

Should we omit the practical aspects in modeling the server clusters?

Thanh-Binh V. Lam

Abstract: This paper investigates the impact of some practical aspects which are simplified on modeling the server clusters. Three considered aspects are S1) the distribution of the switching times, S2) the avoidance of turning off the servers in setup process, and S3) the presence of the shutdown time.

The results demonstrate that the accuracy of the calculation for energy consumption metrics in the server clusters might be improved unless we omit these practical aspects. In particular, both S2 and S3 expose considerable impact on the accuracy of the calculation for the average energy consumption. Taking these aspects into account might support the operators in planning and calculating the energy consumption for their server clusters.

Keywords: Server clusters, setup time, shutdown time, analytical model, practical aspects.

Introduction

The accuracy of the calculation for performance measures and energy consumption metrics is one of the important factors in managing server clusters of the IT service operators. Precisely, it helps them in planning as well as controlling the operation of their system. When modeling server clusters, however, the authors usually omit some practical aspects in the operation of servers, such as, setup time, shutdown time, procedures to turn on and turn off a server, etc. As a consequence, the operation of the server cluster does not behave as its natural dynamics. Throughout this paper, we demonstrate that these practical aspects play an important role in improving the accuracy of the calculation for performance measures and energy consumption metrics in the server clusters.

The authors [1, 2, 3, 4, 5] omitted the shutdown phase of a server which is impractical. More precisely, they took the setup time into account only. They assumed that a server can be turned off immediately without any cost in terms of energy consumption. Particularly, the authors [1] proposed to stop servers in setup process. Moreover, the authors [5, 6] assumed that a block of servers can be powered up or powered down simultaneously. From a practical point of views, even the servers are homogeneous and the commands to power them up or down are initiated at the same time, they could not finish their setup/shutdown periods in the same seconds. The authors [7] first took into account these aspects within some different workloads. However, they did not consider the size of server clusters. The question is whether these aspects impact on the accuracy of the calculation for performance measures and energy consumption metrics in server clusters in the case of varying the cluster size?

More precisely, our goals in this paper are

- to study the impact of distribution of switching times to the accuracy of the calculation for performance measures,
- to check whether turning off a server in setup process effects to the calculation for energy consumption in the system,
- to investigate the presence of shutdown time in the accuracy of the calculation for performance measures,
- to review the impact of cluster size on the calculation for performance parameters in the server clusters.

The rest of the paper is organized as follows. Section presents the abstract model which are considered throughout this paper. The simulation experiments are demonstrated in Section . Finally, Section concludes our work.

System modeling

We follow [7] to apply the abstract model, the parameters of the switching times and the numerical formula. In this work, however, we investigate the impact of S1, S2, and S3 on the accuracy of the

calculation for performance measures in server clusters within some different cluster sizes, namely, $K = 10; 50; 200$ servers.

We apply the service time is $1/\mu = 100s$ which is more than three times to the mean setup time and the mean shutdown time. This assumption is reasonable for a server to serve a job corresponding to such switching times. The average active power of the reference server is $P_{\max} = 56.1W$ [8].

We take six following scenarios into comparison:

Scenario	S1 (distribution)		S2	S3
	Setup time	Shutdown time		
U123	Uniform	Uniform	Yes	Yes
E123	Exponential	Exponential	Yes	Yes
U12	Uniform	No	Yes	No
E12	Exponential	No	Yes	No
U1	Uniform	No	No	No
E1	Exponential	No	No	No

Table 1: Configuration of scenarios in simulation

Simulation experiments

We built a software to simulate the abstract model. Simulation runs were performed with the confidence level of 99% and the accuracy (i.e. the ratio of the half-width of the confidence interval and the mean of collected data) is less than 0.0009.

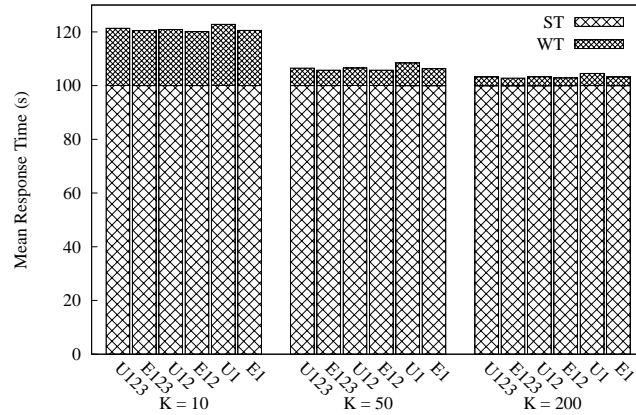


Figure 3: Mean response time vs. cluster size for $1/\mu = 100s$ and $\rho = 0.7$.

Figure 3 presents the mean response time against the cluster size for all scenarios. It is well-known that the mean response time can be decomposed into the mean waiting time (WT) and the mean service time (ST). Apparently, when the number of servers in the cluster increases, the jobs incur less waiting time. Generally, the impact of these aspects on the calculation for performance measures in terms of the mean response time is insignificant regardless the cluster size. Interestingly, the mean response time in the case of scenario U1 is higher than that of times in U12 and U123 for all scenarios. S2 is the dominant aspect which causes this phenomenon. If we do not allow turning off a server in setup process immediately (scenario U123, U12), there is a chance that after finishing its setup period, the server can

serve a waiting job. This job arrives during the times that server is in setup process. In other words, this job does not incur much waiting time. When we omit S2 and S3 (scenario **U1**), obviously, the arrival jobs have to wait for the full setup time which eventually increases the waiting time. Additionally, these practical aspects do not impact on the accuracy of the calculation for performance measures of server clusters under varying of cluster size.

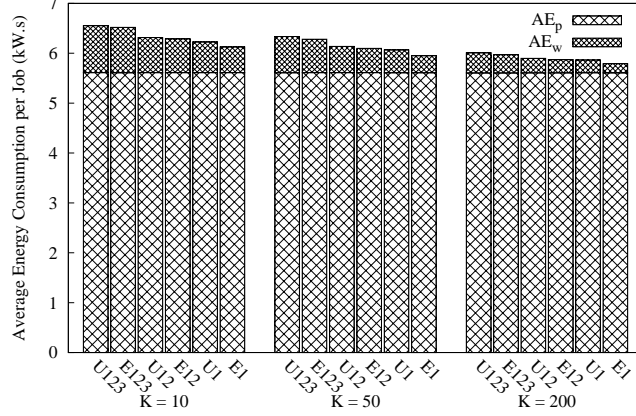


Figure 4: Average energy consumption per job vs. cluster size for $1/\mu = 100s$ and $\rho = 0.7$.

Figure 4 shows the considerable impact of these aspects on the accuracy of the calculation for average energy consumption. We define AE_p is the amount of energy that server consumed to serve job directly. AE_w is the amount of energy that server consumed in switching periods. Obviously, AE_w is wasted because it is not used to serve job at all. The impact of S2 and S3 are outperform their counterpart S1 in all scenarios in terms of the average energy consumption. In the case of small cluster size ($K = 10$), servers have to switch in and out power saving mode frequently which consumes much energy. The deviations between scenario **U123** and scenario **U12** in the case of small cluster size is 0.66 (25%), medium cluster size is 0.05 (27.6%), and large cluster size is 0.03 (27.9%). Visibly, the mismatch between cluster size and workload might causes much wasted energy consumption to the system. Another observation is the average energy consumption in models **U** is larger than that of energy consumption in models **E**. In other words, the server has longer time to consume power if the switching times is distributed uniformly.

Conclusion

In this paper, we demonstrate that:

- Although the impact of the distribution of switching times (S1) is insignificant, we should apply the uniform distribution to the switching times to improve the accuracy of the calculation for performance measures and energy consumption metrics of server clusters.
- The assumption on turning off the servers in setup process (S2) should be avoided.
- The shutdown time (S3) should be taken into account in modeling the dynamics of servers.
- These practical aspects do not impact on the accuracy of the calculation for performance measures and energy consumption metrics in server clusters under varying of cluster size.

The results in this paper might be useful to the operators who desire to evaluate the performance and to calculate the energy consumption of their server clusters. In the future, we would consider the operation policies, i.e. turn on/off a block of servers, then we investigate its impacts within practical aspects to the accuracy of the calculation for performance measures in the server clusters.

References

- [1] Artalejo, J.R., Economou, A., Lopez-Herrero, M.J.: Analysis of a multiserver queue with setup times. *Queueing Syst. Theory Appl.* 51(1-2), 53–76 (Oct 2005), <http://dx.doi.org/10.1007/s11134-005-1740-6>
- [2] Do, T.V., Rotter, C.: Comparison of scheduling schemes for on-demand IaaS requests. *Journal of Systems and Software* 85(6), 1400–1408 (2012),
- [3] Gandhi, A., Harchol-Balter, M.: How data center size impacts the effectiveness of dynamic power management. In: *Communication, Control, and Computing (Allerton)*, 2011 49th Annual Allerton Conference on. pp. 1164–1169 (Sept 2011)
- [4] Tian, N., Zhang, Z.G.: *Vacation Queueing Models: Theory and Applications*. International Series in Operations Research & Management Science, Springer (2006)
- [5] Mitrani, I.: Managing performance and power consumption in a server farm. *Annals of Operations Research* 202(1), 121–134 (2013), <http://dx.doi.org/10.1007/s10479-011-0932-1>
- [6] A. Borthakur and G. Choudhury, A multiserver Poisson queue with a general startup time under N-Policy, *Calcutta Statistical Association Bulletin* 49 (1999) 199-213.
- [7] Do N.H., Lam TB.V. (2015) Some Practical Aspects on modeling Server Clusters. In: Le Thi H., Nguyen N., Do T. (eds) *Advanced Computational Methods for Knowledge Engineering. Advances in Intelligent Systems and Computing*, vol 358. Springer, Cham, https://doi.org/10.1007/978-3-319-17996-4_36
- [8] SPEC, Fujitsu FUJITSU Server PRIMERGY RX1330 M3 machine, https://www.spec.org/power_ssj2008/results/res2017q2/power_ssj2008-20170315-00744.html, March 2017