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, Anton Cervin

Background, applications as queuing networks

Model cloud applications as networks of queues

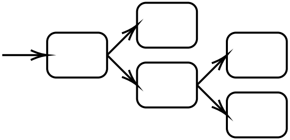

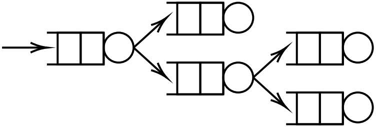

Evaluate queue length PMF using MC simulations, expensive

Instead, approximate important metrics (e.g. mean queue length, response time) using methods such as the **fluid model**

Background, mixed networks

Multiclass queues: Each queue has a set of *classes*, each with its own G_s and routing destination.

$P_{i,j}^{r,s}$ - the probability that a completed request of class r in queue i gets routed to class s in queue j.

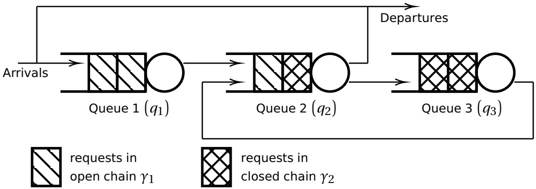

Chains: disjoint request paths over classes, known as *open* if classes allows departures/arrivals of requests, or *closed* if not.

Mixed network: containing both open/closed chains.

Background, the mean-field fluid model

For certain types of queueing networks, possible to derive a fluid model via the **mean-field approximation** of the queue lengths $\boldsymbol{X}(t)$.

Approximate $\mathbb{E}[\boldsymbol{X}(t)]$ with $\boldsymbol{x}(t)$, determined by

$$\dot{\boldsymbol{x}} = F(\boldsymbol{x}), \qquad \text{where } \boldsymbol{x}(0) = \boldsymbol{X}(0),\ F(\boldsymbol{X}) = \sum_{l \in \mathcal{L}} l f(\boldsymbol{X}, l)$$

Converges as "*system size*" goes to infinity.

Closed multiclass networks of PS and INF queues, where each class has a PH distributed service-time fulfills this [1].

Phase-type distribution: model a distribution as the time-to-absorption in a Markov chain of $\mathcal{S}$ transient states and one absorbing state. Parametrized by ζ, Ψ, ψ.

[1] F. Pérez and G. Casale, *Line: Evaluating Software Applications in Unreliable Environments*, IEEE Transactions on Reliability (2017)

Contributions

1. Compact mean-field fluid model for mixed networks.
2. Smoothed mean-field model for increased accuracy at low system sizes.
3. Closed from approximation of response time CDF.

Compact mean-field model for mixed networks

The mean-field fluid model can be extended to mixed networks by adding transition rate functions for arrivals/departures to $\mathcal{L}$.

Can be shown that normalized $\boldsymbol{X}(t) \to \boldsymbol{x}(t)$ when $\{\boldsymbol{\lambda}, \boldsymbol{k}, \boldsymbol{X}(0)\} \to \infty$.

For each queue i, class r and phase state a the drift becomes

$$F_{i,r,a}(\boldsymbol{X}) = \sum_b \Psi_{b,a}^{i,r} \theta_{i,r,b}(\boldsymbol{X}) + \zeta_a^{i,r} \sum_{j,s,b} \psi_b^{j,s} P_{j,i}^{s,r} \theta_{j,s,b}(\boldsymbol{X}) + \zeta_a^{i,r} \lambda^{i,r}$$

where $\theta_{i,r,a}(\boldsymbol{x}) = x_{i,r,a} \cdot g_{i,r,a}(\boldsymbol{x}) = x_{i,r,a} \frac{\min(k_i, \sum_{s,b} x_{i,s,b})}{\sum_{s,b} x_{i,s,b}}$

Compact mean-field model for mixed networks

The drift can be represented in a more manageable form. Using

$$\mathbf{\Psi} = \operatorname{diag}\left(\Psi^{1,1}, \Psi^{1,2}, \Psi^{1,3}, \ldots\right)$$
$$\boldsymbol{A} = \operatorname{diag}\left(\zeta^{1,1}, \zeta^{1,2}, \zeta^{1,3}, \ldots\right)$$
$$\boldsymbol{B} = \operatorname{diag}\left(\psi^{1,1}, \psi^{1,2}, \psi^{1,3}, \ldots\right)$$

$$\boldsymbol{P} = \begin{bmatrix} P^{\cdot,\cdot}_{1,1} & P^{\cdot,\cdot}_{1,2} & \cdots \\ P^{\cdot,\cdot}_{2,1} & P^{\cdot,\cdot}_{2,2} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

we can create $\boldsymbol{W} = \mathbf{\Psi} + \boldsymbol{B}\boldsymbol{P}\boldsymbol{A}^T$ and

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

Product between linear phase-to-phase transitions $\boldsymbol{W}$ and nonlinear effect of queueing discipline $\theta(\boldsymbol{X})$.

Improving the mean-field fluid model

Mean-field error, want $\boldsymbol{x}$ to approximate $\mathbb{E}(\boldsymbol{X})$ but

$$\frac{d}{dt}\mathbb{E}[\boldsymbol{X}] = \mathbb{E}[F(\boldsymbol{X})] \neq F(\mathbb{E}[\boldsymbol{X}]) = \frac{d}{dt}\boldsymbol{x}$$

The queuing network case

$$\mathbb{E}\left[\boldsymbol{W}^T\theta(\boldsymbol{X}) + \boldsymbol{A}\boldsymbol{\lambda}\right] = \boldsymbol{W}^T\mathbb{E}[\theta(\boldsymbol{X})] + \boldsymbol{A}\boldsymbol{\lambda} \neq \boldsymbol{W}^T\theta(\mathbb{E}[\boldsymbol{X}]) + \boldsymbol{A}\boldsymbol{\lambda}$$

Can we find another $\hat{\theta}(\mathbb{E}[\boldsymbol{X}])$ that improves accuracy?

Improving the mean-field fluid model

Problem,

$$\mathbb{E}\left[\theta_{i,r,a}(\boldsymbol{X})\right] = \sum_{\boldsymbol{z}} \mathbb{P}\left(\boldsymbol{X} = \boldsymbol{z}\right) z_{i,r,a} \frac{\min(k_i, \sum_{s,b} z_{i,s,b})}{\sum_{s,b} z_{i,r,a}}$$

$$\theta_{i,r,a}\left(\mathbb{E}\left[\boldsymbol{X}\right]\right) = \mathbb{E}\left[X_{i,r,a}\right] \frac{\min\left(k_i, \sum_{s,b} \mathbb{E}\left[X_{i,s,b}\right]\right)}{\sum_{s,b} \mathbb{E}\left[X_{i,s,b}\right]}$$

Let $\hat{\theta}_{i,r,a}\left(\mathbb{E}\left[\boldsymbol{X}\right]\right) = \mathbb{E}\left[X_{i,r,a}\right] \hat{g}_{i,r,a}\left(\mathbb{E}\left[\boldsymbol{X}\right]\right)$, then by summing over all states/classes in queue i

$$\hat{g}_i\left(\mathbb{E}\left[\boldsymbol{X}\right]\right) = \frac{\sum_c \mathbb{P}\left(\sum_{s,b} X_{i,r,a} = c\right) \min\left(k_i, c\right)}{\sum_{s,b} \mathbb{E}\left[X_{i,s,b}\right]} = \frac{k_i \rho_i(\boldsymbol{X})}{\sum_{s,b} \mathbb{E}\left[X_{i,s,b}\right]}$$

Dependence on the PMF of $\boldsymbol{X}$, we need to allow $\hat{g}$ to change

Improving the mean-field fluid model

One such possible function is the inverse p-norm

$$\hat{g}_i\left(\boldsymbol{x}, p_i\right) = \frac{1}{\left(1 + \left(k_i^{-1} \sum_{s,b} \boldsymbol{x}_{i,s,b}\right)^{p_i}\right)^{1/p_i}}$$

Some nice properties

- $\exists p_i^o$ s.t. equality holds
- $p_i \to \infty$ gives back $g_i(\boldsymbol{X})$
- $\hat{g}_i\left(\boldsymbol{x}, p_i\right)$ monotonic in both $\sum \boldsymbol{x}_i$ and $p_i > 0$

  -> can find $\boldsymbol{p}^o$ for stationary systems in $\mathcal{O}\left(\log n\right)$

  -> better approx. than $g_i(\boldsymbol{X})$ if $p_i^o \leq p_i$

Closed-form approximation of RT distribution

Response time percentiles are often of interest to obtain.

The probability of finding a request in a certain state evolves as

$$\dot{\boldsymbol{\pi}}(t) = \left(\Psi^{i,r}\right)^T \frac{\min\left(k_i, \sum \boldsymbol{X}_i(t)\right)}{\sum \boldsymbol{X}_i(t)} \boldsymbol{\pi}(t), \qquad \boldsymbol{\pi}(0) = \zeta^{i,r}$$

Cumbersome, instead assuming that the request receives the mean processor share yields the closed form solution

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

Remaining in (i, r) at t is the complement of leaving before t, hence

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

Closed-form approximation of RT distribution

Possible to extend to cover (almost) any subset $\mathcal{C}_R$ of classes

Let $\boldsymbol{P}_R := \{\forall\, r, s \in \mathcal{C} \;\; (\boldsymbol{P}_R)_{r,s} = \boldsymbol{P}_{r,s}$ if $r, s \in \mathcal{C}_R$ otherwise $(\boldsymbol{P}_R)_{r,s} = 0\}$, create $\boldsymbol{W}_R = \boldsymbol{\Psi} + \boldsymbol{B}\boldsymbol{P}_R\boldsymbol{A}^T$ and let $\beta\left(\boldsymbol{x}^*, \boldsymbol{p}^*\right)$ be the entrance prob. vector to $\mathcal{C}_R$ of an arbitrary request at stationarity. Then,

$$\Theta_{\mathcal{C}_R}\left(t \mid \boldsymbol{p}\right) = 1 - \hat{\beta}\left(\boldsymbol{x}^*, \boldsymbol{p}^*\right)^T \boldsymbol{A}^T \exp\left[D^{\hat{g}\left(\boldsymbol{x}^*, \boldsymbol{p}^*\right)} \boldsymbol{W}_R\, t\right] \mathbb{1}.$$

Approximation is better the less stochastic the queueing system is

-> unfortunately worse the closer to $\rho = 1$ we are.

Simulation experiment

Test the smoothed fluid model, its robustness in $\boldsymbol{p}$ and the response time approximation on the more advanced two-tier model

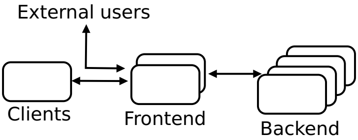

Translate into a mixed PS network with the two chains

$$\text{open chain:} \quad e \to f_1^i \to b_1^j \to f_2^k \to e \qquad i,k \in [1,2],\ j \in [1,4]$$

$$\text{closed chain:} \quad c \to f_3^i \to b_2^j \to f_4^k \to c \qquad i,k \in [1,2],\ j \in [1,4]$$

Fit $\boldsymbol{p}^*$ on data, and see how well it is possible to predict the system metrics after performing a random perturbation on **Workload**, **Horizontal** or **Vertical** scaling.

Simulation experiment

Draw 1500 random perturbations, calculate the relative error (**RE**) on mean class population and 95th percentile [2].

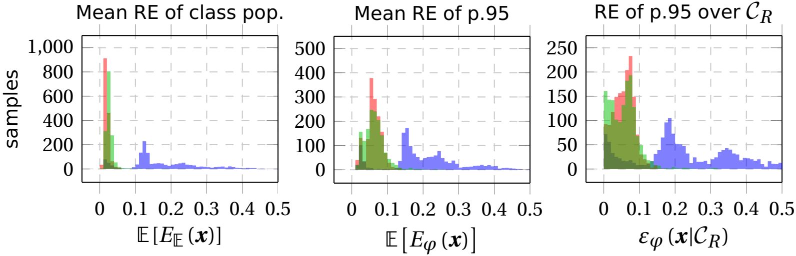

(blue) the standard mean-field model.

(red) the smoothed model with $\hat{p}^*$ re-fitted for each perturbation.

(green) the smoothed model with $\hat{p}$ fitted to the baseline system.

[2] The code: `https://github.com/JohanRuuskanen/Performance2021_code`

Conclusion

1. Compact mean-field fluid model for mixed networks.
2. Smoothed mean-field model for increased accuracy at low system sizes.
3. Closed from approximation of response time CDF.

Thank you for listening!