 : T_n > t\}$, and thus C_τ simply has the distribution of a first cycle given that its length is larger than t:

$$P(X^* \in \cdot \mid T^e = t) = P(X(t) \in \cdot \mid T > t) = \frac{P(X(t) \in \cdot\,,\ T > t)}{\bar{F}(t)}.$$

□

Simulating the stationary distr. of a reg. proc. (4)

Proof.

Since T^e has density $f_e(t) = \lambda \bar{F}(t)$, we obtain

$$
\begin{aligned}
P(X^* \in \cdot) &= \int_0^\infty \frac{P(X(t) \in \cdot\,,\; T > t)}{\bar{F}(t)} \lambda \bar{F}(t) dt \\
&= \lambda \int_0^\infty P(X(t) \in \cdot\,,\; T > t) dt \\
&= \lambda E \int_0^T I\{X(t) \in \cdot\} dt \\
&= \pi(\cdot).
\end{aligned}
$$

□

Simulating the stationary distr. of a reg. proc. (5)

Proposition 1 remains valid in a discrete-time setting too in which case the density of T^e is replaced by the probability mass function $P(T^e = n) = \lambda P(T \geq n)$ on the positive integers $n \geq 1$.

Random Assignment (RA) model (1)

Given a $c-$server queueing model, the *random assignment model (RA)* is the case when each of the c servers forms its own FIFO single-server queue, and each arrival to the system, independent of the past, randomly chooses queue i to join with probability $1/c,\ i \in \{1, 2, \ldots c\}$. In the M/G/c case, we refer to this as the RA M/G/c model.

Random Assignment (RA) model (2)

The following is a special case of Lemma 1.3, Page 342 in [1]. (Such results and others even more general are based on the early work (1979, 1980) of S. Foss and R. Wolff.)

Lemma

Let $Q_F(t)$ denote total number of customers in system at time $t \geq 0$ for the FIFO M/G/c queue, and let $Q_{RA}(t)$ denote total number of customers in system at time t for the corresponding RA M/G/c model in which both models are initially empty and fed exactly the same input of Poisson arrivals $\{t_n\}$ and iid service times $\{S_n\}$. Assume further that for both models the service times are used by the servers in the order in which service initiations occur (S_n is the service time used for the n^{th} such initiation). Then

$$P(Q_F(t) \leq Q_{RA}(t), \text{ for all } t \geq 0) = 1. \tag{6}$$

Random Assignment (RA) model (3)

The importance of Lemma 2 is that it allows us to jointly simulate versions of the two stochastic processes $\{Q_F(t) : t \geq 0\}$ and $\{Q_{RA}(t) : t \geq 0\}$ while achieving a coupling such that (6) holds. In particular, whenever an arrival finds the RA model empty, the FIFO model is found empty as well. These consecutive epochs in time constitute regeneration points (for both models) due to the iid assumptions on the input. We explain how to use these facts to our advantage next.

Random Assignment (RA) model (4)

$$\mathbf{W}_{n+1} = R(\mathbf{W}_n + S_n\mathbf{e} - A_n\mathbf{f})^+, \ n \geq 0, \tag{7}$$

for the FIFO model defines a Markov chain and for the stable M/G/c case, visits to the empty state **0** form positive recurrent regeneration points. Sure, we can simulate iid cycles, starting with $\mathbf{W}_0 = \mathbf{0}$, but we do not know how to simulate a copy of T^e, equilibrium cycle length. So we can not directly use the Proposition with such regeneration points.

But the RA model too regenerates each time an arrival to it finds an empty system, and since $Q_F(t) \leq Q_{RA}(t), \ t \geq 0,$ these RA regeneration points also serve as regeneration points for the FIFO model.

Random Assignment (RA) model (5)

Letting $\mathbf{Q}_n = (Q_{1,n}, \ldots, Q_{c,n}) = \mathbf{Q}(t_n-)$ denotes the number in system (at each node) as found by the n^{th} arrival to the RA model, we set $\mathbf{Q}_0 = 0$ and define

$$T = \min\{n \geq 1 : \mathbf{Q}_n = 0\}.$$

Random Assignment (RA) model (6)

Moreover, for the RA model, we indeed can simulate an equilibrium cycle length T^e. This is because (as we shall next see), we know how to simulate exactly from the stationary distribution of the RA model.

Random Assignment (RA) model (7)

Letting $\mathbf{V}_n = (V_n(1), \dots, V_n(i))$ denote workload (at each node) as found by the n^{th} arriving customer to the RA model, we have, for each node $i \in \{1, 2, \dots, c\}$,

$$V_{n+1}(i) = (V_n(i) + S_n I\{U_n = i\} - A_n)^+, \; n \geq 0, \tag{8}$$

where here, S_n is an iid service time of the n^{th} (Poisson rate λ) arriving customer, and independently $\{U_n : n \geq 0\}$ denotes an iid sequence of random variables with the discrete uniform distribution over $\{1, 2, \dots, c\}$.

Random Assignment (RA) model (8)

From PASTA, the stationary distribution of $\mathbf{V}_n$ is in fact of the form

$$(D(1), \ldots, D(c)), \tag{9}$$

where the $D(i)$ here are iid distributed as D from Pollaczek-Kintchine:

$$D = \sum_{j=1}^{L} Y_j.$$

Algorithm for simulating T^e

1. Initialize $\mathbf{V}_0 = (D(1), \ldots, D(c))$.
2. Simulate sequentially $\{\mathbf{V}_n : n \geq 1\}$ using the recursion in (8) until time

$$T^e = \min\{n \geq 1 : \mathbf{V}_n = 0\}.$$

Final Algorithm

1. Simulate one copy of T^e.
2. Independently simulate a copy of a first cycle (number in system for RA, coupled with the FIFO model) with corresponding cycle length T.
3. If $T < T^e$, then go back to step (2).
4. Construct the FIFO cycle $C = \{\mathbf{W}_1, \ldots, \mathbf{W}_T\}$. Set $\mathbf{W} = \mathbf{W}_{T^e}$.
5. Output $\mathbf{W}$.

References

- Asmussen, S. (2003). *Applied Probability and Queues* (2nd Ed.). Springer-Verlag, New York.
- S. Asmussen, P. W. Glynn (2007) *Stochastic Simulation*, Springer-Verlag, New York.
- S. Asmussen, P. W. Glynn, & H. Thorisson (1992). Stationary detection in the initial transient problem. *ACM TOMACS*, **2**, 130-157.
- J. Blanchet and K. Sigman (2011) On exact sampling of stochastic perpetuities. *Journal of Applied Probability*. (To appear.)
- W.S. Kendall (2004). Geometric ergodicity and perfect simulation. *Electron. Comm. Probab.* **9**, 140-151.
- J. G. Propp and D.G. Wilson (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, **9**, 223-253.

References

- S. G. Foss (1980). Approximation of mutichannel queueing systems. *Siberian Math. J.*, **21**, 132-40.
- S. G. Foss and N. I. Chernova (2001). On optimality of the FCFS discipline in mutiserver queueing systems and networks. *Siberian Math. J.*, **42**, 372-385.
- K. Sigman (2011). Exact simulation of the stationary distribution of the FIFO M/G/c queue. *Journal of Applied Probability*. **Special Volume 48A** (To appear.)
- Wolff, R.W. (1987). Upper bounds on work in system for multi-channel queues. *J. Appl. Probab.* **14**, 547-551.