Fragility Risks of Low Latency Dynamic Queuing in Large-Scale Clouds: Complex System Perspective

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

Abstract— **Economic, reliability, and low latency benefits of the cloud computing model are the result of a high level of dynamic resource sharing, made possible by a high degree of system interconnectivity. This paper suggests a more nuanced view of the effect of interconnectivity and resource sharing since interconnectivity may contribute to system fragility, which is associated with a possibility of abrupt/discontinuous system transition to a persistent overloaded regime with unacceptably high delays. This possibility is a result of the same system interconnectivity, which enables not only desirable but also undesirable load transfer throughout the networked system. Our results indicate that benefits of dynamic queuing disciplines, including low latency, can be realized provided that the relevant systemic risks of undesirable contagion are eliminated or mitigated. Due to the intractability of conventional performance models of large-scale interconnected systems, we use the "Complex Systems" methodology, e.g., a mean-field approximate performance model.**

Keywords-cloud computing model, dynamic queuing, low latency, fragility.

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 cloud computing model is just an example of the current trend towards interconnectivity, which is enabled by technological advances and driven by economics [2]. The major economic, reliability, and convenience advantages of interconnected service systems, as compared to the model of dedicated resources, are due to benefits of dynamic resource sharing. These benefits include the ability to accommodate "small" local demand/capacity imbalances by dynamic routing, i.e., by dynamically redirecting load to distant network portions with available resources, as well as elimination of the fixed cost and reduction of the marginal cost for users of cloud computing infrastructure due to the economy of scale [1].

The benefits of dynamic resource sharing are especially significant in the uncertain, volatile, and hostile environment of highly variable demand, hardware/software failures, and possible malicious attempts to disrupt services. These benefits can be naturally quantified by an increase in the system operational region for a given set of operational scenarios. Various mechanisms designed for taking advantage of dynamic resource sharing, including statistical multiplexing, dynamic routing [3], replication [4], etc., have been investigated. Broadly speaking, these mechanisms have been evaluated and compared on their ability to enlarge the system operational region, given the quality of service requirements, e.g., with respect to delay, reliability, etc.

ISBN 978-3-901882-94-4 © 2017 IFIP

This paper suggests that the same economic forces which brought the cloud computing model into life may also be responsible for the inherent "fragility" of this model. Indeed, economic pressures force the system to operate close to the boundary of the operational region where system resources are almost fully utilized. Due to inherent uncertainties, demand variability, and hardware/software failures, the system will occasionally breach the boundary of the stability region. This raises a question about system performance deterioration and system ability to recover from this deterioration.

While the notion of "fragility" has been receiving significant attention lately [5]-[6], this paper takes probably the most basic view of system fragility. According to this view, fragile systems suffer abrupt/discontinuous performance deterioration as the boundary of the operational region is breached. This basic view, which is consistent with intuitive notion of "fragility," has well defined topological foundations in terms of bifurcation theory, or in a potential case, in terms of catastrophe theory. From a practical perspective, abrupt/discontinuous instabilities are more dangerous than gradual for the following three interrelated reasons. First, abrupt/discontinuous instabilities are more likely to result in unacceptably high performance deterioration as the system breaches the boundary of the operational region. Second, these instabilities are often associated with the existence of metastable, i.e., persistent, states, inside the operational region. This creates risk of the system abruptly transitioning to an undesirable metastable state as a result of sufficiently large fluctuation. Finally, the third disadvantage of abrupt/discontinuous instabilities is that their prediction may be more difficult than for gradual transitions [7].

This paper suggests that dynamic resource sharing may create a form of "robust yet fragile" [5] phenomenon since dynamic resource sharing, while making the system robust to "small" demand/capacity imbalances, may make the system "fragile" to sufficiently large demand/capacity imbalances. Indeed, the same dynamic load redistribution, which allows system to accommodate "small" demand/capacity imbalances, also creates a contagion mechanism for systemic congestion. The possibility of undesirable contagion necessitates control and mitigation of the relevant contagion risks at the cost of reduction in the system capacity/operational region, and thus a reduction in economic efficiency. While the tradeoff between economic efficiency and risk of overload is well known for stand-alone resources, we argue that in networked systems this tradeoff may be more acute due to a possibility of discontinuous/abrupt transition to a persistent congestion in a sizable portion of the system.

The paper is organized as follows. Section II describes an operational model of a system of dynamically shared resources. Section III proposes an approximate mean-field performance model for this system. Section IV analyses this performance model in a case of symmetric system. Finally, Section V discusses the potential fragility of system performance and outlines directions of future research.

II. System of Shared Resources

Subsection A describes an operational model of a system of dynamically shared resources, and subsection B introduces an approximate mean-field performance model of this system.

A. General Model

Following [3] consider a system with I classes of jobs (requests) and J service groups, where group $j = 1,..,J$ includes N_j servers and an infinite buffer. 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.

The static routing strategy $S(q_{ij})$, which 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 the 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 of the entire demand:

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

and the system has almost no spare capacity:

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

Although conditions (1)-(2) appear to be restrictive, it can be shown that in market economy they arise naturally as a result of market pressures. Assuming that the demand is elastic and service provider controls demand through service pricing in an attempt to maximize the revenue, conditions (1)-(2) are the result of this revenue maximization, which also determines routing probabilities q_{ij} [8].

In practice, due to variability of the exogenous demand and limited system reliability, the system may not have sufficient resources to accommodate occasional resource demand/supply imbalances, e.g., because delay requirements may limit buffer sizes. These imbalances can be mitigated with dynamic resource sharing made possible by system interconnectivity. In this paper we consider a dynamic routing strategy $D(q_{ij})$, which generalizes strategy [9] by allowing for load balancing.

Strategy $D(q_{ij})$ is determined by probabilities $q_{ij} \ge 0$, where $q_{ij} > 0$ for $j \in J_i$ and $q_{ij} = 0$ for $j \notin J_i$, and $J_i \subseteq \{1,..,J\}$. Introduce vector $\delta = (\delta_j, j = 1,..,J)$, where $\delta_j = 0$ if at least one server from group j is available and $\delta_j = 1$ otherwise. According to strategy $D(q_{ij})$, a request of class i arriving when at least one server in groups $j \in J_i$ is available: $\sum_{j \in J_i}(1-\delta_k)q_{ik} > 0$, immediately occupies an available server $j \in J_i$ with probability

$$\pi_{ij}(\delta) = (1-\delta_j)q_{ij} \Big/ \sum_{k \in J_i}(1-\delta_k)q_{ik} \tag{3}$$

Otherwise, i.e., if group $j \in J_i$ has no available servers: $\sum_{j \in J_i}(1-\delta_k)q_{ik} = 0$, then $\pi_{ij}(\delta) = 0$, and the arriving request with probability q_j joins the queue to server group $j \in J_i$.

B. System with Native Services

In this paper we only consider 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":

$$I = J, \ \mu_{ij} \le \mu_{ii}, \ i, j = 1,..,I, \ i \ne j, \tag{4}$$

It can be shown that under (4) and some natural assumptions on the model parameters, provider revenue maximization yields a routing which mostly allocates requests to native servers:

$$q_{ii} = 1 - \varepsilon, \ \sum_{j \ne i} q_{ij} = \varepsilon, \ \varepsilon \to 0 \tag{5}$$

Strategy $D(q_{ij})$, where probabilities q_{ij} satisfy (5), first attempts to route an arriving request to a "native" server if one is available. Otherwise, the arriving request is routed to an available feasible server $j \in J_i \setminus i$ with probabilities proportional to q_{ij}. If all servers in groups $j \in J_i$ are occupied, the request joins the queue to server group i:

$$\pi_{ii}(\delta)=\begin{cases}1 & if \quad \delta_i=0\wedge\prod_{j\in J_i}\delta_j=1\\ 0 & otherwise\end{cases} \tag{6}$$

and

$$\pi_{ij}(\delta)=\begin{cases}\dfrac{\delta_i(1-\delta_j)q_{ij}}{\sum_{k\in J_i\setminus i}(1-\delta_k)q_{ik}} & if \quad \sum_{k\neq i}(1-\delta_k)q_{ik}=0\\ 0 & otherwise\end{cases} \tag{7}$$

Further in this paper we consider a particular case of symmetric, fully interconnected system with "native services," where $\Lambda_i=\Lambda$, $\mu_{ii}=\mu$, $\mu_{ij}=\mu/(1+\chi)$; $i,j=1,..,I$, $i\neq j$, and parameter $\chi\geq 0$ characterizes inefficiency of a non-native service as compared to the native service. We assume that dynamic resource sharing is characterized by probabilities $q_{ii}=1-\varepsilon, q_{ij}=\varepsilon/K, \varepsilon\to 0$. For $i=1,..,I$, $j\in J_i$, set J_i is comprised of K randomly selected service groups from $\{1,..,I\}\setminus i$. Parameter K characterizes the level of resource sharing: case $K=0$ corresponds to a system without dynamic resource sharing while case $K=I-1$ corresponds to the highest level of resource sharing. Note that despite numerous simplifications, our model is more realistic than model [9] in the following two practically important aspects. First, our model includes overhead χ associated with non-native services, and second, it allows for a non-complete resource sharing for $K<I-1$.

III. Mean-Field Performance Model

In this paper we are interested in performance of large-scale resource sharing systems with a large number of server groups: $I>>1$, and a large degree of resource sharing: $|J_i|>>1$, $i=1,..,I$, where $|J_i|$ is the size of set J_i. Using the methodology of complex systems [10]-[11], we analyze performance of this system under mean-field approximation, which is based on several simplifying assumptions. The basic simplifying assumption neglects correlations between queues at different servers. Specifically, in the steady state, random variables δ_i are approximately jointly statistically independent for $i=1,..,I$: $P(\delta)\approx\widetilde{P}(\delta|\bar{\delta})$, where

$$\widetilde{P}(\delta|\bar{\delta})\approx\prod_i[\bar{\delta}_i^{\delta_i}(1-\delta_i)^{1-\bar{\delta}_i}] \tag{8}$$

and $\bar{\delta}_i:=E[\delta_i]$. Subsection A proposes a mean-field approximation for the case of a single native server for each class of requests. Subsection B extends this approximation to a case of multiple native servers for some classes of requests.

A. Single Native Server for Each Class of Requests

In this subsection we consider a case when each server group $i=1,..,I$ is comprised of a single server: $N_i=1$. In addition to the basic assumption (8), we also assume that the probabilities of server $j=1,..,I$ being available, $p_j^{(0)}$, and the probability of this server being occupied and having $k+1$ queued requests, $p_j^{(k)}$, are as follows: $p_j^{(0)}\approx\widetilde{p}_j^{(0)}$ and $p_j^{(k)}\approx p_j^{(k)}$, where

$$\widetilde{p}_j^{(1)}=\widetilde{p}_j^{(0)}\widetilde{\rho}_{j0},\ \widetilde{p}_j^{(k+1)}=p_j^{(k)}\widetilde{\rho}_{j1},\ k=1,.., \tag{9}$$

and probabilities (9) satisfy normalization conditions

$$\sum_{k\geq 1}\widetilde{p}_j^{(k)}=1 \tag{10}$$

Parameters $\widetilde{\rho}_j$, $\widetilde{\rho}_{j1}$, $j=1,..,J$ in (9) play the roles of "state-dependent effective loads," which account for the additional load due to dynamic resource sharing. Under approximation (8)-(9), self-consistency conditions result in the following expressions for the effective loads:

$$\widetilde{\rho}_{j0}(\bar{\delta})=\sum_i\frac{\Lambda_i\pi_{ij}(\bar{\delta})}{\mu_{ij}}\sum_{\delta_{-j}}\prod_{k\neq j}[\bar{\delta}_k^{\bar{\delta}_k}(1-\delta_k)^{1-\bar{\delta}_k}], \tag{11}$$

where $\delta_{-j}:=(\delta_i, i\neq j)$, and

$$\widetilde{\rho}_{j1}(\bar{\delta})=\frac{\Lambda_j\pi_{jj}(\bar{\delta})}{\mu_{jj}}\sum_{\delta:\delta_j=1}\prod_{k\neq j}[\delta_k^{\bar{\delta}_k}(1-\delta_k)^{1-\bar{\delta}_k}]. \tag{12}$$

Combining (9)-(12) we obtain the following expression for the approximate probability that server j is available:

$$\widetilde{p}_j^{(0)}=\frac{1-\widetilde{\rho}_{j1}(\bar{\delta})}{1-\widetilde{\rho}_{j1}(\bar{\delta})+\widetilde{\rho}_{j0}(\bar{\delta})}. \tag{13}$$

Since $\bar{\delta}_j=1-\widetilde{p}_j^{(0)}$, relations (13) produce the following closed system of fixed-point equations for approximation $\widetilde{\delta}\approx\bar{\delta}$:

$$\widetilde{\delta}_j=\frac{\widetilde{\rho}_{j0}(\widetilde{\delta})}{1-\widetilde{\rho}_{j1}(\widetilde{\delta})+\widetilde{\rho}_{j0}(\widetilde{\delta})}. \tag{14}$$

After solving system (14) one can evaluate the average number of requests queued at server j, $\widetilde{L}_j=\sum_{k\geq 1}k\widetilde{p}_j^{(k)}$:

$$\widetilde{L}_j=\frac{\widetilde{\rho}_{j0}(\widetilde{\delta})}{[1-\widetilde{\rho}_{j1}(\widetilde{\delta})][1-\widetilde{\rho}_{j1}(\widetilde{\delta})+\widetilde{\rho}_{j0}(\widetilde{\delta})]}. \tag{15}$$

B. Multiple Native Servers for Some Classes of Requests

We call a system lightly or heavily loaded if the probability for a class i to be backlogged is small:

$$E\left[\prod_{j\in J_i}\bar{\delta}_j\right]<<1, \tag{16}$$

or, respectively, large:

$$E\left[\prod_{j\in J_i}\bar{\delta}_j\right]\approx 1, \quad (17)$$

$i=1,..,I$. Kolmogorov's zero-one law implies that a large-scale system: $I>>1$ with a large degree of resource sharing: $|J_i|>>1$, $i=1,..,I$, under approximation of independence (8) can be either in a lightly loaded (16) or a heavily loaded (17) regime. Thus, as exogenous loads increase, one may expect abrupt/discontinuous transition from a lightly loaded regime with almost empty queues to a heavily loaded regime with long queues. In the next section we quantify this statement.

In a heavily loaded regime, since servers are almost always occupied, the system operates as a system with a single server for each class of requests, where this single server has the capacity of the entire server group. Thus, in a heavily loaded regime, average queue sizes can be approximated by (15), where vector $\tilde{\delta}=(\tilde{\delta}_j)$ is determined by fixed-point system (14) with effective loads

$$\tilde{\rho}_{j0}(\bar{\delta})=\frac{1}{N_j}\sum_i\frac{\Lambda_i\pi_{ij}(\bar{\delta})}{\mu_{ij}}\sum_{\delta_{-j}}\prod_{k\neq j}[\bar{\delta}_k^{\bar{\delta}_k}(1-\delta_k)^{1-\bar{\delta}_k}], \quad (18)$$

and

$$\tilde{\rho}_{j1}(\bar{\delta})=\frac{\Lambda_j\pi_{jj}(\bar{\delta})}{N_j\mu_{jj}}\sum_{\delta:\delta_j=1}\prod_{k\neq j}[\delta_k^{\bar{\delta}_k}(1-\delta_k)^{1-\bar{\delta}_k}]. \quad (19)$$

In a case of "light load," when requests do not wait for the service, the system evolution is described by a Markov process $X(t)=(x_{ij}(t))_{i,j=1}^{I,J}$, where $x_{ij}(t)$ is the number of class i requests occupying class j servers. Process $X(t)$ has a unique steady-state distribution $P(X)$, which is a solution of the corresponding system of Kolmogorov equations. Since an astronomically high dimension of this Kolmogorov system makes direct solution computationally infeasible, we consider a mean-field approximation (8), which leads to the following expression for the overflow probabilities $\delta_j\approx\tilde{\delta}_j$:

$$\tilde{\delta}_j=\frac{1}{Z_j(\tilde{\delta})}\sum_{n_{1j}+..n_{Ij}=N_j}\prod_{i=1}^{I}\frac{[\tilde{\rho}_{ij}(\tilde{\delta})]^{n_{ij}}}{n_{ij}!}, \quad (20)$$

In (20) the normalization constants are

$$Z_j(\tilde{\delta})=\sum_{n_{1j}+..n_{Ij}\leq N_j}\prod_{i=1}^{I}\frac{[\tilde{\rho}_{ij}(\tilde{\delta})]^{n_{ij}}}{n_{ij}!}, \quad (21)$$

and the "effective" average loads are

$$\tilde{\rho}_{ij}(\tilde{\delta})=\frac{\Lambda_i}{\mu_{ij}}\sum_{\delta}\pi_{ij}(\delta)\prod_{k\neq i,j}[\delta_k^{\tilde{\delta}_k}(1-\delta_k)^{1-\tilde{\delta}_k}]. \quad (22)$$

The closed system of J non-linear mean-field equations (20)-(22) for the vector of overflow probabilities $\tilde{\delta}=(\tilde{\delta}_j)$ has the form of a fixed-point system. In the case of a large number of group j servers $N_j\to\infty$, equation (20) takes the following form:

$$\tilde{\delta}_j=\max\left(0,\ 1-1\Big/\sum_i\tilde{\rho}_{ij}(\tilde{\delta})\right). \quad (23)$$

If $N_j\to\infty$ for $j=1,..,J$, relations (20)-(22) form a closed system of mean-field equations under fluid approximation. These equations also have the form of a fixed-point system.

Under approximation (23), condition (16) for existence of a light load regime is as follows:

$$\prod_{j\in J_i}\tilde{\delta}_j<<1,\ i=1,..,I, \quad (24)$$

where vector $\tilde{\delta}=(\tilde{\delta}_j)$ is determined by fixed-point system (22)-(23). Under fluid approximation (22)-(23), existence of solution $\tilde{\delta}_j=0$ is a sufficient condition for existence of a light load regime. Since condition $\rho_j<0$ guarantees existence of solution $\tilde{\delta}_j=0$, this condition also guarantees existence of a light load regime. Under fluid approximation (22)-(23), a necessary and sufficient condition for existence of a light load regime can be formulated in terms of the Perron-Frobenius eigenvalue of linearized system (22)-(23) [9]-[10].

IV. SYMMETRIC SYSTEM

This section discusses the performance of a symmetric system of dynamically shared resources under the mean-field approximation proposed in the previous section. Subsection A considers system with single native server for each class of requests. Subsection B considers system with multiple native servers for some classes of requests.

A. Single Server of Each Class

For a symmetric system, equations (11)-(12) take, respectively, the following forms:

$$\tilde{\rho}^{(0)}(\tilde{\delta})=\left(1+(1+\chi)\tilde{\delta}\frac{1-\tilde{\delta}^K}{1-\tilde{\delta}}\right)\rho, \quad (25)$$

$$\tilde{\rho}^{(1)}(\tilde{\delta})=\tilde{\delta}^K\rho, \quad (26)$$

where $\rho=\Lambda/\mu$, and thus system (14) simplifies to the following single fixed-point equation:

$$\tilde{\delta}=\frac{1-\tilde{\delta}+(1+\chi)\tilde{\delta}(1-\tilde{\delta}^K)}{1-\tilde{\delta}+(1+\chi\tilde{\delta})(1-\tilde{\delta}^K)\rho}\rho \quad (27)$$

In this paper we consider the particular case of a "large-scale" symmetric system, $I\to\infty$, with high degree of resource sharing $K\to\infty$. In this particular case the fixed-point equation (18) takes the following form:

$$\tilde{\delta} = \frac{1+\chi\tilde{\delta}}{1+\rho+(\rho\chi-1)\tilde{\delta}}\rho, \tag{28}$$

and the expression for the average queue size (15) becomes:

$$\tilde{L} = \begin{cases} \dfrac{1+\chi\tilde{\delta}}{(1+\rho)(1-\tilde{\delta})+(1+\chi)\tilde{\delta}}\rho & if \quad \tilde{\delta}<1 \\ 1/(1-\rho) & if \quad \tilde{\delta}=1 \end{cases}. \tag{29}$$

Figure 1 shows the solution to fixed-point equation (28).

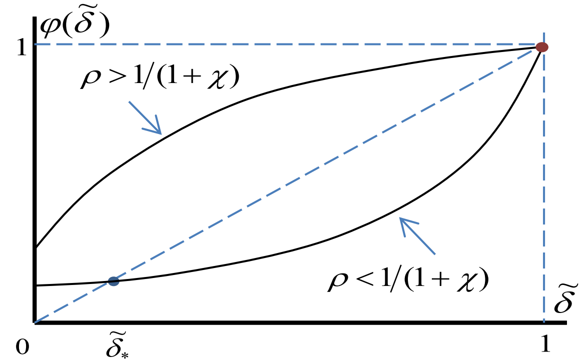

Figure 1. Solution to fixed-point equation (19).

For $\rho < 1/(1+\chi)$ equation (28) has a unique stable solution $\tilde{\delta}_* = \rho/(1-\rho\chi)$ in [0,1). For $\rho > 1/(1+\chi)$ equation (28) has unique stable solution $\tilde{\delta}_* = 1$.

Figure 2 sketches the average queue size in a symmetric system for a case without resource sharing $\tilde{L} = \rho/(1-\rho)$, and for a case of complete resource sharing vs. exogenous load.

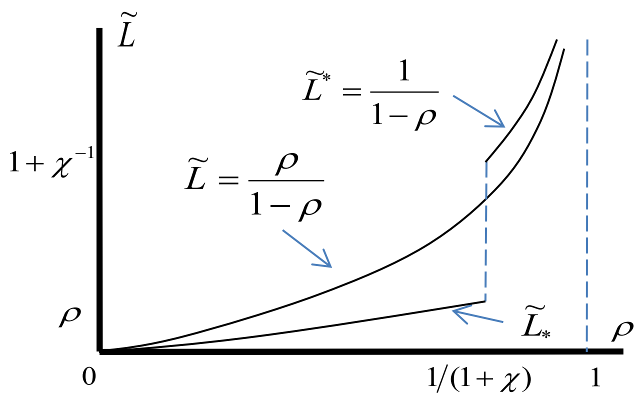

Figure 2. Average queue size: single server of each class.

Expression (29) yields $\tilde{L} = \tilde{L}_*$, where

$$\tilde{L}_* = \frac{\rho}{1+\rho[1-\rho(1+\chi)]} \tag{30}$$

for $\rho < 1/(1+\chi)$, and $\tilde{L} = \tilde{L}^*$, where $\tilde{L}^* = 1/(1-\rho)$, for $\rho > 1/(1+\chi)$. Our analysis indicates that (a) while dynamic resource sharing is beneficial for sufficiently light exogenous load, it may create systemic overload as the exogenous load increases, and (b) the transition from a "normal/operational" to a "abnormal/overloaded" regime may occur abruptly/discontinuously. Although we analyzed a highly idealized system and our analysis was based on a number of approximations, our analysis is consistent with simulation results of various distributed systems with dynamic resource sharing [12] and some analytical results [3].

B. *Multiple Servers of Each Class*

In a symmetric case, the fixed-point system (20)-(22) takes the form of the following single fixed-point equation

$$\tilde{\delta} = \frac{1}{Z(\tilde{\delta})} \sum_{n_1+n_2=N} \frac{\rho^{n_1}\beta^{n_2}(\tilde{\delta})}{n_1!n_2!}, \tag{31}$$

where $\rho = \Lambda/\mu$, the additional load due to dynamic resource sharing is

$$\beta = (1+\chi)\rho\frac{\tilde{\delta}}{1-\tilde{\delta}}, \tag{32}$$

and the normalization constant is

$$Z(\tilde{\delta}) = \sum_{n_1+n_2 \le N} \frac{\rho^{n_1}\beta^{n_2}(\tilde{\delta})}{n_1!n_2!}. \tag{33}$$

Under fluid approximation, equations (31)-(33) take the form of the following equation:

$$\tilde{\delta} = \max\left(0,\ 1-\frac{1-\tilde{\delta}}{(1+\chi\tilde{\delta})\rho}\right) \tag{34}$$

Figure 3 sketches the solution to fixed-point equation (31)-(33) for different exogenous loads ρ.

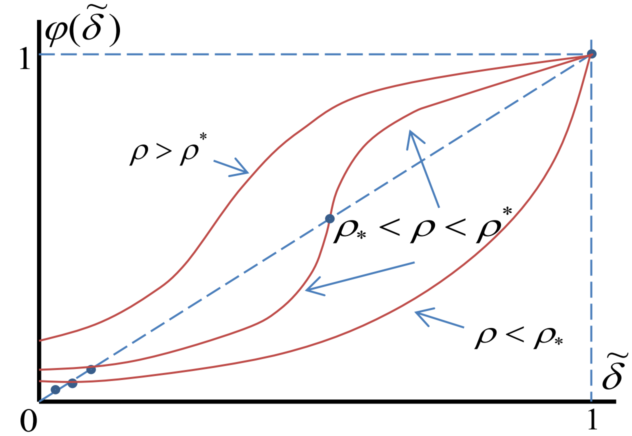

Figure 3. Solution to fixed-point equation (19).

For sufficiently low and high exogenous loads $\rho < \rho_*$ and $\rho > \rho^*$ equation (31)-(33) has a unique globally stable equilibria δ_* and $\delta = 1$ respectively. For intermediate loads $\rho_* < \rho < \rho^*$, sufficiently large number of servers in each class N, and sufficiently high level of resource sharing, equilibria δ_* and $\delta = 1$ coexist as locally stable, and are separated by an unstable equilibrium. Critical loads ρ_* and ρ^* depend on the system parameters N, χ. In particular, $\rho_*|_{N\to\infty} = 1/(1+\chi)$ and $\rho^*|_{N\to\infty} = 1$. Following accepted practice, we interpret globally stable solutions as describing stable system equilibria, and locally stable solutions as describing metastable system equilibria. Note that the metastability is the result of the positive feedback in the effective load $\tilde{\rho}$ due to dynamic resource sharing.

Figure 4 sketches the average number of backlogged requests at a server group as a function of "slowly" changing exogenous load ρ.

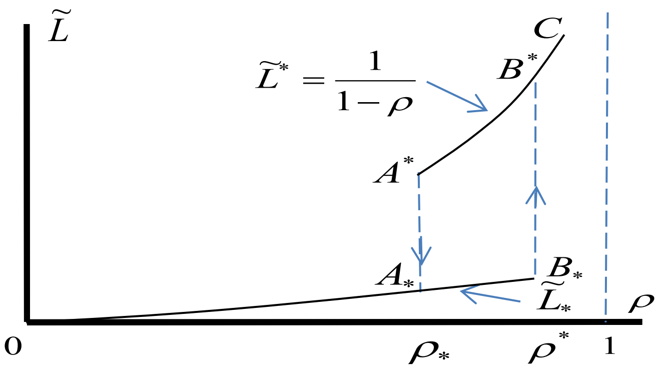

Figure 4. Average queue size: multiple servers of each class.

As load ρ "slowly" increases, the average number of backlogged requests at a server group $\tilde{L}$ follows curve $0A_*B_*B^*C$. As ρ "slowly" decreases, $\tilde{L}$ follows curve $CB^*A^*A_*0$. Curves $0A_*$ and B^*C correspond to the globally stable "normal" and "congested" system equilibria in cases of light: $\rho < \rho_*$, and heavy: $\rho > \rho^*$ loads, respectively. Branches A_*B_* and A^*B^* correspond to the coexisting "normal" and "congested" metastable system equilibria respectively, in a case of intermediate load: $\rho_* < \rho < \rho^*$. Note that discontinuities at the critical loads ρ_* and ρ^* as well as hysteresis loop $A_*B_*B^*A^*A_*$ indicate discontinuous, i.e., the first order phase transition.

V. Discussion & Future Research

A possibility of discontinuous emergence of systemic congestion as a result of dynamic resource sharing suggests a paradigm shift in the design of highly reliable, low delay systems. System designers/operators should combine taking advantage of dynamic resource sharing with controlling risk of abrupt systemic overload/failure due to higher risks associated with abrupt/discontinuous overload/failure than with gradual/continuous ones. The higher risk is not only a result of higher performance deterioration, but also of the unexpected nature of abrupt/discontinuous systemic events. Indeed, gradual/continuous systemic events are typically accompanied by some observable indicators [7], and thus can be predicted and mitigated by initiating the appropriate control actions. This may not be the case for abrupt/discontinuous systemic events since economic pressures drive system designers/operators towards the point of instability, where system resources are fully utilized, increasing the risk of instability.

A major economic driver of the current trend for interconnectivity is the ability of dynamic resource sharing to reduce delays and raise reliability with much higher resource utilization, and thus much higher economic efficiency. This paper suggests that designers/operators of distributed systems, which take advantage of dynamic resource sharing, should mitigate and control the risk of abrupt systemic overload on the boundary of the normal/operational region.

It is straightforward to extend our analysis in order to incorporate limited reliability of system components. Our preliminary results indicate that system component failures make both positive and negative effects of dynamic resource sharing more pronounced. Indeed, while dynamic resource sharing allows system designers/operators to mitigate the negative effect of component failures by reallocating load to the operational components, this reallocation may persist even after the failed components recover. Ultimately, we hope to extend our analysis to more realistic models of distributed interconnected systems for the purpose of quantification and managing of the relevant risk/benefit tradeoffs.

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] 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.

[4] P. Harrison, Z. Qiu, J. Perez, Tackling Latency via Replication in Distributed Systems, ACM/SPEC International Conference on Performance Engineering (ICPE 2016).

[5] J.C. Doyle, D.L. Alderson, L.Li, Steven Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, "The ''robust yet fragile'' nature of the Internet," PNAS, vol. 102, no. 41 October 11, 2005, 14497–14502.

[6] N. Taleb and R. Douady, Mathematical Definition, Mapping, and Detection of (Anti)Fragility, 2013.

[7] 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, 2009, 53-59.

[8] D. Mitra, R. Ramakrishnan, Q. Wang, "Combined economic modeling and traffic engineering," ITC-17, Amsterdam, Nerthelands, 2001.

[9] D. Gamarnik, J.N. Tsitsiklis, and M. Zubeldia, "Delay, Memory, and Messaging Tradeoffs in Distributed Service Systems," Sigmetrics 2016.

[10] V. Marbukh, "Towards unified Perron-Frobenius framework for managing systemic risk in networked systems," European Safety and Reliability Conference (ESREL 2015), Zurich, Switzerland, 2015.

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

[12] P. Echenique, J. Gomez-Gardenes, and Y. Moreno, "Dynamics of jamming transitions in complex networks," 2004, EPL (Europhysics Letters), Volume 71, Number 2.