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

INFORMATION TECHNOLOGY LABORATORY

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

Vladimir Marbukh

FIT 2017

Outline

- Empirical observations & modeling perspectives
- Markov model and approximations of systemic risk
- Cloud models
- Gradual vs. abrupt instabilities
- Implications for Internet transport
- Conclusion, future research

Complex/Networked Systems: Empirical Observations

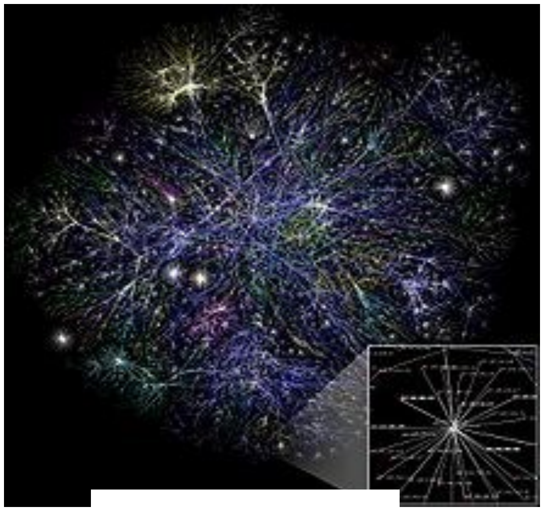

The Internet

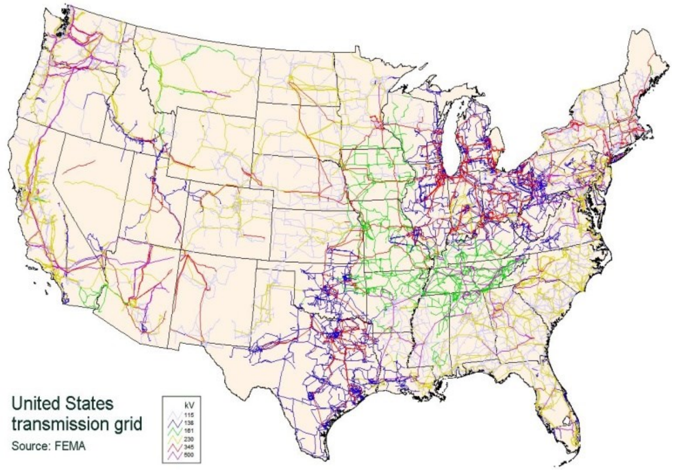

The Power Grid

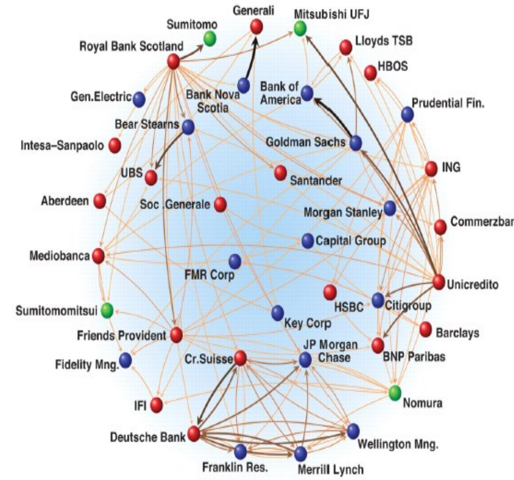

The Financial Network

Inherent connectivity systemic benefit/risk tradeoff
Connectivity is economically driven (rich gets richer, economy of scale, risk sharing, etc.)
Economics fail to address systemic risks of: (cyber)security, cascading failures, etc.

Conventional Risk Management: use historical data to extrapolate, i.e., “fight the last war”.

Challenge: unexpected consequences due to
- *externalities* due to strategic selfish or malicious (cybersecurity, terrorism) components
- non-linear component interactions, randomness, e.g., stochastic resonance

Ultimate Goal: systemic risk/benefit control through combination of regulations/incentives

Markov Micro-description

Markov process with locally interacting components [R. Dobrushin, 1971]

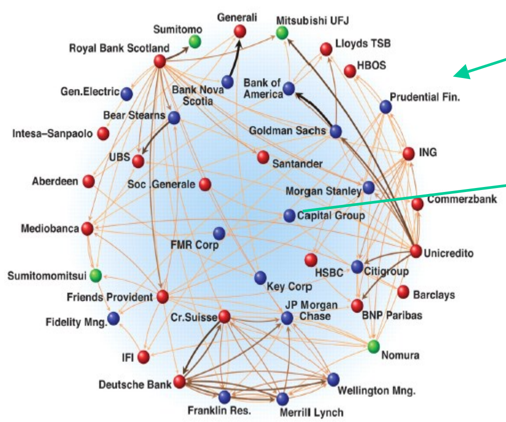

Graph: nodes=components, (directed) links=interactions

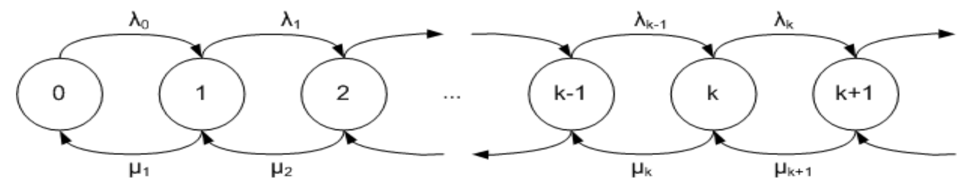

Internal node dynamics $x_n(t)$ Markov process with transition rates dependent on internal states of neighbors

System microstate: $X(t) = (x_1(t),.., x_N(t))$

Non-steady and steady probabilities $P(t, X) = \Pr(X(t) = X), \quad P(X) = \lim_{t\to\infty} P(t, X)$
are solutions to the corresponding Kolmogorov equations.

Kolmogorov system's dimension ~ exp(N) => solution intractable, metastability

In "very particular case" of time reversible Markov process, P(X) ~ exp[U(X)]
Local minima of potential U(X) = metastable states (Landau theory of phase transitions)

In a general case we use mean-field approximation based on "hypothesis of chaos propagation":

$$P(t, X) \approx \prod_{n=1}^{N} P(t, x_n)$$

Individual & Systemic Risks

Desirable states

$$\{x_n\} \setminus \{x_n\}^* : \delta_n = 0$$

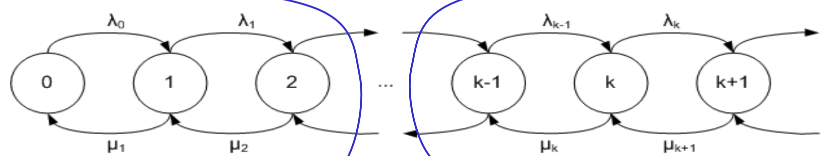

Undesirable states

$$\{x_n\}^* : \delta_n = 1$$

Negative externalities: $$\delta_{-n}^1 \le \delta_{-n}^2 \Rightarrow E[\delta_n | \delta_{-n}^1] \le E[\delta_n | \delta_{-n}^2]$$

where $$\delta_{-n} = (\delta_i, i \ne n)$$

Individual risk: $$s_n = E[\delta_n \chi_n(\delta_{-n})]$$

where $\chi_n(\delta_{-n}) = 0(1)$ depending whether Individual risk can (can't) be transferred to the neighboring components

Example: $$s_n = E\left[\delta_n \prod_{i \in J_n} \delta_j\right]$$ when $$\chi_n(\delta_{-n}) = \prod_{i \in J_n} \delta_i$$

Lorenz, J., Battiston, S., and Schweitzer, F. 2009

Systemic risk: $$S = \left(\sum_n w_n s_n\right) \Big/ \left(\sum_n w_n\right)$$

Cloud: Operational Model

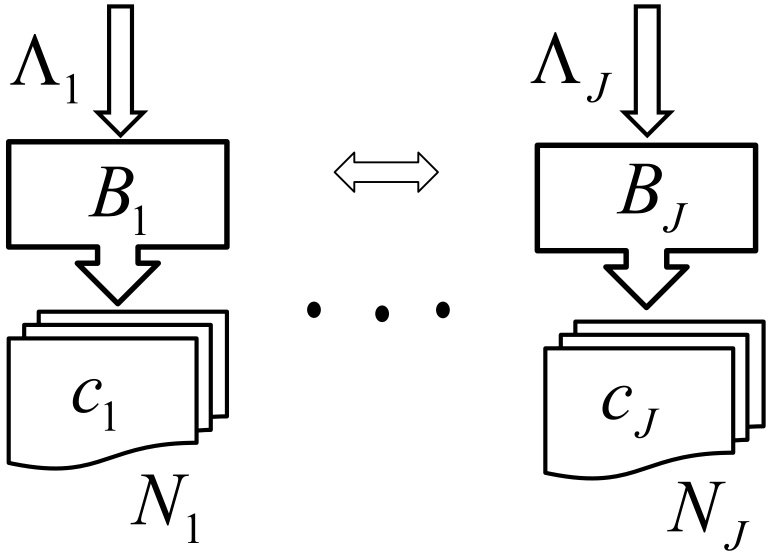

Server group j:
operational with prob. $1-f_j$
non-operational with prob. f_j

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

Cloud: Markov Model

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

$$\Omega(\omega) = \prod_{i=1}^{I} [f_i^{\omega_i} (1 - f_i)^{1-\omega_i}],$$ where

$\omega_i = 0, (1)$ if server group i is operational (non-operational)

Loss probability for class i jobs is:

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

$\delta_i = 0,1$ if server group i is, or respectively, is not available

q_i probability that class i job is admitted to the native server group

α_i probability that class i job attempts for non-native service if $\delta_i = 1$

J_i characterizes system topology

Markov description is intractable even for moderate size systems since it requires solving ~2^{I+N+B} Kolmogorov equations for 2^I vectors (ω_i)

Cloud: Mean-field & Fluid Approximations

$$E\left[\prod_{i\in\{i\}}\delta_i \middle| \omega\right] \approx \prod_{i\in\{i\}}\bar{\delta}_i(\omega_i),$$ where

$\delta_i = 0,(1)$ if server group i has (does not have) available resources

$$\bar{\delta}_i(\omega_i) = \omega_i + (1-\omega_i)\tilde{\delta}_i, \qquad \tilde{\delta}_i \approx \frac{\dfrac{(N_i\tilde{\rho}_i)^{N_i+B_i}}{N_i!N_i^{B_j}}}{\sum_{k=0}^{N_i}\dfrac{(N_i\tilde{\rho}_i)^i}{k!} + \dfrac{(N_i\tilde{\rho}_i)^{N_i}}{N_i!}\dfrac{1-\tilde{\rho}_i^{B_i+1}}{1-\tilde{\rho}_i}}$$

Informally: utilizations of different server group are jointly statistically independent and described by Erlang distribution with loads determined by self-consistency conditions, i.e., mean-field equations:

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

In a case of large server groups: $N_i + B_i \to \infty$, fluctuations are negligible: $\tilde{\delta}_i = \max(0,\ 1-1/\tilde{\rho}_i)$, resulting in fluid approximation.

Symmetric Cloud: Loss Model

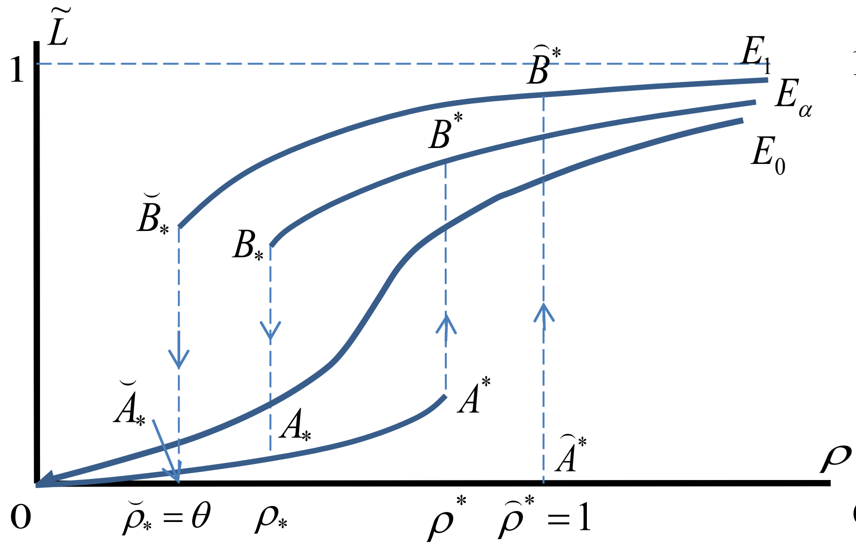

Revenue loss vs. exogenous load for different levels of resource sharing

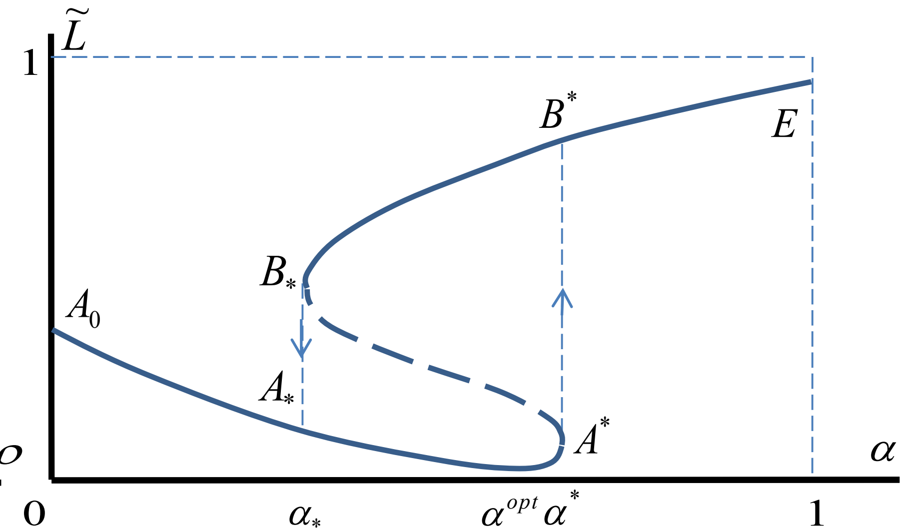

Revenue loss vs. resource sharing level for medium exogenous load

Implications:

- for sufficiently low level of resource sharing – no metastability
- as resource sharing level increases, metastability emerges
- performance in the "normal" ("congested") metastable state gets better (worse)
- ***economics drives system operator towards stability boundary***

Symmetric Cloud: Queuing Model

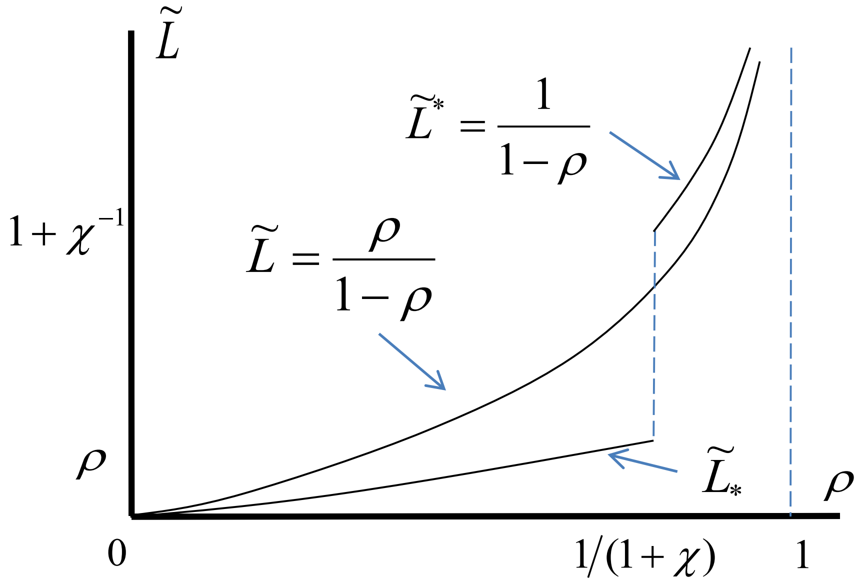

Small service groups: discontinuity in queue size vs. exogenous load for sufficient level of resource sharing

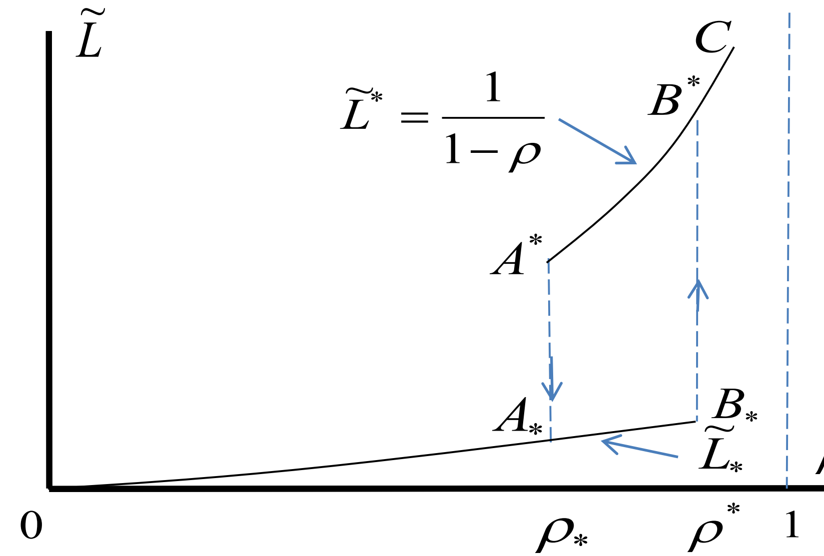

Large service groups: discontinuity & metastability in queue size vs. exogenous load for sufficient resource sharing

Implications:

- for sufficiently low level of resource sharing – no discontinuous instability
- as resource sharing level increases, discontinuous instability emerges
- performance in the “normal” (“congested”) metastable state gets better (worse)
- ***economics drives system operator towards stability boundary***

Resource Sharing Drivers

Generic: economy of scale
Specific: multiplexing gain due to mitigating local imbalances

We propose to quantify benefits of resource sharing by operational region increase

Inefficiency of accommodating component i's individual risk/load by component j

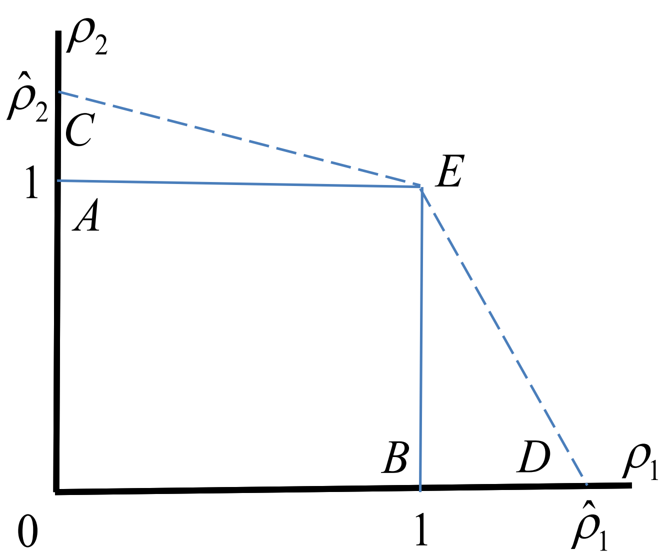

$$\chi_{ij} > \chi_{ii} = 1, \; i \neq j$$

System operational region without: risk sharing OAEBO:

$$\rho_i \leq 1$$

System operational region with complete risk sharing OACEDBO:

$$\rho_i + \chi_{ji}(C_j / C_i)[\rho_j - 1]^+ \leq 1$$

where: $[x]^+ := \max(0, x)$

Operational Region Boundary: Gradual/Abrupt Instability

Loss system under fluid approximation with risk amplification

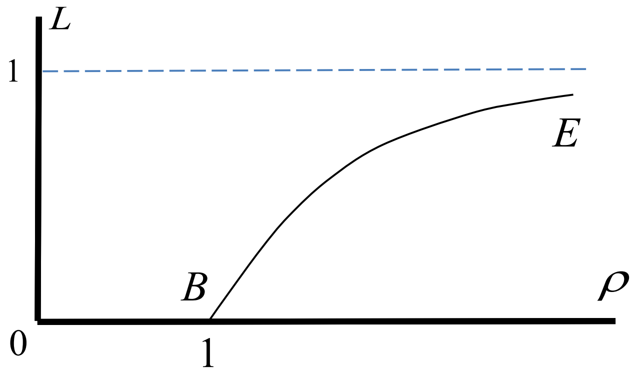

Low level of resource sharing

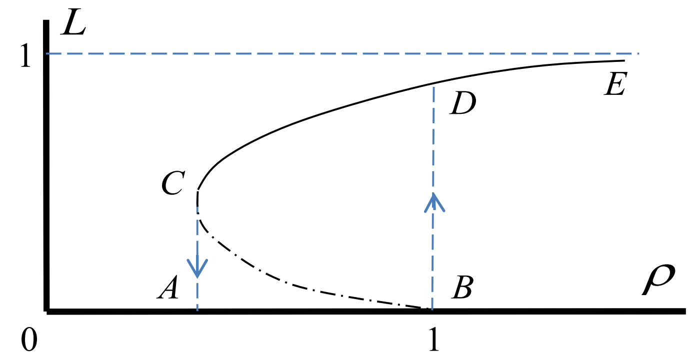

High level of resource sharing

Thesis: since instabilities are unavoidable due to exogenous demand variability, hardware break downs, etc., systemic risk management should favor *gradual* rather than *abrupt* instability on the boundary of the operational region.

Motivation:

- Gradual instabilities may be signaled by critical slowdown, anomalous fluctuations, etc. [M. Scheffer, et al., Early-warning signals for critical transitions, *Nature*, 2009].
- Abrupt/discontinuous instabilities may cause unacceptably high performance deterioration as system gets outside operational region.
- Abrupt/discontinuous instabilities are typically associated with undesirable metastable states inside operational region.

Perron-Frobenius Measure of Systemic Risk

Mean-field equations: $\tilde{\delta}_i = \varphi_i(\tilde{\delta}),\ i = 1,..,N$

Key features of these equations linearized about "normal" equilibrium:
• have a form of fixed-point system
• inside operational region have low systemic risk (normal) solution: $S \approx 0$
• non-negative due to negative externalities: local overload overflows to neighboring components.

Since "normal" equilibrium loses stability as **Perron-Frobenius eigenvalue** of the linearized system $\tilde{\delta} = A\tilde{\delta}$ crosses point $\gamma(A) = 1$ from below, **system stability margin and risk of cascading overload can be quantified** by

$$\Delta(A) = 1 - \gamma(A)$$

In particular, condition of gradual instability on the boundary of operational region in terms of Perron-Frobenius eigenvalue of the linearized mean-field system under fluid approximation just outside operational region:

$$\gamma(A) < 1 \quad \Longleftrightarrow \quad \Delta(A) > 0$$

This in effect condition that the **boundary** of operational region is "**safe**."

Feasible and Safe Parameter Regions

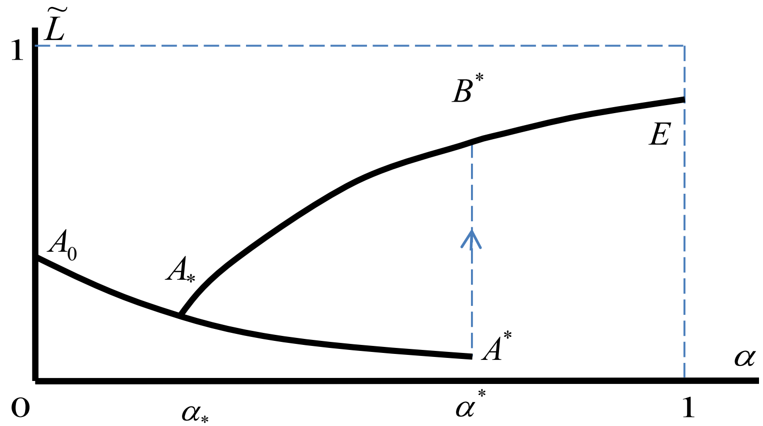

Performance loss vs. resource sharing.

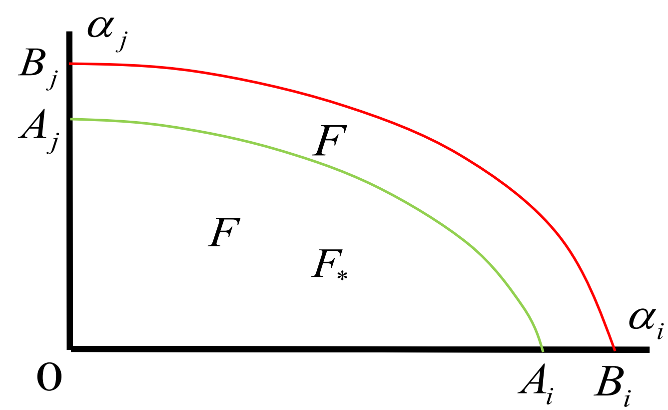

Feasible and safe regions.

Revenue loss at the operational regime boundary: $L = [\gamma L + bL^2 + cL^3]^+$

Feasible parameter region: $F = \{\alpha : \gamma(\alpha) < 1\}$

Safe parameter region: $F_* = \{\alpha : \gamma(\alpha) < 1,\ b(\alpha) < 0\}$

Systemic risk of abrupt/discontinuous instability: $R := 1 - \Pr(\gamma < 1, b < 0)$

Effect of Bounded Rationality

Consider bounded rationality due to uncertain exogenous demand ρ

We assume ρ to be a normal random variable $\rho := N(\tilde{\rho}, \sigma)$

Fixed-point equation: $\tilde{\delta} = E_{\rho}\left\{ \left[1 - \frac{1}{[1+(1+\chi)q\tilde{\delta}]\rho} \right]^{+} \right\}$

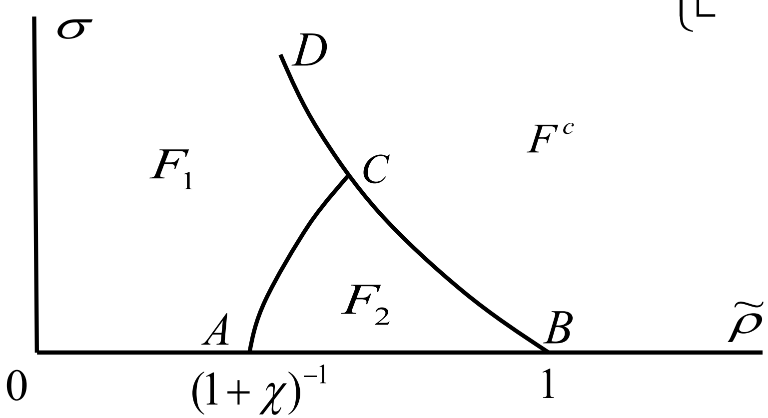

<- Phase diagram

- $F := F_1 \bigcup F_2$: operational equilibrium $\tilde{\delta} = 0$ stable
- F_1 (F_2) : operational equilibrium $\tilde{\delta} = 0$ globally (locally) stable
- F^c : operational equilibrium $\tilde{\delta} = 0$ unstable

Implication: bounded rationality may increase global stability region (C)

Implications for Internet Transport: TCP + Congestion-aware Routing => Instability

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

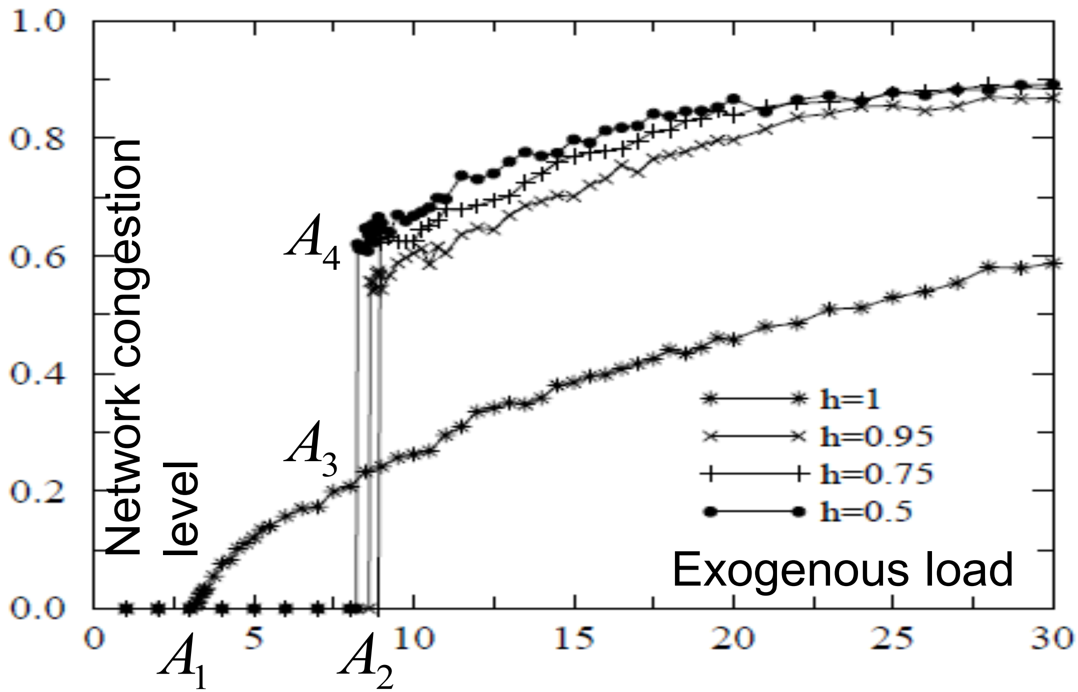

h=1: congestion oblivious (minimum hop count) routing

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.

User Defined Routing: Braess Paradox

Braess paradox, (1969): infrastructure expansion/redundancy may do harm

4000 selfish travelers choose minimum cost/delay route

Without link AB:
Delay=2000/100+45=65

After adding link AB:
Delay=4000/100+4000/100=80

Price of Anarchy (PoA) = 80/65

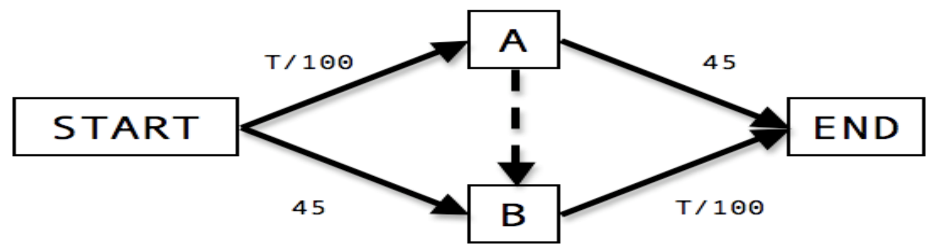

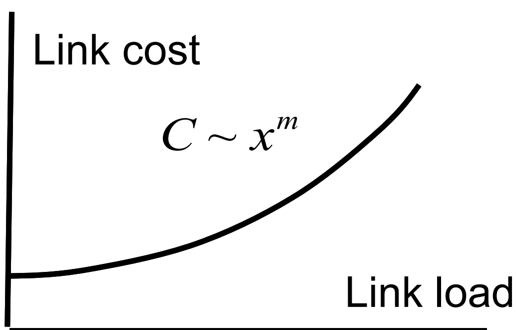

Externalities depend on m:
m=0, no externalities, PoA=1
m>0: negative externalities, PoA>1

Upper bound for PoA independent of network topology (T. Roughgargen, 2002)
m=1: PoA ~ 1.333
m=2: PoA ~ 1.626
m=3: PoA ~ 1.896
m: PoA ~ m/ln(m)

Randomness may cause abrupt deterioration of user defined routing performance due to discontinuous instability (word of caution for SDN)

Congestion-aware Routing: Analytical Modeling

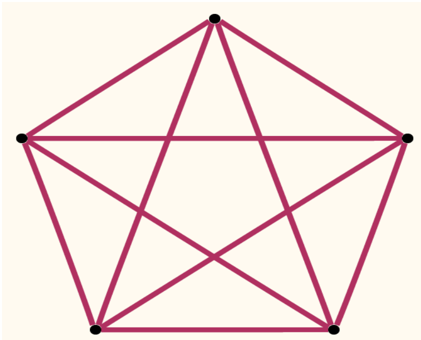

Arriving request is routed directly if possible, otherwise an available 2-link transit route.
Performance: request loss rate *L*.

Risk amplification: load increase → more transit routes → load increase .. Result: cascading overload

Mitigation technique: trunk reservation

Simulation [F. Kelly, 2010]

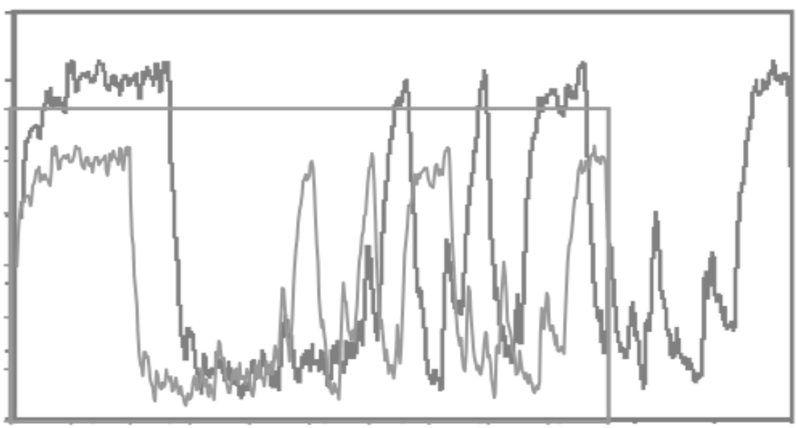

Mean-field approximation: $\tilde{\delta} = \varphi(\tilde{\delta})$

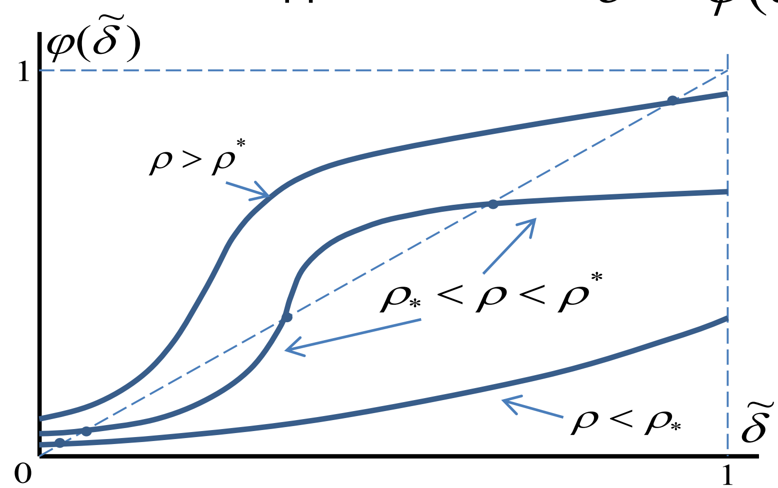

Initial results: randomness may cause abrupt instability for TCP with congestion-aware routing and Multi-Path TCP, fairness mitigates

Conclusions & Future Research

Conclusions:

- Since systemic instabilities are unavoidable, system designers/ operators should avoid abrupt in favor of gradual systemic instabilities
- Existence of inherent tradeoff between economic efficiency under normal conditions and risks of cascading overload/failure resulting in abrupt transition to persistent undesirable state.
- Due to negative externalities, operational equilibrium loses stability in a single dimension determined by the P-F eigenvector, and stability margin is determined by the P-F eigenvalue.

Future research:

- Verification/validation mean-field approximation through simulations, measurements and rigorous analysis (doubtful).
- Possibility of online measurement of the P-F eigenvalue as a basis for “early warning system.”
- Possibility of controlling Networked Systems through a combination of regulations and pricing, based on the P-F eigenvalue.

Thank you!