Performance Analysis for Data Centers: A Novel Heavy-Traffic Approach

Daniela Hurtado, Siva Theja Maguluri
Industrial and Systems Engineering, Georgia Institute of Technology

Cloud computing and data centers play an important role in powering the era of big data and machine learning. Today's data centers are huge, and most of the jobs need quick response times. Therefore, there are several challenges in optimal design and operation of data centers. Queueing theory is a powerful tool in designing and analyzing performance of various resource allocation algorithms. The goal is to characterize the delay experiences by jobs in a data center. Given that it is challenging to characterize the delay or queue lengths in general, heavy traffic theory has been developed to study queueing systems in the asymptotic regime when they are loaded very close to their capacity. In this work, we propose a novel view of the heavy-traffic theory based on the moment-generating function of the queue lengths. We will use this view to study the load-balancing problem in data centers. Future work includes studying resource allocation in data center networks.

Single server queue

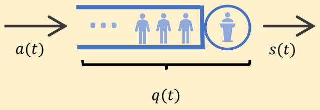

$q(t)$: Number of customers in system at the beginning of time slot t

$a(t)$: Arrivals in time slot t

$s(t)$: Potential service in time slot t

$u(t)$: Unused service in time slot t

$$q(t+1) = [q(t) + a(t) - s(t)]^+ = q(t) + a(t) - s(t) + u(t)$$

$$q(t+1) * u(t) = 0$$

If for $t \geq 0$:

- $a(1), a(2), \dots, a(t)$ are i.i.d. , $s(1), s(2), \dots, s(t)$ are i.i.d.
- $a(k)$ and $s(k)$ are independent of each other and of $q(k)$ for all $k \leq t$
- $\epsilon = E[s(1)] - E[a(1)] > 0$
- MGF of $a(1)$ and $s(1)$ exist

Then, the steady state distribution of the scaled queue length $\epsilon * \overline{q}$ in heavy traffic is:

$$Expo\left(\frac{\sigma_a^2 + \sigma_s^2}{2}\right)$$

Our proof: For any real number θ

$$\left(e^{\epsilon\theta q(t+1)} - 1\right)\left(e^{-\epsilon\theta u(t)} - 1\right) = 0$$

$$E\left[e^{\epsilon\theta\overline{q}(t+1)}\right] = E\left[e^{\epsilon\theta\overline{q}(t)}\right]$$

$$E\left[e^{\epsilon\theta\overline{q}}\right] = \frac{1 - E\left[e^{-\epsilon\theta u}\right]}{1 - E\left[e^{-\epsilon\theta(s-a)}\right]}$$

Taylor approximation

$$\lim_{\epsilon\to 0} E\left[e^{\epsilon\theta\overline{q}}\right] = \frac{1}{1 - \theta\left(\frac{\sigma_a^2 + \sigma_s^2}{2}\right)}$$

MGF of $Expo(\lambda)$ is $\frac{1}{1-\lambda\theta}$

Load balancing problem

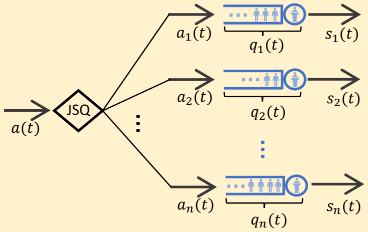

For $i \in \{1, \dots, n\}$, define:

$q_i(t)$: Number of customers in system i at the beginning of time slot t

$a_i(t)$: Arrivals to system i in time slot t

$s_i(t)$: Potential service in system i in time slot t

$u_i(t)$: Unused service in system i in time slot t

$\boldsymbol{q}(t) = (q_1(t), \dots, q_n(t))$

$\boldsymbol{s}(t) = (s_1(t), \dots, s_n(t))$

$\boldsymbol{u}(t) = (u_1(t), \dots, u_n(t)) \quad a(t) = \sum_{i=1}^{n} a_i(t)$

For each $i \in \{1, \dots, n\}$:

$$q_i(t+1) = q_i(t) + a_i(t) - s_i(t) + u_i(t)$$
$$\& \quad q_i(t+1) * u_i(t+1) = 0$$

Notation:

$$\mu_\Sigma = \sum_{i=1}^{n} E[s_i(1)] \qquad \lambda = E[a(1)]$$

$$\boldsymbol{q}_\parallel(t) = \frac{\mathbf{1}}{n}\sum_{i=1}^{n} q_i(t) \quad \& \quad \boldsymbol{q}_\perp(t) = \boldsymbol{q}(t) - \boldsymbol{q}_\parallel(t)$$

Assumptions: For $t \geq 0$:

- $a(t)$ are i.i.d.
- $s_i(1), s_i(2), \dots, s_i(t)$ are i.i.d. for each i
- $a(k)$ and $s_1(k), \dots, s_n(k)$ are independent of each other and of $q(k)$ for all $k \leq t$
- $\epsilon = \mu_\Sigma - \lambda > 0$

State space collapse [4]:
Let $\mu_{min} = \min_i E[s_i(1)]$ and $\delta \in (0, \mu_{min})$ and $\epsilon \in (0, (\mu_{min} - \delta)n)$. Then, for each $r = 1,2, \dots$ there exist a finite constant N_r such that $E[\|\boldsymbol{q}_\perp\|^r] \leq N_r$

$$E[\boldsymbol{q}_i] \approx \frac{1}{n}\sum_{j=1}^{n} q_j \quad \text{as } \epsilon \to 0$$

If the assumptions hold and for $t \geq 0$:

- There exists $A_{max} < \infty$ such that $a(t) \leq A_{max}$ with probability 1
- There exists $S_{max} < \infty$ such that $s_i(t) \leq S_{max}$ with probability 1 for each i

Then, the steady state distribution of the scaled sum of queue lengths $\epsilon * \|\overline{\boldsymbol{q}}\|_1$ in heavy traffic is:

$$Expo\left(\frac{\sigma_a^2 + \sum_{i=1}^{n}\sigma_{s_i}^2}{2}\right)$$

Our proof:

Part 1: For any real number θ

$$E\left[\left(e^{\epsilon\theta\|\boldsymbol{q}^+\|_1} - 1\right)\left(e^{\epsilon\theta\|\boldsymbol{u}^+\|_1} - 1\right)\right]$$
is $o(\epsilon^2)$

$$\& \ E\left[e^{\epsilon\theta\|\overline{\boldsymbol{q}}^+\|_1}\right] = E\left[e^{\epsilon\theta\|\overline{\boldsymbol{q}}\|_1}\right]$$

$$E\left[e^{\epsilon\theta\|\overline{\boldsymbol{q}}^+\|_1}\right] = \frac{1 - E\left[e^{-\epsilon\theta\|\boldsymbol{u}\|_1}\right] + o(\epsilon^2)}{1 - E\left[e^{-\epsilon\theta\sum_i(s_i - a_i)}\right]}$$

As $\epsilon \to 0$: $\epsilon * \left\|\overline{\boldsymbol{q}}_\parallel\right\|_1 \sim Expo\left(\frac{\sigma_a^2 + \sum_{i=1}^{n}\sigma_{s_i}^2}{2}\right)$

Part 2: For any real vector $\boldsymbol{\theta}$

$$\langle\boldsymbol{\theta}, \boldsymbol{q}\rangle = \sum_{i=1}^{n}\theta_i q_i = \frac{\|\boldsymbol{\theta}\|_1}{n}\|\boldsymbol{q}_\parallel\|_1 + \langle\boldsymbol{\theta}, \boldsymbol{q}_\perp\rangle$$

$$+$$

$\lim_{\epsilon\to 0} \epsilon * q_{\perp i} = 0$ w.p. 1

Slutsky's theorem

$$\lim_{\epsilon\to 0} E\left[e^{\epsilon\langle\boldsymbol{\theta}, \overline{\boldsymbol{q}}\rangle}\right] = \frac{1}{1 - \frac{\|\boldsymbol{\theta}\|_1}{n}\left(\sigma_a^2 + \sum_i \sigma_{s_i}^2\right)}$$

Future work:

Resource allocation problem

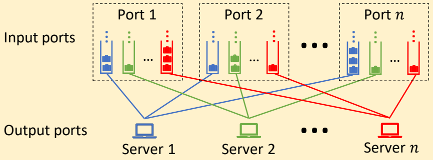

- Resource allocation? MaxWeight
- Distribution of queue lengths? Future work

References

[1] R.Srikant, L.Ying. *Communication networks: an optimization, control and stochastic networks perspective*. Cambridge University Press, 2013.

[2] A.Mood, F.Graybill, and C.Duane. D. boes (1974). Introduction to theory of statistics

[3] S.T.Maguluri, R.Srikant, et al. Heavy traffic queue length behavior in a switch under maxweight algorithm. *Stochastic systems*, 6(1):211-250, 2016.

[4] A. Eryilmaz and R. Srikant. Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queuing systems*, 72(3-4):311-359, 2012

[5] B.Hajek. *Random processes for engineers*. Cambridge University Press, 2015.

[6] A.Gut. *Probability: a graduate course*, volume 75. Springer Science & Business Media, 2012.