

Improving the Mean-Field Fluid Model of Processor Sharing Queueing Networks for Dynamic Performance Models in Cloud Computing

Johan Ruuskanen, Tommi Berner, Karl-Erik Årzén, and Anton Cervin
Department of Automatic Control, Lund University, Sweden

{johan.ruuskanen, tommi.nylander, karl-erik.arzen, anton}@control.lth.se

ABSTRACT

Resource management in cloud computing is a difficult problem, as one needs to balance between adequate service to clients and cost minimization in a dynamic environment of interconnected components. To make correct decisions in such an environment, good performance models are necessary. A common modeling methodology is to use networks of queues, but as these are prohibitively expensive to evaluate for many applications, approximation methods for key metrics are frequently employed. One such method—that provides both transient solutions and short, scalable computation times—is the fluid model, which approximates the dynamics of the mean queue lengths using a system of ordinary differential equations. However, finding a fluid model that can adequately approximate an arbitrary queueing network is in general difficult. In the paper, we extend the state of the art with the following three contributions. First, we show that for any mixed multiclass queueing network of processor sharing and delay queues with phase-type service time distributions, such a fluid model can be found via the mean-field approximation. Furthermore, we propose an improved model based on smoothing of the processor share function that improves the performance of certain systems. Finally, using the smoothed mean-field model, we introduce an accurate closed-form approximation of the response time CDF over any subset of classes and queues.

Keywords

Queueing network; Processor sharing; Mean-field approximation; Fluid model; Response time approximation.

Acknowledgement

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The authors are members of the ELLIIT Excellence Center.

1. INTRODUCTION

Multiclass queueing networks in which only some request paths of classes are subjected to external arrivals are known as *mixed queueing networks* [Baskett et al. 1975]. These networks are of interest to study in performance modeling of

cloud systems, as they can be used to model advanced behaviour such as interactions between synchronous and asynchronous calls to software layers. As exact analysis of queueing networks is often intractable and also prohibitively expensive to evaluate by simulation for many real-time applications, different approximation techniques of relevant performance metrics are often employed.

In certain cases, it is possible to translate the queueing network to a CTMC from which a fluid model can be obtained via the mean-field approximation,

$$\frac{d}{dt} \mathbf{x} = F(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{X}(0) \quad (1)$$

where \mathbf{X} is the request population vector in the network, $\mathbf{x}(t)$ an approximation of $\mathbb{E}[\mathbf{X}(t)]$ and $F(\mathbf{X})$ the drift function of the corresponding CTMC. We will refer to such a fluid model as a *mean-field fluid model*. Considering multiclass networks, in [Pérez and Casale 2017] it was shown that for closed networks, Kurtz’s theorem holds for any combination of multiclass *processor-sharing* (PS) or *delay* (INF) queues with general service times in the form of *phase-type* (PH) distributions.

This paper is an extended abstract of [Ruuskanen et al. 2021]. We refer to the original paper for proofs, more related work, further motivations, additional contributions and multiple experimental evaluations.

2. MIXED PS QUEUEING NETWORK

We built on the results of [Pérez and Casale 2017] and show that the results regarding Kurtz’s theorem can be extended to mixed networks. Further, a compact nonlinear matrix form of the corresponding drift function is developed.

Before continuing, some definitions are needed. Let \mathcal{Q} be the set of queues, where each queue i has \mathcal{C}_i classes and where each class $r \in \mathcal{C}_i$ has a PH distributed service time with $\mathcal{S}_{i,r}$ phase states. Further, let \mathcal{C} and \mathcal{S} be the sets of all classes and phase states in the network. The PH distributions are parametrized by $\Psi^{i,r} \in \mathbb{R}^{|\mathcal{S}_{i,r}| \times |\mathcal{S}_{i,r}|}$, $\psi^{i,r} \in \mathbb{R}_+^{|\mathcal{S}_{i,r}| \times 1}$ and $\zeta^{i,r} \in \mathbb{R}_+^{|\mathcal{S}_{i,r}| \times 1}$ describing; the transition rates between the transient states, from the transient states to the absorbing state, and the entrance probabilities respectively. Further, let $P \in \mathbb{R}^{|\mathcal{C}| \times |\mathcal{C}|}$ be a substochastic matrix where $P_{i,j}^{r,s}$ defines the transition probability from $(i,r) \in (\mathcal{Q}, \mathcal{C}_i)$ to $(j,s) \in (\mathcal{Q}, \mathcal{C}_j)$. If $\sum_{s,j} P_{i,j}^{r,s} < 1$ then there is a possibility for a request to depart the network once completed at (i,r) . Finally, let each class be subjected to the possibility of Poisson arrivals, whose rates are given by $\lambda \in \mathbb{R}^{|\mathcal{C}| \times 1}$.

If we assume that $\mathbf{X} \in \mathbb{Z}_+^{|\mathcal{S}|}$ is a vector of the population in each phase state, and stack the parameter matrices for the phase-type distributions into the following block diagonals

$$\begin{aligned}\mathbf{\Psi} (\in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}) &= \text{diag}(\Psi^{1,1}, \Psi^{1,2}, \Psi^{1,3}, \dots) \\ \mathbf{B} (\in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{C}|}) &= \text{diag}(\psi^{1,1}, \psi^{1,2}, \psi^{1,3}, \dots) \\ \mathbf{A} (\in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{C}|}) &= \text{diag}(\zeta^{1,1}, \zeta^{1,2}, \zeta^{1,3}, \dots)\end{aligned}$$

and form

$$\mathbf{P} = \begin{bmatrix} P_{1,1}^{\cdot,\cdot} & P_{1,2}^{\cdot,\cdot} & \dots \\ P_{2,2}^{\cdot,\cdot} & P_{2,2}^{\cdot,\cdot} & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} \quad (2)$$

obtaining the drift function can be simplified as follows.

THEOREM 1. *The drift function of the mixed queueing network of PS and INF queues can be expressed as*

$$\mathbf{F}(\mathbf{X}) = \mathbf{W}^T \theta(\mathbf{X}) + \mathbf{A} \boldsymbol{\lambda} \quad (3)$$

PROOF. See [Ruuskanen et al. 2021, Th. 1]. \square

where $\mathbf{W} = \mathbf{\Psi} + \mathbf{BPA}^T$, and $\theta(\mathbf{x})$ a function s.t. $\forall (i, r, a) \in (\mathcal{Q}, \mathcal{C}_i, \mathcal{S}_{i,r})$

$$\begin{aligned}\theta_{i,r,a}(\mathbf{X}) &= X_{i,r,a} g_{i,r,a}(\mathbf{X}) \\ g_{i,r,a}(\mathbf{X}) &= \frac{\min(k_i, \sum_{s \in \mathcal{C}_i} \sum_{b \in \mathcal{S}_{i,s}} X_{i,s,b})}{\sum_{s \in \mathcal{C}_i} \sum_{b \in \mathcal{S}_{i,s}} X_{i,s,b}}\end{aligned} \quad (4)$$

where $g_{i,r,a}(\mathbf{X})$ is known as the processor share.

Furthermore, the corresponding solution of the mean-field fluid model $\mathbf{x}(t)$ can be shown to be the limit of a scaled version of the mixed network. Introduce the sequence $\{\mathbf{X}^{(v)}\}_{v \geq 1}$, where $\mathbf{X}^{(v)}(0) = v\mathbf{X}(0)$, $\mathbf{k}^{(v)} = v\mathbf{k}$, and $\boldsymbol{\lambda}^{(v)} = v\boldsymbol{\lambda}$, then

THEOREM 2. *For any $\delta > 0$,*

$$\lim_{v \rightarrow +\infty} \mathbb{P} \left\{ \sup_{t \leq T} \left| \frac{\mathbf{X}^{(v)}(t)}{v} - \mathbf{x}(t) \right| > \delta \right\} = 0 \quad (5)$$

PROOF. See [Ruuskanen et al. 2021, Th. 3]. \square

3. SMOOTHED MEAN-FIELD MODEL

The convergence result in Theorem 2 does not state how accurate the approximation is, and mean-field approximations are in general known to suffer from approximation errors when the system size is small. To improve accuracy in these cases, we created a novel computationally cheap improvement for mean-field fluid models of mixed PS networks.

The main problem for these mean-field models lies in that

$$\mathbb{E}[\theta_a(\mathbf{X})] \neq \theta_a(\mathbb{E}[\mathbf{X}]) \quad (6)$$

due to the minimum function and the fraction in $g_a(\mathbf{X})$. To improve similarity between the quantities in (6), we created a smoothing over the minimum using the inverse p-norm

$$\hat{g}_a(\mathbf{x}, p_{Q(a)}) = \frac{1}{\left(1 + \left(k_{Q(a)}^{-1} \sum \mathbf{x}_{Q(a)}\right)^{p_{Q(a)}}\right)^{1/p_{Q(a)}}} \quad (7)$$

$$p_{Q(a)} > 0, \forall a \in \mathcal{S}$$

which gives the following smoothed mean-field fluid model

$$\begin{aligned}\frac{d\mathbf{x}}{dt} &= \mathbf{W}^T \hat{\theta}(\mathbf{x}, \mathbf{p}) + \mathbf{A} \boldsymbol{\lambda} \\ \mathbf{x}(0) &= \mathbf{X}(0), \quad \mathbf{p}(t) > \mathbf{0}\end{aligned} \quad (8)$$

The smoothing is dependent on the parameter \mathbf{p} , but as (7) can be shown to be monotone in $p_{Q(a)}$, the optimal \mathbf{p}^* for a stationary system can be quickly found from data.

4. CLOSED-FORM APPROXIMATION OF RESPONSE TIME CDF

In performance models of cloud systems, different response time percentiles are often of interest to capture. Retrieving estimates from a fluid model are however difficult, and current techniques rely on either Chebyshev bounds, or by solving an extended fluid model. However, we showed that it is possible to create accurate closed-form approximations.

The main idea is that we can obtain the time-dynamics of the probability of finding a specific request at a certain phase state in $(i, r) \in (\mathcal{Q}, \mathcal{C}_i)$ via the Chapman-Kolmogorov equation. The resulting equation however becomes intractable, but assuming that every request receives the mean processor share, and approximating the mean processor share using $\hat{g}(\mathbf{x}, \mathbf{p})$ we end up with the following linear ODE

$$\mathbb{E}[\dot{\boldsymbol{\pi}}(t)] \approx \left(\Psi^{i,r}\right)^T \hat{g}_i[\mathbf{x}^*(\mathbf{p}), \mathbf{p}] \mathbb{E}[\boldsymbol{\pi}(t)], \quad \boldsymbol{\pi}(0) = \zeta^{i,r} \quad (9)$$

which yields a closed-form approximative solution to $\boldsymbol{\pi}(t)$. Here $\mathbf{x}^*(\mathbf{p})$ is the stationary solution to (8) given \mathbf{p} . As the response time CDF can be shown to be given as the complement of the probability that the request remains in (i, r) , an approximation can be obtained as

$$\Phi_{i,r}(t | \mathbf{p}) \approx 1 - \boldsymbol{\pi}(0)^T \exp\left[\hat{g}_i[\mathbf{x}^*(\mathbf{p}), \mathbf{p}] \Psi^{i,r} t\right] \mathbb{1} \quad (10)$$

We further approximate the response time CDF over almost any subset of classes $\mathcal{C}_R \subset \mathcal{C}$. Let $\beta := \mathbb{R}_+^{|\mathcal{C}| \times 1}$ be the probability of entering specific states in \mathcal{C}_R at $t = 0$, i.e., $\sum \beta = 1$, $\beta_r \geq 0$ if $r \in \mathcal{C}_R$ else $\beta_r = 0$, and introduce the matrix $\mathbf{P}_R := \mathbb{R}_+^{|\mathcal{C}| \times |\mathcal{C}|}$ which only includes the class-to-class transition probabilities between classes in \mathcal{C}_R , i.e. $\forall r, s \in \mathcal{C}$ $(\mathbf{P}_R)_{r,s} = \mathbf{P}_{r,s}$ if $r, s \in \mathcal{C}_R$ otherwise $(\mathbf{P}_R)_{r,s} = 0$. Then,

$$\Phi_{\mathcal{C}_R}(t | \beta) \approx 1 - \beta^T \mathbf{A}^T \exp\left[D^{\hat{g}[\mathbf{x}^*(\mathbf{p}), \mathbf{p}]} \mathbf{W}_R t\right] \mathbb{1} \quad (11)$$

where $\mathbf{W}_R = \mathbf{\Psi} + \mathbf{BPA}^T$ and $D^{\mathbf{x}}$ the diagonal matrix s.t. $D_{aa}^{\mathbf{x}} = x_a \forall a \in \mathcal{S}$. Further,

THEOREM 3. $\Phi_{\mathcal{C}_R}(t | \beta)$ is a valid CDF for all valid choices of \mathcal{C}_R and flow matrices defined for the drift in Theorem 1.

PROOF. See [Ruuskanen et al. 2021, Th. 7] \square

5. REFERENCES

- F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. 1975. Open, closed, and mixed networks of queues with different classes of customers. *J. ACM* 22, 2 (April 1975), 248–260. <https://doi.org/10.1145/321879.321887>
- Juan F. Pérez and Giuliano Casale. 2017. Line: Evaluating Software Applications in Unreliable Environments. *IEEE Transactions on Reliability* 66, 3 (2017), 837–853.
- Johan Ruuskanen, Tommi Berner, Karl-Erik Årzén, and Anton Cervin. 2021. Improving the mean-field fluid model of processor sharing queueing networks for dynamic performance models in cloud computing. *Performance Evaluation* 151 (2021), 102231. <https://doi.org/10.1016/j.peva.2021.102231>