

STOCHASTIC PERFORMANCE ANALYSIS OF UNRELIABLE IAAS CLOUDS

[YongBuo Wang, HaiChuan Zhang, KunYin Guo, Yunni Xia*, WeiLing Li]
[School of computers, Chongqing University, China 40014]
[xiayunni@hotmail.com]

Abstract

Performance evaluation of cloud data-centers has drawn considerable attention from academy and industry. In this study, we present an analytical approach to the performance analysis of Infrastructure-as-a-Service cloud data-centers with unreliable task executions and resubmissions of unsuccessful tasks. Several performance metrics are considered and analyzed under variable load intensities, failure frequencies, multiplexing abilities, and service intensities. We also conduct a case study based on a real-world cloud data-center and employ a confidence interval check to validate the correctness of the proposed model.

Keywords: IaaS cloud, performance, queuing networks

1. INTRODUCTION

Cloud data-centers are key enablers for the scalability of the cloud platform. Cloud computing relies on data-centers to deliver expected services. The widespread adoption of the cloud computing paradigm mandates the exponential growth in the data-centers' computational, network, and storage resources. Managing the computational resources to deliver specified performance [1] is among the key challenges.

Expected request response time which decides system responsiveness and request rejection rate which determines users' satisfaction are usually considered as the most important performance metrics to evaluate a service system. As will be discussed later in this paper, cloud data-centers are usually subject to errors/faults, and error/failure-handling activities could have strong impact on final performance. Due to the difficulties with building monolithic models capable of capturing related factors, measurement-based approaches are frequently used [2-4]. However, these approaches are intractable due to exhaustive experimentations. Therefore, their value is limited. Comprehensive analytical performance models are more preferable in this situation. Although some other analytical models [5-8] are proposed, those works are limited mainly because they assume the system failure-free to simplify the performance/QoS calculation.

This paper focuses on analytical performance analysis of IaaS cloud data-centers with request rejection and resubmission. For this purpose, a stochastic model is proposed and product-form expressions of multiple performance metrics are derived. We validate the model by experiment based on an actual IaaS cloud and the results indicate our model is trustable.

IaaS cloud is a form of cloud computing that provides virtualized computing resources over the Internet. The cloud management unit of an IaaS data-center maintains a request buffer for consecutively-arrived requests, which can be

usually described by an arrival rate, λ . The capacity of such buffer, denoted by c , can be specified before using (e.g., the capacity limit can be specified through the *FRAME_SIZE* property in OpenStack). Requests arrived either leave by rate ϑ or are resubmitted by rate $1-\vartheta$ when the capacity limit is reached. For the performance evaluation purpose, we are interested in knowing request response time, i.e., the expected interval time between request arrival and the corresponding VM ready for execution (e.g., the time of *INSTANCE_SPAWNED* defined in OpenStack).

As shown by Fig. 1, VM instantiation requires multiple steps and interferences with various services and components. Averaged speed (or process rate) of the cloud management unit to spawn a VM, denoted by μ , can be obtained by the reciprocal of averaged instantiation times, e.g., intervals between *INSTANCE_BUILDING* times and *INSTANCE_SPAWNED* times in OpenStack. With the help of VM multiplexing [9] mechanism supported by today's multi-core/multi-threading technologies, multiple VMs can be instantiated on a same PM. The maximum number of VMs that can be instantiated on a PM, denoted by m , is usually bounded. Note that high multiplexing level is not always welcomed because VM interference may cause performance and reliability degradation.

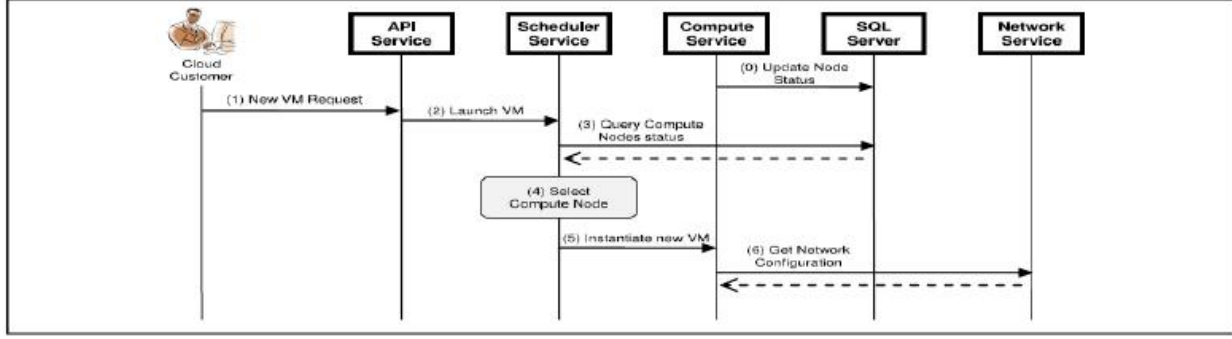


Fig. 1: Sequence chart of VM instantiation on OpenStack

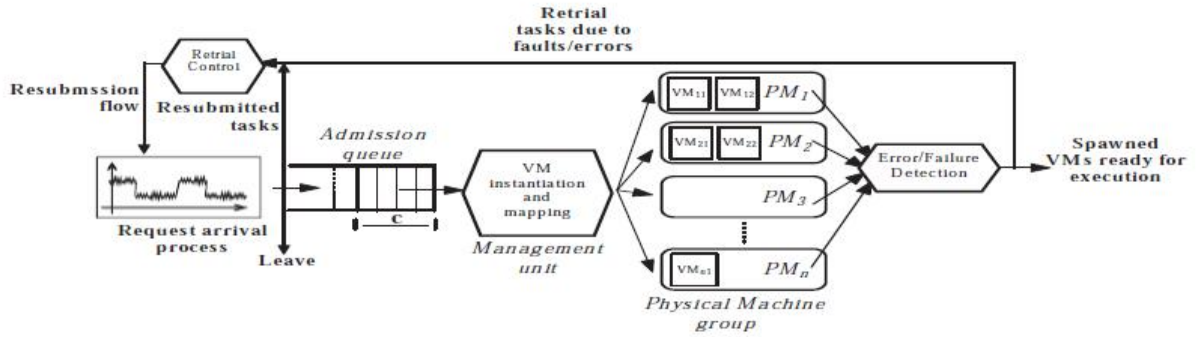


Fig. 2: The cloud provisioning control-flow of an IaaS data-center with task errors/faults

Moreover, Fig. 1 suggests potential of unsuccessful VM instantiation because interactions with local or remote services and components are often error/failure-prone.

Errors/failures can strongly impact cloud performance due to the overhead needed to conduct compensation/transactional-rollback activities and re-instantiate the faulty request. Based on the above discussions, an abstract control-flow model of VM instantiation on unreliable IaaS cloud data-center is illustrated in Fig. 2. It abstracts implementation details of IaaS cloud paradigm while preserving the control-flow contents useful for performance analysis in a context of queueing-networks. Its objective is to derive the quantitative effects of varying request arrival rates, VM instantiation rates, resource scale, and error intensity on cloud performance. The system under study is consequently mapped into an instance of queueing network problems solved in the following section.

2. STOCHASTIC ANALYSIS

Let $N(t)=n$ mean that the number of tasks waiting or being instantiated is n at time t , $M(t)=m$ mean that the number of requests being resubmitted m , and

$X(t)=(N(t), M(t))$ denote the system state at time t , the resulting state space is therefore $E \in \{0, 1, \dots, k\} \times \{0, 1, \dots, \infty\}$. Since the inter-arrival time, VM instantiation time, and resubmission processing time are all exponentially distributed, $X(t)$ is a Markovian process on state space E .

Based on the state transition chart, the corresponding transition-rate matrix, Q , can be derived as:

$$Q = \begin{bmatrix} A_0 & C & & & & \\ B_1 & A_1 & C & & & \\ & B_2 & A_2 & C & & \\ & & \ddots & \ddots & \ddots & \\ & & & B_{k-1} & A_{k-1} & \tilde{C} \\ & & & & \tilde{B}_k & \tilde{A}_k \end{bmatrix} \quad (1)$$

It is easy to see that $X(t)$ is irreducible and non-periodical. Let $\pi_{i,j}(t)$ denote the probability that the Markovian process is at state (k, j) and $\pi_{i,j} = \lim_{t \rightarrow \infty} \pi_{i,j}(t)$, we have that $\pi_{i,j}$ can be calculated as below if the stationary distribution exists:

$$\pi_{k,j} = \pi_{k,j-1} \prod_{l=g+1}^j \rho_l \quad (2)$$

where

$$\rho_l = \frac{e \times f \times \mu + (1 - \theta)\lambda}{l \times \theta \times \mu' + e \times (1 - f)\mu} \quad (3)$$

It is easy to see that ρ_l decreases with l . Consequently, there exists $u \in \mathbb{N}^+$ such that $\rho_u < l$ and

$$\begin{aligned} & \sum_{j=g+1}^{\infty} \left\{ \prod_{l=g+1}^j \left(\frac{e \times f \times \mu + (1 - \theta)\lambda}{l \times \theta \times \mu' + e \times (1 - f)\mu} \right) \right\} \\ &= \sum_{j=g+1}^{\infty} \left\{ \prod_{l=g+1}^j \rho_l \right\} < \sum_{j=g+1}^u (\rho_{g+1})^{j-g} + \sum_{j=u+1}^{\infty} \left(\prod_{l=g+1}^j \rho_l \right) \\ &< \sum_{j=g+1}^u (\rho_{g+1})^{j-g} + (\rho_{g+1})^{u-g} \sum_{j=u+1}^{\infty} \left(\prod_{l=u+1}^j \rho_l \right) \\ &< \sum_{j=g+1}^u (\rho_{g+1})^{j-g} + (\rho_{g+1})^{u-g} \sum_{j=u+1}^{\infty} (\rho_l)^{j-u} \\ &= \sum_{j=g+1}^u (\rho_{g+1})^{j-g} + (\rho_{g+1})^{u-g} \frac{\rho_u}{1 - \rho_u} < \infty \end{aligned} \quad (4)$$

The above derivation leads to

$$\sum_{j=g+1}^{\infty} \left\{ \prod_{l=g+1}^j \left(\frac{e \times f \times \mu + (1 - \theta)\lambda}{l \times \theta \times \mu' + e \times (1 - f)\mu} \right) \right\} < \infty \quad (5)$$

and therefore the stationary distribution exists according to the limit theorems of birth-death processes.

Since the stationary distribution exists, we have the steady-state probabilities of each state as:

$$\pi Q = 0, \sum_{i=0}^{k-1} \sum_{j=0}^g \pi_{i,j} + \sum_{j=0}^{\infty} \pi_{k,j} = 1 \quad (6)$$

Since the stationary distribution exists, we also have:

$$\begin{aligned} & \pi_{k-1,g} \lambda + (1 - \theta) \pi_{k,g-1} \lambda + e(1 - f) \mu \\ &= \pi_{k,g} (e \times \mu + \lambda(1 - \theta) + \theta \times g \times \mu') \\ & \quad - \pi_{k,g+1} (g + 1) \times \theta \times \mu' \end{aligned} \quad (7)$$

From (2) we have:

$$\pi_{k,g+1} = \pi_{k,g} \frac{\lambda(1 - \theta) + e \times f \mu}{(g + 1) \theta \times \mu' + e \times (1 - f) \mu} \quad (8)$$

Combining the above equation with (7), we have:

$$\pi_{k,g-1} \lambda(1 - \theta) = \pi_{k,g} (e \times (1 - f) \mu + g \times \theta \times \mu') \quad (9)$$

which suggests that $A_k = A_k$ and $B_k = B_k$.

According to (6) and (9), we have:

$$\begin{aligned} & T_0 \frac{1}{W} A_0 + T_1 \frac{1}{W} B_1 = 0 \\ & T_i \frac{1}{W} C + T_{i+1} \frac{1}{W} A_{i+1} + T_{i+2} \frac{1}{W} A_{i+2} = 0, 0 \leq i \leq k - 2 \\ & T_{k-1} \frac{1}{W} C + T_k \frac{1}{W} A_k = 0 \end{aligned} \quad (10)$$

where T_0 is a basic solution of $T_0(V_k A_k + V_{k-1} C) = 0$ and $\pi_{i,j}$ is subject to:

$$\begin{aligned} & (\pi_{0,0}, \dots, \pi_{0,g}) = T_0 \frac{1}{W} \\ & (\pi_{i,0}, \dots, \pi_{i,g}) = (\pi_{0,0}, \dots, \pi_{0,g}) V_i, 0 < i \leq k \\ & \pi_{k,j} = T_0 \frac{1}{W} V_k \omega_1 \prod_{l=g+1}^j \frac{e \times f \times \mu + (1 - \theta)\lambda}{l \times \theta \times \mu' + e \times (1 - f)\mu} \end{aligned} \quad (11)$$

where ω_1 is a column vector with its dimension being $g+1$ and is equal to $(0, \dots, 0, 1)^T$. and, V_k is subject to:

$$\begin{aligned} & V_0 = I \\ & V_1 = -A_0(B_1)^{-1} \\ & V_2 = -(V_0 C + V_1 A_1)(B_2)^{-1} \\ & V_i = -(V_{i-2} C + V_{i-1} A_{i-1})(B_i)^{-1}, 2 < i \leq k \end{aligned} \quad (12)$$

W is calculated as:

$$W = T_0 \left(\sum_{i=0}^k V_i \right) \omega + T_0 V_k \omega_1 \sum_{j=g+1}^{\infty} \left(\prod_{l=g+1}^j \frac{e \times f \times \mu + (1 - \theta)\lambda}{l \times \theta \times \mu' + e \times (1 - f)\mu} \right) \quad (13)$$

From (10), we have:

$$T_1 \frac{1}{W} = -T_0 \frac{1}{W} A_0(B_1)^{-1} = T_0 \frac{1}{W} V_1 \quad (14)$$

and similarly:

$$T_2 \frac{1}{W} = -T_0 \frac{1}{W} (V_1 A_1 + C)(B_2)^{-1} = T_0 \frac{1}{W} V_2 \quad (15)$$

Consequently, we can finally have:

$$T_i = -T_0 V_i \quad (16)$$

and the product form solution of T_i can be obtained by using this equation.

Combining (16) with (2), we can obtain the solutions of steady-state probabilities of all states, $\pi_{i,j}$. Note that a similar derivation skill can be found in [10].

3. PERFORMANCE RESULTS

We consider the following as the performance metrics: 1) Expected request response time, T ; and 2) Request rejection rate, R .

As suggested by Fig. 2, T denotes the expected interval between request Arrival and the moment of the corresponding VM being instantiated and ready for execution. Response time is a frequently used measure of efficiency and responsiveness of computer systems. Lower response time also allows for higher system reliability since in an unreliable system failures/errors are more likely to happen when a longer response time is needed.

To analyze T , we first have to calculate the probability that a cloud task enters the resubmission state, P_r :

$$P_r = (1 - \theta) \left(\sum_{j=0}^{\infty} \pi_{k,j} \right) + f \left(1 - \sum_{j=0}^{\infty} \pi_{k,j} \right) \quad (17)$$

The expected number of retrials of a cloud task, N_r , can be obtained as:

$$N_r = \frac{1}{1 - P_r} - 1 \quad (18)$$

The expected time for a task to wait before it is resubmitted to the arrival task flow, T_r , can therefore be calculated as:

$$T_r = \frac{\lambda' / \mu' + \frac{P_0 (\lambda' / (g \times \mu')) (\lambda' / \mu')^g}{g! (1 - \lambda' / (g \times \mu'))}}{\lambda'} \quad (19)$$

where P_0 denotes the probability that no task being resubmitted:

$$P_0 = \frac{1}{\sum_{l=0}^{g-1} \frac{(\lambda' / \mu')^l}{l!} + \frac{(\lambda' / \mu')^g}{g!}} \left(\frac{1}{1 - \lambda' / g \times \mu'} \right) \quad (20)$$

and λ' is the rate of resubmission flow into the arrival task flow:

$$\lambda' = (\lambda + \lambda') \left[\left(\sum_{j=0}^{\infty} \pi_{k,j} \right) (1 - \theta) + \left(1 - \sum_{j=0}^{\infty} \pi_{k,j} \right) f \right] \quad (21)$$

where λ' can be calculated as:

$$\lambda' = \lambda \frac{\sum_{j=0}^{\infty} \pi_{k,j} (1 - \theta) + f - \sum_{j=0}^{\infty} \pi_{k,j} \times f}{1 - \sum_{j=0}^{\infty} \pi_{k,j} (1 - \theta) - f + \sum_{j=0}^{\infty} \pi_{k,j} \times f} \quad (22)$$

The expected total time for a task to spend before its final successful trial on condition that it is not rejected, T_b , can therefore be obtained as:

$$T_b = N_r (T_r \sum_{j=0}^{\infty} \pi_{k,j} + (T_r + T_v) (1 - \sum_{j=0}^{\infty} \pi_{k,j})) \quad (23)$$

where T_v is calculated as:

$$T_v = \frac{1}{\mu} + \sum_{j=0}^g \pi_{0,j} \frac{\rho (\rho \times e)^e}{\lambda'' (1 - \rho)^2 \times (e)!} \quad (24)$$

$\rho = \lambda'' / (e \times \mu)$ and

$$\lambda'' = (\lambda + \lambda') \left(1 - \sum_{j=0}^{\infty} \pi_{k,j} \right) \quad (25)$$

Finally, we have T as:

$$T = T_b + T_v \quad (26)$$

Request rejection rate, R , can be expressed as the ratio of the number of rejected tasks, due to the capacity constraint or PM failures/errors, to the total number of requests submitted to the IaaS data-center.

For user satisfaction, a low rejection rate is always preferable.

$$R = \theta \sum_{j=0}^{\infty} \pi_{k,j} \quad (27)$$

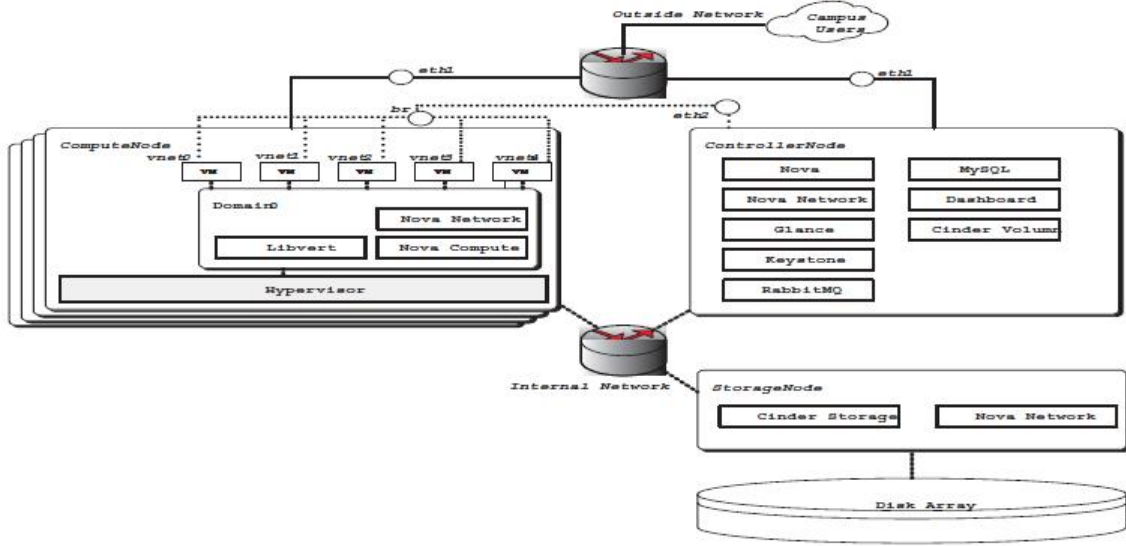


Fig. 3: The architectural view of the Course-Management and Assignment-Submission cloud

4. CASE STUDY AND MODEL VALIDATION

For the model validation purpose, we conduct a case study on a real-world cloud data-center, the Course-Management and Assignment-Submission cloud for undergraduate students of ChongQing University (CQU). Its protocol and architectural views are illustrated in Fig. 3

The cloud system is based on a symmetric server group of 6 Sugon I450 servers (4-CPU Intel Xeon 5506/128G RAM/15TB RAID but only 3-CPU/8G RAM/4TB RAID is assigned as cloud users' space). Each PM can therefore concurrently support no more than 32 VMs. The capacity of the waiting buffer for requests c is 16. The faulty rate f is 0.13-0.79%. The occurrence rate of impatient wait is 11.3% when the waiting buffer is fully occupied, meaning that $\vartheta=0.113$.

As shown in Table 1, the logfile covers time-stamps of each request's arrival and departure time in consecutive periods of 60 minutes from 09:00 to 22:00, Feb. 26, 2016.

For the model validation purpose, we derive 90% confidence intervals from the experimental performance data. By using a normal distribution as the fitting function, we derive the confidence interval of T as:

$$intv(T) = [\bar{ct} - z_{1-\alpha/2} \frac{sdv}{\sqrt{\hat{s}}}, \bar{ct} + z_{1-\alpha/2} \frac{sdv}{\sqrt{\hat{s}}}] \quad (28)$$

where sdv means standard deviation, z the z -distribution, and α the confidence level.

Finally, the confidence interval of R is also based on a Bernoulli distribution as the fitting function:

$$intv(R) = [\bar{r} - z_{1-\alpha/2} \sqrt{\frac{\bar{r} - \bar{r}^2}{\hat{s}}}, \bar{r} + z_{1-\alpha/2} \sqrt{\frac{\bar{r} - \bar{r}^2}{\hat{s}}}] \quad (29)$$

Fig. 4 implies the correctness of the proposed theoretical model.

Fig. 5 illustrates performance changes with variations in request arrival rates when

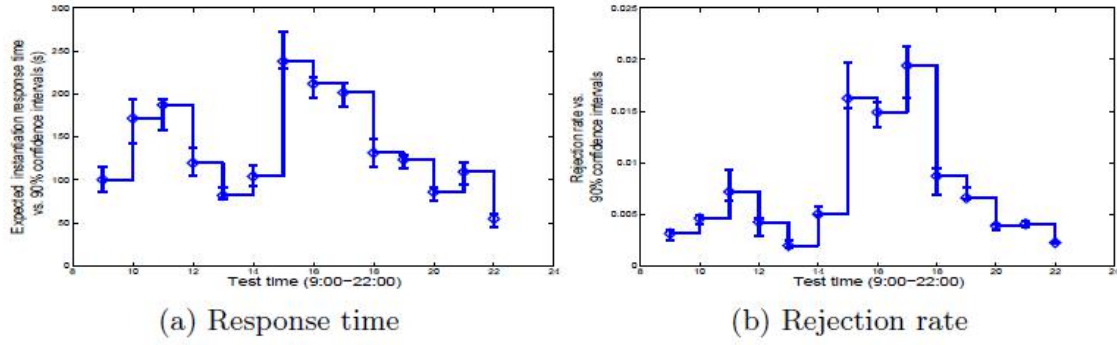


Fig. 4: Validation through confidence interval check

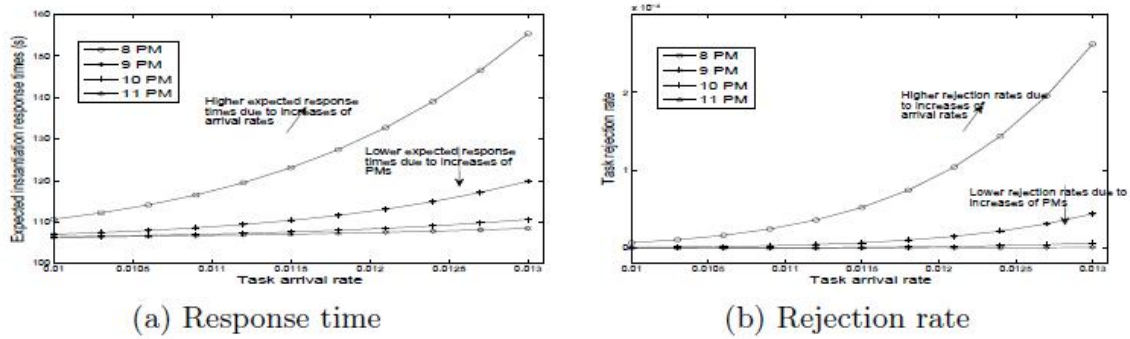


Fig. 5: Analytical performance results vs. arrival rate at different number of PMs

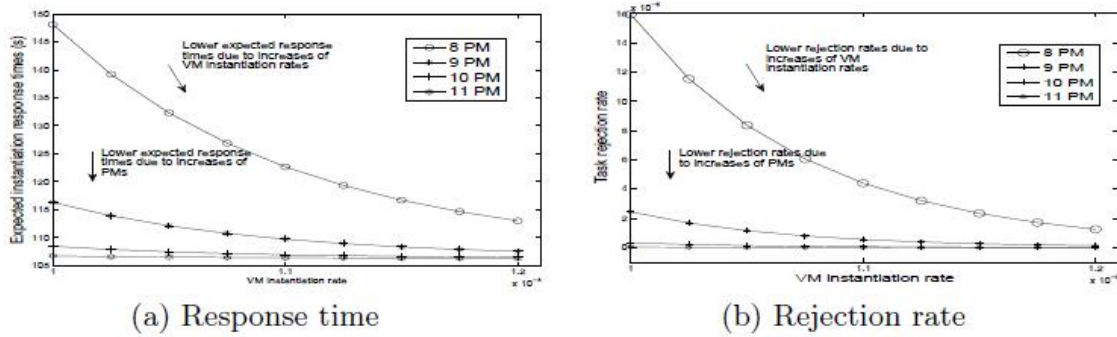


Fig. 6: Analytical performance results vs. VM instantiation rate at different number of PMs

$m=2, c=8, \mu=0.00125, \mu'=0.01, g=4, f=0.08, \vartheta=0.1$. Increasing arrival rate leads to higher expected response time and rejection rate. IaaS cloud maintains a small number of PMs. It can be seen that clouds with more PMs are more resistant to performance loss when arrival rate increases.

Fig. 6 illustrates performance changes with variations in VM instantiation rates when $m=2, c=8, \mu'=0.01, g=4, f=0.08, \vartheta=0.1, \lambda=0.01$. Increasing VM instantiation rate

leads to lower expected response time and rejection rate. It can also be seen that clouds with fewer PMs are more sensitive to performance improvements when VM instantiation rate increases.

Fig. 7 illustrates performance changes with variations in buffer sizes. The growth of such size leads to lower expected response time and rejection rate. Such performance improvements are strong when the size is small. It is also seen that clouds with more PMs always have higher performance.

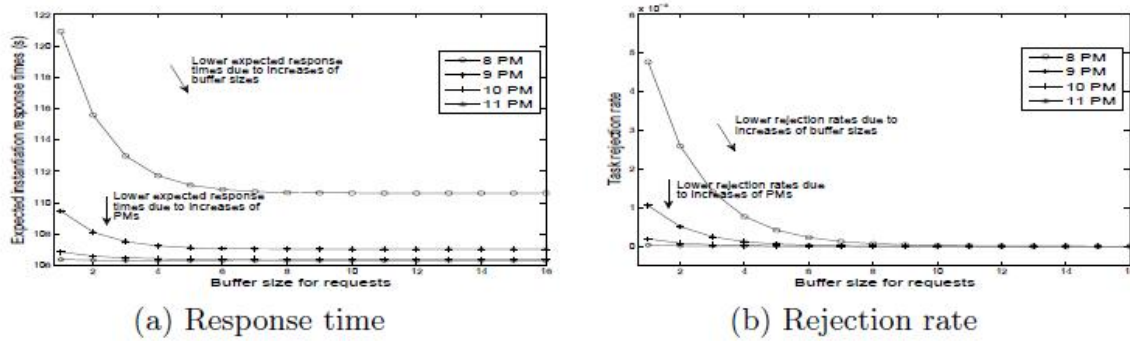


Fig. 7: Analytical performance results vs. request buffer size at different number of PMs

10 R. Zhu, Y. Zhu (2009). Performance Analysis of Call Centers Based on M/M/s/k+G Queue with Retrial, Feedback and Impatience. Proc.Int. Conf. on Grey Sys. and Inte. Services, pp. 1779--1784.

5. CONCLUSIONS AND FURTHER STUDIES

A comprehensive performance-determination model is proposed in this work for failure/error-prone IaaS cloud data-centers with request rejection and resubmission. We consider expected request response time and request rejection as the performance metrics and study the impact of varying system conditions (error intensity, VM instantiation rate, multiplexing ability, request load, etc.) on cloud performance. For the model validation purpose, we conduct a confidence interval check based on performance test results of a real-world cloud application.

6. REFERENCES

- 1 J. Rao, Y. D. Wei, J. Y. Gong, C. Z. Xu (2013). QoS Guarantees and Service Differentiation for Dynamic Cloud Applications, IEEE Trans. Netw. Service Manag., vol. 10, pp. 43--55.
- 2 S. Ostermann, R. Iosup, N. Yigitbasi, T. Fahringer (2009). A performance analysis of EC2 cloud computing services for scientific computing. Proc. Int. Conference on Cloud Computing, pp. 931--945.
- 3 N. Yigitbasi, A. Iosup, D. Epema, S. Ostermann (2009). C-meter: A framework for performance analysis of computing clouds. Proc. Int. Symp. on Cluster Comput. and the Grid, pp. 472--477.
- 4 E. Deelman, G. Singh, M. Livny, J. B. Berriman, J. Good (2008). The cost of doing science on the cloud: the Montage example. Proc. Int. Conf. High Perf. Comput. Netw. Storage and Anal., pp. 1--12.
- 5 Y. N. Xia, M. C. Zhou, X. Luo, S. C. Pang, Q. S. Zhu, J. Li (2015). Stochastic Modeling and Performance Analysis of Migration-enabled and Error-prone Clouds. IEEE Trans. Indus. Info., vol. 11, pp. 495--504.
- 6 K. Xiong, H. Perros (2009). Service performance and analysis in cloud computing. Proc. World Conf. on Services-I, pp. 693--700.
- 7 Y. S. Dai, B. Yang, J. Dongarra, G. Zhang (2009). Cloud service reliability: Modeling and analysis. Proc. IEEE Pacific Rim Int. Symp. On Depend. Comput., pp. 784--789.
- 8 B. Yang, F. Tan, Y. S. Dai, S. C. Guo (2009). Performance evaluation of cloud service considering fault recovery. Proc. Int. Conf. on Cloud Comput., pp. 571--576.
- 9 X. Meng (2010). Efficient Resource Provisioning in Compute Clouds via VM Multiplexing. Proc. Int. Conf. Aut. Comput. (ICAC2010), pp. 11--20.