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

INFORMATION TECHNOLOGY LABORATORY

Considering Emergence in Global Information Systems

Kevin Mills

(includes joint work with Jian Yuan of Tsinghua University)

MCSD Seminar
July 27, 2004

"Natural Laws for Manmade Systems?"

Disclaimer*

If you came to this seminar looking for answers, then you came to the wrong place.

If you came to this seminar looking for questions, then you came to the right place.

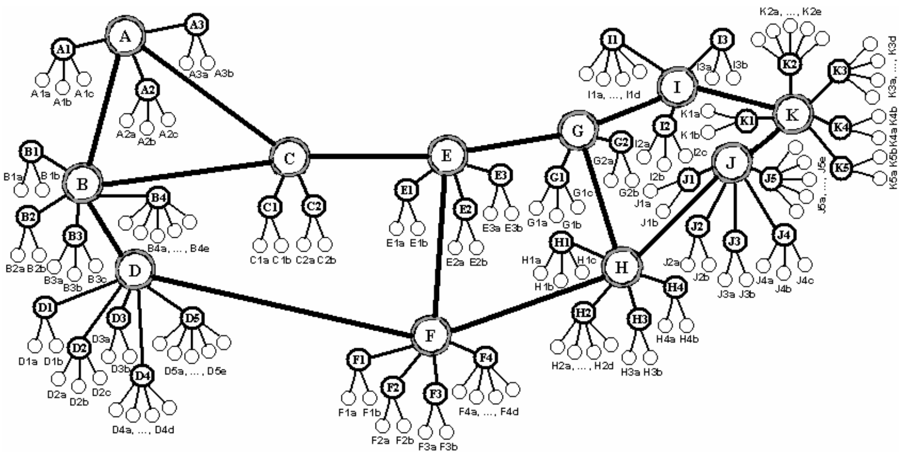

Normalized flow from domain *i* to domain *j*

$$f_{ij} = (x_{ij} - m_{ij}) / \sigma_{ij}$$

Cross-correlation between time-averaged flow vectors

$$C_{(ij)(kl)} = \left\langle f_{ij}(t) f_{kl}(t) \right\rangle$$

Principle Component Analysis of cross-correlation

$$\mathbf{Cw} = \lambda \mathbf{w}$$

Compute the *i*th domain's contribution to *k*th domain

$$\sum_{i,k} (w_{ik}^1)^2 = 1$$

Compute relative strength of flows to *k*th domain

$$S_k = \sum_{i}^{L} (w_{ik}^1)^2$$

*I know you what some math, so I thought I should not make you wait too long (and I decided to omit this).

Seminar Overview

- What is emergence?
- Where (possibly) does emergence arise?
- How (maybe) does emergence arise?
- How (perhaps) can we recognize emergence?
- Emergence by Design vs. Emergence by Nature
- Searching for Emergent Behavior in Large-scale Networks
  - **Flat 2-D homogeneous Cellular Automata (CA) and 1/*f* noise**
  - **Two-tiered homogeneous CA and wavelets**
- Challenges of recognizing emergence in information systems
- Challenges in interpreting, exploiting, eliciting, and controlling emergence

What is emergence?

- **<u>Operational view</u>**: System-wide behavior results – emerges – from interactions among individual elements, rather than from explicit behaviors incorporated into individual elements
  - For example, though each of the 10^{15} cells in a human embryo possess the same DNA, they differentiate (through gene activation and inhibition) into 256 different cell types (e.g., blood, bone, muscle, and neural cells) that organize into the essential systems of the human body
  - The specific role of each cell is not assigned, but rather emerges during embryo development
- **<u>Empirical view</u>**: Systems self-organize into a complex state – poised between predictable cyclic behavior and unpredictable chaos – leading to a statistically predictable distribution of observed changes in system state
  - For example, Earth's tectonic plates exist in a complex state that leads to a distribution of earthquakes with a frequency inversely related to magnitude
  - Such distributions have also been observed in a number of physical and social systems: variations in commodity prices, extinction rates in paleontology, global temperature over time, and frequency of cities by size
  - Measured behaviors lead to a power-law distribution that signifies a system that has self-organized (or emerged) into a complex state

Some traits (possibly) common to emergent systems*

- **Autonomous action** – individual elements act independently without benefit of a master control element
- **Local information** – elements act based on (physically or logically) local information without benefit of a global view
- **Dynamic population** – elements added and deleted naturally without system survival depending on individual elements
- **Collective interaction** – system behavior arises from interactions among many similar independent elements
- **Adaptation** – individual elements can adapt to changing goals, information, or environmental conditions
- **Evolution** – individual elements possess the ability to evolve their behavior over time

*K. N. Lodding, "Hitchhiker's Guide to Biomorphic Software", Queue, June 2004, pp. 66-75.

Some (possible) examples of (operational) emergence

insect colonies

Benard systems

© M.F. Schatz and J.L.Rogers 1998

economies

slime molds

cities

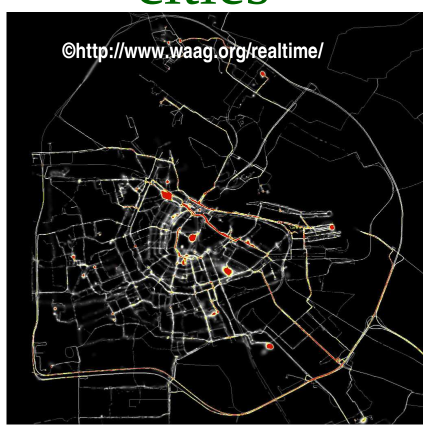

embryo development

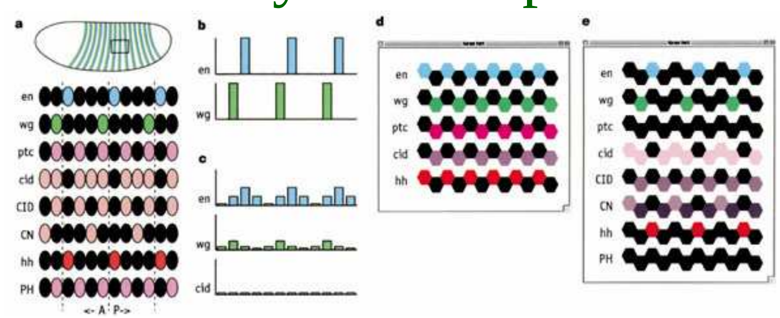

©http://emergent.brynmawr.edu 2003

Some (possible) examples of (empirical) emergence

forest fires

earthquakes

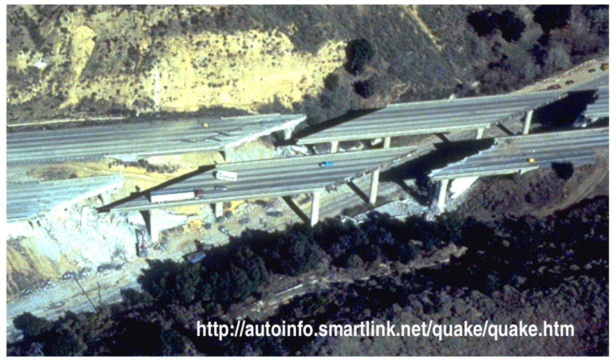

highway traffic flows

species evolution

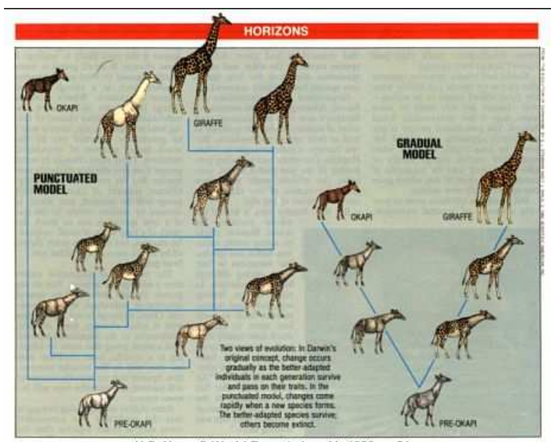

U.S. News & World Report, Jan. 11, 1988, p. 51.

avalanches

information networks

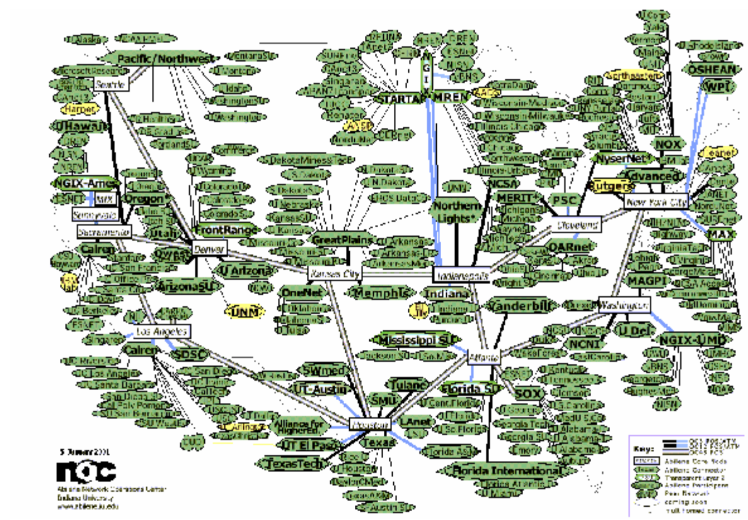

http://www.ics.uci.edu/relations/develop/rs2001/teitelbaum/sld012.htm

How (maybe) does emergence arise?

- **Scale** – requires critical mass in the number of system elements (order emerges from many interactions over space and time)
- **Simplicity** – requires that each element behave rather simply (difficult to construct elements to act on complete information)
- **Locality** – requires interaction among “neighbors” (limits speed of information dissemination)
- **Randomness** – requires chance interactions among elements (increases degree of information dissemination)
- **Feedback –** requires ability to sense environmental conditions (allows some estimation of global state)
- **Adaptation –** requires that each element can vary its behavior (allows system state to change with time)

How (perhaps) can we recognize emergence?

self-similarity

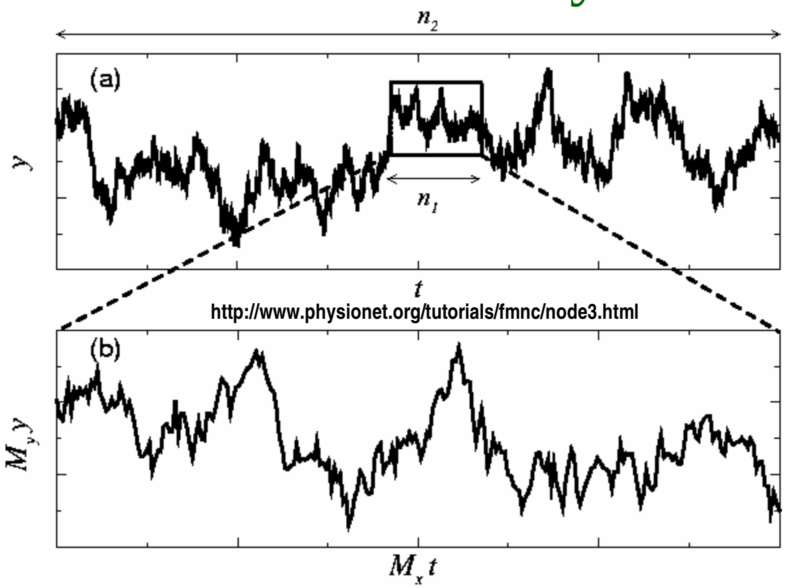

fractal patterns

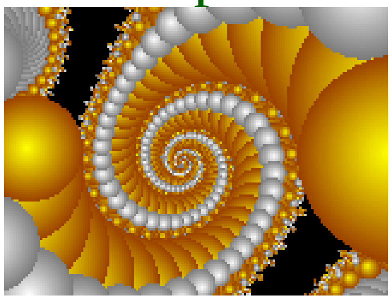

http://www.mbfractals.com/usergal/dougowen.html

linear wavelets

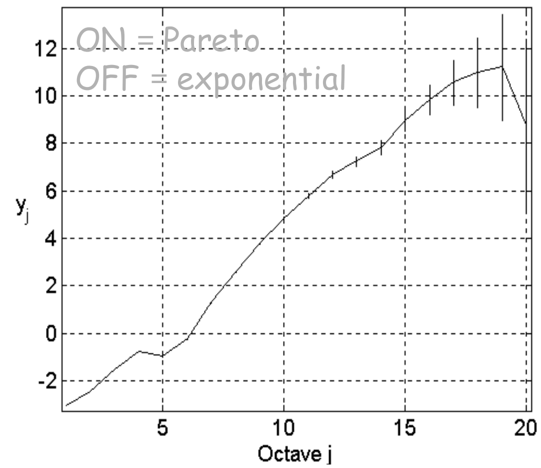

1/*f* noise

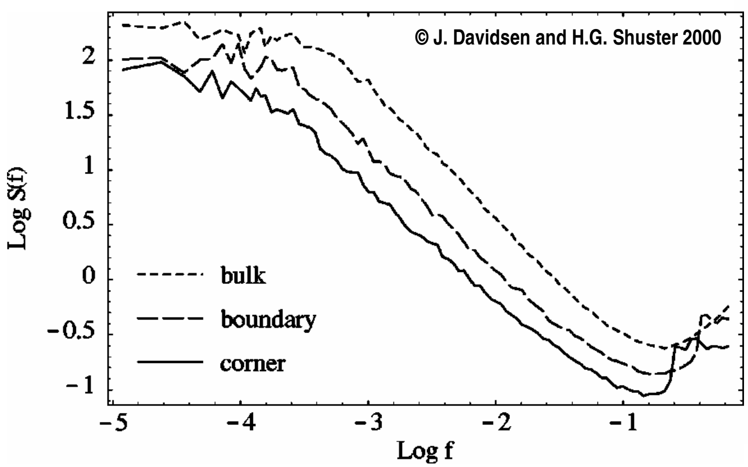

power laws

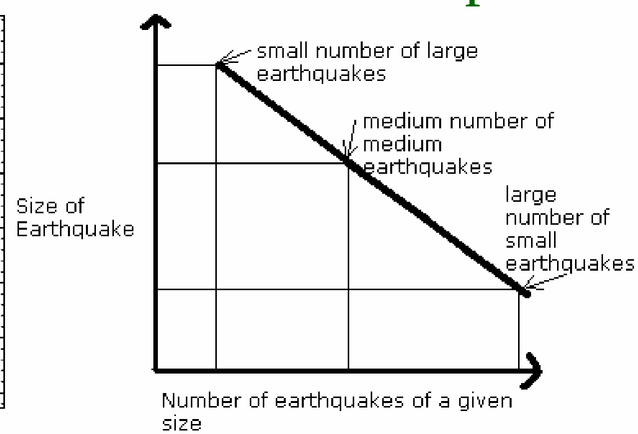

http://complexity.orcon.net.nz/powerlaw.html

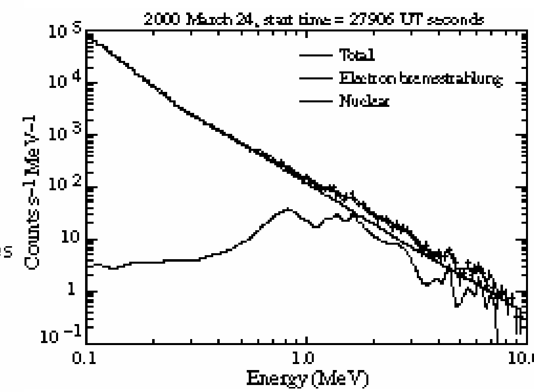

http://heseweb.nrl.navy.mil/gamma/solarflare/24mar00.htm

Emergence by Design vs. Emergence by Nature

- **By Design** – some researchers view emergence as a property that is "designed" into systems
  - Inspires research into techniques to generate desired emergent behaviors
- **By Nature** – some researchers view emergence as an "innate" property of natural systems
  - Inspires research to discover and explain emergent behaviors

- **Possible implications for information systems**
  - Some researchers think we should investigate models (such as artificial life, cellular automata, swarms, biomorphic software, and intelligent agents) to generate emergent behavior in information systems
  - Some researchers suspect that large-scale information systems inherently exhibit emergent properties

Motivation for Our Work

- **Urgency** – growing dependence on large-scale information systems (e.g., Internet, Web, Grid)
- **Suspicion** – that inherently exhibit emergent properties
- **Fear** – that we do not now understand at a macroscopic level
- **Hope** – that we can eventually understand, predict, and control macroscopic behavior in large-scale systems

Our Research Agenda

- Do large-scale information systems (Internet, Web, Grid) inherently exhibit emergent behaviors?
  - If so, are the behaviors desirable, undesirable, or mixed?
  - If so, can we explain, predict, and exploit the behaviors?
- Can we devise effective decentralized mechanisms to elicit desired emergent properties in large-scale information systems?

Examples of Our Research

- Exploring implications of space and time in communications networks
  - Using a flat homogeneous Cellular Automata (CA) and 1/*f* noise
- Investigating current understanding of Internet behavior
  - Using a two-tiered homogeneous CA and wavelets
- Investigating techniques for spatial-temporal traffic analysis in the Internet (omitted from talk – but references provided)
  - Using a three-tiered heterogeneous CA and principle components analysis of cross-correlation matrices

Example #1 – Exploring Implications of Space and Time

- **Goal** – characterize correlation in congestion at different network sizes and time granularities of observation
- **Method** – collect and analyze data from simulation of a homogeneous 2-D CA model that can employ three different means of feedback control
  - (1) open-loop (no control)
  - (2) connection-admission control (CAC) and
  - (3) transmission-control protocol (TCP) flow control
- **Analysis Methods** – log-log plots of power spectral density vs. frequency (i.e., 1/*f* noise) from time-series of
  - Node throughput
  - System congestion state

2-D Homogeneous CA Model of a Network

L **x** ***L*** **grid of interconnected nodes**

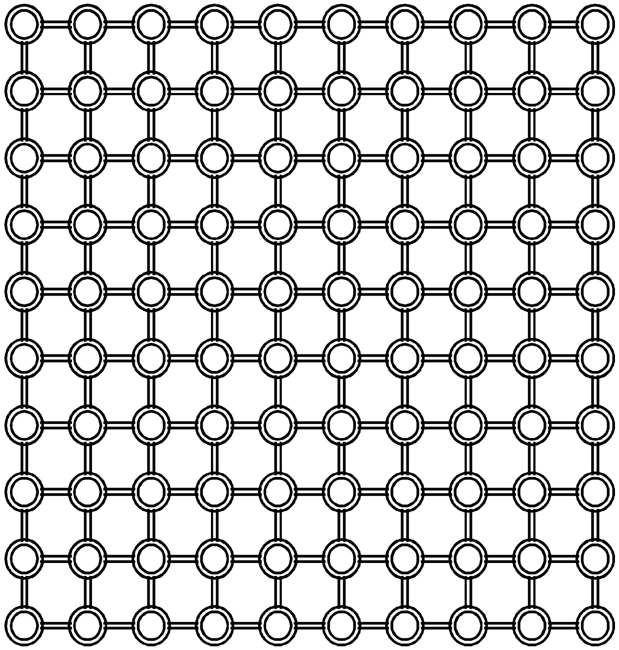

Each node interconnected to four neighbors (boundary nodes interconnected as needed)

- **Nodes** – **generate** source packets (subject to flow restrictions), maintain unlimited length queue, **forward** packets on to neighbors, **consume** packets if node is destination
- **Generation process** – each node has an on-off process:
  - at each time step, generate a packet if on and congestion control permits
  - do not generate packet if *off* or if congestion control forbids
  - duration of on and off periods **exponentially distributed** with means λon and λoff, respectively
- **Congestion control algorithms** (*explained soon*)
  - Open Loop
  - Connection admission control (CAC)
  - TCP Flow Control
- **Routing** – next hop selected **nearest neighbor** (random selection when equidistant)
- **System State**
  - *Xout* is number of packets received by a selected destination node during a time interval *T*
  - *Nr* is the number of packets in router *r*

Behavior of a Node in 2-D CA Model

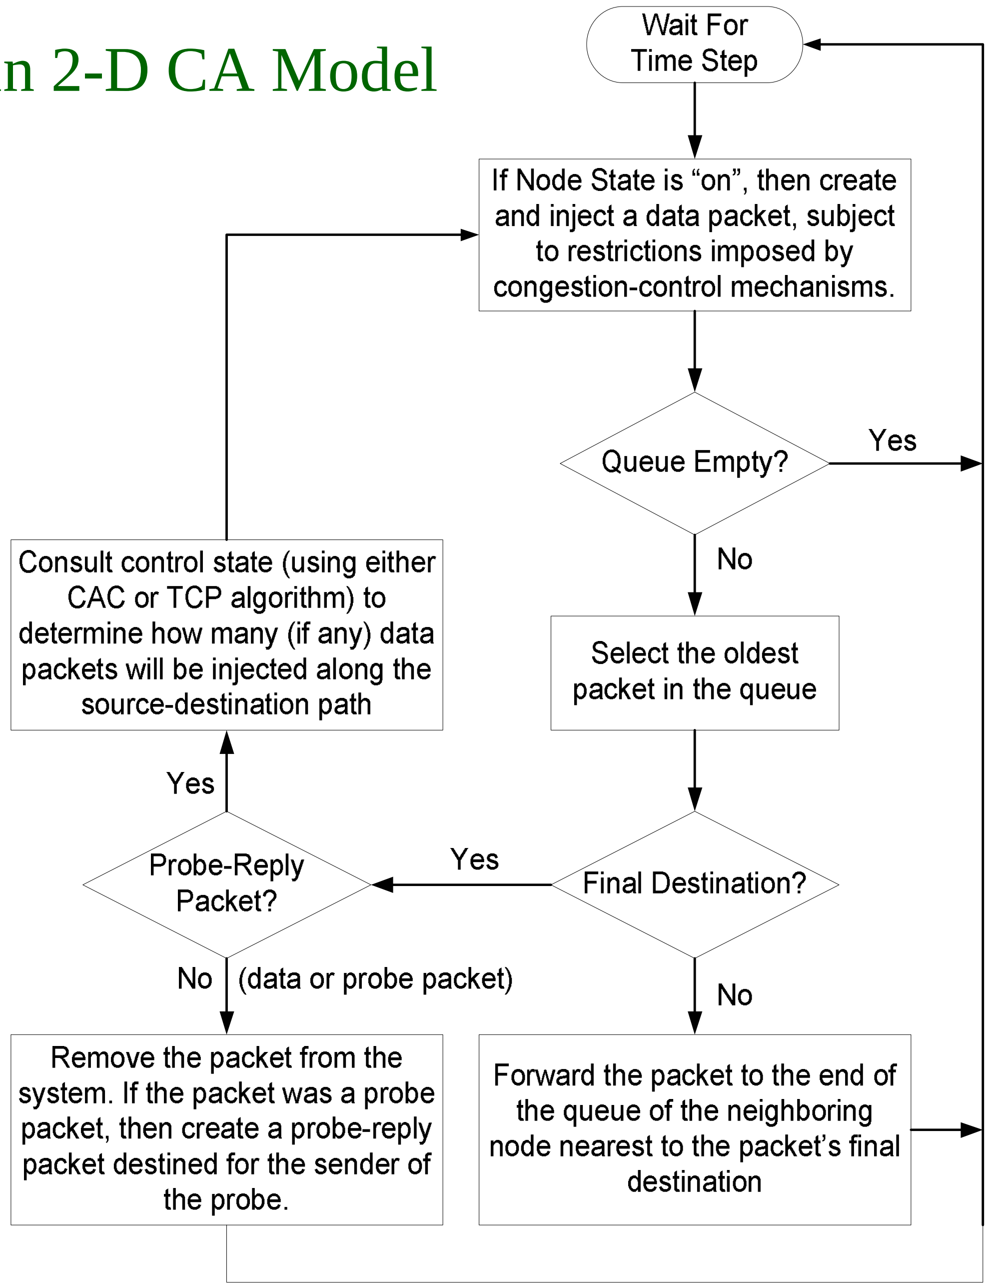

Alternate Congestion-Control Algorithms

- **Open-Loop –** send a packet at each time step when on
- **CAC –** source sends probe packet at beginning of each on period and destination returns probe reply from which source can compute a round-trip time normalized (*Nrtt*) by distance between source and destination
  - If *Nrtt* < some threshold (*Drtt*), then send a data packet at each time step of on period; else send probe packet at next time step
- **TCP –** source sends data packets and destination sends acknowledgment packet for each data packet. Source computes *Nrtt* for each data-ack pair and uses *Nrtt* and *Drtt* in a TCP-like congestion control algorithm. For each ack received the source does:
  - If *Nrtt* > *Drtt*, set slow-start threshold to ½ congestion window; otherwise, if congestion window < slow-start threshold, then congestion window++
  - If congestion window $\geq$ slow-start threshold, then in congestion window = congestion window + 1/congestion window
  - At each time step source generates a packet, but can only have as many packets in transmission as the congestion window allows

Some Time-Series from 2-D CA Model Using TCP Congestion Control

The total time shown on these, and similar graphs, is equal to *T* x *t*, the sample interval size (*T*) multiplied by the number of sample intervals (*t*).

Unless otherwise indicated λ*on* = 100, λ*off* = 500, *Drtt* = 50

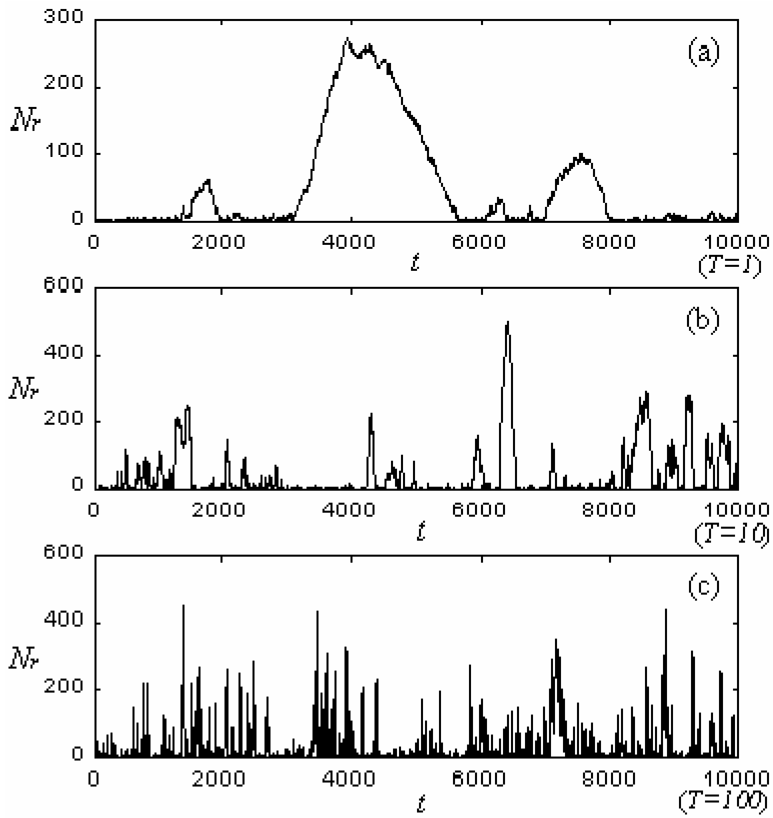

Time series of queue length (*Nr*) at three time granularities *T* = 1, 10, and 100.

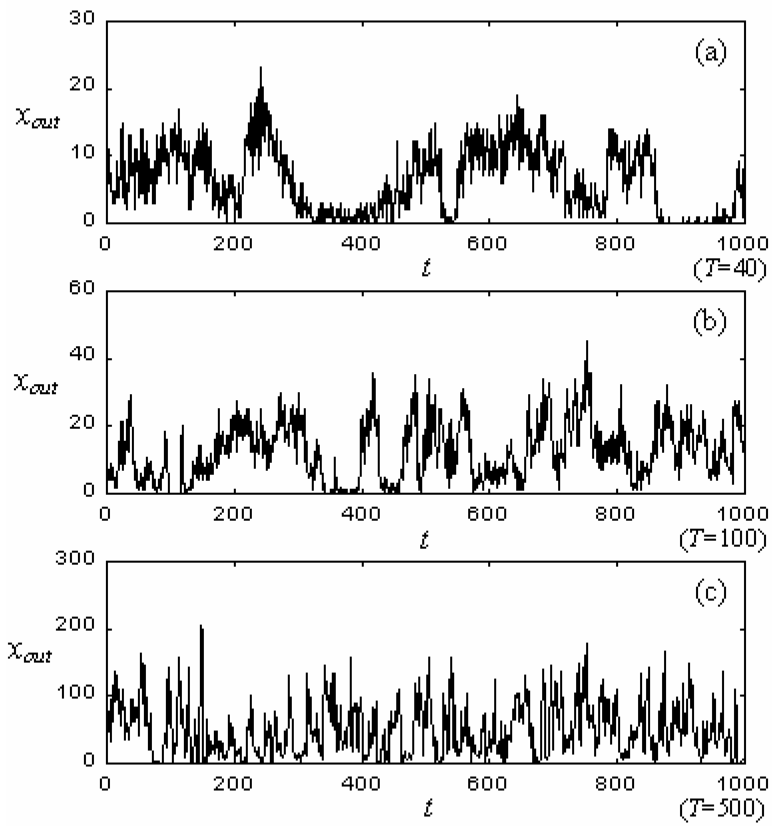

Time series of *Xout* at three time granularities *T* = 40, 100, and 500.

Log-Log Plots of Power Spectra vs. Frequency for *Xout*

***L* = 8 for *T* = 80 and 400**

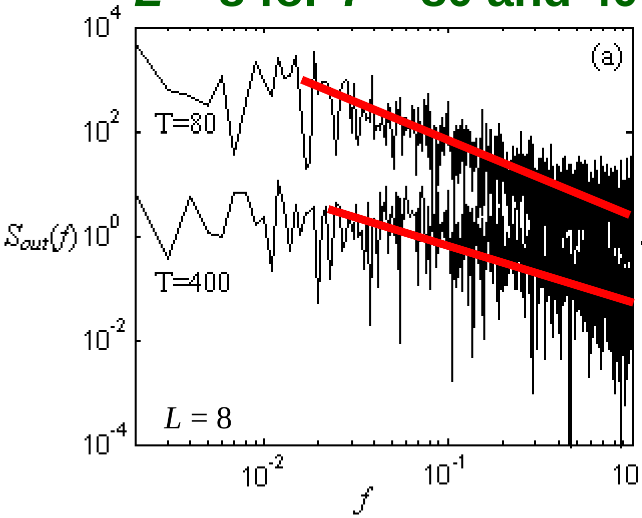

***L* = 32 for *T* = 400 and 1000**

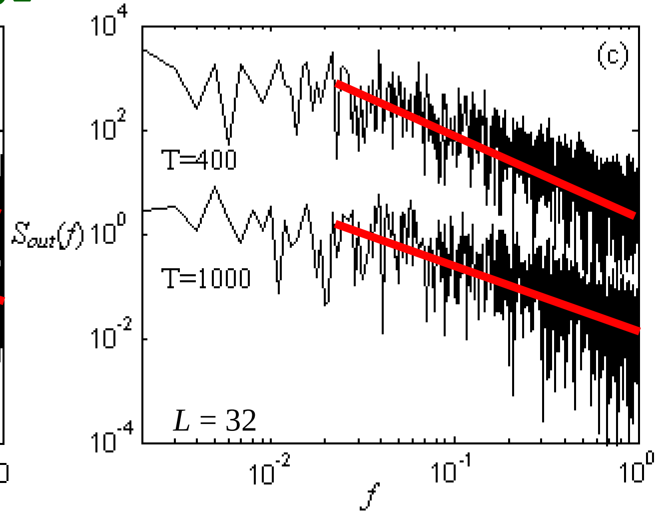

Presence of 1/*f* noise suggests evidence of collective effect

***L* = 16 for *T* = 200 and 600**

HOMOGENEOUS MODEL EXHIBITS SIMILAR BEHAVIOR EVERYWHERE, SO WE CAN SAMPLE ANY NODE AS REPRESENTATIVE OF ALL NODES

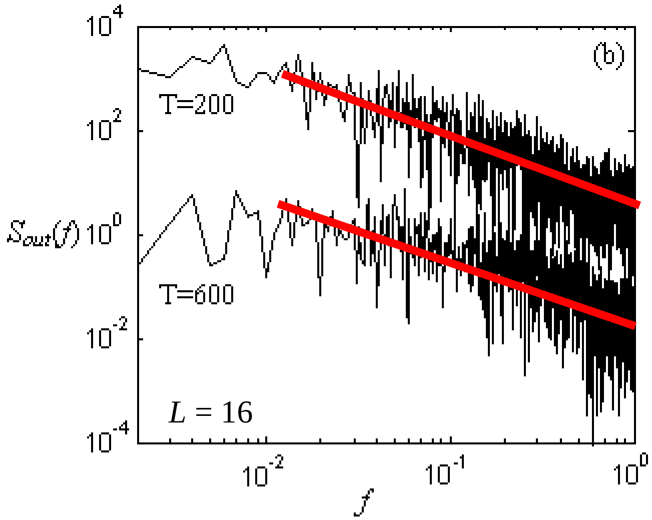

Self-similarity decays for same system size as *T* increases

Self-similarity holds for same *T* as system size increases

Suggests that correlation in congestion increases with system size

NEXT WE MONITOR NETWORK-WIDE CONGESTION

Technique to Monitor Network-Wide Congestion

Red nodes are congested

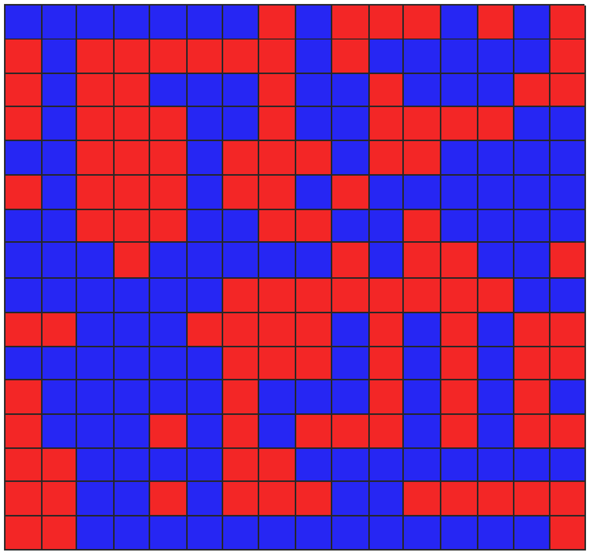

Blue nodes are not congested

- Define threshold Y such that if $Nr > Y$, node r is congested ($Y = 5$ here)
- At any given time granularity T, count the number y of congested nodes

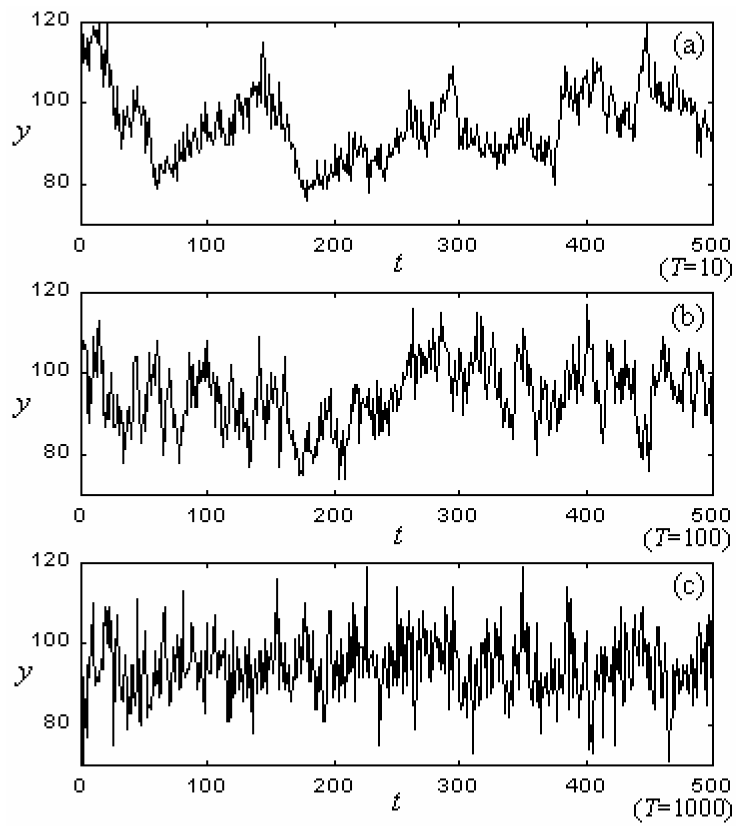

Time-series of number of congested nodes for various time granularities

Log-Log Plots of Power Spectra vs. Frequency for *y*

L = 8 for T = 50 and 400

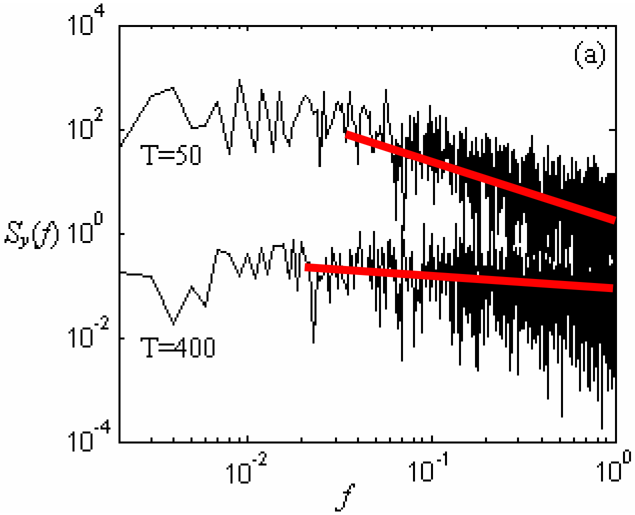

L = 32 for T = 400 and 1000

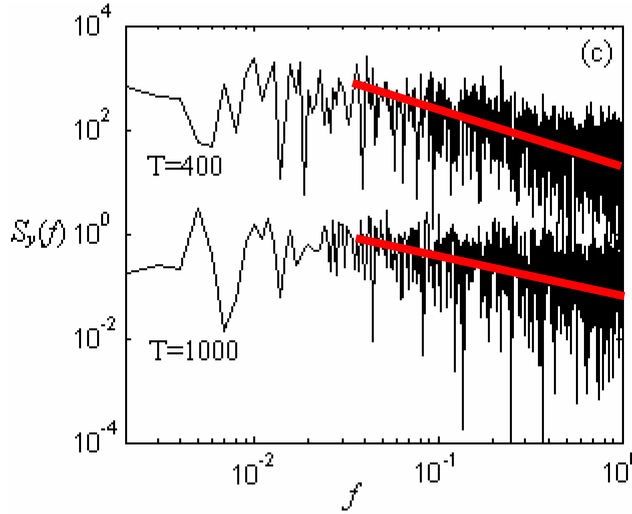

Suggests that some time scale exists for a given network size where the most evident 1/*f* noise (and collective effect) exists and where the network behavior will be most coherent

L = 16 for T = 200 and

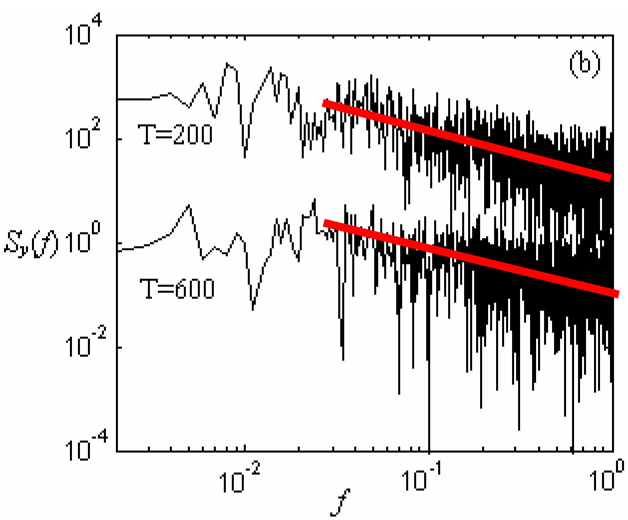

Self-similarity decays for same system size as *T* increases

Self-similarity holds for same *T* as system size increases

Suggests that collective behavior in a large network holds more profound influence on congestion and predictability

NEXT WE CONSIDER INFLUENCE OF NETWORK SIZE

Embedded Subsets of a Network vs. Network of Same Size

Red nodes are

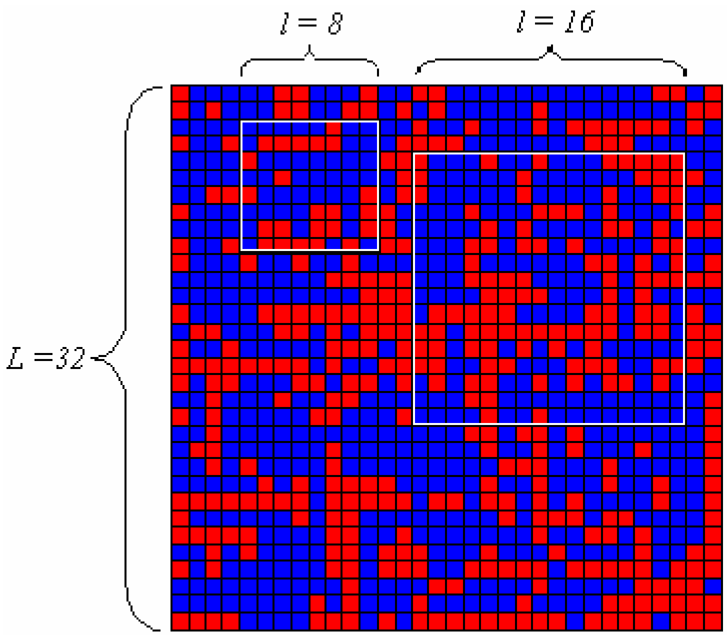

- Response in a network sub-area might have different characteristics than response in a network of the same size as the sub-area

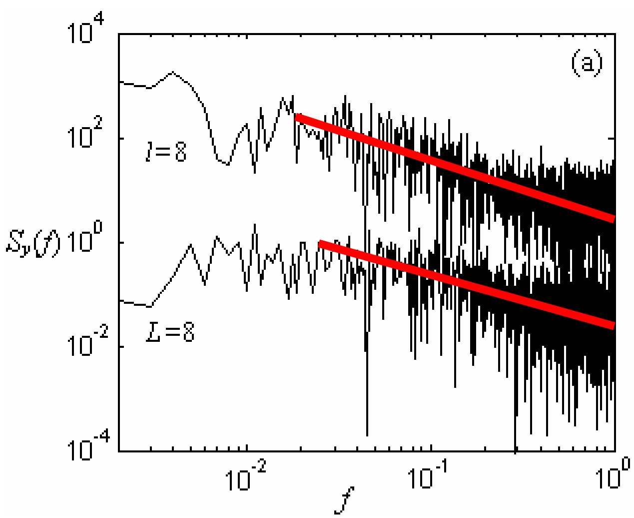

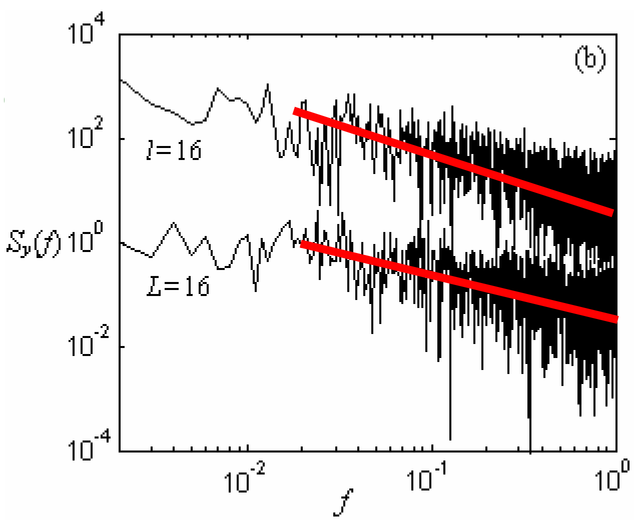

Suggest that network sub-areas exhibit stronger correlation in congestion when compared at the same time scale with a network of the same size as the sub-area

Kevin Mills

Example #2 –Investigating current understanding of Internet behavior

- **Goal** – improve current understanding of correlation structure of network traffic by identifying and studying fundamental causalities arising from multiple protocol layers operating in a sufficiently large network
  - What is the role of user behavior?
  - What is the role of transmission dynamics?
  - What is the role of network structure?
- **Method** – collect and analyze data from simulation of a homogeneous two-tiered (router tier and host tier) CA that represents different protocol layers
  - (1) Application layer – on-off periods (exponential and heavy-tailed distributions)
  - (2) Transport layer – TCP flow control and TCP Friendly Rate Control
  - (3) Network layer –number of hosts per router and capacity of router links
- **Analysis Methods** – wavelet analysis of router throughput over ranges of timescales

Analysis approach based on wavelets

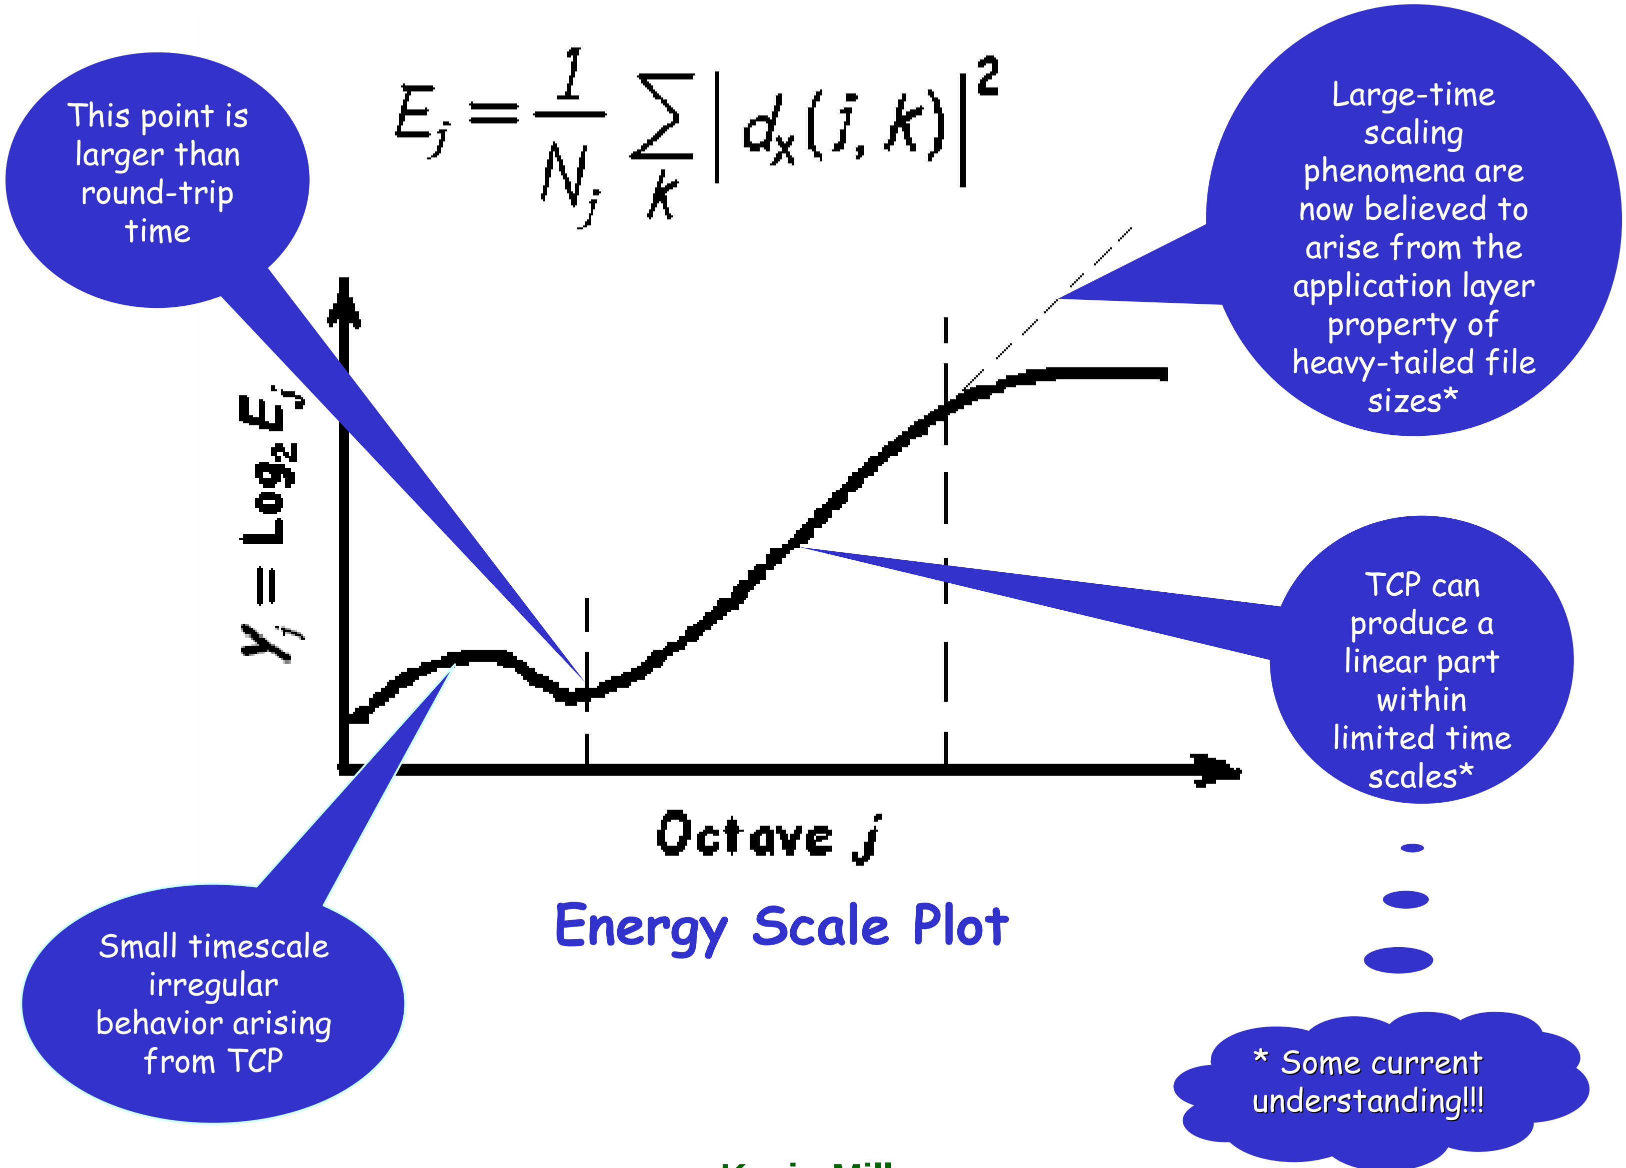

Two-Tiered Homogeneous CA Model of a Network

L x ***L*** **grid of routers connected by links of capacity n_l packets per time step**

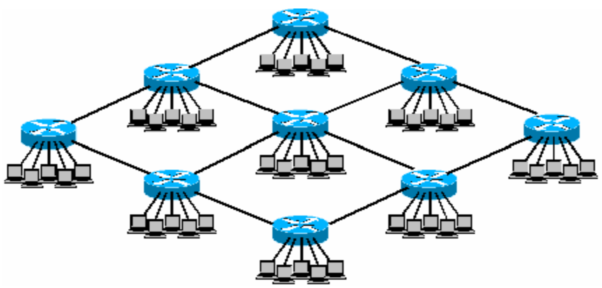

Equal number n_s of sources attached to each router with a variable number ($\leq 2n_s$) of receivers attached to each router

- **Sources** – **generate** packets (subject to congestion control algorithms)
- **Generation process** – each source has **on-off process**:
  - at beginning of each on period randomly select a receiver
  - at each time step, generate a packet if permitted
  - duration of on and off periods **exponentially distributed** (with means λon and λoff) or **Pareto distributed** with means .24 λon /1.2 and .24 λoff /1.2
- **Congestion control algorithms** (*explained soon*)
  - TCP Flow Control
  - TCP Friendly Rate Control (TFRC)
- **Routers** – maintain **limited length queue** (50 packets here) and **forward** packets on to neighbors
  - next hop selected **nearest neighbor to the left** (so that packets between source-destination pairs are split among the two equidistant routes)
- **Receivers** – **consume** packets
- **System State**
  - number of packets consumed and forwarded by a selected destination router during each time step

Alternate Congestion-Control Mechanisms

- **TCP** – source sends data packets and expects destination to send ack for each data packet.
  - If ack indicates missed data packet, set slow-start threshold to ½ congestion window
  - If ack indicates no data packet missed:
    - if congestion window < slow-start threshold, then congestion window++
    - else congestion window = congestion window + 1/congestion window
  - At each time step source generates a packet, but can only have as many packets in transmission as the congestion window allows
- **TFRC** – receiver computes packet loss rate and feeds that back to sender, which estimates round-trip time (RTT)
  - Source inputs packet loss rate and estimated RTT into a TFRC throughput equation to learn when to transmit the next packet (i.e., what should be the interval between packet transmissions)

Investigating Effects of Application Layer

L = 3, n_s = 10, n_l = 5, TCP – note n_l that determines granularity of observation

Exponential λ_{on} = 200, λ_{off} = 2000

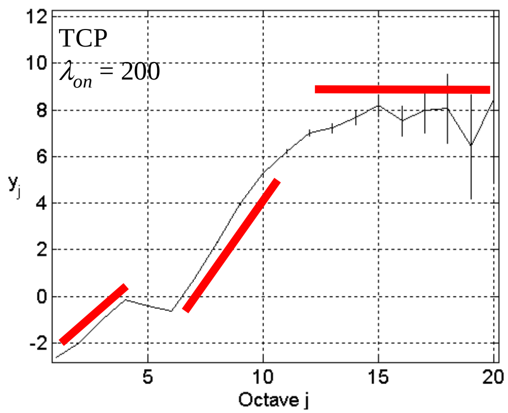

Heavy-tailed distribution of file sizes (modeled as *on* periods) leads to a pronounced autocorrelation in traffic over a range of about 11 octaves

Exponential distribution in file sizes leads to a more limited autocorrelation in traffic over a range of about 6 octaves

These results are consistent with the results of others; thus, raising confidence in our model

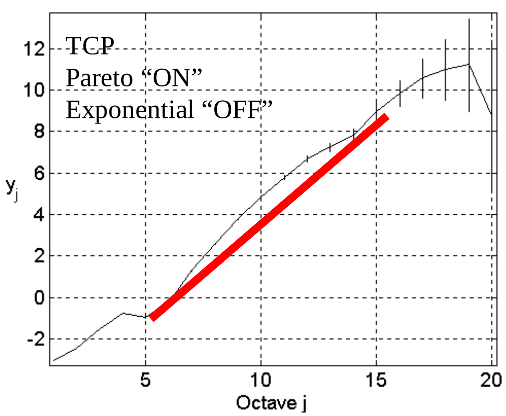

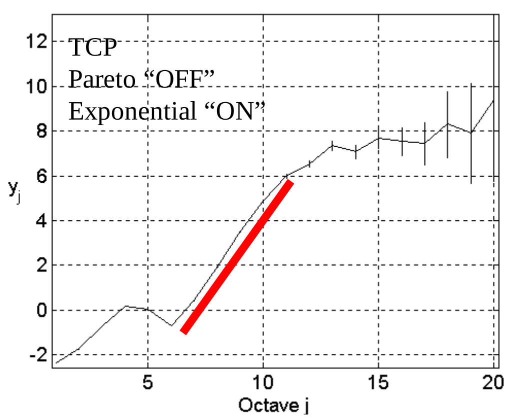

Investigating Effects of Transport Layer

Heavy-tailed distribution of file sizes appears to give rise to long-range dependence regardless the transport mechanism used; however lowering the link capacity destroys
correlation structure for TFRC but not TCP

$L = 3, n_s = 10, \lambda_{on} = 200, \lambda_{off} = 2000$

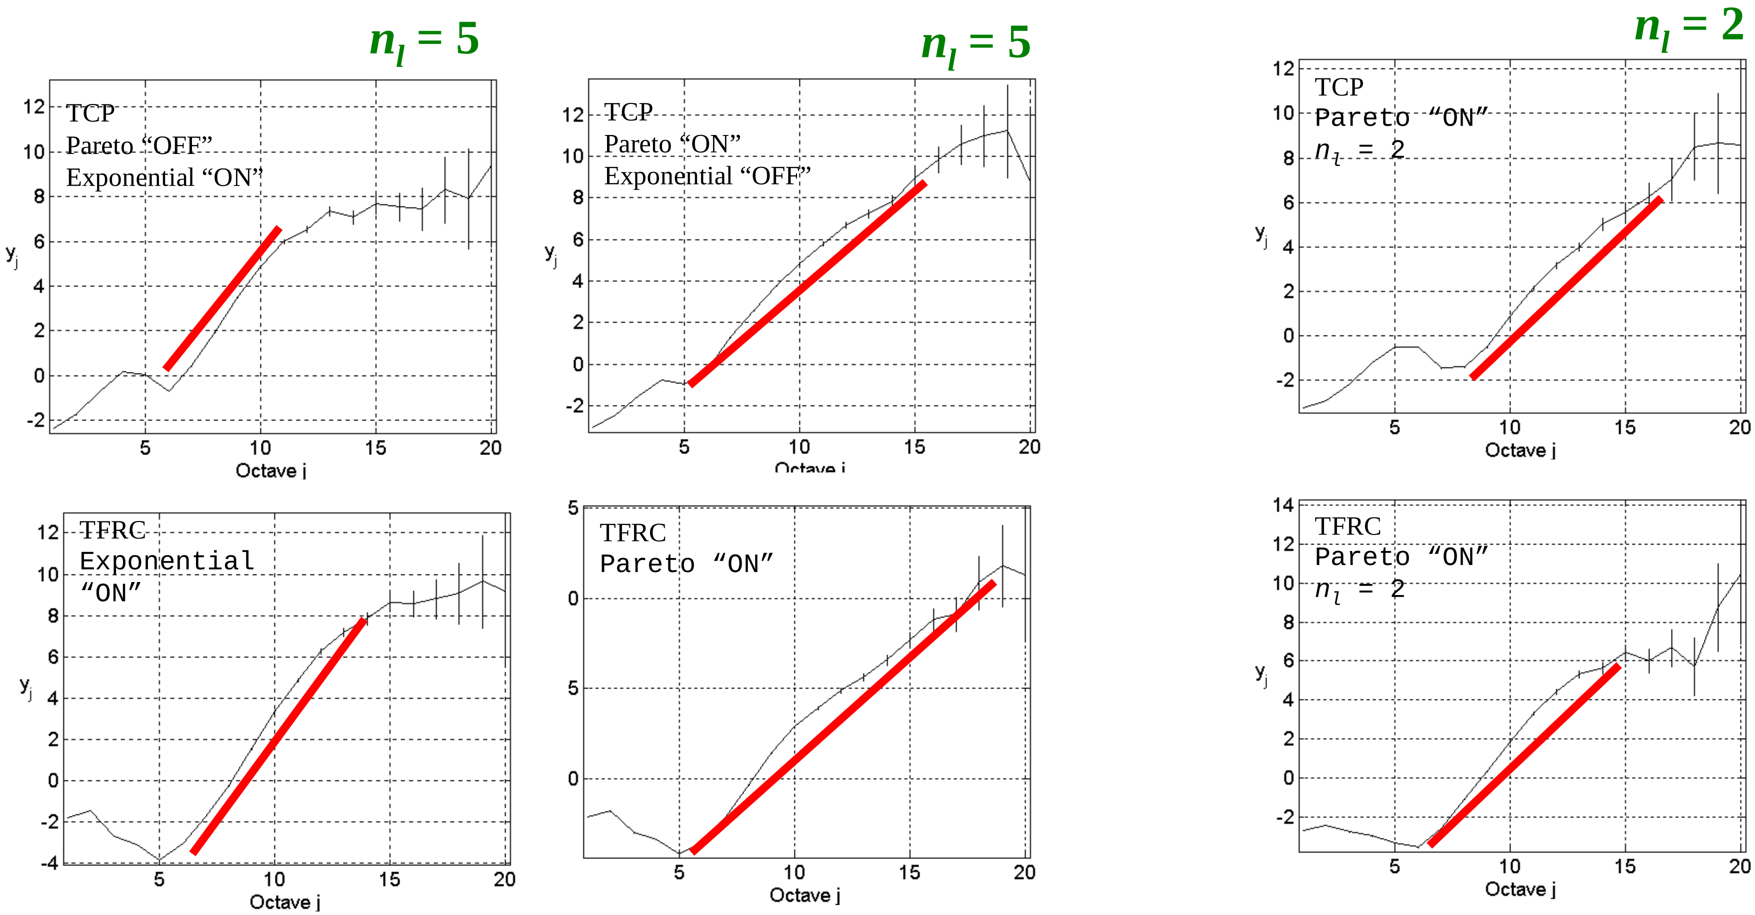

Kevin Mills

Investigating Effects of Link Capacity

Restricting network capacity appears to strengthen correlation structure, while expanding network capacity appears to weaken correlation structure

***L* = 3, Exponential λ_{on} = 200, λ_{off} = 2000**

***n_l* = 2**

***n_l* = 20**

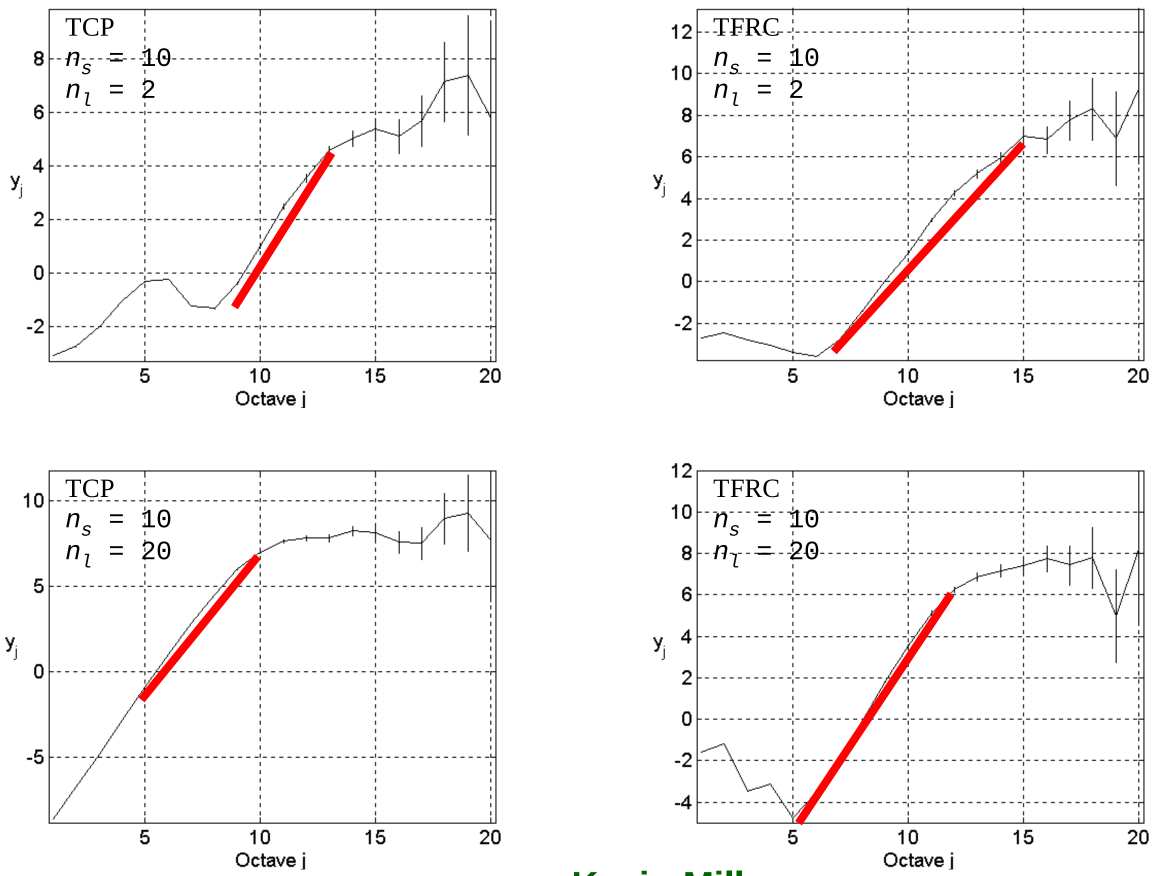

Kevin Mills

Investigating Effects of Traffic Demands

Independent of transport mechanism, increasing the traffic demand for a fixed network capacity increases correlation, while increasing network capacity for a fixed traffic demand weakens correlation **L = 3, Exponential λ_{on} = 200, λ_{off} = 2000**

n_s = 10
n_l = 5

n_s = 40
n_l = 5

n_s = 40
n_l = 20

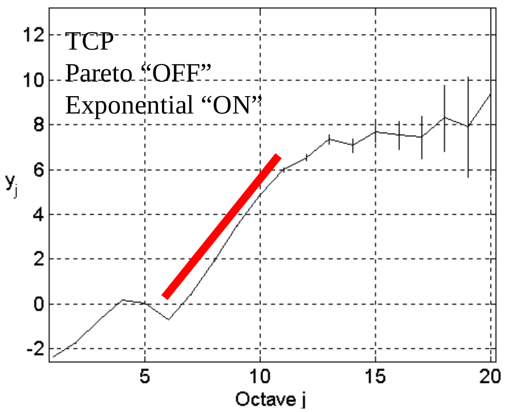

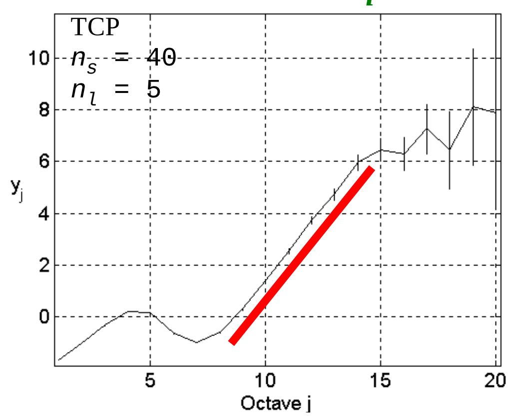

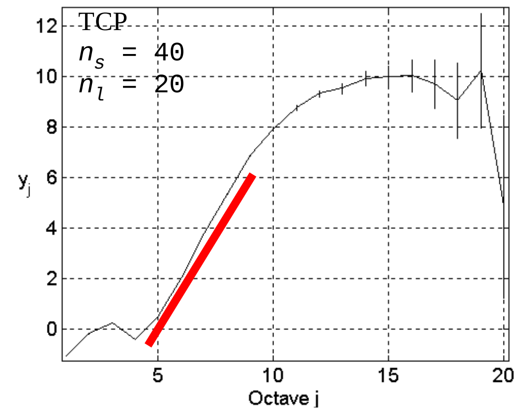

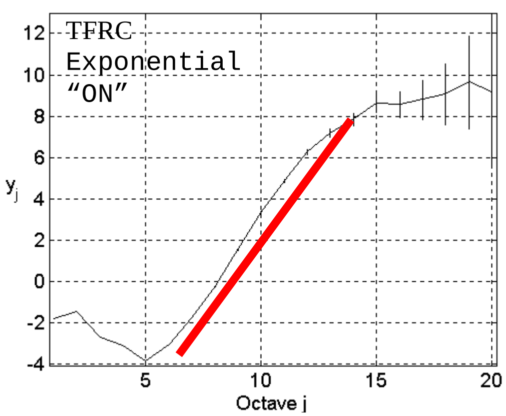

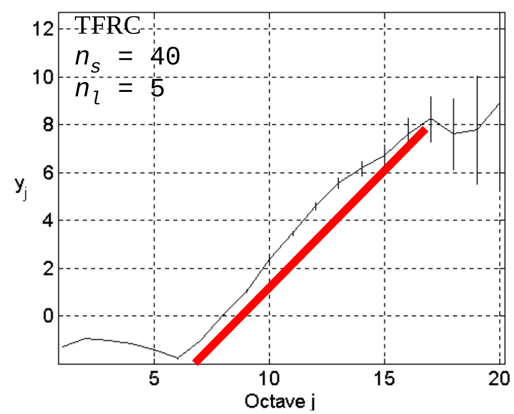

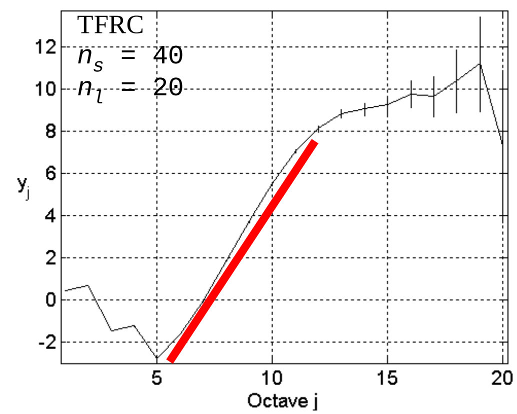

Investigating Effects of Network Size

Independent of transport mechanism, increasing network size increases correlation structure, given the same traffic demand and network capacity

n_s = 10, n_l = 1, Exponential λ_{on} = 200, λ_{off} = 2000

***L* = 3**

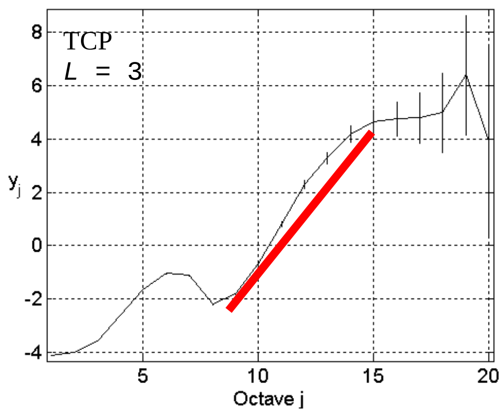

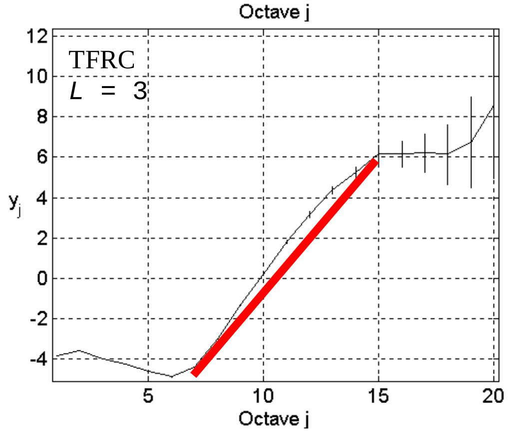

***L* = 9**

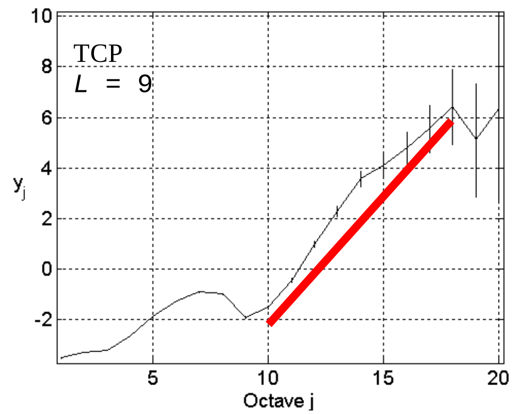

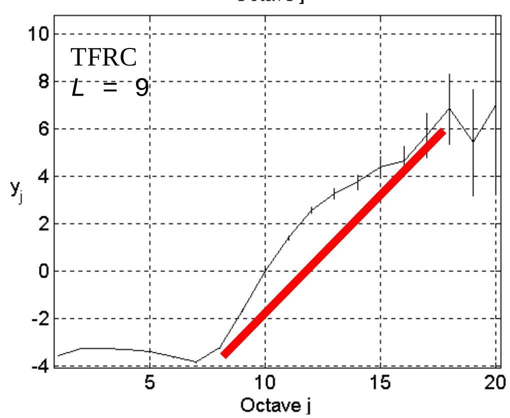

***L* = 27**

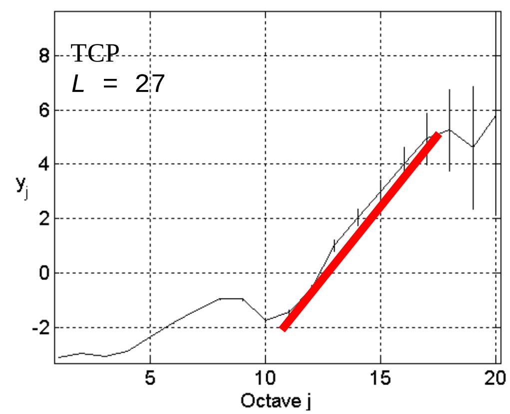

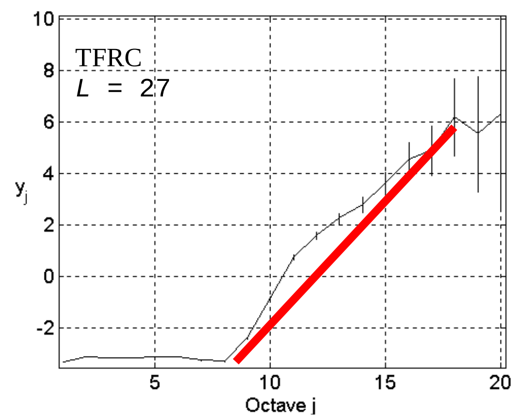

Challenges in Recognizing Emergence

- What data should be collected?
- How much data should be collected?
- At what time granularity should data be collected?
- How should collected data be analyzed?
- How can data be collected and analyzed in real-time throughout a large-scale network (which is not homogeneous)?

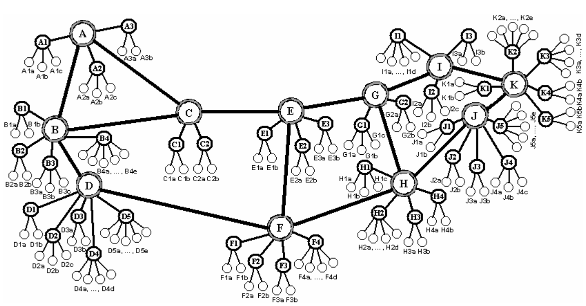

Normalized flow from domain *i* to domain *j*

$$f_{ij} = (x_{ij} - m_{ij}) / \sigma_{ij}$$

Cross-correlation between time-averaged flow vectors

$$C_{(ij)(kl)} = \left\langle f_{ij}(t) f_{kl}(t) \right\rangle$$

Principle Component Analysis of cross-correlation

$$\boldsymbol{Cw} = \lambda \boldsymbol{w}$$

Compute the *i*th domain's contribution to *k*th domain

$$\sum_{i,k} (w_{ik}^1)^2 = 1$$

Compute relative strength of flows to *k*th domain

$$S_k = \sum_{i}^{L} (w_{ik}^1)^2$$

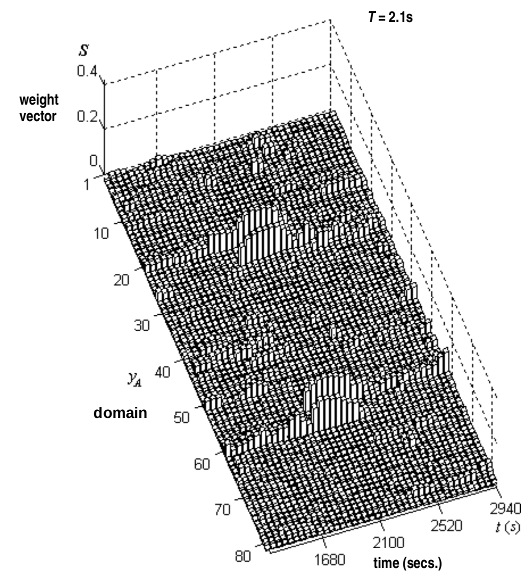

Kevin Mills

Challenges in Interpreting, Exploiting, Eliciting, and Controlling Emergence

- Can evidence of emergent behavior be attributed to appropriate cause(s)?
- Can coherent behavior be recognized and acted upon in time to effect control?
- Can decentralized feedback and adaptation be applied effectively to elicit a desired coherent state?
- Does a system that self-organizes to a “critical” (coherent state) imply a substantial probability of exhibiting chaotic behavior?
- If so, then can a system operating at a critical state be prevented from exhibiting chaotic behavior?

Conclusions

I wish we knew more about these questions because I suspect that our large-scale information systems (e.g., Internet, Web, and Grid) will exhibit emergent properties long before we are able to understand what is happening and why, or to do anything about it.

For Further Reading on Emergence

M. Schroeder, Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise, W. H. Freeman, New York, 1991.

P. Bak, How Nature Works: the science of self-organized criticality, Copernicus, New York, 1996.

Y. Bar-Yam, The Dynamics of Complex Systems: Studies in Nonlinearity, Perseus Books, Cambridge, Mass., August 1997.

J. H. Holland, Emergence: From Order to Chaos, Perseus Books, Cambridge, Mass., 1998.

S. Johnson, Emergence: The Connected Lives of Ants, Brains, Cities, and Software, Scribner, New York, 2001.

K. N. Lodding, "Hitchhiker's Guide to Biomorphic Software", *QUEUE*, June 2004, pp. 66-76.

For Further Reading on Network Modeling and Analysis

I. Csabai, "1/f Noise in Computer Network Traffic", J. Phys., A 27 (12), L417-L421 (1994).

M. Takayasu, H. Takayasu, T. Sato, "Critical behaviors and 1/f noise in information traffic", Physica A 233, 824-834 (1996).

T. Ohira and R. Sawatari, "Phase transition in computer network traffic model", Phys. Rev. E 58, 193-195 (1998).

V. Paxson, "End-to-End Internet Packet Dynamics", *IEEE/ACM Transactions on Networking* 7 (3), 277-292, 1999.

B. Liu, Y. Guo, J. Kurose, D. Towsley, and W. Gong, "Fluid simulation of large scale networks: issues and tradeoffs", *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, 1999.

J. Cowie, H. Liu, J. Liu, D. Nicol and A. Ogielski, "Towards Realistic Million-Node Internet Simulations", *Proceedings of the 1999 International Conference on Parallel and Distributed Processing Techniques and Applications*, 1999.

K. Fukuda, M. Takayasu, and H. Takayasu, "Spatial and temporal behavior of congestion in Internet traffic", Fractals 7, 23-31 (1999).

A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control", *Proc. ACM SIGCOMM*, 301-313, 1999.

R. V. Sole and S. Valverde, "Information transfer and phase transitions in a model of Internet traffic", Physica A 289 (3-4), 595-605, 2001.

For Further Reading on Related Work by Yuan and Mills

J. Yuan, Y. Ren, and X. Shan, "Self-Organized Criticality in a Computer Network Model", Physics Review, E 61(1), 1067-1071, 2000.

J. Yuan, Y. Ren, F. Liu, X. Shan, "Phase transition and collective correlation behavior in the complex computer network", Acta Physica Sinicia, 50 (7), 1221-1225, July 2001.

J. Yuan and K. Mills, "Exploring Collective Dynamics in Communication Networks", *The NIST Journal of Research*, Volume 107, No. 2, March-April 2002. http://w3.antd.nist.gov/~mills/papers/YuanMills2002.pdf

J. Yuan and K. Mills, "Macroscopic Dynamics in Large-Scale Data Networks", chapter in upcoming book Complex Dynamics in Communication Networks, edited by Ljupco Kocarev and Gábor Vattay, to be published by Springer, in press. http://w3.antd.nist.gov/~mills/papers/BookChapterYuanMills.pdf

J. Yuan and K. Mills, "A Cross-Correlation Based Method for Spatial-Temporal Traffic Analysis", accepted for publication by the journal *Performance Evaluation*. http://w3.antd.nist.gov/~mills/unpublished/LRDHT-13revised.pdf

J. Yuan and K. Mills, "Monitoring the Macroscopic Effects of DDoS Flooding Attacks", under review by the *IEEE Transactions on Dependable and Secure Computing*. http://w3.antd.nist.gov/~mills/unpublished/ReformattedTDSC-0026-0204.pdf