2016 IEEE 9th International Conference on Cloud Computing

Systemic Risks in the Cloud Computing Model: Complex Systems Perspective

V. Marbukh
National Institute of Standards and Technology
100 Bureau Drive, Stop 8910
Gaithersburg, MD 20899-8910
E-mail: marbukh@nist.gov

Abstract—**This paper reports on quantification and management of inherent systemic risk/performance tradeoff in the cloud computing model. We view Cloud as a Complex System and associate the systemic risks with a possibility of system phase transition to the undesirable persistent state. Our analysis under mean-field and fluid approximations suggests a shift in cloud architecture design and operation paradigm from maximizing the economic benefits to management and optimization of the inherent systemic risk/benefit tradeoffs. We argue that economics makes this tradeoff more pronounced by driving Cloud service providers towards the boundary of the operational regime, and thus increasing risk of overload when the system does not have sufficient capacity for sustaining the exogenous demand.**

Keywords-cloud computing model, dynamic resource sharing, overload, systemic risk.

I. INTRODUCTION

The NIST definition lists the following five essential characteristics of the cloud computing model: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service [1]. While these traits are oriented towards maximization of the economic and user convenience benefits, they may also be a source of risks and drawbacks due to misuse of the allowed flexibility in resource provisioning [2]. This paper reports on quantifying and managing some of these systemic risk/performance tradeoffs of dynamic resource sharing made possible by the high degree of resource interconnectivity. These tradeoffs are due to the benefits of accommodating occasional resource demand/supply imbalances being inherently associated with risks of local overload spreading over a sizable portion of the system.

Due to intractability of the conventional performance models of a realistic size Cloud, we employ methodology of Complex Systems [2]. As opposed to the conventional view of systemic overload as being continuous with respect to the exogenous utilization, our analysis under mean-field and fluid approximations [3]-[4] indicates a possibility of abrupt/discontinuous systemic overload, which results in the system transitioning to a persistent congested mode through cascades of local overloads. This suggests quantification of the systemic risk of overload by taking into account not only "likelihood" of sustaining the exogenous demand, but also the gradual/continuous or abrupt/discontinuous nature of the overload if the exogenous demand becomes unsustainable.

The paper is organized as follows. Section II introduces operational and performance models of a system of shared resources. Section III analyzes a symmetric system under mean-field approximation and proposes a Perron-Frobenius based approach to analysis of non-symmetric systems under fluid approximation. Finally, Section IV concludes and outlines directions of future research.

II. RESOURCE SHARING MODELS

Subsection A introduces a performance model of statically shared resources, and demonstrates that under broad assumptions economic pressures drive system towards full utilization. Subsection B introduces a performance model of dynamically shared resources.

A. Operational Models

Consider a system with I classes of jobs (requests) and J service groups, where group $j=1,..,J$ includes N_j servers and a buffer capable of holding up to B_j jobs. Jobs of class $i=1,..,I$ arrive following a Poisson process of rate Λ_i, and have an exponentially distributed service time with average $1/\mu_{ij}$ on a class $j=1,..,I$ server. We assume a service strategy which either rejects or accepts an arriving job. In the latter case the job stays until service is completed. We also assume a work-conserving service discipline which does not allow an idle server in a group with at least one buffered job.

Static routing strategy is characterized by probabilities q_{ij} that an arriving request of class i is routed to server group j, where $\sum_j q_{ij} \le 1$ and rejection probabilities $q_{i0} := 1-\sum_j q_{ij}$ characterize admission strategy. We assume that on average, demand for the resources and supply of these resources are matched, i.e., the system is capable of accommodating the entire demand:

$$\sum_j q_{ij} = 1,\ i=1,..,I, \tag{1}$$

and system has almost no spare capacity:

$$\rho_j := (1/N_j)\sum_i q_{ij}\Lambda_i/\mu_{ij} \approx 1,\ j=1,..,J \tag{2}$$

2159-6190/16 $31.00 © 2016 IEEE
DOI 10.1109/CLOUD.2016.136

Despite conditions (1)-(2) appear to be restrictive, it can be shown that in market economy they arises naturally as a result of market pressures. Assuming that service provider controls demand through service pricing in attempt to maximize the revenue, conditions (1)-(2) are the result of this revenue maximization, which also determines routing probabilities q_{ij}. In an important particular case of one-to-one correspondence between request classes and service groups, when "native" service is at least as efficient as "non-native":

$$I = J\,,\ \mu_{ij} \le \mu_{ii}\,,\ i, j = 1,..,I\,,\ i \ne j\ , \tag{3}$$

it can be shown that under some natural assumptions on the model parameters, provider revenue maximization yields routing which allocates requests to native servers only:

$$q_{ij} = 1 \text{ if } i = j \text{, and } q_{ij} = 0 \text{ otherwise.} \tag{4}$$

In practice, due to variability of the exogenous demand and limited system reliability, system may not have sufficient resources to accommodate occasional resource demand/supply imbalances, e.g., because delay requirements may limit buffer sizes. Cloud computing model is expected to mitigate these imbalances with dynamic resource sharing made possible by high degree of resource interconnectivity. In this short paper we consider a generic model of dynamic resource sharing in a case of "small" demand/supply imbalances, when (1)-(2) hold and probability that service group $j = 1,..,J$ has neither available servers nor buffering space, is small.

Introduce vector $\delta = (\delta_j, j \in J)$, where $\delta_j = 0$ if server group j has available resources, i.e., a server, or buffering space, or both. If however, server group j has neither a server nor buffering space available, then $\delta_j = 1$. Since according to our assumptions $\bar{\delta}_j := E[\delta_j] << 1$, $j = 1,..,J$, the main effect of dynamic resource sharing can be described by conditional rerouting probabilities q_{ijk} that a class i request initially routed to server group j is immediately rerouted to server group k in an unlikely case $\delta_j = 1$. The effect of dynamic resource sharing of next order of magnitude with respect to $\bar{\delta}_j$ can be described by the corresponding "second attempt rerouting probabilities," etc.

For a dynamic resource sharing discipline allowing a single rerouting attempt with probabilities q_{ijk}, loss probability for an arriving request of class i is

$$L_i = \sum_j q_{ij} \bar{\delta}_j \sum_{k \ne j} q_{ijk} E[\delta_k | \delta_j = 1], \tag{5}$$

where $E[\delta_k | \delta_j = 1]$ is the conditional expectation of δ_k, given $\delta_j = 1$.

B. Performance Models

Our analysis is based on a mean-field type approximation, which neglects correlations between blockings in different service groups: $E[\delta_k | \delta_j = 1] \approx E[\delta_k]$, and thus allows us to approximate loss (5) as follows:

$$L_i \approx \widetilde{L}_i = \sum_j q_{ij} \bar{\delta}_j \sum_{k \ne j} q_{ijk} \bar{\delta}_k\ , \tag{6}$$

where $\bar{\delta}_k := E[\delta_k]$. Dynamic resource sharing results in additional load due to allowing requests a second attempt to obtain service. The corresponding additional utilization for server group j is

$$\beta_j = \frac{1}{N_j} \sum_i \frac{\Lambda_i}{\mu_{ij}} \sum_{k \ne j} q_{ik} q_{ikj} \bar{\delta}_k\ . \tag{7}$$

We propose to approximate probabilities $\bar{\delta}_j$ by the Erlang formula with S_j servers, buffer size B_j, and utilization $\rho_j + \beta_j$ [5]:

$$\bar{\delta}_j \approx Erl(\rho_j + \beta_j, S_j, B_j), \tag{8}$$

where

$$Erl(\rho, S, B) = \frac{\dfrac{(S\rho)^{S+B}}{S!S^B}}{\displaystyle\sum_{k=0}^{S} \frac{(S\rho)^k}{k!} + \frac{(S\rho)^{S+B}}{S!} \frac{1-\rho^{B+1}}{1-\rho}} \tag{9}$$

Substituting (7) into (8) we obtain a closed system of J non-linear, fixed-point equations for $\widetilde{\delta}_j$, $j = 1,..,J$.

Fluid approximation describes a limit of large service groups $S_j + B_j >> 1$, when equation (9) takes the following form [6]:

$$\widetilde{\delta}_j = [1 - 1/(\rho_j + \beta_j)]^+ \tag{10}$$

since

$$(1 - 1/\rho)^+ = \min_{S+B \to \infty} Erl(\rho, S, B)\,, \tag{11}$$

where $(x)^+ := \max(0, x)$. Combining (10) with (7) we obtain a closed system of fixed-point equations for $\widetilde{\delta}_j$, $j = 1,..,J$.

Following conventional interpretation of mean-field and fluid approximations [7], we interpret multiple solutions of the corresponding approximate systems as describing the metastable, i.e., persistent equilibria of the original Markov model. When $\rho_j \le 1$, $j = 1,..,J$, system (7), (10) has trivial solution $\widetilde{\delta}_j = 0$, which describes the "normal/operational" system regime without losses: $L_i = 0$, $i = 1,..,I$.

In the rest of this paper we discuss the related questions of systemic overload due to short-term and long-term resource demand/supply imbalances. The first question is related to existence of non-trivial solutions to (7), (10) for $\rho_j \le 1$, $j = 1,..,J$. The existence of such a solution indicates a possibility of the system transitioning to a persistently congested mode despite long-term resource demand/supply is balanced. This transition occurs through cascades of local congestion spreading to other parts of the system due to dynamic resource sharing. The second question is related to whether overload emerges gradually/continuously or abruptly/discontinuously as long-term exogenous loads ρ_j exit the operational region $\rho_j \le 1$, $j = 1,..,J$.

III. Systemic Risk/Performance Tradeoff

Subsection A analyzes symmetric system under mean-field approximation (7), (8)-(9). Subsection B discusses stability of the operational equilibrium under fluid approximation (7), (10).

A. Symmetric System under Mean-Field Approximation

Consider a particular case of symmetric system with native services (3)-(4), where $N_i = N$, $\Lambda_i = \Lambda$, $\mu_{ii} = \mu$, $\mu_{ij} = \mu/(1+\chi)$; $i, j = 1,..,I$, $i \ne j$, parameter $\chi \ge 0$ characterizes inefficiency of a non-native service as compared to the native service, and dynamic resource sharing is characterized by probabilities $q_{ijk} = q$, $k \ne i, j$.

Figure 1 sketches the persistent loss $\tilde{L}$ vs. exogenous system load $\rho := \Lambda/(N\mu)$ under mean-field approximation.

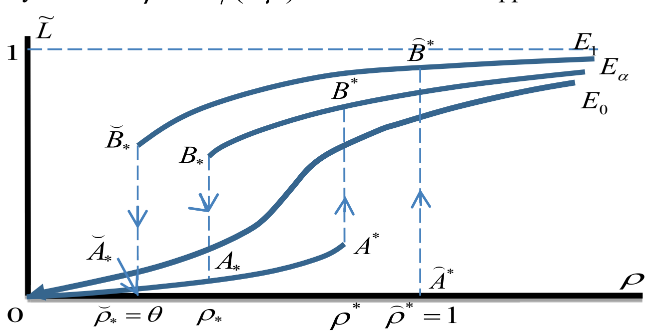

Figure 1. Persistent loss vs. exogenous utilization.

Curve $0E_0$ sketches loss rate $\tilde{L}$ for sufficiently low level of resource sharing α, when mean-field equation has unique solution $\tilde{\delta}$ for all ρ. For sufficiently large resource sharing level α, large service groups: $N + B >> 1$, and sufficiently inefficient non-native service, i.e., small θ, mean-field equation has two stable solutions $\tilde{\delta}_*$ and $\tilde{\delta}^*$ for intermediate load $\rho_*(\alpha) < \rho < \rho^*(\alpha)$. As load ρ "slowly" increases, the loss $\tilde{L}$ follows curve $0A_*(\alpha)A^*(\alpha)B^*(\alpha)E_\alpha$. As ρ "slowly" decreases, the loss $\tilde{L}$ follows curve $E_\alpha B^*(\alpha)B_*(\alpha)A_*(\alpha)0$. Curves $0A_*(\alpha)$ and $E_\alpha B^*(\alpha)$ correspond to the globally stable "normal" and "congested" system equilibria respectively. Branches $A_*(\alpha)A^*(\alpha)$ and $B_*(\alpha)B^*(\alpha)$ correspond to the coexisting "normal" and "congested" metastable system equilibria respectively for intermediate load $\rho_*(\alpha) < \rho < \rho^*(\alpha)$.

Discontinuities at the critical loads $\rho_*(\alpha)$ and $\rho^*(\alpha)$ as well as the hysteresis loop $A_*(\alpha)A^*(\alpha)B^*(\alpha)B_*(\alpha)$ indicate discontinuous transition. Curves $0\breve{A}_*\hat{A}^*$ and $\breve{B}_*\hat{B}^*E_1$ represent loss $\tilde{L}$ in a case of complete resource sharing $\alpha = 1$, when $\rho_{**} := \rho_*(1) = \theta$ and $\rho^{**} := \rho^*(1) = 1$. Increase in the resource sharing increases "spread" between the normal and congested metastable regimes by reducing loss in the normal regime and increasing loss in the congested regime.

Figure 2, which sketches persistent loss vs. "slowly" changing level of resource sharing α, indicates a combination of positive and negative effects of the dynamic resource sharing on the system performance.

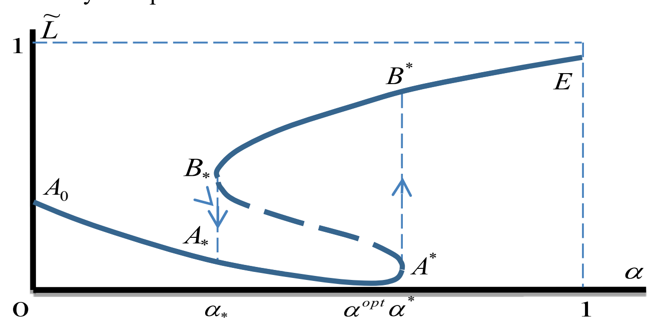

Figure 2. Persistent loss vs. level of resource sharing.

As resource sharing α "slowly" increases (decreases), loss rate $\tilde{L}$ follows curve $A_0 A_* A^* B^* E$ $(EB^* B_* A_* A_0)$.

B. General System under Fluid Approximation

Figure 3 sketches persistent loss vs. exogenous utilization in a symmetric system under fluid approximation for $(1+\chi)q < 1$

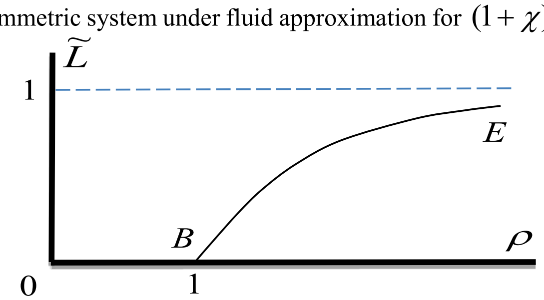

Figure 3. Persistent loss: fluid approximation, $(1+\chi)q < 1$.

Figure 4 sketches persistent loss vs. exogenous utilization in a symmetric system under fluid approximation for $(1+\chi)q>1$

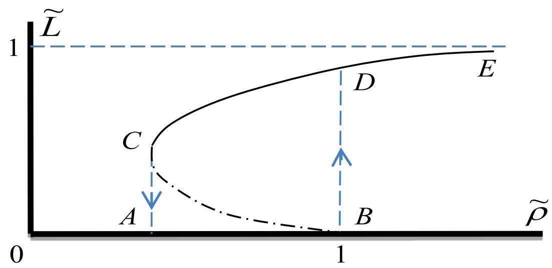

Figure 4. Persistent loss: fluid approximation, $(1+\chi)q>1$.

For $\rho \le 1$ solution $\widetilde{\delta}=0$ corresponds to the desirable "operational" system equilibrium with loss rate $\widetilde{L}=0$. Solution $\widetilde{\delta}=\widetilde{\delta}^*>0$, represented by curve CDE, corresponds to the "congested" system equilibrium with positive loss $L^*=(1-q+q\widetilde{\delta}^*)\widetilde{\delta}^*$. Curve CB represents unstable equilibrium fixed point separating stable equilibrium points $\widetilde{\delta}=0$ and $\widetilde{\delta}=\widetilde{\delta}^*$. While the "operational" and "congested" equilibria are globally stable for "light" and "heavy" exogenous loads respectively, these equilibria are metastable for "intermediate" exogenous load. As exogenous load ρ "slowly" increases from 0 to ∞, system follows curve $0ABDE$. As exogenous load ρ "slowly" decreases from ∞ to 0, system follows curve $EDCA0$. Hysteresis loop $BDCAB$ is indicative of the discontinuous "phase transition."

While for a symmetric system, we characterized continuous/discontinuous nature of overload by analyzing the entire system phase diagram, for a general, realistic-size system, high dimension of the mean-field and fluid approximations make quantitative evaluation of the system phase diagram computationally infeasible. In the rest of this subsection we suggest a computationally tractable criterion of continuous vs. discontinuous overload, which is based on type of bifurcation of the fluid approximation (7), (10) on the boundary of the operational regime $\rho_i \le 1$, $i=1,..,I$.

Consider linearized fixed-point system (7), (10) as $\rho_j \downarrow 1$

$$\widetilde{\delta}_i=\rho_i-1+\frac{1}{N_i}\sum_{j\neq i}\left(\sum_k \frac{\Lambda_k}{\mu_{ki}} q_{kj} q_{kji}\right)\widetilde{\delta}_j\,, \tag{12}$$

$i=1,..,J$. Linear system (12) has solution $\widetilde{\delta}\to 0$ as $\rho_i \downarrow 1$ if $\gamma(A)<1$, and does not have such a solution if $\gamma(A)>1$, where $\gamma(A)$ is Perron-Frobenius eigenvalue of matrix $A=(a_{ij})_{i,j=1}^J$ with non-negative elements

$$a_{ij}=\frac{1}{N_i}\sum_k \frac{\Lambda_k}{\mu_{ki}} q_{kj} q_{kji}\,, \tag{13}$$

and matrix A is assumed irreducible [8].

Thus, condition

$$\gamma(A)<1 \tag{14}$$

ensures gradual/continuous overload as exogenous utilizations ρ_j, $j=1,..,J$ cross unity from below, and is a generalization of the corresponding criterion $(1+\chi)q<1$ for a particular case of a symmetric system. Due to inherent uncertainties in the system parameters, closeness $\gamma(A)$ to unity is one of indications of systemic risk of overload.

IV. Conclusion & Future Research

This paper has suggested that the economic benefits of dynamic resource sharing are inherently associated with systemic risk of overload, which may be either gradual/continuous or abrupt/discontinuous. The conventional economic efficiency maximization framework should be replaced with constrained optimization, which controls the systemic risk of overload. Future work should address the practicality of the proposed Perron-Frobenius based optimization framework at the system design and operational stages. Of particular interest is a potential ability of online measurements of the corresponding Perron-Frobenius eigenvalue to provide "early warning signals" of the system approaching the instability/breaking point [9] for the purpose of initiating appropriate control actions. Since abrupt/discontinuous overload takes the form of "congestion collapse," this constrained optimization should give higher priority to mitigation of the systemic risk of congestion collapse identified with abrupt/discontinuous systemic overload.

References

[1] The NIST Definition of Cloud Computing NIST Special Publication 800-145, http://csrc.nist.gov/publications/PubsSPs.html#800-145.

[2] D. Helbing, Globally networked risks and how to respond, *Nature*. 497, 51-59, (02 May 2013).

[3] V. Marbukh, "On systemic risk in the cloud computing model," 26th International Teletraffic Congress (ITC), Karlskrona, Sweden, 2014.

[4] V. Marbukh, "Perron-Frobenius measure of systemic risk of cascading overload in complex clouds: work in progress," *IFIP/IEEE International Symposium on Integrated Network Management*, Gent, Belgium, 2013.

[5] L. Kleinrock, Queueing Systems Volume 1: Theory, John Wiley & Sons, 1975.

[6] A.L. Stolyar and E.Yudovina, Systems with large flexible server pools: Instability of "natural load balancing," Annals of Applied Probability, Vol. 23, No. 5, 2013, pp. 2099-2138.

[7] Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi, "Stochastic networks with multiple stable points," Annals of Probability, 36(1):255-278, 2008.

[8] S.U. Pillai, T. Suel, and S. Cha, "The Perron-Frobenius theorem: some of its applications," IEEE Signal Processing Magazine, , Vol. 22, Issue 2, March 2005, pp. 62-75.

[9] M. Scheffer, J. Bascompte, W.A. Brock, V. Brovkin, S.R. Carpenter, V. Dakos, H. Held, E.H. van Nes, M. Rietkerk and G. Sugihara, Early-warning signals for critical transitions, *Nature* 461:3, 53-59.