National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce

INFORMATION TECHNOLOGY LABORATORY

SDN: Systemic Risks due to Dynamic Load Balancing

Vladimir Marbukh

IRTF SDN

Abstract

SDN facilitates dynamic load balancing

Systemic benefits of dynamic load balancing:

- economic: higher resource utilization, higher revenue,..
- resilience/robustness to failures, demand variability,..

Systemic risks of dynamic load balancing:

- robust to small yet fragile to large-scale failures/overload
- possibility of abrupt cascading overload
- persistent/metastable systemically congested states

Necessity to manage SDN systemic risk/benefit tradeoff

Congestion-aware Routing in a Delay Network

P. Echenique, J. Gomez-Gardenes, and Y. Moreno, "Dynamics of jamming transitions in complex networks," 2004.

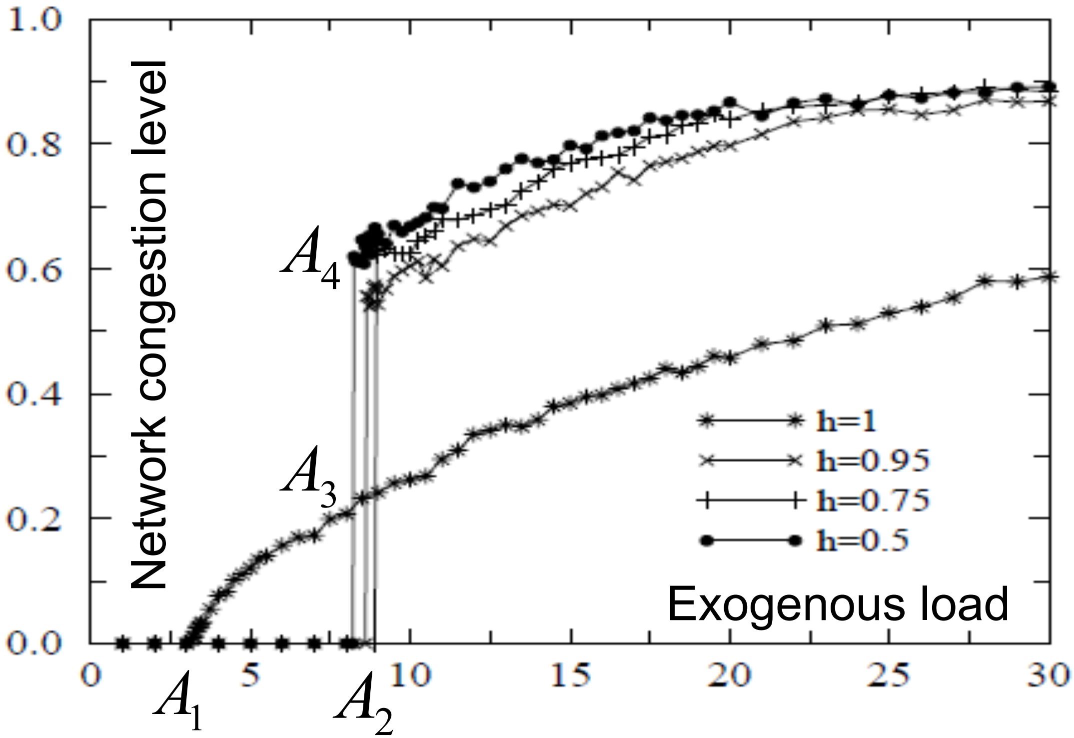

h=1: congestion oblivious (minimum hop count) routing

h=0: congestion aware routing

Minimum-cost routing
Route cost:

$$C_i = hd_i + (1-h)q_i$$

d_i # hops from node i to the destination

q_i queue length at node i

Congestion-aware routing *robust* to small *yet fragile* to large-scale congestion
Benefit: lower network congestion for medium exogenous load from A1 to A2
Risk: hard/severe network overload (discontinuous phase transition) at A2
Economics drives system to the stability boundary A2.

Congestion-aware Routing in Loss Network

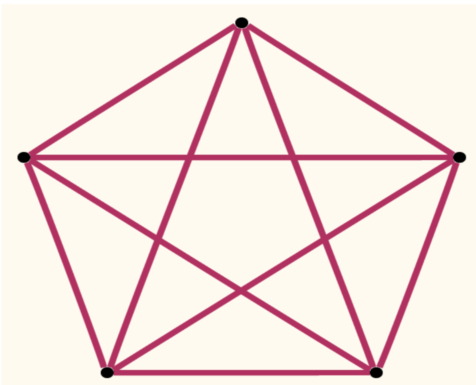

Fully connected network

Arriving request is routed directly if possible, otherwise an available 2-link transit route. Performance: request loss rate *L*.

Positive feedback: load increase → more transit routes → load increase .. = Cascading overload

Combination of selfish requests & variable demand => emergence of congested metastable (persistent) state => robust (to local) yet fragile (to large-scale congestion)

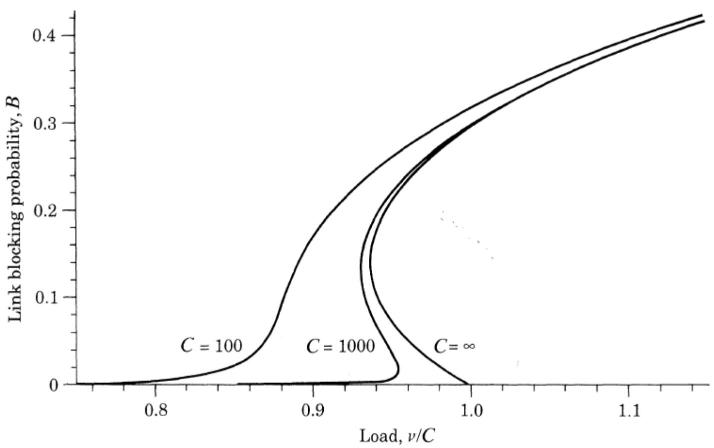

Loss under mean-field approximation [F. Kelly]

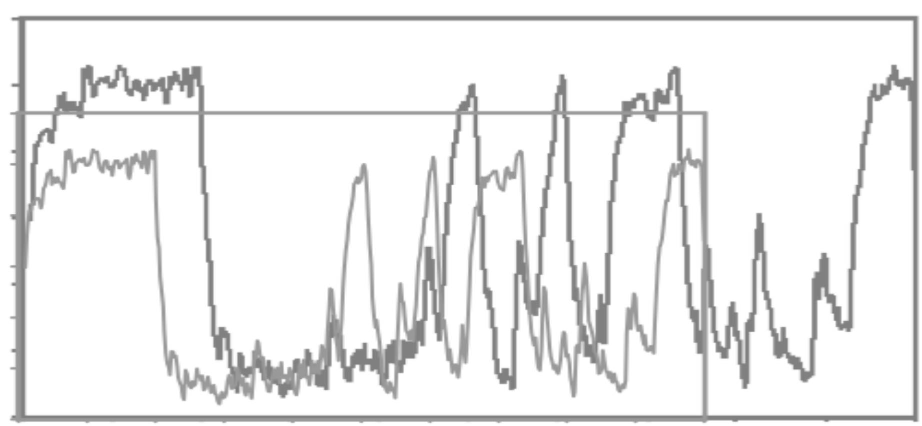

Metastability/Cascading overload [F. Kelly]

Cloud with Dynamic Load Balancing

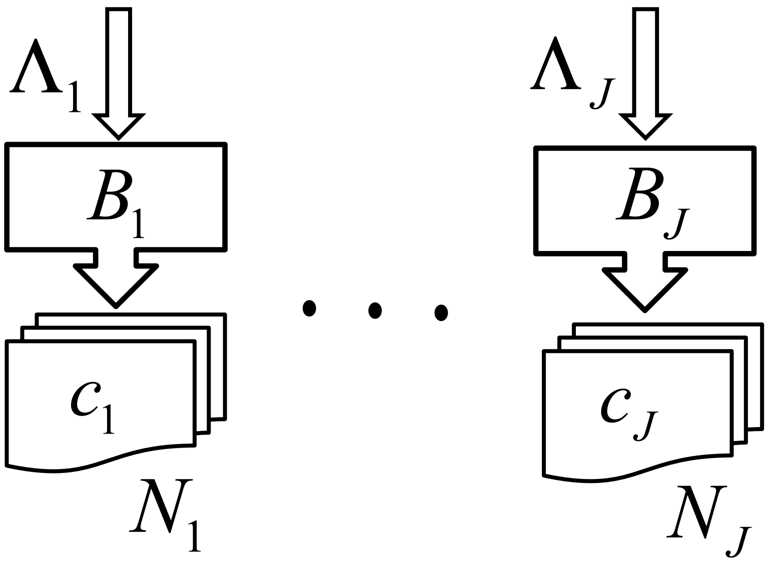

Server group j:
operational with prob. $1 - f_j$
non-operational with prob. f_j

Failures/recoveries on much slower time scale than job arrivals/departures

Static load balancing is possible if:

$$f_j = 0, \quad \rho_j = 1 - O(N_j^{-1/2+\alpha})$$

where utilization is $\rho_j = \Lambda_j / (N_j c_j)$ and $\alpha \geq 0, N_j \to \infty$

Problems: $f_j > 0$, exogenous load uncertain, other uncertainties.
Possible solution: dynamic load balancing based on dynamic utilization, e.g., numbers of occupied servers, queue sizes, etc.

Problem: serving non-native requests is less efficient: $c_{ij} < c_i, \; i \neq j$

and according to A.L. Stolyar and E. Yudovina (2013) this may cause instability of “natural” dynamic load balancing

Dynamic Load Balancing in Cloud [V. Marbukh, 2014]

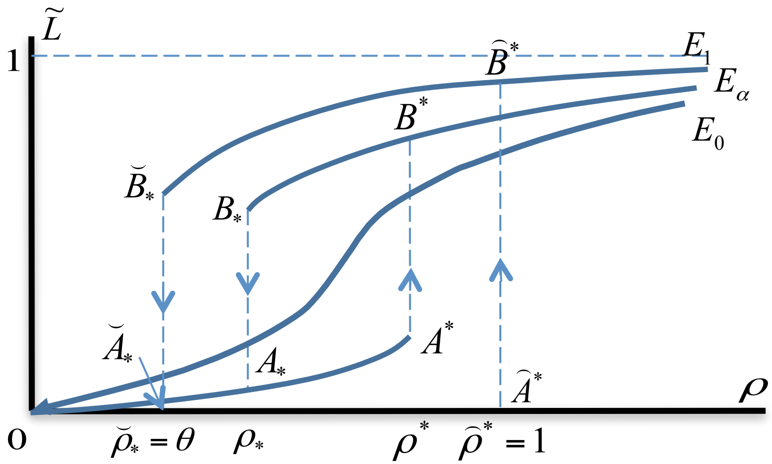

Figure. Lost revenue vs. exogenous load for different levels of resource sharing

(a) As level of resource sharing exceeds certain threshold, metastable/persistent congested equilibrium emerges, making **Cloud *robust*** to local overload ***yet fragile*** to large-scale overload
(b) With further increase in resource sharing, performance of the normal metastable equilibrium improves, while of the congested metastable equilibrium worsens.

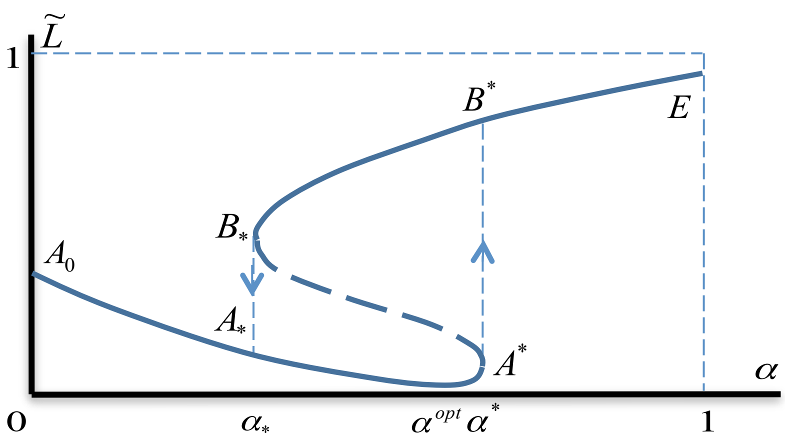

Figure. Provider perspective: lost revenue vs. resource sharing level.

(a) Economics of the “normal” equilibrium drives Cloud from robust to fragile and eventually to stability boundary of the normal equilibrium.
(b) This creates inherent tradeoff between lost revenue:

$$SysLoss(\alpha) = \widetilde{L}_*(\alpha) - \widetilde{L}_*(\alpha^{opt})$$

and systemic risk of large scale overload

$$SysRisk(\alpha) = [\widetilde{L}^*(\alpha) - \widetilde{L}_*(\alpha)]P(\alpha^* \le \alpha)$$

Systemic Performance/Risk Tradeoff in Cloud

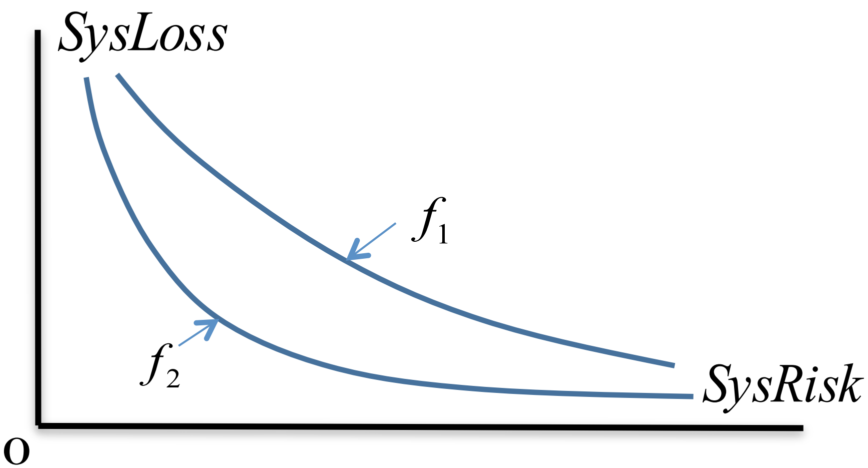

Figure. Risk/Performance tradeoff: $f_1 > f_2$

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

$$SysLoss(\alpha) \approx (\tilde{\alpha}^* - \alpha) f$$

$$SysRisk \sim \exp\left[-\frac{1}{2}\left(\frac{SysLoss}{\sigma f}\right)^2\right]$$

Implication: Uncertainty makes systemic Risk/Performance tradeoff essential

Question: How can one-dimensional analysis describe a heterogeneous Cloud?
Answer: Perron-Frobenius theory due to congestion dynamics being non-negative

Since "normal" equilibrium loses stability as **Perron-Frobenius eigenvalue** of the linearized system crosses point $\gamma = 1$ from below, it is natural to quantify the **system stability margin and risk of cascading overload** by

$$\Delta = 1 - \gamma$$

Word of caution: the above results are obtain under mean-field approximation.

Future Research

- Verification/validation results obtained under mean-field approximation through simulations, measurements on networks and rigorous analysis (doubtful).
- Possibility of online measurement of the Perron-Frobenius eigenvalue for the purpose of using it as a basis for "early warning system."
- Possibility of controlling networks, especially through pricing, based on the Perron-Frobenius eigenvalue.

Thank you!