2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems

Towards Systemic Risk Aware Engineering of Large-Scale Networks: 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 proposes a framework for systemic risk aware engineering of large-scale networks. Economics drives two major evolutionary trends in networked system design/operation: (a) towards the boundary of the system capacity/operational region, where all system resources are fully utilized, and (b) an increase in the system interconnectivity allowing enlargement of this region due to dynamic resource sharing. However, numerous recent systemic failures in various performance oriented networked systems demonstrated systemic risks associated with these trends. Using the "Complex System" perspective, we associate the systemic risks with a possibility of system phase transition to an undesirable persistent state. Specifically, we argue that while growing interconnectivity allows a system to accommodate "small" demand/capacity imbalances, sufficiently large imbalances may result in abrupt/discontinuous instabilities, which is a form of "robust yet fragile" phenomenon. Motivated by higher performance losses and lower predictability of abrupt/discontinuous vs. gradual/continuous systemic overload, the proposed "systemic risk aware performance engineering," intends to enlarge the system capacity/operational region subject to ensuring gradual/continuous system performance deterioration once the boundary of the system operational region is breached.**

Keywords- dynamic resource sharing, performance engineering, systemic risk, complex system.

I. INTRODUCTION

Economics drives two major evolutionary trends in networked system design/operation. The first trend is towards the boundary of the system capacity/operational region, where all system resources are fully utilized, by matching expected demand with available resources through demand pricing and resource provisioning. The second trend is an increase in the system interconnectivity allowing for enlargement of the capacity/operational region through dynamic resource sharing [1]. This trend is driven by the incentive to mitigate both unavoidable variability of the exogenous demand and limited reliability of system components. Indeed, dynamic resource sharing allows a system to accommodate "small" local demand/capacity imbalances, i.e., local congestion, by dynamic routing, i.e., by dynamically redirecting load to a distant network portion with available resources.

However, numerous recent systemic failures in various performance oriented networked systems [1] highlighted systemic risks associated with these trends due to misuse of the allowed flexibility in the dynamic resource sharing [1]. This paper discusses a specific tradeoff between economic efficiency and systemic risk of interconnectivity, which is a form of "robust yet fragile" phenomenon [2]. We argue that dynamic resource sharing, enabled by system interconnectivity, while making system robust to a "small" demand/capacity imbalances, may make the system "fragile" to sufficiently large demand/capacity imbalances. Indeed, the same dynamic load redistribution, which allows for local congestion alleviation, also makes systemic congestion spreading possible. Employing the "Complex Systems" methodology, we associate this "fragility" with a discontinuous/abrupt transition to a persistent congestion in a sizable portion of the system.

Motivated by higher performance losses and lower predictability of abrupt/discontinuous vs. gradual/continuous systemic instabilities [3], we propose "systemic risk aware performance engineering," which intends to enlarge the system capacity/operational region subject to ensuring gradual/continuous system performance deterioration once the boundary of the system capacity/operational region is breached due to exogenous demand variability or limited reliability of the system components [4]-[5]. The paper is organized as follows. Section II describes generic models of dynamic resource sharing in networked systems, and due to intractability of the Markov model of even moderate size resource sharing systems, introduces approximate mean-field performance models. Section III proposes and illustrates a concept of "systemic risk-aware performance engineering."

II. RESOURCE SHARING

Following [6] 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. Different service groups may include geographically distant resources and different types of resources, e.g., memory, CPUs, communication resources, etc. A class i job can be serviced on one of several resource sets. Service times are distributed exponentially with service group specific averages. We assume a service strategy which either

2375-0227 2016
U.S. Government Work Not Protected by U.S. Copyright
DOI 10.1109/MASCOTS.2016.20

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.

An arriving class $i=1,..,I$ job has exponentially distributed service time with average $1/\mu_{ij}$ on a class $j=1,..,I$ server. A static routing strategy is characterized by probabilities q_{ij} that an arriving request of class i is routed to server group j, where rejection probabilities $q_{i0}:=1-\sum_j q_{ij}$ characterize the admission strategy. We assume that on a "slow time scale" the average demand for the resources and supply of these resources are balanced: $\sum_j q_{ij}=1,\ i=1,..,I$, and the system has no spare capacity: $\rho_j:=(1/N_j)\sum_i q_{ij}\Lambda_i/\mu_{ij}\approx 1,\quad j=1,..,J$. It can be shown that in a market economy these two assumptions satisfied as a result of market pressures [7]. Also, in a market economy, the service provider controls routing probabilities q_{ij} in an attempt to maximize revenue brought by serviced demand. We assume one-to-one correspondence between request classes and service groups, and "native" service being at least as efficient as "non-native": $I=J,\ \mu_{ij}\le\mu_{ii},\ i,j=1,..,I$, $i\ne j$. Under these and some other natural assumptions, provider revenue maximization yields routing which allocates requests to native servers: $q_{ij}=1$ if $i=j$, and $q_{ij}=0$ otherwise.

In real systems, due to variability of the exogenous demand and limited system reliability, a system may not have sufficient resources to accommodate the demand bursts, e.g., because delay requirements may limit buffer sizes. Interconnectivity makes dynamic resource sharing possible, and thus allows for mitigation of occasional demand/supply imbalances. To quantify positive and negative effects of dynamic resource sharing, consider a generic model of a "small" demand/supply imbalances on a "fast time scale", when small probability exists that static routing fails due to insufficient resources.

Introduce vector $\delta=(\delta_j)$, where $\delta_j=0$ if server group $j=1,..,I$ has available resources, i.e., a server, or buffering space, or both. Otherwise $\delta_j=1$. Since according to our assumptions $\bar{\delta}_j:=E[\delta_j]<<1$, $j=1,..,I$, 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 $k=1,..,I$ in an unlikely case $\delta_j=1$. The levels of resource sharing are determined by $\alpha_{ij}:=\sum_k q_{ijk}$. The loss probability for an arriving request of class i is

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

where $E[\delta_k|\delta_j=1]$ is the expectation of δ_k conditioned on $\delta_j=1$. Due to intractability of the Markov performance model of a realistic size resource sharing system, our analysis is based on a mean-field approximation [4]-[5], which neglects correlations between blockings in different service groups: $E[\delta_k|\delta_j=1]\approx E[\delta_k]$:

$$L_i\approx\widetilde{L}_i=\sum_j q_{ij}\bar{\delta}_j\sum_{k\ne j} q_{ijk}\bar{\delta}_k\ , \tag{2}$$

where $\bar{\delta}_k:=E[\delta_k]$. Dynamic resource sharing results in additional utilization for server group j equal

$$\beta_j=\frac{1}{N_j}\sum_i\frac{\Lambda_i}{\mu_{ij}}\sum_{k\ne j} q_{ik}q_{ikj}\bar{\delta}_k\ , \tag{3}$$

due to allowing requests a second attempt to obtain service.

Following [4]-[5] we approximate:

$$\bar{\delta}_j\approx Erl(\rho_j+\beta_j,N_j,B_j), \tag{4}$$

where the Erlang distribution

$$Erl(\rho,N,B)=\frac{\dfrac{(N\rho)^{N+B}}{N!N^B}}{\sum_{k=0}^{S}\dfrac{(N\rho)^k}{k!}+\dfrac{(N\rho)^{S+B}}{N!}\dfrac{1-\rho^{B+1}}{1-\rho}} \tag{5}$$

and "effective utilization" $\rho=\rho_j+\beta_j$ accounts for both "native" and "non-native" demand. In the spirit of mean-field approximation, distribution (5) is chosen as being exact in a case when only "native" services are allowed: $q_{ijk}=0, k\ne j$. Substituting (4)-(5) into (3) we obtain a closed system of I non-linear, fixed-point equations for $\widetilde{\delta}_j$, $j=1,..,I$.

III. Risk-Aware Engineering

Figure 1 plots the persistent portion of lost revenue $\widetilde{L}$ vs. the exogenous system load ρ in a symmetric system [4]-[5].

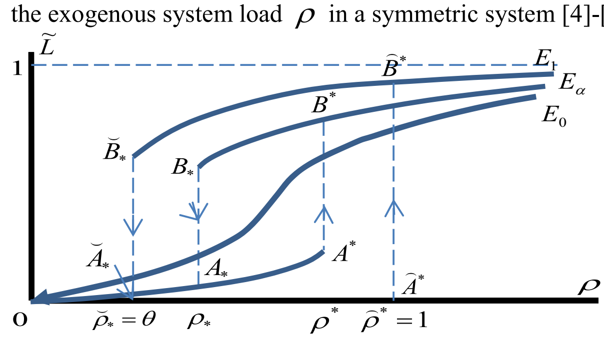

Figure 1. Persistent loss probability vs. exogenous load.

Curve $0E_0$ corresponds to sufficiently low level of resource sharing α, when mean-field equations (3)-(5) have unique solution $\tilde{\delta}$ for all ρ. For sufficiently high resource sharing level α and sufficiently inefficient non-native service, mean-field equations (3)-(5) have two stable solutions $\tilde{\delta}_*$ and $\tilde{\delta}^*$ for intermediate load $\rho_* < \rho < \rho^*$. As load ρ "slowly" increases (decreases), the loss $\tilde{L}$ follows curve $0A_*A^*B^*E$. ($EB^*B_*A_*0$). Curves $0A_*$ and EB^* correspond to the globally stable "normal" and "congested" system equilibria respectively. Branches A_*A^* and B_*B^* correspond to the coexisting "normal" and "congested" metastable system equilibria respectively for intermediate load $\rho_* < \rho < \rho^*$. Discontinuities at the critical loads ρ_* and ρ^* as well as the hysteresis loop $A_*A^*B^*B_*$ indicate discontinuous instability.

Figure 2, which sketches persistent portion of lost revenue rate as a function of "slowly" changing level of resource sharing α, indicates a combination of positive and negative effects of the dynamic resource sharing on the system performance.

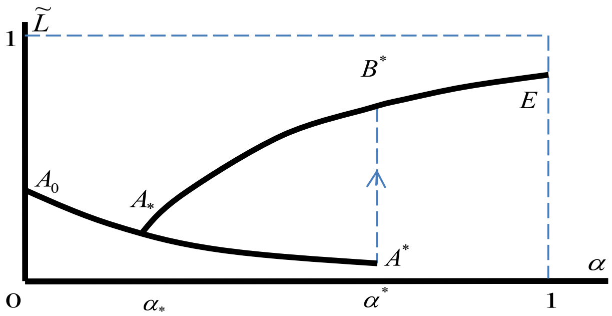

Figure 2. Systemic performance loss vs. resource sharing.

As resource sharing α "slowly" increases (decreases), persistent loss rate $\tilde{L}$ follows curve $A_0A_*A^*B^*E$ $(EB^*A_*A_0)$, and thus $[0,\alpha^*)$ and $[0,\alpha_*)$ represent the "feasible region" and the "safe region" for α. Hysteresis loop $A_*A^*B^*A_*$ in Figure 2 is consistent with the loop in Figure 1. Economic pressures drive network towards resource sharing level $\alpha = \alpha^*$ from below. However, since α^* depends on uncontrollable exogenous parameters, system is likely to transition to undesirable operating point B^*. Lowering level of the resource sharing level to $\alpha = \alpha_*$ eliminates the possibility of abrupt/discontinuous overload at the cost of the corresponding reduction in the expected performance. The expected performance vs. overload risk tradeoff can be managed by controlling resource sharing level $\alpha \in [\alpha_*, \alpha^*]$.

Figure 3 sketches "typical" feasible region $F = OB_iB_jO$ and safe region $F_* = OA_iA_jO$ with respect to engineering parameters (α_i, α_j) in a heterogeneous system.

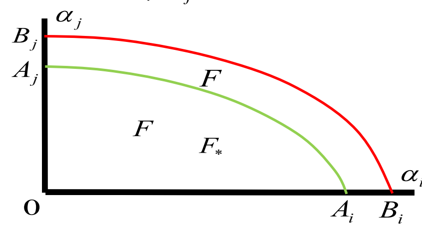

Figure 3. Feasible and safe regions for system parameters.

Since performance oriented systems are driven towards boundary of the operational region, where resources are fully utilized, risk aware system design should mitigate likelihood of abrupt performance deterioration when the boundary is breached due to unavoidable uncertainties.

Our initial results suggest the following "macroscopic fluid description" of networked systems at the onset of systemic instability of the operational equilibrium $L=0$: $L = [\gamma L + bL^2 + cL^3]^+$, where $[x]^+ := \max(0,x)$, γ is the Perron-Frobenius eigenvalue of linearized network evolutionary operator, b and $c<0$ are some constants. According to Perron-Frobenius theory, operational equilibrium is asymptotically stable if $\gamma < 1$, and is unstable if $\gamma > 1$. Thus feasible region in Figure 3 is $F = \{\alpha : \gamma(\alpha) < 1\}$, and safe region is $F_* = \{\alpha : \gamma(\alpha) < 1,\ b(\alpha) < 0\}$. Since "macro-parameters" γ, b, and c are subject to numerous uncertainties, probabilistic characterization of this uncertainty leads to the following quantification of the systemic risk of abrupt instability: $R := 1 - \Pr(\gamma < 1, b < 0)$. An intriguing question to be investigated is feasibility of an "early warning system" of forthcoming systemic instabilities, which relies on online measurements of the system "macro-parameters" for the purpose of initiating appropriate control actions [3].

References

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

[2] 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.

[3] 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.

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

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

[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 and K. Mills, "Demand pricing & resource allocation in market based compute grids: a model and initial results," ICN 2008.