A Large-Scale Network with Moving Servers

Arnab Ganguly
Louisiana State University
aganguly@lsu.edu

Kavita Ramanan
Brown University
Kavita_Ramanan@Brown.edu

Philippe Robert
INRIA Paris
Philippe.Robert@inria.fr

Wen Sun
INRIA Paris
Wen.Sun@inria.fr

1. INTRODUCTION

In this note we analyze a queueing network of N queues and S_N servers. The arrival process to each queue is assumed to be Poisson with parameter λ, and the service times of jobs are independent and exponentially distributed with rate μ. The S_N servers circulate among the queues. When a server has completed the service of a job at a queue, it chooses at random *a non-empty* queue *without* a server if one exists. If all non-empty queues have a server, then the server remains idle. Idle servers are used as soon as an empty queue receive a new job.

This model is a special example of a polling system, for which there is already a huge literature. The main difference with classical models of polling systems is that several servers move around in the network and do not follow a fixed routing schedule among nodes. Furthermore, queues are not located on a circular network. The polling system with moving servers that we study has been used as an idealized model to analyze the performance of Ethernet Passive Optical Networks (EPONs). In this case, queues represent end-users and servers represent the different wavelengths available; see Antunes et al. [1], Boon et al. [2] and Robert and Roberts [9] for example.

As a polling system this special class has received little attention. There are some partial results concerning stability of these systems, by Fricker and Jaïbi [5] and Down [4], for example. In this work, we consider the case of a large network, that is, when N goes to infinity while S_N/N converges to some constant.

Stability Condition and Large Scale Behavior

It is assumed that the number of servers S_N is of the order of αN, i.e.

$$\lim_{N\to+\infty} S_N/N = \alpha \in (0,1). \tag{1}$$

Clearly a necessary stability condition is that $\lambda N{<}\mu S_N$. Therefore, it will be assumed that the condition $\rho\doteq\lambda/\mu{<}\alpha$ holds. We will prove that $\rho{<}\alpha$ is also a necessary stability condition for all N sufficiently large.

If one considers the total amount of work in the system, as long as there are more than S_N non-empty queues, this system could be seen as an $M/M/S_N$ queue for which the stability condition is well known to be $\rho<\alpha$. The problem in our model is to control the state of the network when the system is congested (that is, has large queues) but works below its maximal capacity, that is, there are less than S_N non-empty queues and thus idle servers. We will use a coupling argument and a quadratic Lyapounov function to resolve this problem.

The stability condition can be (roughly) rephrased as saying that the effective service capacity of each file is $\alpha{<}1$. We will show that, under a suitable condition on the initial state, the network behaves on finite time intervals as a set of independent $M/M/1$ queues, each of whose service capacity is 1. A consequence is that, asymptotically, there is a large number of idle servers, of the order of N on any finite time interval. As a result, the sharing of servers among queues increases the effective capacity of each queue of the network. A particular case of this result has been proved in Antunes et al. [1] when the initial states of queues are i.i.d., having a geometric distribution with parameter ρ.

2. STOCHASTIC MODEL

Notations. In what follows, for $y>0$, $(\mathcal{E}_y^n)_{n\in\mathbb{N}}$ denotes a sequence of i.i.d. exponential random variables with parameter y, and $(\mathcal{N}_i^y)_{i\in\mathbb{N}}$ denotes an i.i.d. sequence of Poisson processes with rate y.

For $1{\le}i{\le}N$ and $t{\ge}0$, the state of the system at time t can be represented by $X^N(t)=\left\{(L_i^N(t),U_i^N(t)),1{\le}i{\le}N\right\}$, where

- $L_i^N(t)$ is the number of jobs in queue i at time t;
- $U_i^N(t){=}1$ if a server is at queue i and $U_i^N(t){=}0$ if not.

Clearly X^N is an irreducible Markov process in the state space

$$\mathfrak{S}_N\doteq\left\{\{(\ell_i,u_i)\}{\in}(\mathbb{N}{\times}\{0,1\})^N : \sum_{i=1}^N u_i{=}\sum_{i=1}^N u_i\mathbb{1}_{\{\ell_i>0\}}{\le}S_N\right\}$$

For $(a,b){\in}\mathbb{N}{\times}\{0,1\}$ and $1{\le}i{\le}N$, let $e_i(a,b)$ denote the element of $(\mathbb{N}{\times}\{0,1\})^N$ whose ith coordinate is (a,b) and all other coordinates are $(0,0)$. One defines, for $x{=}\{(\ell_i,u_i)\}{\in}\mathfrak{S}_N$,

$$W^N(x)\doteq\sum_{i=1}^N \ell_i \text{ and } A^N(x)\doteq\sum_{i=1}^N \mathbb{1}_{\{\ell_i>0\}}. \tag{2}$$

Note that $A^N(x)$ is the number of non-empty queues and $W^N(x)$ is the total number of jobs, in state x. For this

Markov process, the transitions from a state $x=\{(\ell_i,u_i)\}\in\mathfrak{S}_N$ are as follows: for $i=1,\dots,N$,

- Arrivals.

$$x\mapsto\begin{cases}x{+}e_i(1,1), \text{ at rate } \lambda & \text{if } A^N(x){<}S_N \text{ and } \ell_i{=}0,\\ x{+}e_i(1,0), \text{ at rate } \lambda & \text{otherwise.}\end{cases}$$

- Departures.
  - If $u_i{=}1$, $\ell_i{>}1$, then $x{\mapsto}x{+}e_i(-1,0)$ occurs at rate μ if $A_N(x){\le}S_N$ and rate $\mu/(A_N(x){-}S_N{+}1)$ if $A_N(x){>}S_N$;
  - If $u_i{=}1$, $\ell_i{=}1$ and $A^N(x){\le}S_N$, $x{\mapsto}x{+}e_i(-1,-1)$ occurs at rate μ;
  - If $u_i{=}1$ and $A^N(x){>}S_N$, then for every j with $u_j{=}0$, $\ell_j{>}0$, $x{\mapsto}x{+}e_i(-1,-1){+}e_j(0,1)$ with rate $\mu/(A_N(x){-}S_N{+}\mathbb{1}_{\{\ell_i>1\}})$.

Let $D_i[0,t]$ denote the number of departures from queue i in the time interval $[0,t]$, and let $(s_i^n)_{n\in\mathbb{N}}$ denote the departure instants. Also, let $\xi_N \doteq \mu S_N/(N{-}S_N{+}1)$.

Proposition 1. *For $i{\in}\{1,\dots,N\}$, there exists a version X^N of the Markov process with initial state $x{=}\{(\ell_j,u_j)\}$ $\in\mathfrak{S}_N$ such that $s_i^n{\le}t_i^n$ for all $n{<}\ell_i$, where*

$$t_i^n =: \mathcal{E}^1_{\xi_N}\mathbb{1}_{\{u_i=0\}}{+}\mathcal{E}^1_\mu{+}\sum_{k=2}^n\left(B_N^k\mathcal{E}^k_{\xi_N}{+}\mathcal{E}^k_\mu\right), \tag{3}$$

with (B_N^k) denoting i.i.d. Bernoulli random variables with

$$\mathbb{P}(B_N^k{=}0)=\frac{1}{N{-}S_N{+}1}.$$

Proof. We proceed by induction. Assume $\ell_i{>}1$ and let $n{=}1$. If $u_i{=}1$, then the departure from queue i occurs after an exponential time $\mathcal{E}^1_\mu$; if $u_i{=}0$, then all S_N servers must be busy at other nodes, and one has additionally to wait for a server to move to queue i. For this, one of the busy servers has to complete service, which happens at rate μS_N, and then choose to move to node i, which will happen with probability $1/(N-S_N+1)$. This proves (3) when $n=1$.

Now, assume $s_i^n{\le}t_i^n$, $n{<}\ell_i{-}1$. If $A^N(X^N(s_i^n)){\le}S_N$ then, since queue i is not empty and has a server, $s_i^{n+1}{=}s_i^n{+}\mathcal{E}^{n+1}_\mu$. Otherwise, one has $A^N(X^N(s_i^n)){>}S_N$, either the server leaving node 1 returns immediately to this node, i.e., $B_N^{n+1}{=}0$, or one has to wait for a time $\mathcal{E}^{n+1}_{\xi_N}$ for the next server to come to node 1. One concludes that $s_i^{n+1}{\le}t_i^{n+1}$. $\square$

Proposition 2. *Under the condition $\rho{<}\alpha$, for N sufficiently large, the Markov process X^N is ergodic.*

Proof. For $t{>}0$ and $1{\le}i{\le}N$, denote by $\widetilde{D}_i(0,t)$ the counting process associated to the sequence $(t_i^n)_{n\in\mathbb{N}}$ of Proposition 1. The renewal theorem, (see, e.g., Grimmett and Stirzaker [6]) and condition (1) imply

$$\begin{aligned}\lim_{t\to+\infty}&\frac{1}{t}\mathbb{E}[\widetilde{D}_i(0,t)]\\ &=1\bigg/\left(\frac{1}{\mu}{+}\left(1{-}\frac{1}{N{-}S_N{+}1}\right)\frac{N{-}S_N{+}1}{\mu S_N}\right)\sim\mu\alpha,\end{aligned}$$

for large N. The condition $\rho{<}\alpha$ gives the existence of $\kappa{>}0$, N_0 and $T_N<\infty$ such that, for $N{\ge}N_0$, the relation

$$\frac{1}{T_N}\mathbb{E}\left[\mathcal{N}_i^\lambda(T_N)-\widetilde{D}_i(0,T_N)\right]\le-\kappa$$

holds for all $i{\in}\{1,\dots,N\}$, where the rate λ Poisson process $\mathcal{N}_i^\lambda$ represents the arrival process at queue i.

Then, given any initial condition $X^N(0){=}\{(L_i^N(0),U_i^N(0))\}$ $\in\mathfrak{S}_N$, Proposition 1 implies

$$\begin{aligned}L_i^N(t)&{\le}\left(L_i^N(0){-}D_i(0,t)\right)^+{+}\mathcal{N}_i^\lambda(t)\\ &\le\left(L_i^N(0){-}\widetilde{D}_i(0,t)\right)^+{+}\mathcal{N}_i^\lambda(t),\end{aligned}$$

Therefore, $L_i^N(t)^2{-}L_i^N(0)^2$ is less than or equal to

$$\begin{aligned}&2(\mathcal{N}_i^\lambda(t){-}\widetilde{D}_i(0,t))L_i^N(0){+}\left((L_i^N(0){-}\widetilde{D}_i(0,t))^+{+}\mathcal{N}_i^\lambda(t)\right)^2\\ &\quad{+}(\mathcal{N}_i^\lambda(t){-}\widetilde{D}_i(0,t))^2{-}\left(L_i^N(0){+}\mathcal{N}_i^\lambda(t){-}\widetilde{D}_i(0,t)\right)^2\\ &\quad\le 2(\mathcal{N}_i^\lambda(t){-}\widetilde{D}_i(0,t))L_i^N(0)+3\mathcal{N}_i^\lambda(t)^2+3\widetilde{D}_i(0,t)^2.\end{aligned}$$

Define, for $x{=}\{(\ell_i,u_i)\}\in\mathfrak{S}_N$, $G(x){=}{:}\ell_1^2{+}\ell_2^2{+}\cdots{+}\ell_N^2$. Then, for $t{\ge}0$, we have

$$\begin{aligned}\mathbb{E}[G(X^N(t))-G(X^N(0)]&\le(\lambda t-\mathbb{E}[\widetilde{D}_1(0,t)])\sum_{i=1}^N L_i^N(0)\\ &+3N\mathbb{E}[\mathcal{N}_1^\lambda(t)^2]+3\sum_{i=1}^N\mathbb{E}[\widetilde{D}_i(0,t)^2].\end{aligned}\tag{4}$$

From relation (4), one gets that if $G(X^N(0))\ge K_N^2$, where

$$K_N{\doteq}\frac{1}{2\kappa T_N}\left[1{+}3N\mathbb{E}\left[\mathcal{N}_1^\lambda(0,T_N)^2\right]{+}3\sum_{i=1}^N\mathbb{E}\left[\widetilde{D}_i(0,T_N)^2\right)\right],$$

then

$$\mathbb{E}[G(X^N(T_N))-G(X^N(0)]\le-1.$$

Thus, the function G is a Lyapounov function for the Markov process X^N. One concludes that, for $N{\ge}N_0$, the Markov process X^N is ergodic (see, e.g., [8, Theorem 8.13]). $\square$

3. LARGE SCALE ASYMPTOTICS

We start with a technical result on the duration of time when there is an idle server.

Lemma 1. *If $\rho{<}\alpha$ and the initial state $X^N(0){=}x_N\in\mathfrak{S}_N$ satisfies $A^N(x_N){\ge}S_N$ and*

$$\limsup_{N\to+\infty}W^N(x_N)/N{=}C_1<+\infty,$$

where $A^N(\cdot)$ and $W^N(\cdot)$ are defined by (2), then, if

$$\tau_N{\doteq}\inf\left\{t{>}0{:}A^N(t){<}S_N\right\},$$

there exists $t_0{>}0$ such that

$$\lim_{N\to+\infty}\mathbb{P}(\tau_N{\le}t_0)=1. \tag{5}$$

Proof. On the time interval $[0,\tau_N]$, one has

$$\left(W^N(X^N(\cdot))\right)\stackrel{\text{dist.}}{=}\left(W^N(x_N){+}\mathcal{N}_1^{\lambda N}(\cdot){-}\mathcal{N}_1^{\mu S_N}(\cdot)\right)$$

holds. In particular,

$$W^N(x_N){+}\mathcal{N}_1^{\lambda N}(\tau_N){-}\mathcal{N}_1^{\mu S_N}(\tau_N)\ge0.$$

The law of large numbers for Poisson processes then shows that relation (5) holds with $t_0{=}2C_1/(\mu\alpha{-}\lambda)$. $\square$

The next proposition gives a mean-field result. Specifically, under a stronger condition on ρ, we show that the system behaves as a system of independent $M/M/1$ queues in which all non-empty queues have a server. Let the processes $\{Q_i\}_{i\in\mathbb{N}}$ represent the dynamics of i.i.d. $M/M/1$ queues with arrival rate λ and service rate μ. Then the stability condition $\rho{<}\alpha{<}1$ implies that each Q_i is ergodic.

PROPOSITION 3. *If the empirical distribution of the $\mathbb{N}$-valued sequence $\{L_i^N(0)\}_{i=1,\dots,N}$ converges in distribution to ν as $N\to\infty$ and if*

$$\rho < \frac{\alpha-\nu(\mathbb{N}^*)}{1-\nu(\mathbb{N}^*)}, \tag{6}$$

with $\mathbb{N}^*\doteq\mathbb{N}\backslash\{0\}$, then as $N\to+\infty$,

$$\left(\frac{1}{N}\sum_{i=1}^N \delta_{(L_i^N(\cdot),U_i^N(\cdot))}\right)$$
$$\Rightarrow \left(\mathbb{P}_\nu(Q_1(\cdot){=}0)\delta_{(0,0)}+\sum_{m=1}^{+\infty}\mathbb{P}_\nu(Q_1(\cdot){=}m)\delta_{(m,1)}\right), \tag{7}$$

where $\delta_{(x,y)}$ denotes the Dirac mass at (x,y), $\mathbb{P}_\nu$ the distribution conditioned on the initial distribution ν and $\Rightarrow$ denotes convergence in distribution of measure-valued processes, (see, e.g., [3, 7]).

PROOF. For $1{\le}i{\le}N$, set $Q_i(0){=}L_i^N(0)$. Due to the independence of the dynamics of the $M/M/1$ queues, one has

$$\left(\frac{1}{N}\sum_{i=1}^N \delta_{Q_i(\cdot)}\right) \Rightarrow \left(\sum_{m=0}^{+\infty}\mathbb{P}_\nu(Q_1(\cdot){=}m)\delta_m\right),$$

as $N\to+\infty$. In particular, this implies that as $N\to+\infty$,

$$\left(\frac{1}{N}\sum_{i=1}^N \mathbb{1}_{\{Q_i(\cdot)>0\}}\right) \Rightarrow \mathbb{P}_\nu(Q_1(\cdot){>}0), \tag{8}$$

where the convergence is of $[0,1]$-valued stochastic processes. Note that

$$\mathbb{P}_\nu(Q_1(t){>}0) \le \nu(\{0\})\mathbb{P}_0(Q_1(t){>}0)+\nu(\mathbb{N}^*) < \alpha,$$

where the last inequality uses the well known property of the $M/M/1$ queue that $t\mapsto Q_1(t)$ is stochastically increasing when $Q_1(0)=0$ and the fact that ρ is the probability that the queue is not empty at equilibrium, and relation (6). Consequently, for $T{>}0$, if

$$\mathcal{H}_T^N \doteq \left\{\sup_{0\le t\le T}\frac{1}{N}\sum_{i=1}^N \mathbb{1}_{\{Q_i(t)>0\}} \le \frac{S_N}{N}\right\},$$

relation (8) shows that $\lim_{N\to\infty}\mathbb{P}(\mathcal{H}_T^N)=1$. To conclude the proof, observe that on the event $\mathcal{H}_T^N$, there are always idle servers up to time T, and so the relation

$$X^N=\left\{L_i^N,U_i^N\right\}_{i=1,\dots,N} \stackrel{\text{dist.}}{=} \left\{Q_i^\nu(\cdot),\mathbb{1}_{\{Q_i^\nu(\cdot)>0\}}\right\}_{i=1,\dots,N}.$$

holds on the time interval $[0,T]$. $\Box$

We now obtain a convergence result under a condition on ρ that does not involve the initial state.

PROPOSITION 4. *For any sequence of initial states $\{x_N\}$ that satisfy*

$$\limsup_{N\to+\infty}\frac{W^N(x_N)}{N} < +\infty, \tag{9}$$

where $W^N(\cdot)$ is defined by relation (2), and if

$$\rho{<}1-\sqrt{1-\alpha}, \tag{10}$$

then there exists some $t_0{>}0$ such that the convergence in (7) holds on any finite time interval after time t_0.

PROOF. We use the results and the notations of Lemma 1. There exists some t_0 such that the event $F_N=\{\tau_N{\le}t_0\}$ has probability close to 1 as N gets large. Let $x_N{=}(\ell_i^N,u_i^N)$, then by Condition (9), the sequence of empirical distributions of $(\ell_i, i{=}1,\dots,N)$ is tight. Since, on the event F_N,

$$L_i^N(\tau_N)\le \ell_i+\mathcal{N}_i^\lambda([0,t_0]),$$

the sequence of empirical distributions of $(L_i^N(\tau_N), i{=}1,\dots,N)$ is tight and, by definition of τ_N, any limit point ν satisfies the relation $\nu(\mathbb{N}^*){\le}\rho$. The strong Markov property of X^N applied at time τ_N and Proposition 3 give that the convergence in (7) holds on any finite time interval after time t_0 if the relation

$$\rho < \frac{\alpha-\rho}{1-\rho} \le \frac{\alpha-\nu(\mathbb{N}^*)}{1-\nu(\mathbb{N}^*)}$$

holds, but this is precisely our condition on ρ. $\Box$

Note that Condition (10) is stronger that the condition $\rho{<}\alpha$.

4. EXTENSIONS

The main extension currently under study is the model where the servers do not return immediately to the queues but wait for an exponential time with parameter γ. The "natural" stability condition in this case is $\lambda(1/\mu{+}1/\gamma)<\alpha$. The analogue of the mean-field result seen in this note would be the statement that, under similar initial conditions, after some finite time, a queue in this system behaves as an $M/M/1$ queue with arrival rate λ and service rate μ. More interestingly, one would like to show that at the level of the whole network, almost surely there is a finite number of non-empty queues without a server.

5. REFERENCES

[1] N. Antunes, C. Fricker, P. Robert, and J. Roberts, *Traffic capacity of large WDM passive optical networks*, 22th ITC Specialist Seminar (Amsterdam), 2010.

[2] M. Boon, R. van der Mei, and E. Winands, *Applications of polling systems*, Surveys in Operations Research and Management Science **16** (2011), 67–82.

[3] Donald A. Dawson, *Measure-valued Markov processes*, École d'Été de Probabilités de Saint-Flour XXI—1991, Springer, Berlin, 1993, pp. 1–260.

[4] D. Down, *On the stability of polling models with multiple server*, J. Appl. Probab. **35** (1998), 925–935.

[5] C. Fricker and M. R. Jaïbi, *Stability of multi-server polling models*, INRIA Research report, No. 3347, 1998.

[6] G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, 2nd ed., OUP, New York, 1992.

[7] H. Kaspi and K. Ramanan, *Law of large numbers limits for many-server queues*, Ann. Appl. Probab. **21** (2011), 33–114.

[8] P. Robert, *Stochastic networks and queues*, Stochastic Modelling and Applied Probability Series, vol. 52, Springer, New-York, 2003.

[9] P. Robert and J. Roberts, *A flow-aware MAC protocol for a passive optical metropolitan area network*, 23th ITC Specialist Seminar (San Francisco), 2011.