CMU SCS

15-826: Multimedia Databases and Data Mining

Graph mining - part 1

Christos Faloutsos

CMU SCS

Thanks

- Deepayan Chakrabarti (CMU)
- Soumen Chakrabarti (*IIT-Bombay*)
- Michalis Faloutsos (UCR)
- George Siganos (UCR)




CMU SCS

Introduction


Internet Map
[lumeta.com]

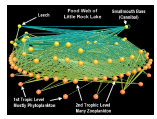
Food Web
[Martinez '91]

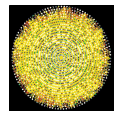
Protein Interactions
[genomebiology.com]

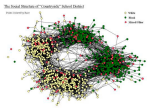
Friendship Network
[Moody '01]

Graphs are everywhere!

CMU SCS

Graph structures

- Physical networks
- Physical Internet
- Telephone lines
- Commodity distribution networks

CMU SCS

Networks derived from "behavior"

- Telephone call patterns
- Email, Blogs, Web, Databases, XML
- Language processing
- Web of trust, epinions

CMU SCS

Outline

Part 1: Topology, 'laws' and generators

Part 2: PageRank, HITS and eigenvalues

Part 3: influence, virus propagation, communities

Motivating questions:

CMU SCS

Part 1: Topology and generators

- What do real graphs look like?
- What properties of nodes, edges are important to model?
- What local and global properties are important to measure?
- How to model and generate realistic graphs?

CMU SCS

Part 2: PageRank, HITS and eigenvalues

- How important is a node?
- Who is the best customer to advertise to?

CMU SCS

Part 3: influence and communities

- How do nodes influence their neighbors?
- How do viruses propagate?

CMU SCS

PART 1: Topology, laws and generators

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Influence, viruses, communities

CMU SCS

Motivating questions

- What do real graphs look like?
  - What properties of nodes, edges are important to model?
  - What local and global properties are important to measure?
- How to generate realistic graphs?

CMU SCS

Motivating questions

Given a graph:

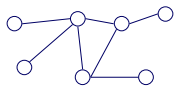

- Are there un-natural sub-graphs? (criminals' rings or terrorist cells)?
- How do P2P networks evolve?

CMU SCS

Why should we care?

Internet Map [lumeta.com]

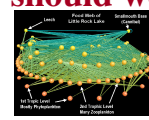

Food Web [Martinez '91]

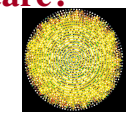

Protein Interactions [genomebiology.com]

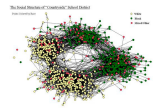

Friendship Network [Moody '01]

CMU SCS

Why should we care?

- **A1: extrapolations**: how will the Internet/Web look like next year?
- **A2: algorithm design**: what is a realistic network topology,
  - to try a new routing protocol?
  - to study virus/rumor propagation, and immunization?

CMU SCS

Why should we care? (cont'd)

- **A3: Sampling**: How to get a 'good' sample of a network?
- **A4: Abnormalities**: is this sub-graph / sub-community / sub-network 'normal'? (what **is** normal?)

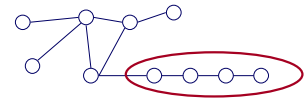

CMU SCS

Outline

Part 1: Topology, 'laws' and generators

- 'Laws' and patterns
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Pairs, influence, communities

CMU SCS

Topology

How does the Internet look like? Any rules?

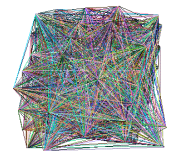

(Looks random – right?)

CMU SCS

Are real graphs random?

- random (Erdos-Renyi) graph – 100 nodes, avg degree = 2
- before layout
- after layout
- No obvious patterns

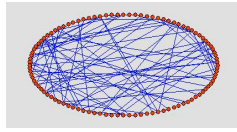

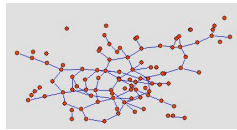

(generated with: pajek
http://vlado.fmf.uni-lj.si/pub/networks/pajek/
)

CMU SCS

Laws and patterns

Real graphs are NOT random!!

- Diameter
- in- and out- degree distributions
- other (surprising) patterns

CMU SCS

Laws – degree distributions

- Q: avg degree is ~2 - what is the most probable degree?

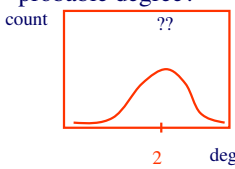

CMU SCS

Laws – degree distributions

- Q: avg degree is ~2 - what is the most proba

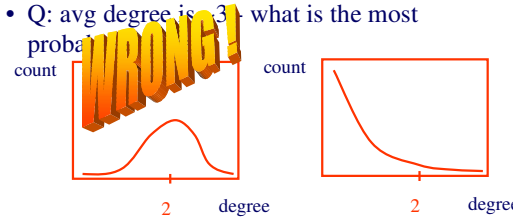

CMU SCS

I.Power-law: outdegree *O*

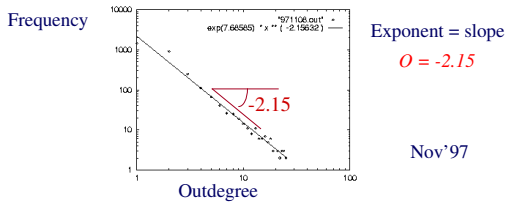

Exponent = slope
O = -2.15

Nov'97

The plot is linear in log-log scale [FFF'99]

freq = degree (-2.15)

CMU SCS

II.Power-law: rank *R*

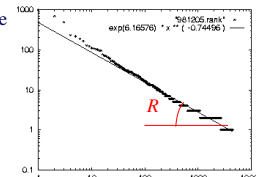

Exponent = slope
R = -0.74

Dec'98

- The plot is a line in log-log scale

CMU SCS

III. Eigenvalues

- Let A be the adjacency matrix of graph
- The eigenvalue λ is:
  - $A \underline{v} = \lambda \underline{v}$, where $\underline{v}$ some vector
- Eigenvalues are strongly related to graph topology

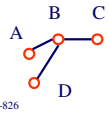

	A	B	C	D
A		1		
B	1		1	1
C		1		
D		1		

CMU SCS

III.Power-law: eigen *E*

Eigenvalue

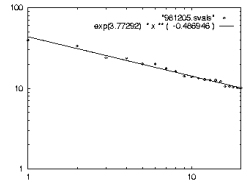

Exponent = slope

E = -0.48

Dec'98

Rank of decreasing eigenvalue

- Eigenvalues in decreasing order (first 20)
- [Mihail+, 02]: *R = 2 * E*

CMU SCS

IV. The Node Neighborhood

- *N(h)* = # of pairs of nodes within *h* hops

CMU SCS

IV. The Node Neighborhood

- Q: average degree = 3 - how many neighbors should I expect within 1,2,… *h* hops?
- Potential answer:

  1 hop -> 3 neighbors

  2 hops -> 3 * 3

  …

  h hops -> 3^h

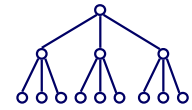

CMU SCS

IV. The Node Neighborhood

- Q: average degree = 3 - how many neighbors should I expect within 1,2,… *h* hops?
- Potential answer:

  1 hop -> 3 neighbors

  2 hops -> 3 * 3

  …

  h hops -> 3^h

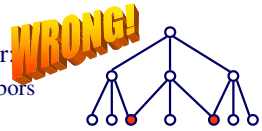

WE HAVE DUPLICATES!

CMU SCS

IV. The Node Neighborhood

- Q: average degree = 3 - how many neighbors should I expect within 1,2,… *h* hops?
- Potential answer:

  1 hop -> 3 neighbors

  2 hops -> 3 * 3

  …

  h hops -> 3^h

'avg' degree: meaningless!

CMU SCS

IV. Power-law: hopplot *H*

of Pairs *H = 4.86*

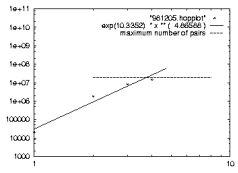

Hops Dec 98

# of Pairs *H = 2.83*

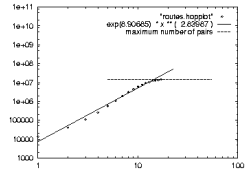

Hops Router level '95

Pairs of nodes as a function of hops $N(h) = h^H$

CMU SCS

Observation

- Q: Intuition behind 'hop exponent'?
- A: 'intrinsic=fractal dimensionality' of the network

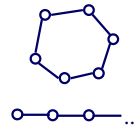

$N(h) \sim h^1$

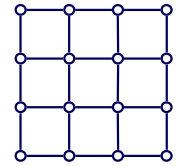

$N(h) \sim h^2$

CMU SCS

But:

- Q1: How about graphs from other domains?
- Q2: How about temporal evolution?

CMU SCS

The Peer-to-Peer Topology

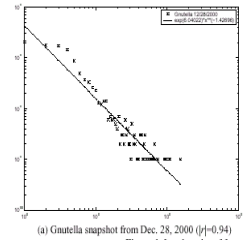

(a) Gnutella snapshot from Dec. 28, 2000 ($|r|$=0.94)

[Jovanovic+]

- Frequency versus degree
- Number of adjacent peers follows a power-law

CMU SCS

More Power laws

- Also hold for other web graphs [Barabasi+, '99], [Kumar+, '99] with additional 'rules' (bi-partite cores follow power laws)

CMU SCS

Time Evolution: rank *R*

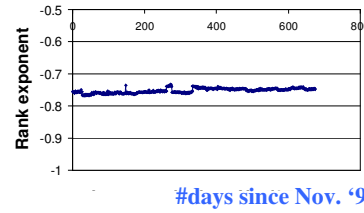

Domain level

The rank exponent has not changed! [Siganos+, '03]

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
  - Power laws for degree, eigenvalues, hop-plot
  - other rules???
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Pairs, influence, communities

CMU SCS

Any other 'laws'?

Yes!

CMU SCS

Any other 'laws'?

Yes!
- Small diameter (~ constant!) –
  - six degrees of separation / 'Kevin Bacon'
  - small worlds [Watts and Strogatz]

CMU SCS

Any other 'laws'?

- Bow-tie, for the web [Kumar+ '99]
- IN, SCC, OUT, 'tendrils'
- disconnected components

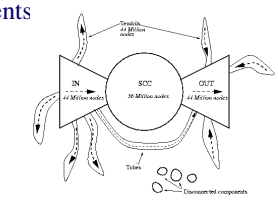

CMU SCS

Any other 'laws'?

- power-laws in communities (bi-partite cores) [Kumar+, '99]

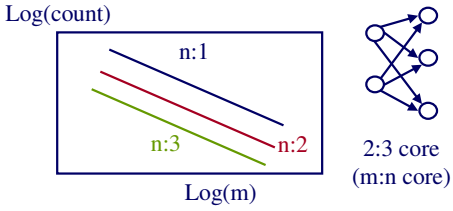

2:3 core
(m:n core)

CMU SCS

Any other 'laws'?

- "Jellyfish" for Internet [Tauro+ '01]
- core: ~clique
- ~5 concentric layers
- many 1-degree nodes

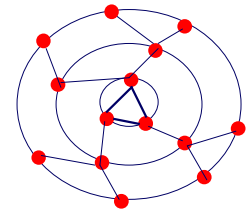

CMU SCS

Summary of 'laws'

- Power laws for degree distributions
- …….... for eigenvalues, bi-partite cores
- Small diameter ('6 degrees')
- 'Bow-tie' for web; 'jelly-fish' for internet

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
  - Power laws for degree, eigenvalues, hop-plot
  - bow-tie, small diameter, ...
  - time evolution
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Pairs, influence, communities

CMU SCS

Evolution of diameter?

- Prior analysis, on power-law-like graphs, hints that

  diameter ~ O(log(N)) or

  diameter ~ O(log(log(N)))
- i.e.., slowly increasing with network size
- Q: What is happening, in reality?

CMU SCS

Evolution of diameter?

- Prior analysis, on power-law-like graphs, hints that

  diameter ~ O(log(N)) or

  diameter ~ O(log(log(N)))
- i.e.., slowly increasing with network size
- Q: What is happening, in reality?
- A: It **shrinks**(!!), towards a constant value

CMU SCS

Shrinking diameter

[Leskovec+05a]
- Citations among physics papers
- 11yrs; @ 2003:
  - 29,555 papers
  - 352,807 citations
- For each month *M*, create a graph of all citations up to month *M*

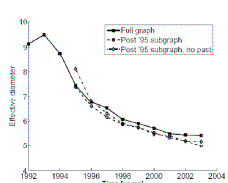

(a) arXiv citation graph

CMU SCS

Shrinking diameter

- Authors & publications
- 1992
  - 318 nodes
  - 272 edges
- 2002
  - 60,000 nodes
    - 20,000 authors
    - 38,000 papers
  - 133,000 edges

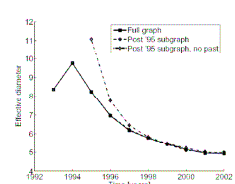

(b) Affiliation network

CMU SCS

Shrinking diameter

- Patents & citations
- 1975
  - 334,000 nodes
  - 676,000 edges
- 1999
  - 2.9 million nodes
  - 16.5 million edges
- Each year is a datapoint

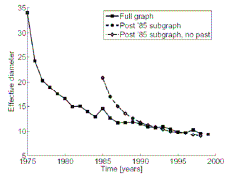

(c) Patents

CMU SCS

Shrinking diameter

- Autonomous systems
- 1997
  - 3,000 nodes
  - 10,000 edges
- 2000
  - 6,000 nodes
  - 26,000 edges
- One graph per day

diameter

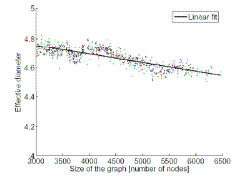

(d) AS

N

CMU SCS

Temporal evolution of graphs

- N(t) nodes; E(t) edges at time t
- suppose that

  N(t+1) = 2 * N(t)
- Q: what is your guess for

  E(t+1) =? 2 * E(t)

CMU SCS

Temporal evolution of graphs

- N(t) nodes; E(t) edges at time t
- suppose that

  N(t+1) = 2 * N(t)
- Q: what is your guess for

  E(t+1) =? X * E(t)
- A: over-doubled!

CMU SCS

Temporal evolution of graphs

- A: over-doubled - but obeying:

  $E(t) \sim N(t)^a$ for all t

  where $1<a<2$

CMU SCS

Densification Power Law

ArXiv: Physics papers and their citations

E(t)

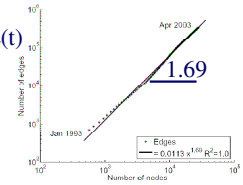

(a) arXiv

N(t)

CMU SCS

Densification Power Law

ArXiv: Physics papers and their citations

E(t)

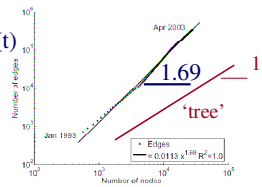

(a) arXiv

N(t)

CMU SCS

Densification Power Law

ArXiv: Physics papers and their citations

E(t)

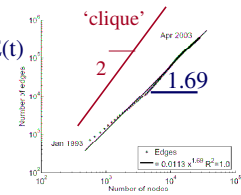

(a) arXiv

N(t)

CMU SCS

Densification Power Law

U.S. Patents, citing each other

E(t)

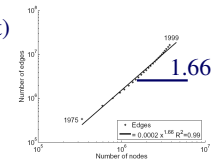

(b) Patents

N(t)

CMU SCS

Densification Power Law

Autonomous Systems

E(t)

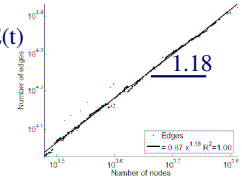

(c) Autonomous Systems

N(t)

CMU SCS

Densification Power Law

ArXiv: authors & papers

E(t)

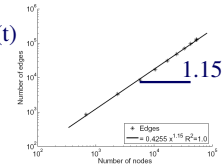

(d) Affiliation network

N(t)

CMU SCS

Summary of 'laws'

- Power laws for degree distributions
- …….. for eigenvalues, bi-partite cores
- Small & **shrinking** diameter ('6 degrees')
- 'Bow-tie' for web; 'jelly-fish' for internet
- ``Densification Power Law'', over time

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Pairs, influence, communities

CMU SCS

Generators

- How to generate random, realistic graphs?
  - Erdos-Renyi model: beautiful, but unrealistic
  - degree-based generators
  - process-based generators
  - recursive/self-similar generators

CMU SCS

Erdos-Renyi

- random graph – 100 nodes, avg degree = 2
- Fascinating properties (phase transition)
- But: unrealistic (Poisson degree distribution != power law)

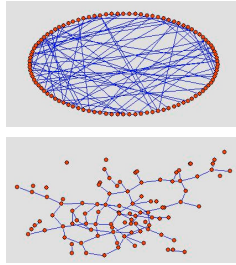

CMU SCS

skip

E-R model & Phase transition

- vary avg degree D
- watch Pc =
  Prob(there is a giant connected component)
- How do you expect it to be?

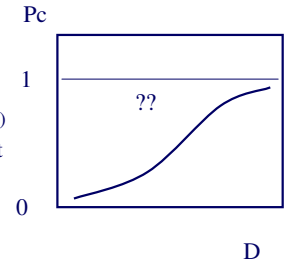

CMU SCS

skip

E-R model & Phase transition

- vary avg degree D
- watch Pc =
  Prob(there is a giant connected component)
- How do you expect it to be?

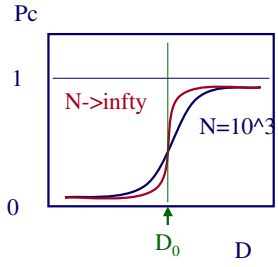

CMU SCS

Degree-based

- Figure out the degree distribution (eg., 'Zipf')
- Assign degrees to nodes
- Put edges, so that they match the original degree distribution

CMU SCS

Process-based

- Barabasi; Barabasi-Albert: Preferential attachment -> power-law tails!
  - 'rich get richer'
- [Kumar+]: preferential attachment + mimick
  - Create 'communities'

CMU SCS

Process-based (cont'd)

- [Fabrikant+, '02]: H.O.T.: connect to closest, high connectivity neighbor
- [Pennock+, '02]: Winner does NOT take all

CMU SCS

Recursive generators

- (RMAT [Chakrabarti+,'04])
- Kronecker product

CMU SCS

Wish list for a generator:

- Power-law-tail in- and out-degrees
- Power-law-tail scree plots
- **shrinking/constant** diameter
- Densification Power Law
- communities-within-communities

CMU SCS

Graph Patterns

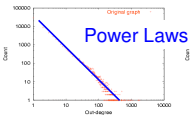

Count vs Indegree

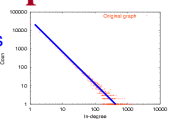

Count vs Outdegree

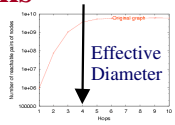

Hop-plot

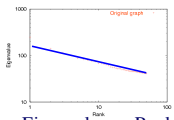

Eigenvalue vs Rank

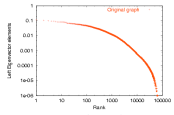

"Network values" vs Rank

CMU SCS

Wish list for a generator:

- Power-law-tail in- and out-degrees
- Power-law-tail scree plots
- **shrinking/constant** diameter
- Densification Power Law
- communities-within-communities

Q: how to achieve all of them?

CMU SCS

Wish list for a generator:

- Power-law-tail in- and out-degrees
- Power-law-tail scree plots
- **shrinking/constant** diameter
- Densification Power Law
- communities-within-communities

Q: how to achieve all of them?

A: Kronecker matrix product [Leskovec+05b]

CMU SCS

Kronecker product

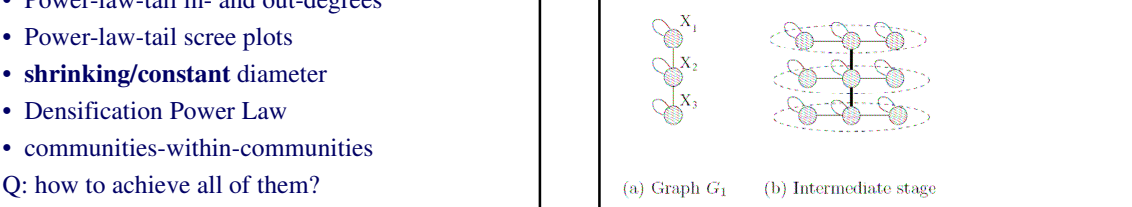

(a) Graph G_1 (b) Intermediate stage

CMU SCS

Kronecker product

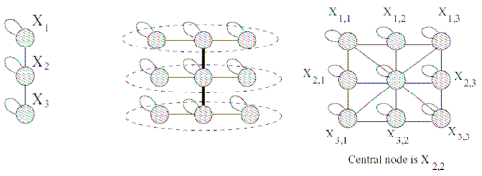

(a) Graph G_1 (b) Intermediate stage (c) Graph $G_2 = G_1 \otimes G_1$

CMU SCS

Kronecker product

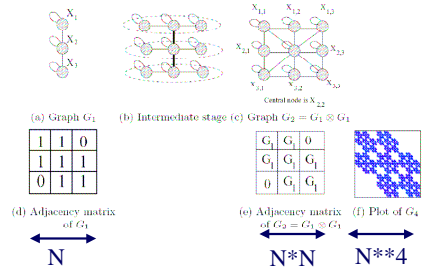

CMU SCS

Properties of Kronecker graphs:

- ✓ Power-law-tail in- and out-degrees
- ✓ Power-law-tail scree plots
- ✓ **constant** diameter
- ✓ perfect Densification Power Law
- ✓ communities-within-communities

CMU SCS

Properties of Kronecker graphs:

- ✓ Power-law-tail in- and out-degrees
- ✓ Power-law-tail scree plots
- ✓ **constant** diameter
- ✓ perfect Densification Power Law
- ✓ communities-within-communities

and we can prove all of the above

(**first and only** generator that does that)

CMU SCS

Kronecker - patents

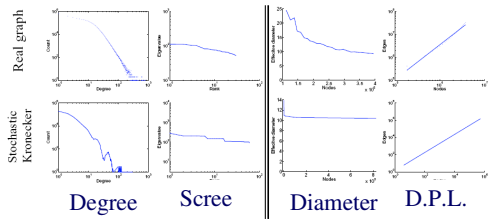

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
- Generators
- Tools

Part 2: PageRank, HITS and eigenvalues

Part 3: Pairs, influence, communities

CMU SCS

Outline

Part 1: Topology, 'laws' and generators
- 'Laws' and patterns
- Generators
- Tools: power laws and fractals
  - Why so many power laws?
  - Self-similarity, power laws, fractal dimension

CMU SCS

Power laws

- Q1: Why so many?
- A1:
- Q2: Are they only in graph-related settings?
- A2:

CMU SCS

Power laws

- Q1: Why so many?
- A1: self-similarity; 'rich-get-richer'
- Q2: Are they only in graph-related settings?
- A2: NO!

CMU SCS

A famous power law: Zipf's law

log(freq)

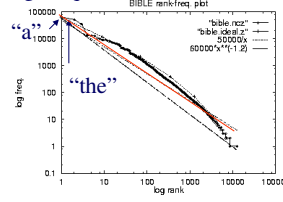

log(rank)

- Bible - **rank** vs **frequency** (log-log)

CMU SCS

Power laws, cont'ed

- web hit counts [Huberman]
- Click-stream data [Montgomery+01]

CMU SCS

Click-stream data

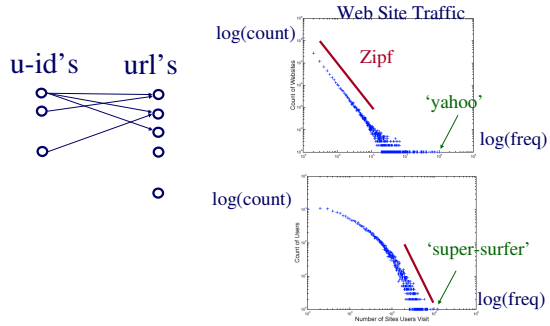

CMU SCS

Swedish sex-web

Albert Laszlo Barabasi
http://www.nd.edu/~networks/Publication%20Categories/04%20Talks/2005-norway-3hours.ppt

Nodes: people (Females; Males)
Links: sexual relationships

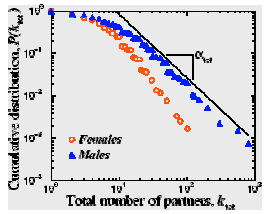

4781 Swedes; 18-74;
59% response rate.

Liljeros et al. *Nature* 2001

CMU SCS

Lotka's law

(Lotka's law of publication count); and citation counts: (*citeseer.nj.nec.com* 6/2001)

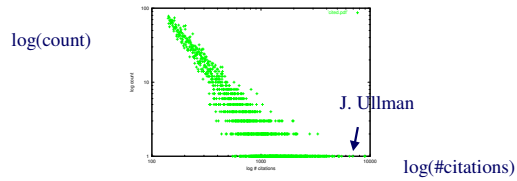

CMU SCS

Power laws

- Q1: Why so many?
- A1: self-similarity; 'rich-get-richer'
- Q2: Are they only in graph-related settings?
- A2: NO!

CMU SCS

Recall: Hop Plot

- Internet routers: how many neighbors within *h* hops? (= **correlation integral**!)

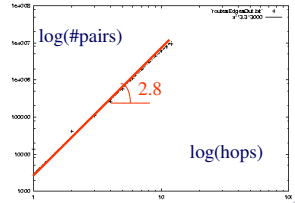

Reachability function: number of neighbors within r hops, vs r (log-log).

Mbone routers, 1995

CMU SCS

Observation

- Q: Intuition behind 'hop exponent'?
- A: 'intrinsic=fractal dimensionality' of the network

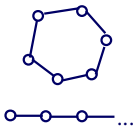

N(h) ~ h^1

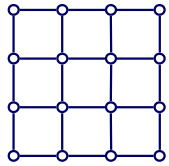

N(h) ~ h^2

CMU SCS

Non-integer dimensionality??

- Q3: How is it possible?
- A3:
- Q4: What does it mean?
- A4:

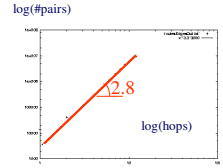

CMU SCS

Non-integer dimensionality

- Q3: How is it possible?
- A3: Through recursion!
- Q4: What does it mean?
- A4: There are groups (quasi-cliques / communities) in every scale

For example, Sierpinski triangle:

CMU SCS

Sierpinski triangle

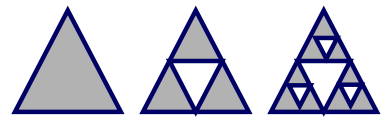

Equivalent graph?

CMU SCS

Sierpinski triangle

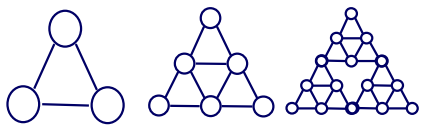

~ Equivalent graph

CMU SCS

Conclusions - Laws

'Laws' and patterns:

- Power laws for degrees, eigenvalues, 'communities'/cores
- Small diameter and shrinking diameter
- Bow-tie; jelly-fish
- densification power law

CMU SCS

Conclusions - Generators

- Preferential attachment (Barabasi)
- Variations
- Recursion – Kronecker graphs

CMU SCS

Conclusions - Tools

- Power laws –
  - rank/frequency plots ~ log-log NCDF
  - log-log PDF
- Self-similarity / recursion / fractals
- 'correlation integral' = hop-plot

CMU SCS

Conclusions - Tools (cont'd)

- Real settings/graphs: skewed distributions
  - 'mean' is meaningless
  - slope of power law, instead

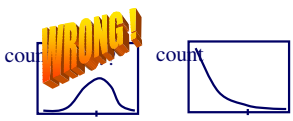

CMU SCS

Conclusions - Tools (cont'd)

- Real settings/graphs: skewed distributions
  - 'mean' is meaningless
  - slope of power law, instead

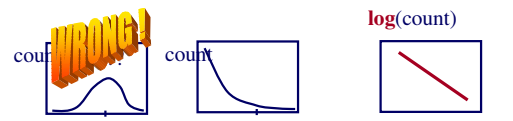

CMU SCS

Conclusions - Tools (cont'd)

- Recursion/self-similarity
  - May reveal non-obvious patterns (e.g., bow-ties within bow-ties within bow-ties) [Dill+, '01]

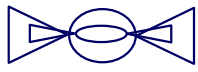

"To iterate is human, to recurse is divine"

CMU SCS

Resources

Generators:

- Kronecker ({christos,jure}@cs.cmu.edu)
- BRITE http://www.cs.bu.edu/brite/
- INET: http://topology.eecs.umich.edu/inet

CMU SCS

Other resources

Visualization - graph algo's:

- Graphviz: http://www.graphviz.org/
- pajek: http://vlado.fmf.uni-lj.si/pub/networks/pajek/

Kevin Bacon web site:
http://www.cs.virginia.edu/oracle/

CMU SCS

References

- [Aiello+, '00] William Aiello, Fan R. K. Chung, Linyuan Lu: *A random graph model for massive graphs*. STOC 2000: 171-180
- [Albert+] Reka Albert, Hawoong Jeong, and Albert-Laszlo Barabasi: *Diameter of the World Wide Web*, Nature 401 130-131 (1999)
- [Barabasi, '03] Albert-Laszlo Barabasi *Linked: How Everything Is Connected to Everything Else and What It Means* (Plume, 2003)

CMU SCS

References, cont'd

- [Barabasi+, '99] Albert-Laszlo Barabasi and Reka Albert. *Emergence of scaling in random networks*. Science, 286:509--512, 1999
- [Broder+, '00] Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. *Graph structure in the web*, WWW, 2000

CMU SCS

References, cont'd

- [Chakrabarti+, '04] *RMAT: A recursive graph generator*, D. Chakrabarti, Y. Zhan, C. Faloutsos, SIAM-DM 2004
- [Dill+, '01] Stephen Dill, Ravi Kumar, Kevin S. McCurley, Sridhar Rajagopalan, D. Sivakumar, Andrew Tomkins: *Self-similarity in the Web*. VLDB 2001: 69-78

CMU SCS

References, cont'd

- [Fabrikant+, '02] A. Fabrikant, E. Koutsoupias, and C.H. Papadimitriou. *Heuristically Optimized Trade-offs: A New Paradigm for Power Laws in the Internet*. ICALP, Malaga, Spain, July 2002
- [FFF, 99] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "*On power-law relationships of the Internet topology*," in SIGCOMM, 1999.

CMU SCS

References, cont'd

- [Jovanovic+, '01] M. Jovanovic, F.S. Annexstein, and K.A. Berman. *Modeling Peer-to-Peer Network Topologies through "Small-World" Models and Power Laws*. In TELFOR, Belgrade, Yugoslavia, November, 2001
- [Kumar+ '99] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, Andrew Tomkins: *Extracting Large-Scale Knowledge Bases from the Web*. VLDB 1999: 639-650

CMU SCS

References, cont'd

- [Leland+, '94] W. E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, *On the Self-Similar Nature of Ethernet Traffic*, IEEE Transactions on Networking, 2, 1, pp 1-15, Feb. 1994.
- [Leskovec+05a] Jure Leskovec, Jon Kleinberg, Christos Faloutsos, *Graphs over Time: Densification Laws, Shrinking Diameters and Possible Explanations* (KDD 2005)

CMU SCS

References, cont'd

- [Leskovec+05b] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos *Realistic, Mathematically Tractable Graph Generation and Evolution, Using Kronecker Multiplication* (ECML/PKDD 2005), Porto, Portugal, 2005.
- [Mihail+, '02] Milena Mihail, Christos H. Papadimitriou: *On the Eigenvalue Power Law*. RANDOM 2002: 254-262

CMU SCS

References, cont'd

- [Milgram '67] Stanley Milgram: *The Small World Problem*, Psychology Today 1(1), 60-67 (1967)
- [Montgomery+, '01] Alan L. Montgomery, Christos Faloutsos: *Identifying Web Browsing Trends and Patterns*. IEEE Computer 34(7): 94-95 (2001)

CMU SCS

References, cont'd

- [Palmer+, '01] Chris Palmer, Georgos Siganos, Michalis Faloutsos, Christos Faloutsos and Phil Gibbons *The connectivity and fault-tolerance of the Internet topology* (NRDM 2001), Santa Barbara, CA, May 25, 2001
- [Pennock+, '02] David M. Pennock, Gary William Flake, Steve Lawrence, Eric J. Glover, C. Lee Giles: *Winners don't take all: Characterizing the competition for links on the web* Proc. Natl. Acad. Sci. USA 99(8): 5207-5211 (2002)

CMU SCS

References, cont'd

- [Schroeder, '91] Manfred Schroeder *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise* W H Freeman & Co., 1991

CMU SCS

References, cont'd

- [Siganos+, '03] G. Siganos, M. Faloutsos, P. Faloutsos, C. Faloutsos *Power-Laws and the AS-level Internet Topology*, Transactions on Networking, August 2003.
- [Watts+ Strogatz, '98] D. J. Watts and S. H. Strogatz *Collective dynamics of 'small-world' networks*, Nature, 393:440-442 (1998)
- [Watts, '03] Duncan J. Watts *Six Degrees: The Science of a Connected Age* W.W. Norton & Company; (February 2003)