CS 525
Advanced Distributed Systems
Spring 2010

Indranil Gupta (Indy)

Structure of Networks

April 29, 2010

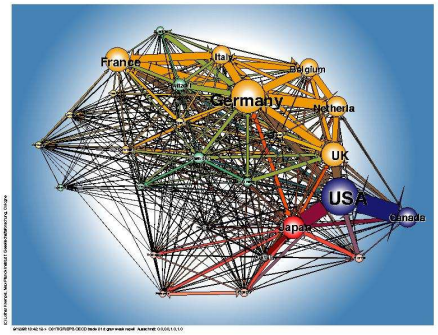

1981

Graph or *Network*

Countries= *nodes*

Treaties= *edges*

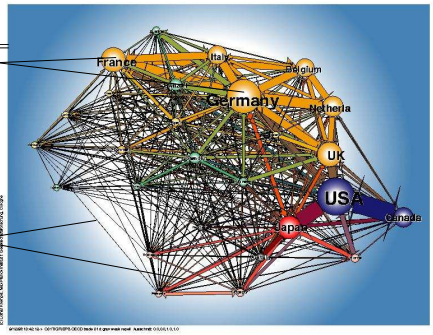

1981

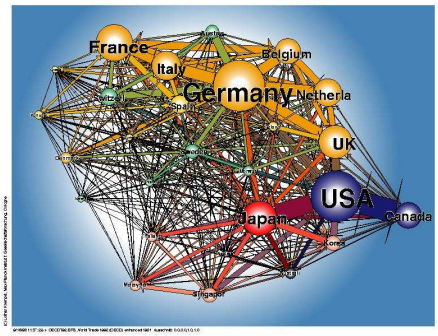

1992

PCs= *nodes*

Connected= *edges*

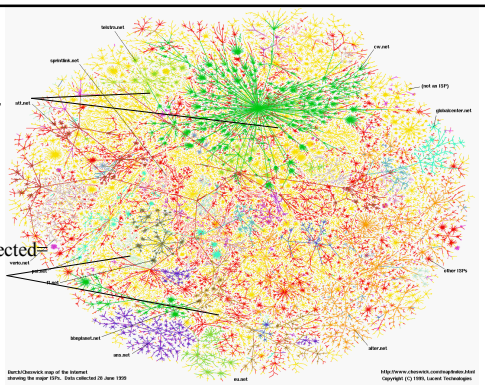

The Internet (Internet Mapping Project, color coded by ISPs)

Food Web of Little Rock Lake, WI

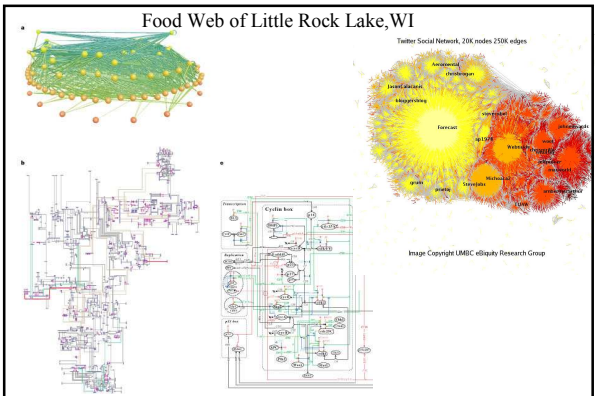

Electric Power Grid

Metabolic reaction network

This Lecture: Common Thread

Networks
- Structure of,
- Dynamics within,

- We'll study networks at three different "levels"

7

Lowermost Level: Basics, Physical Phenomena, and Life

8

Complexity of Networks

- Structural: *human population has 6 B nodes, there are millions of computers on the Internet…*
- Evolution: *people make new friends all the time, ISP's change hands all the time…*
- Diversity: *some people are more popular, some friendships are more important…*
- Node Complexity: *PCs have different CPUs, Windows is a complicated OS…*
- Emergent phenomena: simple end behavior ➔ complex system-wide behavior. *If we understand the basics of climate change, why is the weather so unpredictable?*

9

1. Network Structure

- "Six degrees of Kevin Bacon"
- Milgram's experiment in 1970
- Watts and Strogatz Model
- Kleinberg's algorithmic results

- Recent work on mapping Internet, WWW, p2p overlays, electric power grid, protein networks, co-authorship among scientists
- These networks have "evolved naturally"

10

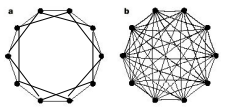

Fully Connected graph

Ring graph

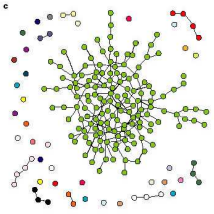

Random graph

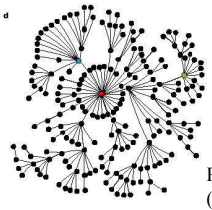

Power Law Graph
(Degenerate: tree)

11

A Scientist's Perspective

- Two important metrics
  - Clustering Coefficient: CC
    - Pr(A-B edge, given an A-C edge and a C-B edge)
  - Path Length of shortest path
- Ring graph: high CC, long paths
- Random graph: low CC, short paths
- Small World Networks: high CC, short paths

12

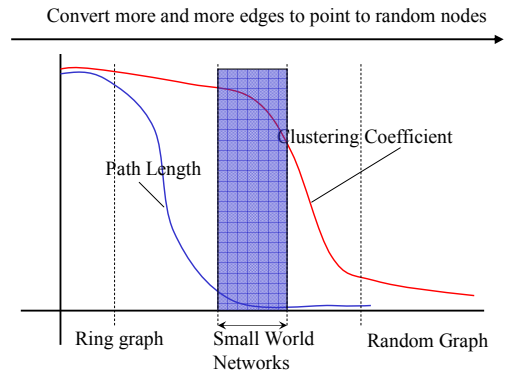

13

Most "natural evolved" networks are (probably) small world

- Network of actors → six degrees of Kevin Bacon
- Network of humans → Milgram's experiment
- Co-authorship network → "Erdos Number"

Many of these networks also "grow incrementally" [Faloutsos and Faloutsos]

14

Another Scientific Viewpoint

That was about "nature of neighbors"; what about *number* of neighbors?

Degree distribution – what is the probability of a given node having k edges (neighbors, friends, …)

- Regular graph: all nodes same degree
- **Gaussian**
- Random graph: **Exponential** $e^{-k.c}$
- **Power law**: $k^{-\alpha}$

15

Basics: The Log-Log Plot

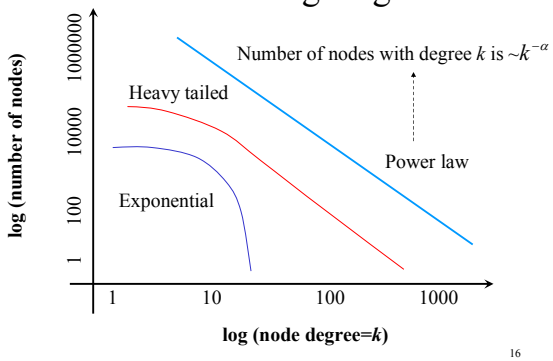

16

WWW is a power law graph $\alpha = 2.1 - 2.4$

NCSTRL co-author graph is power law, with exponential cutoff

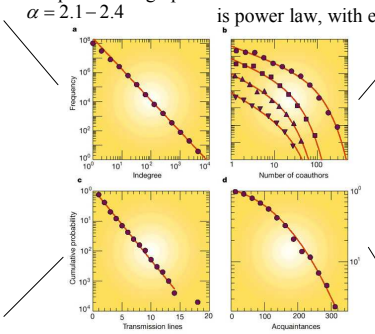

Electric Power Grid graph is exponential

Social network of Utah Mormons is Gaussian

17

Power law vs. Small World

- A lot of small world networks are power law graphs (Internet backbone, telephone call graph, protein networks)
- Not all small world networks are power law (e.g., co-author networks)
- Not all power law networks are small world
- *Preferential Model* for network growth generates power law distributions – special way of incremental growth
  - e.g., Web pages linking to each other
- Power law networks also called scale-free

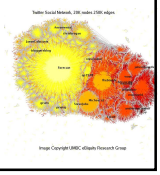

Power law + Small world

Most nodes have small degree, but a few nodes have high degree

Attacks on small world networks
- Killing a large number of randomly chosen nodes does not disconnect graph
- Killing a few high-degree nodes will disconnect graph

"A few (of the many thousand) nutrients are very important to your body"
"The Electric Grid is very vulnerable to terrorists"

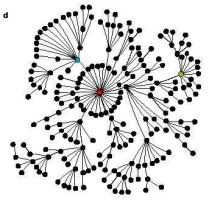

19

2. Network Dynamics

- Strogatz goes on to discuss dynamics of many "natural networks"
- We'll focus on dynamics w.r.t. the Internet and P2P networks in the papers [Akella et al] and [Ripeanu et al]
- But let's just touch a bit on oscillation dynamics in networks…

20

- Networks of coupled dynamical systems
- If each node is a dynamical system, and is affected by its neighbors, what behaviors emerge from the entire network?
- E.g., Social networks, network of neurons in the brain, protein networks, …
- An example of emergent behavior: *self-synchronization*

21

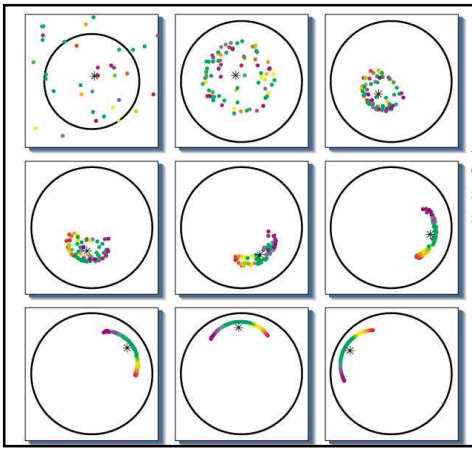

A group of oscillators can self-synchronize

22

Self-Synchronizing Fireflies

- Synchronizing Fireflies of Malaysia
- Each firefly: $\dot{\theta} = \omega$
- Is driven by an external stimulus $\dot{\Phi} = \Psi$

so $\dot{\theta} = \omega + A\sin(\Phi - \theta)$

- Can show self-sync occurs when

$$\omega - \mathrm{A} < \Psi < \omega + \mathrm{A}$$

- For more details see [Strogatz's non-linear dynamics textbook]

23

Why the heart beats by itself

- Consists of a few thousand *sinoatrial cells*
- Each oscillating at its own frequency

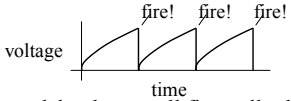

- Peskin's model: when a cell fires, all other cells have a small jump in voltage
- [Why does this self-synchronize?]
  - Think of two sinoatrial cells first
- For more details, see [Strogatz's book "Sync"]

24

Discussion

- What is one problem where a self-synchronizing system could be used to design a distributed protocol?
- Why is the co-authorship network different from the Internet if both follow an incremental / preferential construction?

25

A Level Up: The Internet

26

- [Faloutsos et al] showed that the Internet backbone follows a power law distribution
- [One kind of Dynamism over such a network?]

27

- [Faloutsos et al] showed that the Internet backbone follows a power law distribution
- [One kind of Dynamism over such a network?]
- Routing [Akella et al]

  What is the stress on Internet routers due to
  - Shortest path routing ("efficient")
  - Policy based routing (BGP)

28

Internet is a multi-level topology
At the highest level, it consists of AS's
AS's consist of subnets, then LANS, …
AS-AS routing is done by BGP

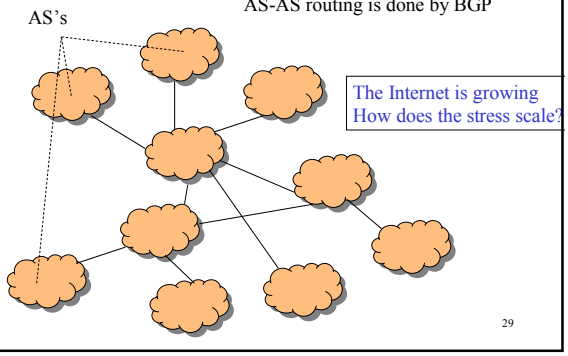

29

Main Result

- Take a power law network (node has degree k with probability $k^{-\alpha}$)
- **Shortest path routing,** with ties broken by higher degree
- With **uniform traffic model** for all pairs of end nodes, maximum edge congestion grows as

$$O(n^{1+\frac{1}{\alpha}})$$ [Theorem 1]

30

Log-log plot again
Power-law and Tree network topologies give superlinear congestion
Exponential Network has sublinear congestion [why?]

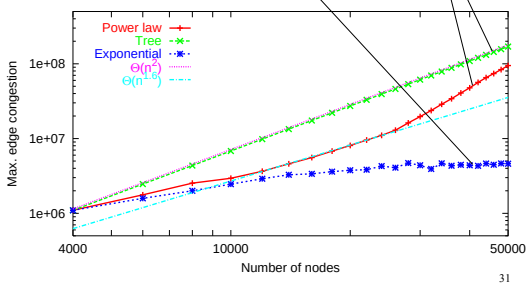

31

Clout Traffic Model: well-connected nodes generate more end to end traffic

+ **Shortest-path routing** has worse max edge congestion than with uniform traffic

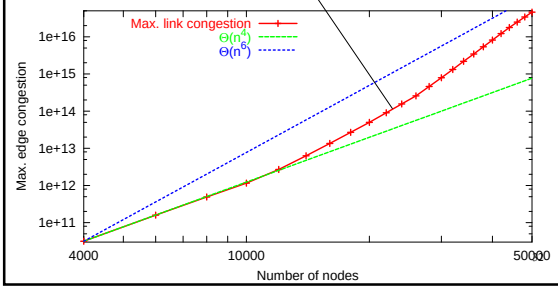

Policy Routing

- Due to ISP to ISP financial contracts, AS to AS edges are either
  - Customer-provider edges, or
  - Peering edges
- Policy routing prefers customer → provider traffic
- Gives "valley free" paths: most edges are customer → provider

33

Clout Traffic Model: well-connected nodes generate more traffic

+ **Policy routing** also gives superlinear growth, *but is better than shortest path routing*

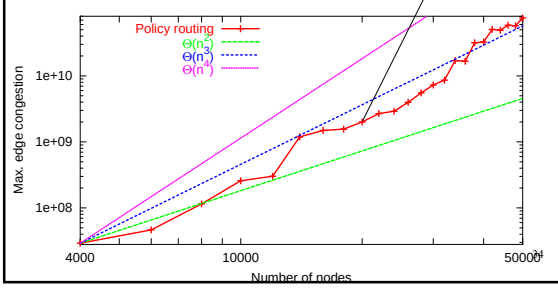

34

Policy-based and Shortest Path Routing give similar edge congestion for the uniform traffic model

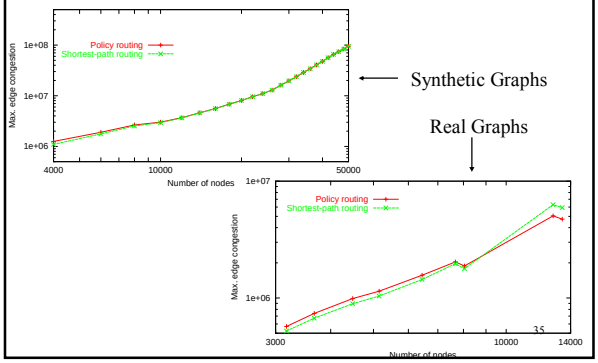

35

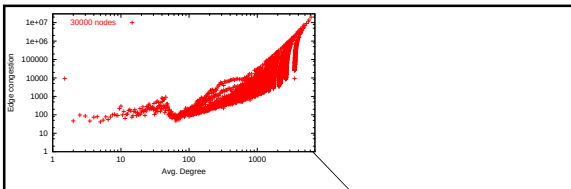

A Solution: Add redundant edges between selective node pairs

Congestion varies linearly with n

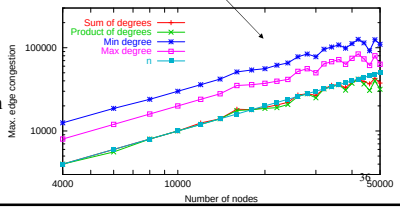

36

Discussion

- Metrics: max edge congestion (why not average?)
  - Instability from single source could spread
    - Think "Electric Power Grid failures"
    - Think "self-synchronizing routers"
- Why is Shortest Path Routing always worse than Policy Routing?
  - Shortest Path Routing is supposed to be "efficient"
  - Outrageous Opinion: Are policies the reason why the Internet stays up and robust? Should the design of Internet be left to non-technicians?

37

Another Level Up: Applications

38

Study of Gnutella

- [Ripeanu et al]
- Gnutella
  - Peer to peer Overlay
  - Users download songs from other users, upload their own songs
  - Each computer host = "peer"
  - Completely decentralized

39

Gnutella Structure

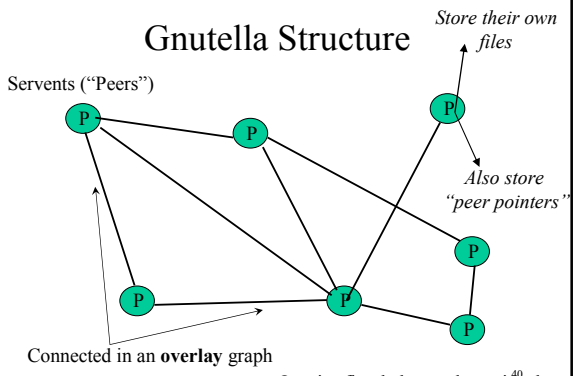

40

Study of Gnutella

- 6 month period 10/00-5/01
  - Before revision of Gnutella protocol (late 2001)
- Characteristics
  - System Size
  - Network Traffic
  - Node Connectivity

41

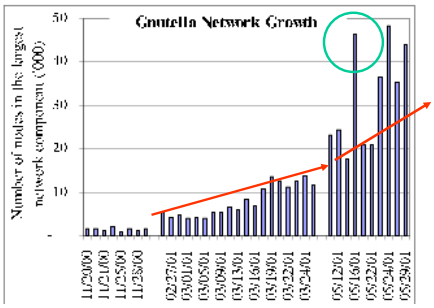

95% of nodes in largest connected component
Quick Growth over time (exponential?)
Spikes

42

Churn Characteristics

- 40% of nodes logged in for less than 4 hours
- 25% nodes alive for more than 24 hours

43

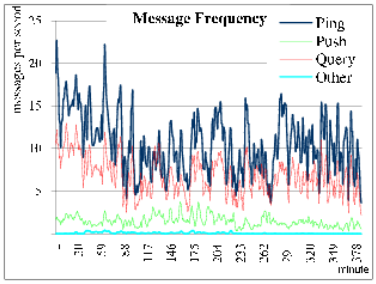

55% ping-pong messages (membership)
36% query messages
Subsequent improvements reduced these to 8%, 92%

95% of nodes less than 7 hops away
[Implication of ttl=7 for query messages?]

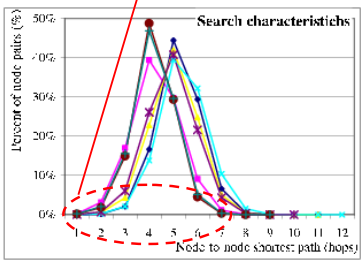

45

Traffic Volume

- 170 K connexns for 50 K node Gnutella
- 6 KBps per connexn ➔1 GBps total
  ➔330 TB/month
- 1.7% of total traffic in Internet Backbone
- Recall: [3/00] 25% UWisc traffic from Napster

46

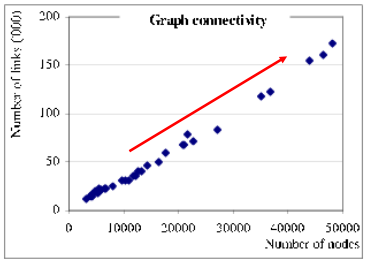

Average degree of node is scale-independent
On average 3.4 edges / node

47

Not power law
But Heavy-tailed

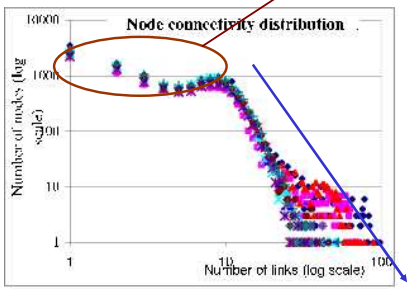

48

Overlay-Network Match

- Does the Gnutella Overlay reflect the underlying Internet structure?
- Entropy technique in paper
  - Nodes identified with their domain names
  - Gnutella graph structure → clustering of nodes
  - Calculate entropy of above clustering and compare with entropy of a random selection of nodes from across domains
  - If same, Gnutella graph is random, otherwise it is more ordered
- Authors find Gnutella clustering entropy to be 8% lower than random clustering entropy
- *Gnutella structure is independent of underlying Internet [hence the term "application overlays"]*

49

Discussion

- Do overlays really reflect the application?
  - Concept of malleable overlays
- Are application-dependent overlays "unfriendly" to the network and other applications?
- What if the overlays are very large? (think PlanetLab)
- What if they are small and there are millions of them? (think overlay hosting services)
- What if they are large *and* many? (think overlay hosting services on top of PlanetLab-style clusters)

50

Another level Up: The Users, Humans, …

51

Summary

- Humans, and the networks connecting them…societal networks, actor networks, co-authorship graphs….
- And we're back full circle!

- We've discussed
  - Network structure
  - Network dynamics
- Many commonalities
  - power law, small world …

  among "naturally evolved" networks
  - social nets, metabolic nets, electric power grid, Internet, WWW,…
- Can look at in awe, but systems design also has to deal with it
- Large opportunity to design distributed protocols inspired by nature (rigorously!) …

52

Discussion

- Questions?

53

Final Report Grading

- Final report will be graded just like a conference reviewer would:
  - Importance of problem
  - Novelty of solution
  - Evaluation of solution
  - Clarity of Presentation
  - Nits (grammar, references, etc.)

54

Semester's Final Lecture

- 3 papers from website + 2 handouts
  - Read as many as you can – you'll enjoy them
  - None of them is technical!
- Next Tuesday's lecture (last lecture) – we'll close the discussion we started in the semester's first lecture

55