On Systemic Risk in the Cloud Computing Model

V. Marbukh

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

Abstract—The major benefits of the emerging cloud computing infrastructure include elimination of the fixed cost and reduction in the marginal cost for the users due to the economy of scale and dynamic resource sharing. This paper argues that these unquestionable economic benefits of dynamic resource sharing are inherently associated with systemic risks and drawbacks. This possibility necessitates shift in cloud architecture design and operation paradigm from maximizing the economic benefits to managing and optimization of the inherent systemic risk/benefit tradeoffs. This paper is a first step towards exposure and quantification of inherent tradeoffs involving system economic efficiency and systemic risk of cascading overload in the cloud model of shared resources.

Keywords—cloud computing model, dynamic resource sharing, economic efficiency, systemic risk, tradeoffs.

I. INTRODUCTION

The NIST definition lists five essential characteristics of the cloud computing model: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service [1]. The economic and convenience advantages of the shared cloud computing architecture as compared to the conventional model of owning computer resources are unquestionable. Among economic advantages are elimination of the fixed cost and reduction of the marginal cost for users of cloud computing infrastructure due to the economy of scale. One can easily identify significant similarities between the current emergence and evolution of the cloud computing model and the earlier emergence and evolution of the Internet. Of course, these similarities are not accidental; they are the result of the same economic pressures. Indeed, convenience for the end users and economic efficiency, which are due to the end-to-end principle and packet switching, resulted in Internet success in competition with telephone network [2].

Assuming that historic parallels between the Internet and Cloud models are rooted in reality, one may attempt to learn major Cloud evolution trends from the historic lessons of the Internet evolution. Probably, the most important lesson is that the core principles of the Internet model, e.g., the end-to-end principle and packet switching, are inherently associated with serious risks and drawbacks, which have led to numerous proposals for modification of the Internet model for the purpose of mitigating and managing of the corresponding risk/benefit tradeoffs [2]. One may expect that the cloud computing model of shared computing and communication resources inherits

both benefits and drawbacks of the Internet model of shared communication resources.

However, current research is mostly concentrated on the benefits of the cloud model [3]-[5]. A notable exception to this research trend is [6]. Due to highly dynamic operating environment, static load balancing, which attempts to match the expected demand to the system resources, may be unable to accommodate unavoidable fluctuations in the exogenous demand even with available buffering. Demand fluctuations can be accommodated with dynamic load balancing allowed by the cloud computing model. However, assuming infinite buffers, [6] rigorously proves that these benefits of dynamic load balancing are achieved at the cost of “much longer” queues and delays as compared to static load balancing. While important, these rigorous results rely on a number of unrealistic assumptions, including infinite buffers.

This paper proposes a mean-field approximation, which overcomes this and some other limitations of model [6] at the cost of losing mathematical rigor. Our results indicate that in a practically relevant parameter region, benefits of dynamic load balancing in the “normal operating regime” come with risks of the system abruptly transitioning to the “persistently congested mode” through the process of cascading overload. A similar phenomenon has been predicted in a rather unrealistic model of a cloud without buffering [7], and has been observed in networks with dynamic routing [8]-[9].

This paper can be viewed as a first step in the direction of paradigm shift in Cloud design and operation: from maximizing economic benefits to identifying and managing the relevant risk/benefit tradeoffs. In particular, we demonstrate that economic pressures push the system to the stability boundary of the normal operating regime, making the tradeoff between “normal system performance” and systemic risk of transitioning to congested regime more acute. We also suggest using the Perron-Frobenius eigenvalue [10] of the “reduced” system performance model as a measure of the systemic risk of cascading overload. Applicability of the Perron-Frobenius theory [10] in systems with dynamic resource sharing is due to the positive feedback in the system component overload, which can spill over to the “neighboring” components.

The paper is organized as follows. Section II introduces a cloud Markov performance model. Since this Markov model is computationally intractable for realistic size clouds, Section III proposes tractable mean-field and fluid approximate performance models. Based on these approximations, Section IV analyses performance of a symmetric cloud, when

dimension of mean-field performance model does not depend on the system size. Section V quantifies systemic risk and risk/performance tradeoff for a symmetric system, and proposes an approach to such a quantification for a general system. Finally, Section VI briefly summarizes and outlines directions of future research.

II. CLOUD MODELS

Subsection A describes cloud operations, including job arrivals, admission, routing and servicing, and Subsection B introduces the corresponding Markov performance model.

A. Service Model

Consider a system with I classes of jobs (requests) and I service groups, where group $i=1,\dots,I$ includes N_i servers and a buffer capable of holding up to B_i jobs. Jobs of class $i=1,\dots,I$ arrive following a Poisson process of rate Λ_i , and have an exponentially distributed service time with average $1/\mu_{ij}$ on a class $j=1,\dots,I$ server. We assume that $\mu_{ij} \leq \mu_{ii}$, $i, j=1,\dots,I$, $i \neq j$, i.e., “native” service is more efficient than “non-native”.

We model limited system reliability by assuming that entire server group $i=1,\dots,I$ is either operational with probability $1-f_i$ or non-operational, i.e., failed, with probability $f_i \in [0,1]$. Operational or failed statuses of different server groups are jointly statistically independent. Note that our failure model implies that failures/recoveries occur on a much slower time scale than request arrivals/departures.

We assume a service strategy which either rejects or accepts an arriving job to some service group where the job stays until service is completed. We also assume work-conserving service disciplines which do not allow idle server(s) in a group which has at least one buffered job. We consider dynamic admission and routing strategies, which are based on the resource availability characterized by the vector $\delta = (\delta_j, j \in J)$ where the binary variable $\delta_i = 0$ if server group i has available resources, i.e., a server, or buffering space, or both. If however server group i has neither a server or a buffering space available, then $\delta_i = 1$.

We consider admission and routing strategies which immediately either reject an arriving class i job with probability q_{i0} , or route this job to a “native” service group i with probability $q_{ii} = 1 - q_{i0}$ if $\delta_i = 0$. If $\delta_i = 1$, the job is immediately either lost with probability $1 - \alpha_i$, or attempts to obtain service at some non-native service group $j \in J_i \subseteq \{1,\dots,I\} \setminus i$ with probability α_i . Thus, if

$\prod_{m \in J_i} \delta_m = 0$, i.e., at least one service group $k \in J_i$ has available resources, the job attempting to obtain a non-native service is routed to service group k with probability

$$g_{ik}(\delta) = (1 - \delta_k) q_{ik} / \sum_{m \in J_i} (1 - \delta_m) q_{im}, \quad (1)$$

where $\sum_{k \in J_i} q_{ik} = 1$, $q_{ik} > 0, \forall k \in J_i$. If however

$\prod_{m \in J_i} \delta_m = 1$, the request is finally lost.

In the above admission and routing strategies, the admission probabilities q_{ii} control the expected exogenous demand in order to match it to the system capacity, sets J_i determine the system “topology”, probabilities α_i determine the level of resource sharing (pooling), and probability matrix $Q = (q_{ij})_{i,j=1}^I$ determines the load balancing/routing strategy.

Assuming that each serviced request of class i brings revenue w_i , the system performance can be characterized by the overall revenue rate

$$R = (1 - L) \sum_i w_i \Lambda_i \quad (2)$$

where the portion of lost revenue is

$$L = \left(\sum_i w_i \Lambda_i \right)^{-1} \left(\sum_i w_i \Lambda_i L_i \right) \quad (3)$$

and the loss probability for a class i job is

$$L_i = 1 - q_{ii} + \left(1 - \alpha_i + \alpha_i E \left[\prod_{j \in J_i} \delta_j \mid \delta_i = 1 \right] \right) q_{ii} E[\delta_i] \quad (4)$$

Our goal is to quantify ability of dynamic resource sharing (resource pooling) to mitigate the mismatch between the uncertain (due to limited reliability) system capacity and fluctuating exogenous demand, assuming that the systematic mismatches have been already eliminated through admission control, capacity planning, pricing, or their combination. In order to isolate the effect of load fluctuations we further assume $q_{ii} = 1$, which allows us to simplify (4) as follows:

$$L_i = \left(1 - \alpha_i + \alpha_i E \left[\prod_{j \in J_i} \delta_j \mid \delta_i = 1 \right] \right) E[\delta_i]. \quad (5)$$

B. Markov Performance Model

Introduce the indicators $\omega_i \in \{0,1\}$, where $\omega_i = 0$ if service group $i=1,\dots,I$ is operational, and $\omega_i = 1$ otherwise. Since different service groups fail independently from each other, the joint probability distribution of $\omega = (\omega_i, i=1,\dots,I)$ is

$$\Omega(\omega) = \prod_{i=1}^I [f_i^{\omega_i} (1 - f_i)^{1-\omega_i}]. \quad (6)$$

Given vector $\omega = (\omega_i)$, cloud evolution is described by a

homogeneous in time Markov process $(n_{ij}(t), m_{ij}(t))$, where $n_{ij}(t)$ is the number of class i requests served by group j , and $m_{ij}(t)$ is the number of class i requests waiting in the group j buffer at moment t .

In a case of finite buffers: $B_j < \infty$, Markov process $(n_{ij}(t), m_{ij}(t))$ has unique steady-state distribution $P(n_{ij}, m_{ij} | \omega) = \lim_{t \rightarrow \infty} \Pr \{n_{ij}(t) = n_{ij}, m_{ij}(t) = m_{ij} | \omega\}$ for any initial condition $(n_{ij}(0), m_{ij}(0))$. This steady-state distribution can be determined from the corresponding steady-state Kolmogorov equations, which we omit here due to limited space.

After solving this steady-state Kolmogorov system, one can calculate conditional loss probabilities

$$L_i(\omega) = \left(1 - \alpha_i + \alpha_i E \left[\prod_{j \in J_i} \delta_j | \delta_i = 1, \omega \right] \right) E[\delta_i | \omega] \quad (7)$$

where

$$\delta_j = \begin{cases} 1 & \text{if } \sum_i n_{ij} = N_j, \sum_i m_{ij} = B_j \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

and mathematical expectations are evaluated with respect to the conditional steady-state distribution $P(n_{ij}, m_{ij} | \omega)$. The unconditional steady-state loss probabilities can be obtained by averaging of the corresponding conditional loss probabilities:

$$L_i = \sum_{\omega} L_i(\omega) \Omega(\omega) \quad (9)$$

where distribution $\Omega(\omega)$ is given by (6).

Unfortunately, the astronomically high dimension of the Kolmogorov system, which has to be solved for each of 2^I possible indicator vectors $\omega \in \{0,1\}^I$, makes the Markov performance model computationally infeasible even for moderate size systems. The source of this difficulty is the correlations between states of different service groups due to the resource sharing allowing a non-native service.

III. APPROXIMATIONS

Subsection A proposes a mean-field approximation, which is computationally tractable for medium-size systems. Subsection B combines this mean-field approximation with fluid approximations in cases of large service groups or large number of service groups and large level of resource sharing.

A. Mean-field Approximation

Steady-state Kolmogorov equations can be explicitly solved in a particular case when only native services are allowed: $\alpha_i = 0$, $i = 1, \dots, I$. In this case, the random variables δ_i conditioned on vector ω are jointly statistically independent

for $i = 1, 2, \dots, I$, and thus for any subset of server groups $\{i\} \subseteq \{1, 2, \dots, I\}$:

$$E \left[\prod_{i \in \{i\}} \delta_i | \omega \right] = \prod_{i \in \{i\}} \bar{\delta}_i(\omega_i), \quad (10)$$

where $\bar{\delta}_i(\omega_i) := E[\delta_i | \omega_i]$ is

$$\bar{\delta}_i(\omega_i) = \omega_i + (1 - \omega_i) \bar{\delta}_i(0), \quad (11)$$

and $\bar{\delta}_i := \bar{\delta}_i(0)$ is

$$\bar{\delta}_i = \frac{(N_i \rho_i)^{N_i + B_i}}{N_i! N_i^{B_i}} \cdot \sum_{k=0}^{N_i} \frac{(N_i \rho_i)^k}{k!} + \frac{(N_i \rho_i)^{N_i}}{N_i!} \frac{1 - \rho_i^{B_i+1}}{1 - \rho_i}. \quad (12)$$

In (12) server group i utilization by native requests is $\rho_{ii} = \Lambda_i / (N_i \mu_{ii})$. Substituting (11)-(12) into (7), we obtain the conditional losses

$$L_i(\omega) = \left(1 - \alpha_i + \alpha_i \prod_{j \in J_i} \bar{\delta}_j(\omega_j) \right) \bar{\delta}_i(\omega_i). \quad (13)$$

Similar to [7], consider a mean-field approximation, which drastically reduces the problem dimension by postulating that (10)-(11) hold approximately:

$$E \left[\prod_{i \in \{i\}} \delta_i | \omega \right] \approx \prod_{i \in \{i\}} \tilde{\delta}_i(\omega_i), \quad (14)$$

where $\tilde{\delta}_i(\omega_i) := E[\delta_i | \omega_i]$ is

$$\tilde{\delta}_i(\omega_i) = \omega_i + (1 - \omega_i) \tilde{\delta}_i, \quad (15)$$

and $\tilde{\delta}_i := \tilde{\delta}_i(0)$ is

$$\tilde{\delta}_i \approx \frac{N_i^{N_i} \tilde{\rho}_i^{N_i + B_i}}{N_i!} \cdot \sum_{k=0}^{N_i} \frac{(N_i \tilde{\rho}_i)^k}{k!} + \frac{(N_i \tilde{\rho}_i)^{N_i}}{N_i!} \frac{1 - \tilde{\rho}_i^{B_i+1}}{1 - \tilde{\rho}_i}. \quad (16)$$

Self-consistency requires “effective utilizations” $\tilde{\rho}_j$ in (16) to satisfy the following “conservation laws”:

$$\tilde{\rho}_j = \rho_{jj} + \sum_{i \neq j} \alpha_i \tilde{\rho}_{ij}(\tilde{\delta}), \quad (17)$$

where the load to be routed, subject to admission strategy, from server group i to server group j is

$$\tilde{\rho}_{ij}(\tilde{\delta}) = \sum_{\omega} \Omega(\omega) \sum_{\delta} \rho_{ij}(\delta | \omega) \prod_j \tilde{p}_j(\delta_j | \omega_j), \quad (18)$$

$\tilde{p}_j(0|0) = 1 - \tilde{\delta}_j$, $\tilde{p}_j(1|0) = \tilde{\delta}_j$, $\tilde{p}_j(0|1) = 0$, and $\tilde{p}_j(1|1) = 1$. Given vector $\omega = (\omega_i)$, the conditional load from server group i to server group j is

$$\rho_{ij}(\delta|\omega) = \rho_{ii} \frac{\mu_{ii}}{\mu_{ji}} g_{ij}(\delta) [\omega_i + (1-\omega_i)\delta_i]. \quad (19)$$

Substituting (19) into (18) and then into (17) and (16), we obtain a closed system of I fixed point equations for I unknowns $\tilde{\delta} = (\tilde{\delta}_i)$:

$$\tilde{\delta}_i = \varphi_i(\tilde{\delta}), \quad i=1, \dots, I. \quad (20)$$

After solving system (20) for $\tilde{\delta} = (\tilde{\delta}_i)$, the conditional losses (7) can be approximated as follows:

$$\tilde{L}_i(\omega) = \left(1 - \alpha_i + \alpha_i \prod_{j \in J_i} \tilde{\delta}_j(\omega_j) \right) \tilde{\delta}_i(\omega_i), \quad (21)$$

where $\tilde{\delta}_i(\omega_i)$ are given by (15). Note that if $\alpha_i = 0$, $i=1, \dots, I$, i.e., resource sharing is not allowed, the mean-field approximation becomes identical to the Markov model. If $\alpha_i \geq 0$, i.e., resource sharing is allowed, the mean-field approximation results in I non-linear algebraic equations (20). Thus, the mean-field approximation has greatly lower dimension than the Markov description at the cost of non-linearity.

B. Fluid Approximation

Fluid approximation, which replaces probability distributions by their averages, can be justified if these averages are much greater than the corresponding standard deviations. In other words, fluid approximation neglects fluctuations of the random variables assuming that these fluctuations are small comparatively to the corresponding averages.

Usually, fluid approximation is associated with a case of “large” service groups $N_j + B_j \rightarrow \infty$, when formula (16) takes the following form:

$$\tilde{\delta}_j = \max(0, 1 - 1/\bar{\rho}_j). \quad (22)$$

Combining (22) with (17)-(19) we obtain a closed system of I fixed-point equations for $\tilde{\delta}_j, j=1, \dots, I$. This system may have stable equilibrium $\tilde{\delta}_j = 0, j=1, \dots, I$ which describes desirable system equilibrium with asymptotically lossless performance: $L_j = 0$ as $N_j + B_j \rightarrow \infty, j=1, \dots, I$. A sufficient condition for existence of this lossless system equilibrium is that the total utilization of each server group $j=1, \dots, I$ by native and non-native jobs is less than one:

$$\rho_{jj} + \sum_{i \neq j} \alpha_i \rho_{ii} f_i \frac{\mu_{ii}}{\mu_{ji}} \sum_{\omega} g_{ij}(\omega) \Omega(\omega) < 1. \quad (23)$$

In a case of a large number of service groups $I \gg 1$ and high level of resource sharing, i.e., large sets $J_i, i=1, \dots, I$, fluid approximation is applicable to the aggregate loads

conditioned on vectors δ and ω . Indeed, in this case the aggregate load on a server group is a sum, i.e., composition, of the native load and a “large” number of “small” non-native loads from other service groups. At any given moment, the aggregate load is the function of the corresponding vectors δ and ω . The fluid approximation replaces average of this function over random vectors δ and ω (17) with the corresponding function of the average vectors $\tilde{\delta}$ and $\tilde{\omega}$:

$$\tilde{\rho}_j(\tilde{\delta}) \approx \rho_{jj} + \sum_{i \neq j} \alpha_i \rho_{ii} \frac{\mu_{ii}}{\mu_{ji}} g_{ij}(\tilde{\delta}) [f_i + (1-f_i)\tilde{\delta}_i]. \quad (24)$$

In a case of both large server groups and large degree of resource sharing, further simplification is possible resulting in the following condition for existence of a “normal” operational state with lossless performance:

$$\rho_{jj} + \sum_{i \neq j} \alpha_i \rho_{ii} f_i \frac{\mu_{ii}}{\mu_{ij}} \sum_{k \in J_i} (1-f_k) q_{ik} < 1. \quad (25)$$

Note that more accurate approximations of (16) than (22) are possible depending on the asymptotic regime as $N_j + B_j \rightarrow \infty$ and $\bar{\rho}_j \rightarrow 1$. For example, if $B_j \rightarrow \infty$ and $\bar{\rho}_j \rightarrow 1$, formula (16) can be simplified as follows:

$$\bar{\delta}_j = \bar{\rho}_j^{B_j} (1 - \bar{\rho}_j) / (1 - \bar{\rho}_j^{B_j+1}). \quad (26)$$

IV. SYMMETRIC SYSTEM

This section analyzes persistent behavior of a symmetric system for which the dimension of the mean-field system is independent of the system size. Subsection A defines a symmetric system and derives the corresponding mean-field and fluid approximations. Subsection B analyses the persistent performance of a symmetric system, which is identified with equilibria of the corresponding mean-field equation. Subsection C discusses the phase diagram of a symmetric system under fluid approximation. A phase diagram separates regions in space of system parameters with different persistent system behavior.

A. Mean-field Equations

Consider a symmetric system, where $N_j = N, B_j = B, \Lambda_j = \Lambda, \mu_{ii} = \mu, f_j = f, \mu_{ij} = \theta \mu, q_{ij} = 1/(I-1)$ for $i, j=1, \dots, I, i \neq j$, and the parameter $\theta \leq 1$ characterizes the relative inefficiency of a non-native service as compared to the native service. In the case of a large number of server groups: $I \rightarrow \infty$, equation (20) takes the form of the following system of fixed point equations

$$\tilde{\delta} = \frac{N^N \tilde{\rho}^{N+B}}{N!} \left/ \left(\sum_{i=0}^N \frac{(N\tilde{\rho})^i}{i!} + \frac{(N\tilde{\rho})^N}{N!} \frac{1 - \tilde{\rho}^{B+1}}{1 - \tilde{\rho}} \right) \right., \quad (27)$$

and

$$\tilde{\rho} = \left(1 + \frac{\alpha}{\theta} \frac{1 - (1-f)(1-\tilde{\delta})}{(1-f)(1-\tilde{\delta})} \right) \rho, \quad (28)$$

where the exogenous load is $\rho = \Lambda/(N\mu)$.

Substituting (28) into the right-hand side of (27) we obtain a single fixed-point equation

$$\tilde{\delta} = \varphi(\tilde{\delta}). \quad (29)$$

Figure 1 shows the solution to the fixed-point equation (29).

Figure 1. Solution of the fixed-point equation (29)

For a sufficiently low and high exogenous load $\rho < \rho_*$ and $\rho > \rho^*$ equation (29) has unique globally stable equilibrium $\tilde{\delta}_*$ and $\tilde{\delta}^*$ respectively. For an intermediate load $\rho_* < \rho < \rho^*$, equilibria $\tilde{\delta}_*$ and $\tilde{\delta}^*$ coexist as locally stable, and are separated by an unstable equilibrium. Critical loads ρ_* and ρ^* depend on the system parameters. Following accepted practice, e.g., see [7]–[9], we interpret *globally stable* equilibria of (29) as describing *stable* system states, and *locally stable* equilibria of (29) as describing *metastable* system states. Note that the metastability is the result of the positive feedback in the effective load $\tilde{\rho}$ due to resource sharing allowing congestion spillover to the neighboring service groups.

B. Persistent Performance

After solving equation (29) the unconditional persistent loss rate can be determined as follows:

$$\tilde{L} = \begin{cases} (1-\alpha)[f + (1-f)\tilde{\delta}] & \text{if } \tilde{\delta} < 1 \\ 1 - (1-f)\theta/\rho & \text{if } \tilde{\delta} = 1 \end{cases} \quad (30)$$

Figure 2 shows the persistent loss rate (30) as a function of “slowly” changing exogenous load ρ for different values of parameter α representing the level of resource sharing.

Figure 2. Symmetric system: loss vs. exogenous load.

Curve $0E_0$ sketches the loss rate $\tilde{L}$ for sufficiently low level of resource sharing α , when equation (29) has unique solution $\tilde{\delta}$ for all ρ . For sufficiently high resource sharing level α , large service groups $N+B$, and sufficiently inefficient non-native service, i.e., small θ , equation (29) has two stable solutions $\tilde{\delta}_*$ and $\tilde{\delta}^*$ for the 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, i.e., the first order phase transition. Curves $0A_\alpha B_\alpha$ and $C_1 D_1 E_1$ represent loss $\tilde{L}$ in a case of complete resource sharing $\alpha=1$, when $\rho_{**} := \rho_*(1) = \theta$ and $\rho^{**} := \rho^*(1) = 1$. Thus, 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 metastable regime.

Figure 3 demonstrates this dual effect of dynamic resource sharing on the system performance by sketching persistent loss rate (30) as a function of the “slowly” changing level of resource sharing α for sufficiently large service groups $N+B$, and sufficiently inefficient non-native service, i.e., small θ .

Figure 3. Symmetric system: loss vs. resource sharing level.

As resource sharing α “slowly” increases (decreases), persistent loss rate $\tilde{L}$ follows curve $A_0 A^* A^* B^* E$ ($EB^* B^* A^* A_0$). Branches $A_0 A^* A^*$ and $EB^* B^*$ in Figure 3 correspond to branches $OA^* A^*$ and $EB^* B^*$ in Figure 2 respectively. Hysteresis loop $A^* A^* B^* B^* A^*$ in Figure 3 corresponds to hysteresis loop $A^* A^* B^* B^* A^*$ in Figure 2.

C. Fluid Approximation

A system phase diagram separates regions in the space of system parameters with different topological structure of mean-field equation (29). This Subsection analyzes phase diagram of a symmetric system under fluid approximation (22), when $N + B \gg 1$:

$$\tilde{\delta} = \max(0, 1 - 1/\tilde{\rho}), \quad (31)$$

introducing normalized exogenous load

$$\beta = \left(1 + \frac{1}{\theta} \frac{f}{1-f}\right) \rho, \quad (32)$$

which is the sum of the native load and non-native load redirected from non-operational service groups.

Solution to the system (28), (31)-(32) yields the following result. In the normal persistent state, losses occur only for requests, which are native for non-operational groups and not redirected to operational groups:

$$\tilde{L}_* = (1 - \alpha)f. \quad (33)$$

The normal persistent state is stable if operational service groups can handle all native and non-native traffic which is rerouted from non-operational service groups:

$$\beta < \frac{(1-f)\theta + f}{(1-f)\theta + \alpha f}. \quad (34)$$

The congested persistent state with loss

$$\tilde{L}^* = 1 - \frac{(1-f)\theta + f}{\beta} > 0 \quad (35)$$

is stable if the operational service groups cannot handle all their

native and non-native traffic once being fully loaded:

$$\beta > \frac{\theta + (1-\theta)f}{\alpha}. \quad (36)$$

The normal and the congested persistent states coexist if

$$\frac{\theta + (1-\theta)f}{\alpha} < \beta < \frac{\theta + (1-\theta)f}{\theta + (\alpha - \theta)f}. \quad (37)$$

Figure 4 shows section of the system phase diagram by plane (β, α) . In Figure 4, $\beta_* = \theta + (1-\theta)f$,

$$\beta^* = 1 + \frac{(1-\theta)f}{(1-f)\theta + f\theta}, \quad (38)$$

and

$$\beta_*^* = \frac{(1-f)\theta + f\theta}{(1-f)\theta}. \quad (39)$$

Inside regions U_* and U^* the system has unique “normal” and “congested” equilibrium with loss (33) and (35) respectively. Inside region U_*^* these two equilibria coexist as metastable.

 Figure 4. Symmetric system: phase diagram (β, α)

The optimal level of resource sharing $\alpha^{opt} = \alpha^{opt}(\beta)$ follows the curve $ABCDE$. For $1 < \beta < \beta^*$, the optimal resource sharing $\alpha^{opt}(\beta)$ lies on the stability boundary of the normal regime CD in Fig. 4.

Figure 5 shows section of the system phase diagram by plane (β, f) .

 Figure 5. Symmetric system: phase diagram (β, f) .

Inside regions U_* and U^* system has unique “normal” and “congested” equilibrium with loss (33) and (35) respectively. Inside region U_* these two equilibria coexist as metastable.

V. SYSTEMIC RISK

Based on the similarity with discontinuous phase transitions, one may expect that system transition from normal to congested persistent mode occurs through the process of cascading overload. This Section discusses and quantifies the corresponding systemic risk and the risk/performance tradeoff. Subsection A analyses a symmetric system using direct solution of the corresponding mean-field equation. Subsection B suggests that Perron-Frobenius theory provides an adequate analytical framework for quantitative evaluation of the systemic risk and risk/performance tradeoff.

A. Symmetric System

Figure 3 indicates a tradeoff between system performance in the persistent “normal” state and risk of system transition to the persistent overloaded state. Indeed, conventional system design aims at the optimization of the system performance in the persistent “normal” state assuming that this state exists and system remains in this state. In practice, due to inherent uncertainties, the optimal and critical levels of resource sharing α^{opt} and α^* respectively, are subject to variability. This variability may result in parameter α exceeding the threshold α^* , which would cause the abrupt system transition from the desirable persistent state represented by branch $A_0A_*A^*$ to the congested persistent mode represented by branch B_*B^*E .

The stability margin of the desirable persistent state can be increased by lowering level of resource sharing α below the estimated optimal level $\tilde{\alpha}^{opt} = E[\alpha^{opt}]$ at the cost of increasing loss rate (33) above the minimum by

$$Loss(\alpha) = \tilde{L}_*(\alpha) - \tilde{L}_*(\tilde{\alpha}^{opt}), \quad (40)$$

where $L_*(\alpha)$ is the system loss rate in the persistent normal state. While parameter α is controlled, stability threshold α^* is a subject to various inherent uncertainties which. We model these uncertainties by assuming that α^* is a random variable, and measure risk by the expected “catastrophic loss”

$$Risk(\alpha) = [\tilde{L}^*(\alpha) - \tilde{L}_*(\alpha)]P(\alpha^* \leq \alpha), \quad (41)$$

where $L^*(\alpha)$ is the system loss rate in the persistent overloaded state. Note that similar analysis holds for other risk measures.

Figure 6 sketches the risk/loss tradeoff (40)-(41), which is the Pareto optimal frontier separating the feasible (upper-right) and infeasible (lower-left) regions on the (Risk, Loss) plane. Since it is impossible to simultaneously reduce Risk and Loss beyond the Pareto optima, the systemic risk management involves two steps: first, reaching the (Risk, Loss) Pareto optimal frontier, and second, maintaining the optimal operating point on this frontier.

Figure 6. Systemic Risk/Loss tradeoff: $f_1 > f_2$.

In the practically important case of large service groups $N + B \gg 1$, $\tilde{\alpha}^{opt} \approx \tilde{\alpha}^*$, and thus loss (40) can be approximated as follows:

$$Loss(\alpha) \approx (\tilde{\alpha}^* - \alpha)f. \quad (42)$$

In a particular case of normally distributed threshold α^* with small standard deviation σ ,

$P(\alpha^* < \alpha) \sim \exp[-(\tilde{\alpha}^* - \alpha)^2 / (2\sigma^2)]$, and thus

$$Risk \sim \exp\left[-\frac{1}{2}\left(\frac{Loss}{\sigma f}\right)^2\right]. \quad (43)$$

Figure 6 indicates that Risk/Loss tradeoff becomes more essential for less reliable systems, i.e., systems with higher probability of service group failures f .

B. General System

Assuming that each arriving class i request brings revenue w_i , the aggregate *unconditional* revenue loss in the normal persistent state L_* is given by (3) with derived from (5) unconditional loss probabilities:

$$L_* = \sum_{\omega} \left(1 - \alpha_i + \alpha_i \prod_{j \in J_i} \tilde{\delta}_{*j}(\omega_j)\right) \tilde{\delta}_{*i}(\omega_i) \Omega(\omega), \quad (44)$$

where

$$\tilde{\delta}_{*i}(\omega_i) = \omega_i + (1 - \omega_i)\tilde{\delta}_{*i}, \quad (45)$$

and $\tilde{\delta}_* = (\tilde{\delta}_{*i})$ is solution to the corresponding mean-field system (20), describing the normal persistent state with $\tilde{\delta}_{*i} < 1$. A conventional approach to cloud optimization is based on the minimization of the revenue loss (44) over vector $\alpha = (\alpha_i)$ characterizing level of resource sharing and matrix $Q = (q_{ij})_{i,j=1}^I$ characterizing the load balancing strategy. However, as the previous Subsection demonstrated, this performance optimization may result in an unacceptably low stability margin of the desirable equilibrium $\tilde{\delta}_* = (\tilde{\delta}_{*i})$.

To estimate stability margin consider the linearized mean-field system (20) around “normal” equilibrium $\tilde{\delta}_* = (\tilde{\delta}_{*i})$:

$$\pi_j = \sum_{i \neq j} a_{ij}(\tilde{\delta}_*) \delta \pi_i, \quad (46)$$

where $a_{ij}(\tilde{\delta}_*) = [\partial \varphi_i(\delta) / \partial \delta_j]_{\delta=\tilde{\delta}_*}$ and $\pi_j = \tilde{\delta}_j - \tilde{\delta}_{*j}$. The key observation is that due to the resource sharing, matrix $A(\tilde{\delta}_*) = (a_{ij}(\tilde{\delta}_*))_{i,j=1}^J$ is non-negative: $a_{ij}(\tilde{\delta}_*) \geq 0$. Indeed, resource sharing enables congestion in one server group to overflow to the “neighboring” server groups. For simplicity we also assume that different groups of servers are sufficiently interconnected so that matrix $A(\tilde{\delta}_*)$ is irreducible.

According to the Perron-Frobenius theory [10], under these assumptions, matrix $A(\tilde{\delta}_*)$ ’s spectral radius γ is an eigenvalue, which has positive left and right eigenvectors. It is known [10] that the trivial solution to (46) $\pi = 0$ is asymptotically stable if $\gamma < 1$, and unstable if $\gamma > 1$. Thus, mean-field system (20) is locally asymptotically stable if $\gamma < 1$, and unstable if $\gamma > 1$. Since “normal” equilibrium $\tilde{\delta}_*$ loses stability as Perron-Frobenius eigenvalue γ crosses point $\gamma = 1$ from below, it is natural to define the system stability margin as follows:

$$\Delta = 1 - \gamma. \quad (47)$$

The margin (47) can be viewed as a measure of the systemic risk of cascading congestion leading to the persistent congested regime.

Risk-aware performance optimization can be naturally formulated as follows:

$$L_{opt}(\Delta) = \min_{\alpha, Q} L_*(\alpha, Q) \quad (48)$$

subject to the maximum allowable risk

$$\gamma(\alpha, Q) \leq 1 - \Delta. \quad (49)$$

After solving this optimization problem, the Loss/Risk tradeoff is given by

$$Loss = L_{opt}(\Delta) - L_{opt}(0), \quad (50)$$

where systemic risk is identified with Δ .

Solutions to risk-aware performance optimization problem (48)-(49) depends on the system parameters. However, unavoidable variability in the system parameters results in variability of the Perron-Frobenius eigenvalue γ around the expected eigenvalue $\tilde{\gamma}^{opt} = 1 - \Delta$. Viewing actual Perron-Frobenius eigenvalue γ as a random variable, it is natural to estimate the systemic risk as follows:

$$Risk = [L^{opt}(\Delta) - L_{opt}(\Delta)]P(\gamma - \tilde{\gamma}^{opt} > \Delta), \quad (51)$$

where $L^{opt}(\Delta)$ is the aggregate revenue loss in the corresponding persistent congested state.

VI. CONCLUSION AND FUTURE RESEARCH

This paper has reported on work in progress on exposure and quantification of the inherent tradeoff between economic efficiency and systemic risk of cascading overload in the cloud computing model. Due to the intractability of the conventional Markov model, we have suggested employing methodologies of “Complex Systems” for the quantitative assessment of the corresponding tradeoffs. Specifically, we have proposed a mean-field approximation, which greatly reduces the complexity of the system description as compared to the conventional Markov model.

We have followed a conventional interpretation of multiplicity of local equilibria of the corresponding mean-field equations as describing the metastable, i.e., persistent, system states. We have identified the stability margin of a metastable system state with the stability margin of the corresponding equilibrium. In particular, our analysis has indicated that dynamic resource sharing allows for mitigation of the limited system reliability. However, this mitigation occurs at the cost of increased risk of cascading overload.

Numerous questions deserve further investigation, e.g., accuracy of the proposed mean-field approximation and nature of the predicted systemic instability. More broadly, future work should address the practicality of the proposed systemic risk measure at the system design and operational stages. Of particular interest is a potential ability of online measurements of the corresponding Perron-Frobenius eigenvalue [11] to provide “early warning signals” of the system approaching the instability/breaking point [12].

REFERENCES

- [1] The NIST Definition of Cloud Computing NIST Special Publication 800-145, <http://csrc.nist.gov/publications/PubsSPs.html#800-145>.
- [2] D.D. Clark, J. Wroclawski, K.R. Sollins, and R. Braden, “Tussle in cyberspace: defining tomorrow’s Internet,” *IEEE/ACM Trans. on Networking*, Vol. 13, No. 3, June 2005, pp. 462-475.
- [3] J. Tsitsiklis and K. Xu, “On the power of (even a little) centralization in distributed processing,” *ACM SIGMETRICS 2011*, San Jose, CA, 2011.
- [4] M. Bramson, Y. Li, and B. Prabhakar, “Randomized load balancing with service time distribution,” *ACM SIGMETRICS 2010*, 2010.
- [5] J. Hu and B. Hajek, The supermarket game, *IEEE ISIT*, 2012.
- [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] 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.
- [8] R. J. Gibbens, P. J. Hunt, and F. P. Kelly, Bistability in communication networks, *Disorder in physical systems*, Oxford Sci. Publ., Oxford Univ. Press, New York, 1990, pp. 113-127.
- [9] Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi, “Stochastic networks with multiple stable points,” *Annals of Probability*, 36(1):255-278, 2008.
- [10] A. Berman, R.J. Plemmons, Nonnegative Matrices in the Mathematical Sciences, 1994, SIAM.
- [11] S. Hanly, Congestion measure of congestion in CDMA networks, *IEEE Transactions on Communications*, Vol. 47, Issue 3, 1999, pp. 426-437.
- [12] 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**, 53-59 (3 September 2009).