


ECS 253 / MAE 253, Lecture 5

April 14, 2025




“Internet measurement and Optimization
approaches to network growth”

Aside on Adjacency Matrix and random walks


Consider undirected edges

$$M_{ij} = \begin{cases} 1 & \text{if edge exists between } i \text{ and } j \\ 0 & \text{otherwise.} \end{cases}$$



$$\begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \end{pmatrix} = M$$

Random walk: State Transition Matrix (Column-normalize the adjacency matrix)



$$M = \begin{pmatrix} 1/4 & 1/3 & 1/2 & 1/4 & 0 \\ 1/4 & 1/3 & 0 & 1/4 & 0 \\ 1/4 & 0 & 1/2 & 0 & 0 \\ 1/4 & 1/3 & 0 & 1/4 & 1/2 \\ 0 & 0 & 0 & 1/4 & 1/2 \end{pmatrix}$$

M will have a basis set of eigenvectors $\{\vec{u}_i\}$ and corresponding eigenvalues λ_i .

Perron-Frobenius Theorem


- Applies to **irreducible, positive, stochastic** matrices.
- “Irreducible” means cannot be block-diagonalized into disjoint pieces. (i.e., network is connected — only one component).
- “Positive” means each entry $M_{ij} > 0$.
- “Stochastic” means column normalized (or row normalized).

Perron-Frobenius Theorem

Leading eigenvalue

- One leading eigenvalue with $\lambda_1 = 1$.
 - The corresponding eigenvector, v_1 , has strictly positive entries and the sum over all the entries, $\sum_i v_1[i] = 1$.
 - This is the stationary distribution of the random walk dynamics.
-
- For **non-negative** matrices ($M_{ij} \geq 0$), similar results, but can't guarantee eigenvectors are positive (in practice, normally still works ... we will come back to this later in the quarter.)

What about networks with directed edges? (e.g., HW1)



- Does it matter whether M_{ij} means an edge from node i to j , or if it means an edge from node j to i ?
- In general, it does not matter. But, sometimes it does matter!
- For the graph pictured, if M_{ij} means an edge from node i to j , then the 5th column will be a vector of all zero's. And there is no way to make it a column-normalized stochastic matrix. (But you can make it row-normalized.)


Last time: “Robust yet fragile”

- “Error and attack tolerance of complex networks”

Random networks with power law degree distribution show:

- Fragility to degree-targeted removal
- Robustness to random node removal

(This is in the context of keeping the full network connected.)



Example histogram of a PA run with N=500 nodes.

Albert, Jeong and Barabasi, *Nature*, 406 (27) 2000




“The Achilles Heel of the Internet”

- “How robust is the Internet?” Yuhai Tu, *Nature* (News and Views) **406** (27) 2000.
- “Scientists spot Achilles heel of the Internet”, CNN, July 26, 2000.

Random vs engineered vs evolved (e.g. biological) systems


Is the Internet really a random power law graph?

- **REDUNDANCY!!!** a key principle in engineering (and evolution?).
- The 'robust yet fragile' nature of the Internet
Doyle, Alderson, Li, Low, Roughan, Shalunov, Tanaka, Willinger, PNAS **102** (4) 2005.




- Degree distribution is not the whole story.

Power law random graph: Robust to random failure, vulnerable to targeting attack



Why did Albert, Jeong and Barabasi find that their sample of the internet topology was vulnerable to degree targeted attack?

What is the Internet?




Internet

- Web of interconnected networks
 - Grows with no central authority
 - Autonomous Systems optimize local communication efficiency
 - The building blocks are engineered and studied in depth
 - Global entity has not been characterized

- Most real world complex-networks have non-trivial properties.


- Global properties can not be inferred from local ones
 - Engineered with large technical diversity
 - Range from local campuses to transcontinental backbone providers



Power Laws in the Internet?

Definition of “node” depends on level of representation

Internet connectivity structures are different at each layer



The Internet hourglass

Applications

Web FTP Mail News Video Audio ping kazaa

Transport protocols

TCP

SCTP

UDP

ICMP

IP

Ethernet 802.11 Power lines ATM Optical Satellit Bluetooth

Link technologies

AS
(BGP routes) →

TraceRoute →

Router level →


(picture from David Alderson)

TCP / IP

- The TCP protocol: a collection of rules for formatting, ordering, and error-checking data sent across a network.
- In 1974, Vincent Cerf and Robert Kahn developed the Transmission Control Protocol (TCP) which was further split into the Internet Protocol (IP) and TCP in 1978.
- In 1982, DoD adopted TCP/IP as the standard protocol in the Internet.
- IP address: a unique 4-byte number to identify each machine

Internet Infrastructure

The IP address



Common top domain names in the US: **.com**, **.mil**, **.edu**, **.org**

Outside of the US, the top-level domain identifies the country:
uk (England), fr (France), cn (China), ...

Internet Infrastructure The Transmission Control Protocol

See also: http://en.wikipedia.org/wiki/Transmission_Control_Protocol


Structure of a TCP/IP packet

Bit offset	Bits 0 - 7		8–15	16–23	24–31
0	Source address <i>Computer sending the packet</i>				
32					
64					
96					
128	Destination address <i>Destination computer</i>				
160					
192					
224					
256	TCP length				<i>Length of the packet</i>
288	Zeros				Next header
320	Source port			Destination port	
352	Sequence number				
384	Acknowledgement number				
416	Data offset	Reserved	Flags	Window	
448	Checksum			Urgent pointer	
480	Options (optional)				
480/512+	<i>Checksum for integrity</i> Data				

Internet Infrastructure

The Transmission Control Protocol

How does the sender know it needs to retransmit:




- TCP a *decentralized* protocol with non-linear ramp-up and random restart.

Autonomous system

A collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators that presents a common, clearly defined routing policy to the Internet

AS-Level Topology

- Nodes = (sets of) entire networks (Autonomous Systems or ASes)
- Links = peering relationships between ASes
- Really a map of economic or business relationships, not of physical connectivity



Internet Measurements

- The Internet is man-made, so why do we need to measure it?
 - Because we still don't really understand it
 - Sometimes things go wrong
 - Malicious users
 - Measurement for network operations
 - Detecting and diagnosing problems
 - What-if analysis of future changes
 - Measurement for scientific discovery
 - Creating accurate models that represent reality
 - Identifying new features and phenomena

How to measure the structure of the Internet?


- **Traceroute** (IP address level) see: unix traceroute command
- **BGP tables** (AS level)
- **“Whois” data** (AS level)

Repositories / public resources (mostly AS level)


- University of Oregon Route Views Project
<http://www.routeviews.org/>
- CAIDA (Cooperative Association for Internet Data Analysis, UCSD)
<http://www.caida.org/home/>

Internet Topology Measurements Probing

■ Direct probing




■ Indirect probing



http://www.caida.org/publications/animations/active_monitoring/traceroute.mpg

Internet Topology Measurement: Background




Problems: Traceroute

- Lakhina, Byers, Crovella, Xie, *INFOCOM*, 2003.
- Achlioptas, Clauset, Kempe, Moore, *STOC*, 2005.
- Achlioptas, Clauset, Kempe, Moore, *J. of ACM*, 56 (4), 2009.

- Build approximately single-source, all-destinations, shortest-path trees. (Union of traceroute samples.)


- Faloutsos³ *SIGCOMM*, 1999.
- Albert, Jeong, Barabasi, *Nature*, 2000.



- Sampling bias

- Nodes close to root sampled more accurately
- High degree nodes sampled more accurately than low degree. (Follow an edge at random, k times as likely to lead to node of degree k than degree 1. See next slide.)

Aside: Edge following probability, q_k



k edges reach node of degree k :

- Let q_k denote the probability of following an edge to a node of degree k .
- q_k is proportional to $k p_k$.
- Precisely, $q_k = \frac{k p_k}{\sum_k k p_k}$


Traceroute sampling bias

- Lakhina, et al *INFOCOM*, 2003: Show empirically that Erdős-Rényi random graphs (Poisson dist) appear to have power law degree distribution.
- Petermann and De Los Rios [2004] and Clauset and Moore [2005]: Even if a power law, the exponent γ is underestimated.
- Achlioptas et al 2005 and 2009: Rigorous proof of bias and consequences.
 - Poisson degree dist
 - d -regular random graphs (all nodes have degree d).
- Recommendation: Traceroute sampling over the union of a very large number of sources more accurate.

AS level topology measurement: Challenges

AS level connections inferred from BGP routing tables .

- AS level does not reflect physical connectivity (geographically distant routers can appear as one AS).
- Hidden subgraphs




- www.sawvis.net
- managed IP and hosting company
- founded 1995
- offering “private IP with ATM at core”

**This “node” is an entire network!
(not just a router)**

<http://www.caida.org/tools/measurement/skitter/>

The Internet?

- Michalis Faloutsos, Petros Faloutsos, Christos Faloutsos, “On power-law relationships of the Internet topology”, ACM SIGCOMM Computer Communication Review Volume 29 , Issue 4 Oct. 1999.
- Only one order of magnitude (even exponential can look power law in a short regime).
- $\gamma \approx 2.1$
- (over 7600 cites)
- Only 6000 nodes
- Consequences for topology generators



(a) Int-11-97


Can there be real Power Laws in data?

- in the WWW sure.
- in a social network ... possible.
- in earthquake magnitude ... yes, but to some cutoff.
- in the Internet?

Why power laws cannot continue: Finite size effects, resource limitations, physical geometric (Internet) vs virtual geometry-free (WWW)....

The “Who-is-Who” network in Budapest


(Analysis by Balázs Szendrői and Gábor Csányi)



Bayesian curve fitting $\rightarrow p(k) = ck^{-\gamma}e^{-\alpha k}$

Another common distribution: power-law with an exponential cutoff

■ $p(x) \sim x^{-a} e^{-x/\kappa}$



but could also be a lognormal or double exponential...

“Power law” → power law with exponential tail

Ubiquitous empirical measurements:

System with: $p(x) \sim x^{-B} \exp(-x/C)$	B	C
Full protein-interaction map of <i>Drosophila</i>	1.20	0.038
High-confidence protein-interaction map of <i>Drosophila</i>	1.26	0.27
Gene-flow/hybridization network of plants as function of spatial distance	0.75	10^5 m
Earthquake magnitude	1.35 - 1.7	$\sim 10^{21}$ Nm
Avalanche size of ferromagnetic materials	1.2 - 1.4	$L^{1.4}$
ArXiv co-author network	1.3	53
MEDLINE co-author network	2.1	~ 5800
PNAS paper citation network	0.49	4.21

(Saturation and PA often put in apriori to explain)

Known Mechanisms for Power Laws

- Phase transitions (singularities)
- Random multiplicative processes (fragmentation)
- Combination of exponentials (e.g. word frequencies)
- **Preferential attachment / Proportional attachment**
(Polya 1923, Yule 1925, Zipf 1949, Simon 1955, Price 1976, Barabási and Albert 1999)

Attractiveness is proportional to size:

$$\frac{ds}{dt} \propto s$$

- Add in **saturation** [Amaral 2000, Börner 2004], get PA with exponential decay .

An alternate view, Mandelbrot, 1953: optimization

(Information theory of the statistical structure of language)


- **Goal:** Optimize information conveyed for unit transmission cost (what probability distribution over words gives most info?)
- Consider an alphabet of d characters, with n distinct words
- Order all possible words by length (A,B,C,....AA,BB,CC.....)
- “Cost” of j -th word, $C_j \sim \log_d j$
- Ave information per word: $H = - \sum p_j \log p_j$
- Ave cost per word: $C = \sum p_j C_j$
- Minimize: $\frac{d}{dp_j} \left(\frac{C}{H} \right) \implies p_j \sim j^{-\alpha}$

Optimization versus Preferential Attachment origin of power laws

Mandelbrot and Simon's heated public exchange

- A series of six letters between 1959-61 in *Information and Control*.
- Optimization on hold for many years, but recently resurfaced:
- Calson and Doyle, HOT, 1999
- Fabrikant, Koutsoupias, and Papadimitriou, 2002
- Solé, 2002

Simon and Mandelbrot's exchange




From Barabasi *Network Science*

FKP (Fabrikant, Koutsoupias, and Papadimitriou, 2002)


An optimization model of internet growth

- Nodes arriving sequentially at random in a unit square.
- Upon arrival, node i connects to an already existing node j that minimizes “cost”:
$$\alpha d_{ij} + h_j$$
- d_{ij} is Euclidean distance between i and j .
 h_j is the hop distance from j to the root node.
- i.e., connect to the closest node that has good network performance



FKP cont

- αd_{ij} introduces a *scale*. The first node to arrive an uninhabited area collects all the subsequent arrivals.
- Eventually get hubs-and-leaf structure, but the hubs grow in degree super-linearly.



Tempered Preferential Attachment

[D'Souza, Borgs, Chayes, Berger, Kleinberg, *PNAS* 2007.]

[Berger, Borgs, Chayes, D'Souza, Kleinberg, *ICALP* 2004.]

[Berger, Borgs, Chayes, D'Souza, Kleinberg, *CPC*, 2005.]

- **Optimization**

Like FKP, start with linear tradeoffs, but consider a scale-free metric. (Plus will result in local model.) Gives rise to:

→ **PA**


→ **Saturation**

→ **Viability**

(Not all children have equal fertility, not all spin-offs equally fit, etc).

Competition-Induced Preferential Attachment

Consider points arriving sequentially, uniformly at random along the unit line:



Each incoming node, t , attaches to an existing node j (where $j < t$), which minimizes the function:

$$F_{tj} = \min_j [\alpha_{tj} d_{tj} + h_j]$$

Where $\alpha_{tj} = \alpha \rho_{tj} = \alpha n_{tj} / d_{tj}$.

The “cost” becomes: $F_{tj} = \min_j [\alpha n_{tj} + h_j]$





$$F_{tj} = \min_j [\alpha n_{tj} + h_j]$$

- $\alpha_{tj} = \alpha \rho_{tj}$ local density, e.g. real estate in Manhattan.
- Reduces to n_{tj} — number of points in the interval between t and j
- “Transit domains” — captures realistic aspects of Internet costs (i.e. AS/ISP-transit requires BGP and peering).
- Like FKP, tradeoff initial connection cost versus usage cost.
- Note cases $\alpha = 0$ and $\alpha > 1$.

The process on the line (for $1/3 < \alpha < 1/2$)

“Border Toll Optimization Problem” (BTOP)


$$F_{tj} = \min_j [\alpha n_{tj} + h_j]$$





(A **local** model – connect either to closest node, or its parent.)

Mapping onto a tree



(equal in distribution to the line)





$t=1$




$t=2$



$t=3$



$t=4$




From line to tree

Integrating out the dependence on interval length from the conditional probability:

$$\begin{aligned} \Pr [x_{t+1} \in I_k | \pi(t)] &= \int \Pr [x_{t+1} \in I_k | \pi(t), \vec{s}(t)] dP(\vec{s}(t)) \\ &= \int s_k(t) dP(\vec{s}(t)) = \frac{1}{t+1}, \end{aligned}$$

i.e., The probability to land in the k -th interval is uniform over all intervals.

Preferential attachment with a cutoff



Let $d_j(t)$ equal the degree of **fertile** node j at time t .

The number of **intervals** contributing to j 's fertility is $\max(d_j(t), A)$.

Probability node $(t + 1)$ attaches to node j is:

$$Pr(t + 1 \rightarrow j) = \max(d_j(t), A) / (t + 1).$$

The process on degree sequence

Let $N_0(t) \equiv$ number of infertile vertices.

Let $N_k(t) \equiv$ number of fertile vertices of degree k
(for $1 \leq k < A$).

Let $N_A(t) \equiv$ number of fertile vertices of degree $k \geq A$
(i.e. $N_A(t) = \sum_{k=A}^{\infty} N_k(t)$ “the tail”)

In terms of $p_k(t)$:

$$p_1(t+1)(t+1) - p_1(t)(t) = Ap_A(t) - p_1(t)$$

$$p_k(t+1)(t+1) - p_k(t)(t) = (k-1)p_{k-1}(t) - kp_k(t), \quad 1 < k < A$$

$$p_A(t+1)(t+1) - p_A(t)(t) = (A-1)p_{A-1}(t).$$

Proposition 1 (Convergence of expectations to stationary distribution): $p_k(t) \rightarrow p_k$.

$$p_1 = Ap_A - p_1$$

$$p_k = (k-1)p_{k-1} - kp_k, \quad 1 < k < A$$

$$p_A = (A-1)p_{A-1}.$$

Proposition (2): (Concentration) (i.e., How big are the fluctuations about $n_k(t)$?) Requires second-moment method.

Recursion relation

$$p_k = (k - 1)p_{k-1}(t) - kp_k(t), \quad 1 < k < A.$$

Implies

$$p_k = \prod_{i=2}^k \left(\frac{i-1}{i+1} \right) p_1, \quad 1 < k < A.$$

Power law for $1 < k < A$

$$\frac{p_k}{p_1} = \prod_{i=2}^k \left(\frac{i-1}{i+1} \right) = \frac{2}{k(k+1)}$$

$$\sim c k^{-2}$$

Exponential decay for $k > A$

Recursion relation: $p_k = A(p_{k-1} - p_k), \quad k \geq A.$

Implies

$$p_k = \left(\frac{A}{A+1}\right)^{k-A} p_A, \quad k \geq A.$$

$$p_k = \left(1 - \frac{1}{A+1}\right)^{k-A} p_A = \left[\left(1 - \frac{1}{A+1}\right)^{A+1}\right]^{(k-A)/(A+1)} p_A$$

$$\sim \exp[-(k-A)/(A+1)] p_A.$$


April 2007

April 10, 2007 | vol. 104 | no. 15 | pp. 6091–6492

PNAS

Proceedings of the National Academy of Sciences of the United States of America • www.pnas.org

Preferential attachment from optimization



Fossil teeth of earliest modern humans

Engineered halogen bonds direct conformation

Ancient origin of lentiviruses

Mutated *Clock* gene induces mania

Linear optimization and transportation networks

(Applying the “FKP” ideas)

We will study these in-depth later


- M. T. Gastner, M.E.J. Newman, “The spatial structure of networks”, cond-mat/0407680, 2004.
- M. T. Gastner, M.E.J. Newman, “Shape and efficiency in spatial distribution networks”, *Journal of Statistical Mechanics*, 2006.
- M. T. Gastner, M.E.J. Newman, “Optimal design of spatial distribution networks”, *Physical Review E*, 74, 016117, 2006.

Optimal networks of optimally located facilities


The optimal network design problem then consists of two parts.

First, we distribute p facilities on the map by solving the p -median problem.

Then we find the network minimizing the total cost C .



Different routing strategies



Summary

- Internet measurement :
 - Traceroute sampling (router level)
 - Peering agreements/ routing tables (AS level)
- Optimization approaches to network growth :
 - FKP (leads to hubs and leaves;
bi-modal not power law degree distribution in $N \rightarrow \infty$ limit)
 - TPA (Pref Attachment with saturation, fertility / viability)
 - Gastner/Newman: FKP approach to transport networks.