Core Models of Complex Networks

Principles of Complex Systems
CSYS/MATH 300, Spring, 2013 | #SpringPoCS2013

Prof. Peter Dodds
@peterdodds

Department of Mathematics & Statistics | Center for Complex Systems | Vermont Advanced Computing Center | University of Vermont

These slides brought to you by:

Outline

Models

Some important models:

1. Generalized random networks;
2. Small-world networks;
3. Generalized affiliation networks;
4. Scale-free networks;
5. Statistical generative models (p^*).

Models

Generalized random networks:

- Arbitrary degree distribution P_k.
- Create (unconnected) nodes with degrees sampled from P_k.
- Wire nodes together randomly.
- Create ensemble to test deviations from randomness.

Building random networks: Stubs

Phase 1:

- Idea: start with a soup of unconnected nodes with stubs (half-edges):

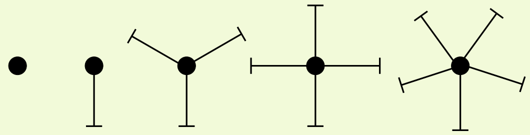

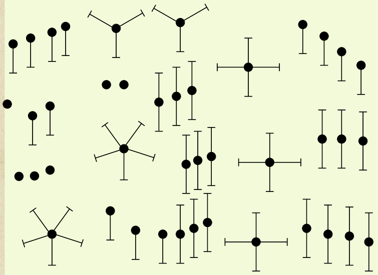

- Randomly select stubs (not nodes!) and connect them.
- Must have an even number of stubs.
- Initially allow self- and repeat connections.

Building random networks: First rewiring

Phase 2:

- Now find any (A) self-loops and (B) repeat edges and randomly rewire them.

(A)

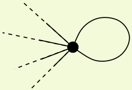

(B)

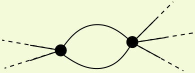

- Being careful: we can't change the degree of any node, so we can't simply move links around.
- Simplest solution: randomly rewire two edges at a time.

General random rewiring algorithm

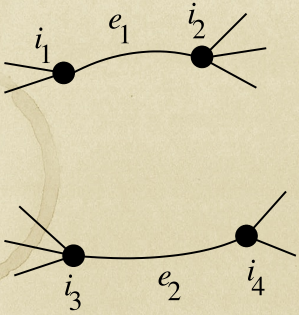

- Randomly choose two edges. (Or choose problem edge and a random edge)
- Check to make sure edges are disjoint.

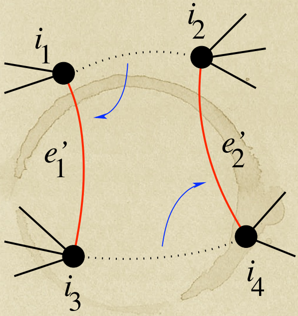

- Rewire one end of each edge.
- Node degrees do not change.
- Works if e_1 is a self-loop or repeated edge.
- Same as finding on/off/on/off 4-cycles. and rotating them.

Sampling random networks

Phase 2:

- Use rewiring algorithm to remove all self and repeat loops.

Phase 3:

- Randomize network wiring by applying rewiring algorithm liberally.
- Rule of thumb: # Rewirings $\simeq 10 \times$ # edges [10].

People thinking about people:

How are social networks structured?

- How do we define and measure connections?
- Methods/issues of self-report and remote sensing.

What about the dynamics of social networks?

- How do social networks/movements begin & evolve?
- How does collective problem solving work?
- How does information move through social networks?
- Which rules give the best ‘game of society?’

Sociotechnical phenomena and algorithms:

- What can people and computers do together? (google)
- Use Play + Crunch to solve problems. Which problems?

Social Search

A small slice of the pie:

- Q. Can people pass messages between distant individuals using only their existing social connections?
- A. Apparently yes...

Milgram's social search experiment (1960s)

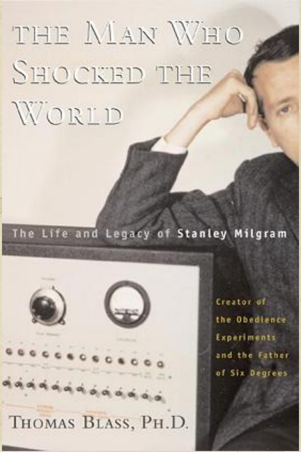

http://www.stanleymilgram.com

- Target person = Boston stockbroker.
- 296 senders from Boston and Omaha.
- 20% of senders reached target.
- chain length $\simeq 6.5$.

Popular terms:

- The Small World Phenomenon;
- "Six Degrees of Separation."

The problem

Lengths of successful chains:

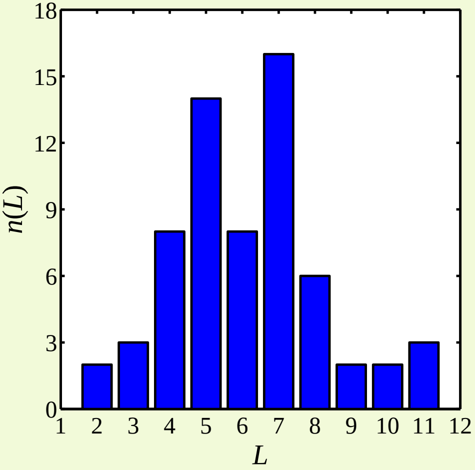

From Travers and Milgram (1969) in Sociometry: [13]
"An Experimental Study of the Small World Problem."

The problem

Two features characterize a social 'Small World':

1. Short paths exist, (= Geometric piece)

   and

2. People are good at finding them. (= Algorithmic piece)

Social Search

Milgram's small world experiment with email:

"An Experimental study of Search in Global Social Networks"
P. S. Dodds, R. Muhamad, and D. J. Watts,
Science, Vol. 301, pp. 827–829, 2003. [6]

Social search—the Columbia experiment

- 60,000+ participants in 166 countries
- 18 targets in 13 countries including
  - a professor at an Ivy League university,
  - an archival inspector in Estonia,
  - a technology consultant in India,
  - a policeman in Australia,
    and
  - a veterinarian in the Norwegian army.
- 24,000+ chains

We were lucky and contagious (more later):

"Using E-Mail to Count Connections" (⊞), Sarah Milstein, New York Times, Circuits Section (December, 2001)

The UNIVERSITY of VERMONT

All targets:

Table S1

Target	*City*	*Country*	*Occupation*	*Gender*	*N*	N_c (%)	r (r_0)	<L>
1	Novosibirsk	Russia	PhD student	F	8234	20(0.24)	64 (76)	4.05
2	New York	USA	Writer	F	6044	31 (0.51)	65 (73)	3.61
3	Bandung	Indonesia	Unemployed	M	8151	0	66 (76)	n/a
4	New York	USA	Journalist	F	5690	44 (0.77)	60 (72)	3.9
5	Ithaca	USA	Professor	M	5855	168 (2.87)	54 (71)	3.84
6	Melbourne	Australia	Travel Consultant	F	5597	20 (0.36)	60 (71)	5.2
7	Bardufoss	Norway	Army veterinarian	M	4343	16 (0.37)	63 (76)	4.25
8	Perth	Australia	Police Officer	M	4485	4 (0.09)	64 (75)	4.5
9	Omaha	USA	Life Insurance Agent	F	4562	2 (0.04)	66 (79)	4.5
10	Welwyn Garden City	UK	Retired	M	6593	1 (0.02)	68 (74)	4
11	Paris	France	Librarian	F	4198	3 (0.07)	65 (75)	5
12	Tallinn	Estonia	Archival Inspector	M	4530	8 (0.18)	63(79)	4
13	Munich	Germany	Journalist	M	4350	32 (0.74)	62 (74)	4.66
14	Split	Croatia	Student	M	6629	0	63 (77)	n/a
15	Gurgaon	India	Technology Consultant	M	4510	12 (0.27)	67 (78)	3.67
16	Managua	Nicaragua	Computer analyst	M	6547	2 (0.03)	68 (78)	5.5
17	Katikati	New Zealand	Potter	M	4091	12 (0.3)	62 (74)	4.33
18	Elderton	USA	Lutheran Pastor	M	4438	9 (0.21)	68 (76)	4.33
Totals					98,847	384 (0.4)	63 (75)	4.05

Social search—the Columbia experiment

- Milgram's participation rate was roughly 75%
- Email version: Approximately 37% participation rate.
- Probability of a chain of length 10 getting through:

$$.37^{10} \simeq 5 \times 10^{-5}$$

- $\Rightarrow$ 384 completed chains (1.6% of all chains).

Social search—the Columbia experiment

- Motivation/Incentives/Perception matter.
- If target *seems* reachable
  $\Rightarrow$ participation more likely.
- Small changes in attrition rates
  $\Rightarrow$ large changes in completion rates
- e.g., $\searrow$ 15% in attrition rate
  $\Rightarrow$ $\nearrow$ 800% in completion rate

Social search—the Columbia experiment

Comparing successful to unsuccessful chains:

- Successful chains used relatively weaker ties:

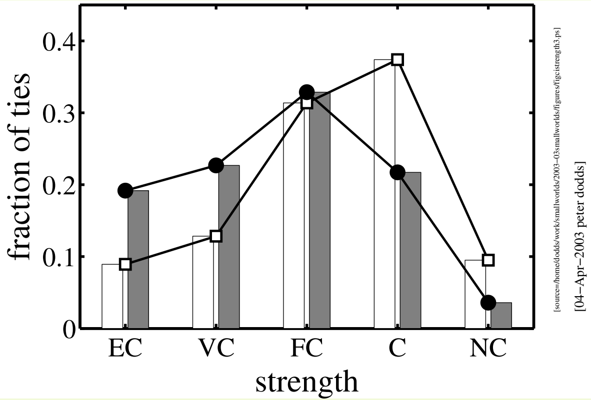

Social search—the Columbia experiment

Successful chains disproportionately used:

- Weak ties, Granovetter [7]
- Professional ties (34% vs. 13%)
- Ties originating at work/college
- Target's work (65% vs. 40%)

. . . and disproportionately avoided

- hubs (8% vs. 1%) (+ no evidence of funnels)
- family/friendship ties (60% vs. 83%)

Geography → Work

Social search—the Columbia experiment

Senders of successful messages showed little absolute dependency on

- age, gender
- country of residence
- income
- religion
- relationship to recipient

Range of completion rates for subpopulations:

30% to 40%

Social search—the Columbia experiment

Nevertheless, some weak discrepencies do exist...

Contrived hypothetical above average connector:

Norwegian, secular male, aged 30-39, earning over $100K, with graduate level education working in mass media or science, who uses relatively weak ties to people they met in college or at work.

Contrived hypothetical below average connector:

Italian, Islamic or Christian female earning less than $2K, with elementary school education and retired, who uses strong ties to family members.

Social search—the Columbia experiment

Mildly bad for continuing chain:

choosing recipients because “they have lots of friends” or because they will “likely continue the chain.”

Why:

- Specificity important
- Successful links used relevant information. (e.g. connecting to someone who shares same profession as target.)

Social search—the Columbia experiment

Basic results:

- $\langle L \rangle = 4.05$ for all completed chains
- L_* = Estimated 'true' median chain length (zero attrition)
- Intra-country chains: $L_* = 5$
- Inter-country chains: $L_* = 7$
- All chains: $L_* = 7$
- Milgram: $L_* \simeq 9$

Usefulness:

Harnessing social search:

- Can distributed social search be used for something big/good?
- What about something evil? (Good idea to check.)
- What about socio-inspired algorithms for information search? (More later.)
- For real social search, we have an incentives problem.
- Which kind of influence mechanisms/algorithms would help propagate search?
- Fun, money, prestige, ... ?
- Must be 'non-gameable.'

Red balloons:

A Grand Challenge:

- 1969: The Internet is born (⊞) (the ARPANET (⊞)—four nodes!).
- Originally funded by DARPA who created a grand Network Challenge (⊞) for the 40th anniversary.
- Saturday December 5, 2009: DARPA puts 10 red weather balloons up during the day.
- Each 8 foot diameter balloon is anchored to the ground somewhere in the United States.
- Challenge: Find the latitude and longitude of each balloon.
- Prize: $40,000.

*DARPA = Defense Advanced Research Projects Agency (⊞).

Where the balloons were:

Finding red balloons:

The winning team and strategy:

- MIT's Media Lab (⊞) won in less than 9 hours. [11]
- Pickard et al. "Time-Critical Social Mobilization," [11] Science Magazine, 2011.
- People were virally recruited online to help out.
- Idea: Want people to both (1) find the balloons, and (2) involve more people.
- Recursive incentive structure with exponentially decaying payout:
  - $2000 for correctly reporting the coordinates of a balloon.
  - $1000 for recruiting a person who finds a balloon.
  - $500 for recruiting a person who recruits the balloon finder, . . .
  - (Not a Ponzi scheme.)
- True victory: Colbert interviews Riley Crane (⊞)

Finding balloons:

Clever scheme:

- Max payout = \$4000 per balloon.
- Individuals have clear incentives to both
  1. involve/source more people (spread), and
  2. find balloons (goal action).
- Gameable?
- Limit to how much money a set of bad actors can extract.

Extra notes:

- MIT's brand helped greatly.
- MIT group first heard about the competition a few days before. Ouch.
- A number of other teams did well (⊞).
- Worthwhile looking at these competing strategies. [11]

The social world appears to be small... why?

Theory: how do we understand the small world property?

- Connected random networks have short average path lengths:

$$\langle d_{AB} \rangle \sim \log(N)$$

  N = population size,

  d_{AB} = distance between nodes A and B.

- But: social networks aren't random...

Simple socialness in a network:

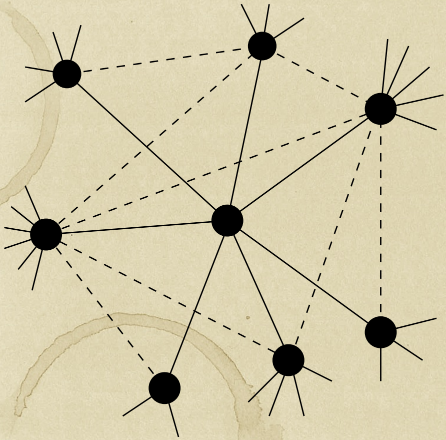

Need "clustering" (your friends are likely to know each other):

Non-randomness gives clustering:

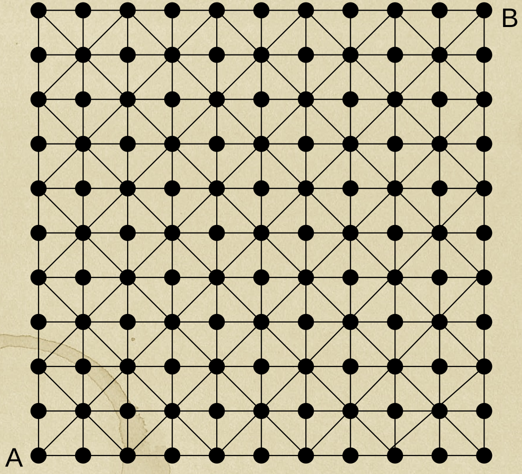

$d_{AB} = 10 \rightarrow$ too many long paths.

Randomness + regularity

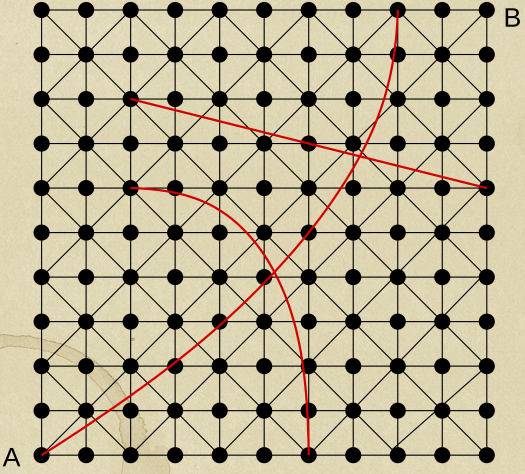

Now have $d_{AB} = 3$

$\langle d \rangle$ decreases overall

Small-world networks

Introduced by Watts and Strogatz (Nature, 1998) [15]
“Collective dynamics of ‘small-world’ networks.”

Small-world networks were found everywhere:

- neural network of C. elegans,
- semantic networks of languages,
- actor collaboration graph,
- food webs,
- social networks of comic book characters,...

Very weak requirements:

- local regularity + random short cuts

Toy model:

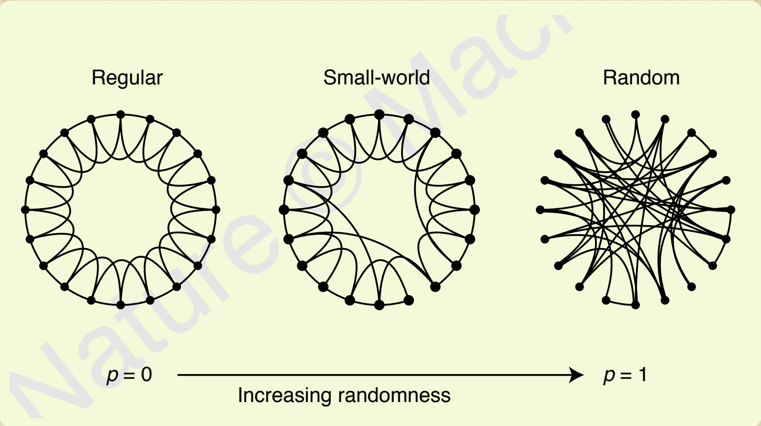
Regular
Small-world
Random
$p = 0$
Increasing randomness
$p = 1$
Nature © Mac

The structural small-world property:

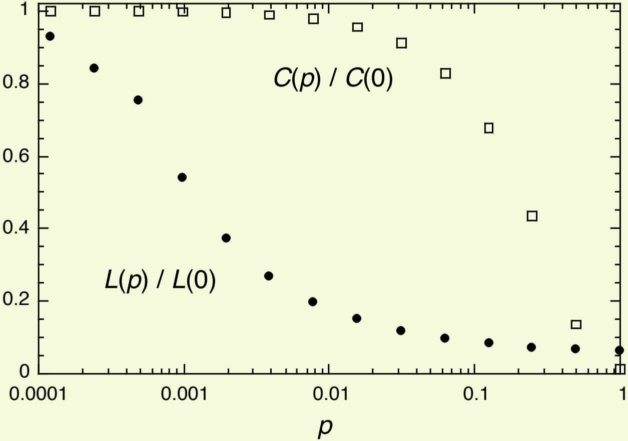

- $L(p)$ = average shortest path length as a function of p
- $C(p)$ = average clustring as a function of p

Previous work—finding short paths

But are these short cuts findable?

Nope. [8]

Nodes cannot find each other quickly
with any local search method.

Need a more sophisticated model...

Previous work—finding short paths

- What can a local search method reasonably use?
- How to find things without a map?
- Need some measure of distance between friends and the target.

Some possible knowledge:

- Target's identity
- Friends' popularity
- Friends' identities
- Where message has been

Previous work—finding short paths

Jon Kleinberg (Nature, 2000) [8]
"Navigation in a small world."

Allowed to vary:

1. local search algorithm

   and

2. network structure.

Previous work—finding short paths

Kleinberg's Network:

1. Start with regular d-dimensional cubic lattice.
2. Add local links so nodes know all nodes within a distance q.
3. Add m short cuts per node.
4. Connect i to j with probability

$$p_{ij} \propto x_{ij}{}^{-\alpha}.$$

- $\alpha = 0$: random connections.
- α large: reinforce local connections.
- $\alpha = d$: connections grow logarithmically in space.

Previous work—finding short paths

Theoretical optimal search:

- "Greedy" algorithm.
- Number of connections grow logarithmically (slowly) in space: $\alpha = d$.
- Social golf.

Search time grows slowly with system size (like $\log^2 N$).

But: social networks aren't lattices plus links.

Previous work—finding short paths

- If networks have hubs can also search well: Adamic et al. (2001) [1]

$$P(k_i) \propto k_i^{-\gamma}$$

where k = degree of node i (number of friends).

- Basic idea: get to hubs first (airline networks).
- But: hubs in social networks are limited.

The problem

If there are no hubs and no underlying lattice, how can search be efficient?

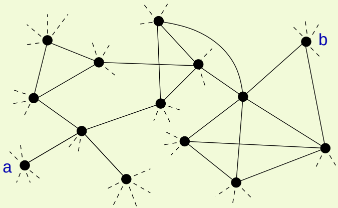

Which friend of a is closest to the target b?

What does 'closest' mean?

What is 'social distance'?

Models

One approach: incorporate identity.

Identity is formed from attributes such as:

- Geographic location
- Type of employment
- Religious beliefs
- Recreational activities.

Groups are formed by people with at least one similar attribute.

Attributes ⇔ Contexts ⇔ Interactions ⇔ Networks.

Social distance—Bipartite affiliation networks

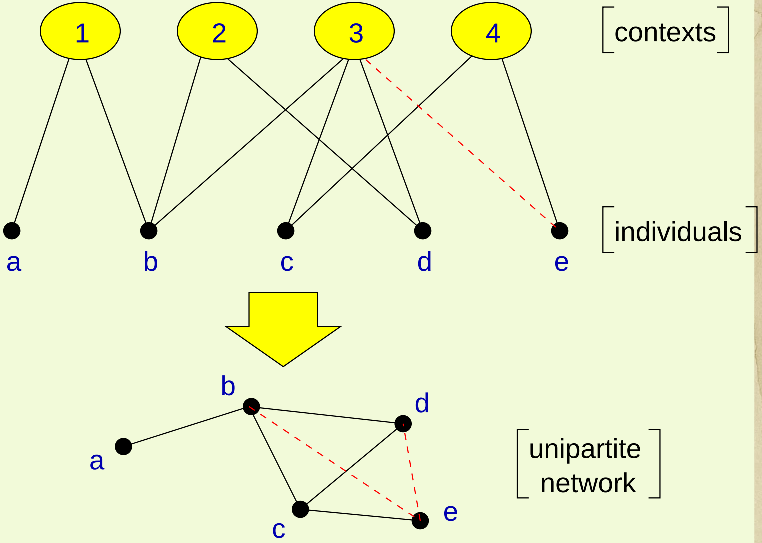

- Bipartite affiliation networks: boards and directors, movies and actors.

Social distance—Context distance

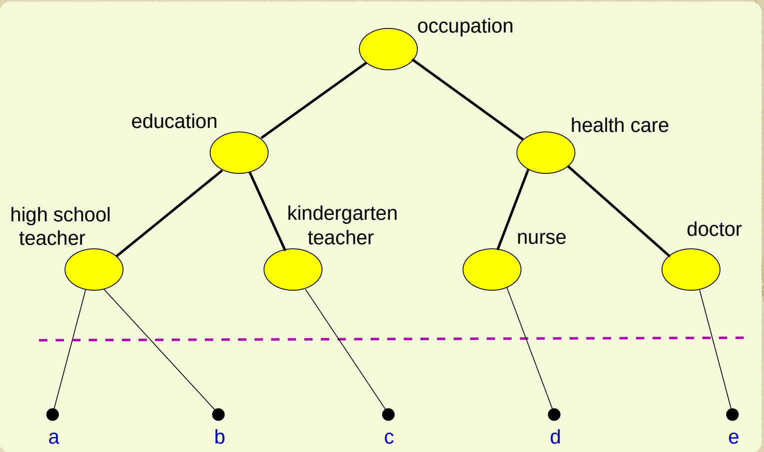

Models

Distance between two individuals x_{ij} is the height of lowest common ancestor.

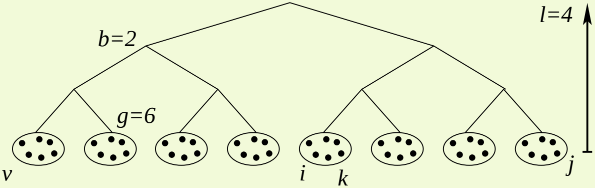

$x_{ij} = 3$, $x_{ik} = 1$, $x_{iv} = 4$.

Models

- Individuals are more likely to know each other the closer they are within a hierarchy.
- Construct z connections for each node using

$$p_{ij} = c \exp\{-\alpha x_{ij}\}.$$

- $\alpha = 0$: random connections.
- α large: local connections.

Models

Generalized affiliation networks

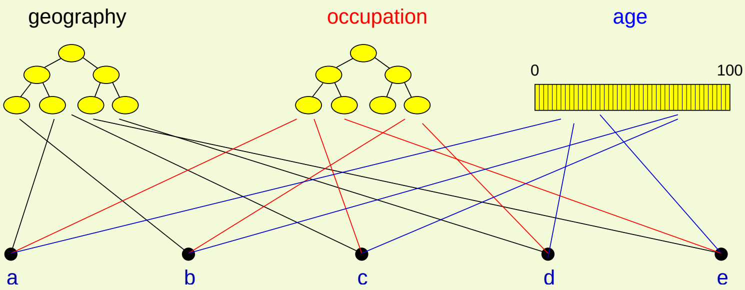

- Blau & Schwartz [4], Simmel [12], Breiger [5], Watts *et al.* [14]

The model

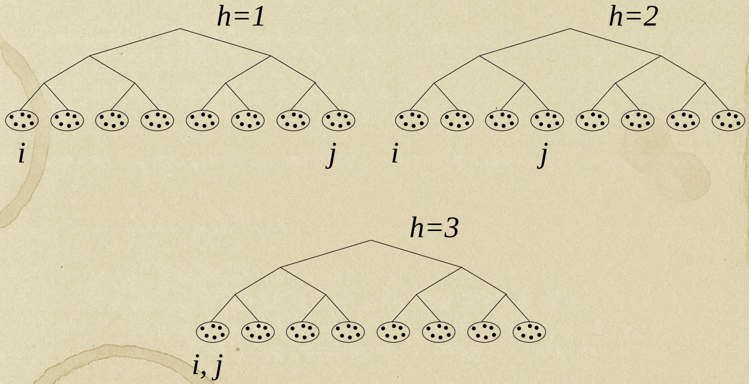

$\vec{v}_i = [1\ 1\ 1]^T$, $\vec{v}_j = [8\ 4\ 1]^T$

$x^1_{ij} = 4,\ \ x^2_{ij} = 3,\ \ x^3_{ij} = 1.$

Social distance:

$$y_{ij} = \min_h x^h_{ij}.$$

The model

Triangle inequality doesn't hold:

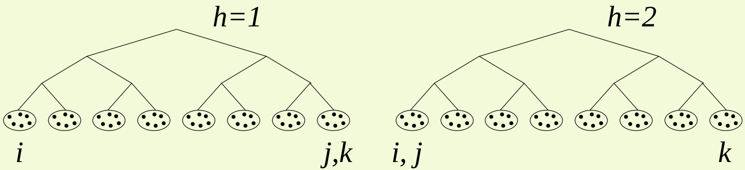

$$y_{ik} = 4 > y_{ij} + y_{jk} = 1 + 1 = 2.$$

The model

- Individuals know the identity vectors of
  1. themselves,
  2. their friends,
     and
  3. the target.
- Individuals can estimate the social distance between their friends and the target.
- Use a greedy algorithm + allow searches to fail randomly.

The model-results—searchable networks

$\alpha = 0$ versus $\alpha = 2$ for $N \simeq 10^5$:

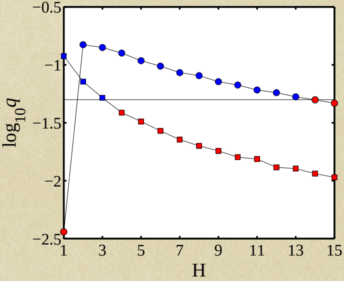

$q \geq r$

$q < r$

$r = 0.05$

q = probability an arbitrary message chain reaches a target.

- A few dimensions help.
- Searchability decreases as population increases.
- Precise form of hierarchy largely doesn't matter.

The model-results

Milgram's Nebraska-Boston data:

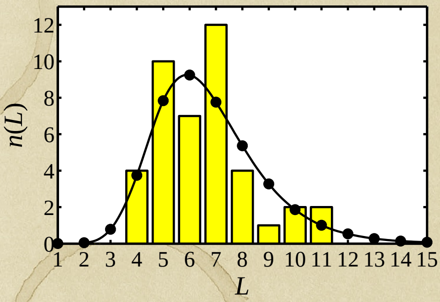

Model parameters:

- $N = 10^8$,
- $z = 300$, $g = 100$,
- $b = 10$,
- $\alpha = 1$, $H = 2$;

- $\langle L_{\rm model} \rangle \simeq 6.7$
- $L_{\rm data} \simeq 6.5$

Social search—Data

Adamic and Adar (2003)

- For HP Labs, found probability of connection as function of organization distance well fit by exponential distribution.
- Probability of connection as function of real distance $\propto 1/r$.

Social Search—Real world uses

- Tags create identities for objects
- Website tagging: `http://bitly.com`
- (e.g., Wikipedia)
- Photo tagging: `http://www.flickr.com`
- Dynamic creation of metadata plus links between information objects.
- Folksonomy: collaborative creation of metadata

Social Search—Real world uses

Recommender systems:

- Amazon uses people's actions to build effective connections between books.
- Conflict between 'expert judgments' and tagging of the hoi polloi.

Nutshell for Small-World Networks:

- Bare networks are typically unsearchable.
- Paths are findable if nodes understand how network is formed.
- Importance of identity (interaction contexts).
- Improved social network models.
- Construction of peer-to-peer networks.
- Construction of searchable information databases.

Scale-free networks

- Networks with power-law degree distributions have become known as scale-free networks.
- Scale-free refers specifically to the degree distribution having a power-law decay in its tail:

$$P_k \sim k^{-\gamma} \text{ for 'large' } k$$

- One of the seminal works in complex networks:
  Laszlo Barabási and Reka Albert, Science, 1999:
  "Emergence of scaling in random networks"[3]
  Google Scholar: Cited $\approx 16,050$ times

  (as of March 18, 2013)
- Somewhat misleading nomenclature...

Scale-free networks

- Scale-free networks are not fractal in any sense.
- Usually talking about networks whose links are abstract, relational, informational, . . . (non-physical)
- Primary example: hyperlink network of the Web
- Much arguing about whether or networks are 'scale-free' or not. . .

Some real data (we are feeling brave):

From Barabási and Albert's original paper [3]:

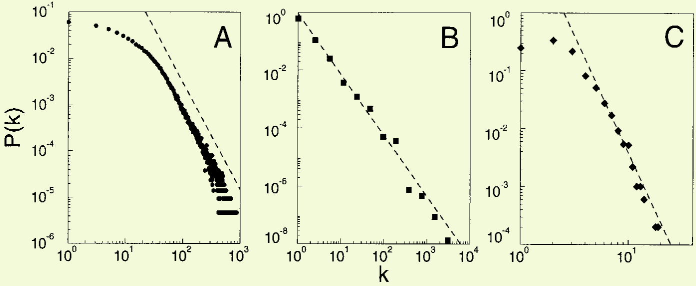

Fig. 1. The distribution function of connectivities for various large networks. (**A**) Actor collaboration graph with $N = 212{,}250$ vertices and average connectivity $\langle k \rangle = 28.78$. (**B**) WWW, $N = 325{,}729$, $\langle k \rangle = 5.46$ (6). (**C**) Power grid data, $N = 4941$, $\langle k \rangle = 2.67$. The dashed lines have slopes (A) $\gamma_{\text{actor}} = 2.3$, (B) $\gamma_{\text{www}} = 2.1$ and (C) $\gamma_{\text{power}} = 4$.

Random networks: largest components

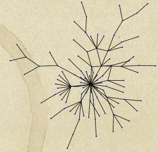

$\gamma = 2.5$
$\langle k \rangle = 1.8$

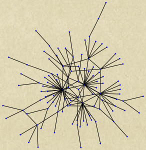

$\gamma = 2.5$
$\langle k \rangle = 2.05333$

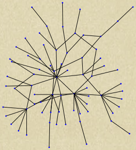

$\gamma = 2.5$
$\langle k \rangle = 1.66667$

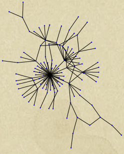

$\gamma = 2.5$
$\langle k \rangle = 1.92$

$\gamma = 2.5$
$\langle k \rangle = 1.6$

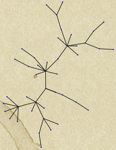

$\gamma = 2.5$
$\langle k \rangle = 1.50667$

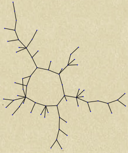

$\gamma = 2.5$
$\langle k \rangle = 1.62667$

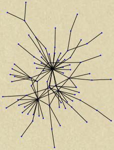

$\gamma = 2.5$
$\langle k \rangle = 1.8$

Scale-free networks

The big deal:

- We move beyond describing networks to finding mechanisms for why certain networks are the way they are.

A big deal for scale-free networks:

- How does the exponent γ depend on the mechanism?
- Do the mechanism details matter?

BA model

- Barabási-Albert model = BA model.
- Key ingredients:
  Growth and Preferential Attachment (PA).
- Step 1: start with m_0 disconnected nodes.
- Step 2:
  1. Growth—a new node appears at each time step $t = 0, 1, 2, \ldots$.
  2. Each new node makes m links to nodes already present.
  3. Preferential attachment—Probability of connecting to ith node is $\propto k_i$.
- In essence, we have a rich-gets-richer scheme.
- Yes, we've seen this all before in Simon's model.

BA model

- Definition: A_k is the attachment kernel for a node with degree k.
- For the original model:

$$A_k = k$$

- Definition: $P_{\text{attach}}(k,t)$ is the attachment probability.
- For the original model:

$$P_{\text{attach}}(\text{node } i, t) = \frac{k_i(t)}{\sum_{j=1}^{N(t)} k_j(t)} = \frac{k_i(t)}{\sum_{k=0}^{k_{\max}(t)} k N_k(t)}$$

where $N(t) = m_0 + t$ is # nodes at time t and $N_k(t)$ is # degree k nodes at time t.

Approximate analysis

- When $(N+1)$th node is added, the expected increase in the degree of node i is

$$E(k_{i,N+1} - k_{i,N}) \simeq m \frac{k_{i,N}}{\sum_{j=1}^{N(t)} k_j(t)}.$$

- Assumes probability of being connected to is small.
- Dispense with Expectation by assuming (hoping) that over longer time frames, degree growth will be smooth and stable.
- Approximate $k_{i,N+1} - k_{i,N}$ with $\frac{\mathrm{d}}{\mathrm{d}t} k_{i,t}$:

$$\frac{\mathrm{d}}{\mathrm{d}t} k_{i,t} = m \frac{k_i(t)}{\sum_{j=1}^{N(t)} k_j(t)}$$

where $t = N(t) - m_0$.

Approximate analysis

- Deal with denominator: each added node brings m new edges.

$$\therefore \sum_{j=1}^{N(t)} k_j(t) = 2tm$$

- The node degree equation now simplifies:

$$\frac{\mathrm{d}}{\mathrm{d}t} k_{i,t} = m \frac{k_i(t)}{\sum_{j=1}^{N(t)} k_j(t)} = m \frac{k_i(t)}{2mt} = \frac{1}{2t} k_i(t)$$

- Rearrange and solve:

$$\frac{\mathrm{d}k_i(t)}{k_i(t)} = \frac{\mathrm{d}t}{2t} \Rightarrow \boxed{k_i(t) = c_i\, t^{1/2}.}$$

- Next find c_i ...

Approximate analysis

- Know ith node appears at time

$$t_{i,\text{start}} = \begin{cases} i - m_0 & \text{for } i > m_0 \\ 0 & \text{for } i \le m_0 \end{cases}$$

- So for $i > m_0$ (exclude initial nodes), we must have

$$k_i(t) = m \left(\frac{t}{t_{i,\text{start}}} \right)^{1/2} \quad \text{for } t \ge t_{i,\text{start}}.$$

- All node degrees grow as $t^{1/2}$ but later nodes have larger $t_{i,\text{start}}$ which flattens out growth curve.
- First-mover advantage: Early nodes do best.

Approximate analysis

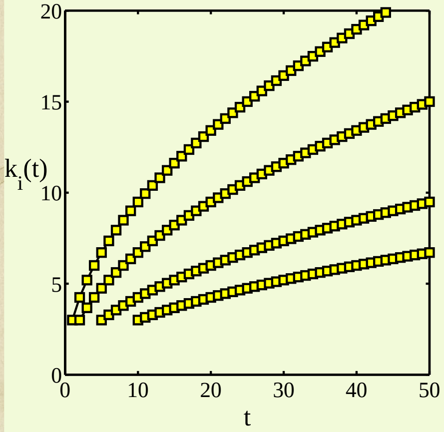

- $m = 3$
- $t_{i,\text{start}} = 1, 2, 5,$ and 10.

Degree distribution

- So what's the degree distribution at time t?
- Use fact that birth time for added nodes is distributed uniformly between time 0 and t:

$$\mathbf{Pr}(t_{i,\text{start}})\mathrm{d}t_{i,\text{start}} \simeq \frac{\mathrm{d}t_{i,\text{start}}}{t}$$

- Also use

$$k_i(t) = m\left(\frac{t}{t_{i,\text{start}}}\right)^{1/2} \Rightarrow t_{i,\text{start}} = \frac{m^2 t}{k_i(t)^2}.$$

Transform variables—Jacobian:

$$\frac{\mathrm{d}t_{i,\text{start}}}{\mathrm{d}k_i} = -2\frac{m^2 t}{k_i(t)^3}.$$

Degree distribution

- $$\mathbf{Pr}(k_i)\mathrm{d}k_i = \mathbf{Pr}(t_{i,\mathrm{start}})\mathrm{d}t_{i,\mathrm{start}}$$
- $$= \mathbf{Pr}(t_{i,\mathrm{start}})\mathrm{d}k_i \left| \frac{\mathrm{d}t_{i,\mathrm{start}}}{\mathrm{d}k_i} \right|$$
- $$= \frac{1}{t}\mathrm{d}k_i \, 2\frac{m^2 t}{k_i(t)^3}$$
- $$= 2\frac{m^2}{k_i(t)^3}\mathrm{d}k_i$$
- $$\propto k_i^{-3}\mathrm{d}k_i \,.$$

Degree distribution

- We thus have a very specific prediction of $\mathbf{Pr}(k) \sim k^{-\gamma}$ with $\gamma = 3$.
- Typical for real networks: $2 < \gamma < 3$.
- Range true more generally for events with size distributions that have power-law tails.
- $2 < \gamma < 3$: finite mean and 'infinite' variance (wild)
- In practice, $\gamma < 3$ means variance is governed by upper cutoff.
- $\gamma > 3$: finite mean and variance (mild)

Back to that real data:

From Barabási and Albert's original paper [3]:

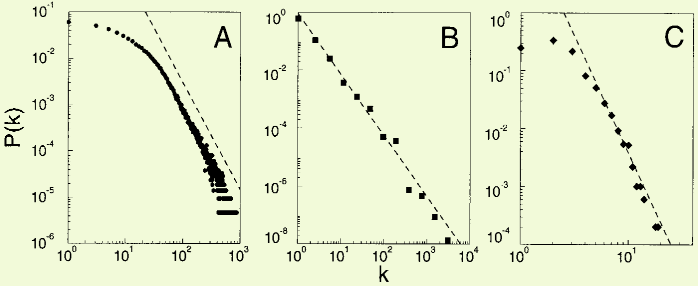

Fig. 1. The distribution function of connectivities for various large networks. (**A**) Actor collaboration graph with $N = 212{,}250$ vertices and average connectivity $\langle k \rangle = 28.78$. (**B**) WWW, $N = 325{,}729$, $\langle k \rangle = 5.46$ (6). (**C**) Power grid data, $N = 4941$, $\langle k \rangle = 2.67$. The dashed lines have slopes (A) $\gamma_{\text{actor}} = 2.3$, (B) $\gamma_{\text{www}} = 2.1$ and (C) $\gamma_{\text{power}} = 4$.

Examples

Web	$\gamma \simeq 2.1$ for in-degree
Web	$\gamma \simeq 2.45$ for out-degree
Movie actors	$\gamma \simeq 2.3$
Words (synonyms)	$\gamma \simeq 2.8$

The Internets is a different business...

Things to do and questions

- Vary attachment kernel.
- Vary mechanisms:
  1. Add edge deletion
  2. Add node deletion
  3. Add edge rewiring
- Deal with directed versus undirected networks.
- Important Q.: Are there distinct universality classes for these networks?
- Q.: How does changing the model affect γ?
- Q.: Do we need preferential attachment and growth?
- Q.: Do model details matter? Maybe . . .

Preferential attachment

- Let's look at preferential attachment (PA) a little more closely.
- PA implies arriving nodes have complete knowledge of the existing network's degree distribution.
- For example: If $P_{\mathrm{attach}}(k) \propto k$, we need to determine the constant of proportionality.
- We need to know what everyone's degree is...
- PA is $\therefore$ an outrageous assumption of node capability.
- But a very simple mechanism saves the day...

Preferential attachment through randomness

- Instead of attaching preferentially, allow new nodes to attach randomly.
- Now add an extra step: new nodes then connect to some of their friends' friends.
- Can also do this at random.
- Assuming the existing network is random, we know probability of a random friend having degree k is

$$Q_k \propto kP_k$$

- So rich-gets-richer scheme can now be seen to work in a natural way.
- Later: we'll see that the nature of Q_k means your friends have more friends that you. #disappointing

Robustness

- Albert et al., Nature, 2000:
  "Error and attack tolerance of complex networks" [2]
- Standard random networks (Erdős-Rényi) versus Scale-free networks:

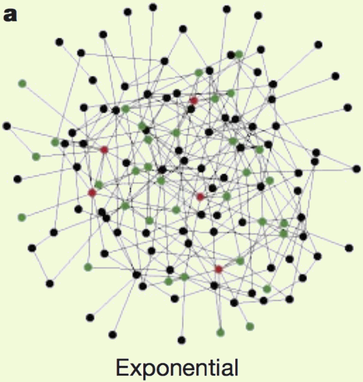

from Albert et al., 2000

Robustness

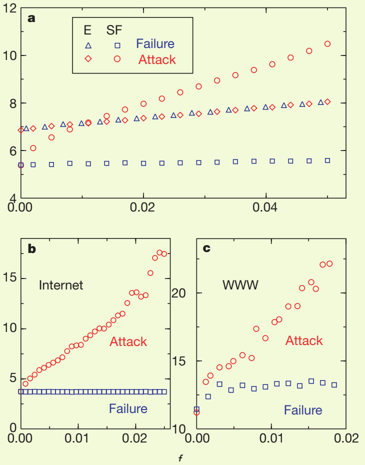

from Albert et al., 2000

- Plots of network diameter as a function of fraction of nodes removed
- Erdős-Rényi versus scale-free networks
- blue symbols = random removal
- red symbols = targeted removal (most connected first)

Robustness

- Scale-free networks are thus robust to random failures yet fragile to targeted ones.
- All very reasonable: Hubs are a big deal.
- But: next issue is whether hubs are vulnerable or not.
- Representing all webpages as the same size node is obviously a stretch (e.g., google vs. a random person's webpage)
- Most connected nodes are either:
  1. Physically larger nodes that may be harder to 'target'
  2. or subnetworks of smaller, normal-sized nodes.
- Need to explore cost of various targeting schemes.

Generalized model

Fooling with the mechanism:

- 2001: Krapivsky & Redner (KR) [9] explored the general attachment kernel:

$$\mathbf{Pr}(\text{attach to node } i) \propto A_k = k_i^{\nu}$$

  where A_k is the attachment kernel and $\nu > 0$.
- KR also looked at changing the details of the attachment kernel.
- We'll follow KR's approach using rate equations (⊞).

Generalized model

- Here's the set up:

$$\frac{\mathrm{d}N_k}{\mathrm{d}t} = \frac{1}{A}\left[A_{k-1}N_{k-1} - A_k N_k\right] + \delta_{k1}$$

where N_k is the number of nodes of degree k.

1. One node with one link is added per unit time.
2. The first term corresponds to degree $k-1$ nodes becoming degree k nodes.
3. The second term corresponds to degree k nodes becoming degree $k-1$ nodes.
4. A is the correct normalization (coming up).
5. Seed with some initial network (e.g., a connected pair)
6. Detail: $A_0 = 0$

Generalized model

- In general, probability of attaching to a specific node of degree k at time t is

$$\textbf{Pr}(\text{attach to node } i) = \frac{A_k}{A(t)}$$

where $A(t) = \sum_{k=1}^{\infty} A_k N_k(t)$.

- E.g., for BA model, $A_k = k$ and $A = \sum_{k=1}^{\infty} k N_k(t)$.
- For $A_k = k$, we have

$$A(t) = \sum_{k'=1}^{\infty} k' N_{k'}(t) = 2t$$

since one edge is being added per unit time.

- Detail: we are ignoring initial seed network's edges.

Generalized model

- So now

$$\frac{\mathrm{d}N_k}{\mathrm{d}t} = \frac{1}{A}\left[A_{k-1}N_{k-1} - A_k N_k\right] + \delta_{k1}$$

becomes

$$\frac{\mathrm{d}N_k}{\mathrm{d}t} = \frac{1}{2t}\left[(k-1)N_{k-1} - kN_k\right] + \delta_{k1}$$

- As for BA method, look for steady-state growing solution: $N_k = n_k t$.
- We replace $\mathrm{d}N_k/\mathrm{d}t$ with $\mathrm{d}n_k t/\mathrm{d}t = n_k$.
- We arrive at a difference equation:

$$n_k = \frac{1}{2\cancel{t}}\left[(k-1)n_{k-1}\cancel{t} - kn_k\cancel{t}\right] + \delta_{k1}$$

Universality?

- Insert question from assignment 7 (⊞)
  As expected, we have the same result as for the BA model:

$$N_k(t) = n_k(t)t \propto k^{-3} \text{ for large } k.$$

- Now: what happens if we start playing around with the attachment kernel A_k?
- Again, we're asking if the result $\gamma = 3$ universal (⊞)?
- KR's natural modification: $A_k = k^\nu$ with $\nu \neq 1$.
- But we'll first explore a more subtle modification of A_k made by Krapivsky/Redner [9]
- Keep A_k linear in k but tweak details.
- Idea: Relax from $A_k = k$ to $A_k \sim k$ as $k \to \infty$.

Universality?

- Recall we used the normalization:

$$A(t) = \sum_{k'=1}^{\infty} k' N_{k'}(t) \simeq 2t \text{ for large } t.$$

- We now have

$$A(t) = \sum_{k'=1}^{\infty} A_{k'} N_{k'}(t)$$

  where we only know the asymptotic behavior of A_k.
- We assume that $A = \mu t$
- We'll find μ later and make sure that our assumption is consistent.
- As before, also assume $N_k(t) = n_k t$.

Universality?

- For $A_k = k$ we had

$$n_k = \frac{1}{2}\left[(k-1)n_{k-1} - kn_k\right] + \delta_{k1}$$

- This now becomes

$$n_k = \frac{1}{\mu}\left[A_{k-1}n_{k-1} - A_k n_k\right] + \delta_{k1}$$

$$\Rightarrow (A_k + \mu)n_k = A_{k-1}n_{k-1} + \mu\delta_{k1}$$

- Again two cases:

$$k = 1 : n_1 = \frac{\mu}{\mu + A_1}; \qquad k > 1 : n_k = n_{k-1}\frac{A_{k-1}}{\mu + A_k}.$$

Universality?

- Time for pure excitement: Find asymptotic behavior of n_k given $A_k \to k$ as $k \to \infty$.
- Insert question from assignment 7 (⊞)
  For large k, we find:

$$n_k = \frac{\mu}{A_k} \prod_{j=1}^{k} \frac{1}{1 + \frac{\mu}{A_j}} \propto k^{-\mu-1}$$

- Since μ depends on A_k, details matter...

Universality?

- Now we need to find μ.
- Our assumption again: $A = \mu t = \sum_{k=1}^{\infty} N_k(t) A_k$
- Since $N_k = n_k t$, we have the simplification $\mu = \sum_{k=1}^{\infty} n_k A_k$
- Now subsitute in our expression for n_k:

$$1 \cancel{\mu} = \sum_{k=1}^{\infty} \frac{\cancel{\mu}}{\cancel{A_k}} \prod_{j=1}^{k} \frac{1}{1 + \frac{\mu}{A_j}} \cancel{A_k}$$

- Closed form expression for μ.
- We can solve for μ in some cases.
- Our assumption that $A = \mu t$ looks to be not too horrible.

Universality?

- Consider tunable $A_1 = \alpha$ and $A_k = k$ for $k \geq 2$.
- Again, we can find $\gamma = \mu + 1$ by finding μ.
- Insert question from assignment 7 (⊞)
  Closed form expression for μ:

$$\frac{\mu}{\alpha} = \sum_{k=2}^{\infty} \frac{\Gamma(k+1)\Gamma(2+\mu)}{\Gamma(k+\mu+1)}$$

  #mathisfun

-
$$\mu(\mu - 1) = 2\alpha \Rightarrow \mu = \frac{1 + \sqrt{1 + 8\alpha}}{2}.$$

- Since $\gamma = \mu + 1$, we have

$$0 \leq \alpha < \infty \Rightarrow 2 \leq \gamma < \infty$$

- Craziness...

Sublinear attachment kernels

- Rich-get-somewhat-richer:

$$A_k \sim k^{\nu} \text{ with } 0 < \nu < 1.$$

- General finding by Krapivsky and Redner:[9]

$$n_k \sim k^{-\nu} e^{-c_1 k^{1-\nu} + \text{correction terms}}.$$

- Stretched exponentials (truncated power laws).
- aka Weibull distributions.
- Universality: now details of kernel do not matter.
- Distribution of degree is universal providing $\nu < 1$.

Sublinear attachment kernels

Details:

- For $1/2 < \nu < 1$:

$$n_k \sim k^{-\nu} e^{-\mu\left(\frac{k^{1-\nu} - 2^{1-\nu}}{1-\nu}\right)}$$

- For $1/3 < \nu < 1/2$:

$$n_k \sim k^{-\nu} e^{-\mu \frac{k^{1-\nu}}{1-\nu} + \frac{\mu^2}{2} \frac{k^{1-2\nu}}{1-2\nu}}$$

- And for $1/(r+1) < \nu < 1/r$, we have r pieces in exponential.

Superlinear attachment kernels

- Rich-get-much-richer:

$$A_k \sim k^{\nu} \text{ with } \nu > 1.$$

- Now a winner-take-all mechanism.
- One single node ends up being connected to almost all other nodes.
- For $\nu > 2$, all but a finite # of nodes connect to one node.

Nutshell:

Overview Key Points for Models of Networks:

- Obvious connections with the vast extant field of graph theory.
- But focus on dynamics is more of a physics/stat-mech/comp-sci flavor.
- Two main areas of focus:
  1. Description: Characterizing very large networks
  2. Explanation: Micro story $\Rightarrow$ Macro features
- Some essential structural aspects are understood: degree distribution, clustering, assortativity, group structure, overall structure,...
- Still much work to be done, especially with respect to dynamics... #excitement

References I

[1] L. Adamic, R. Lukose, A. Puniyani, and B. Huberman.
Search in power-law networks.
Phys. Rev. E, 64:046135, 2001. pdf (⊞)

[2] R. Albert, H. Jeong, and A.-L. Barabási.
Error and attack tolerance of complex networks.
Nature, 406:378–382, 2000. pdf (⊞)

[3] A.-L. Barabási and R. Albert.
Emergence of scaling in random networks.
Science, 286:509–511, 1999. pdf (⊞)

[4] P. M. Blau and J. E. Schwartz.
Crosscutting Social Circles.
Academic Press, Orlando, FL, 1984.

References II

[5] R. L. Breiger.
The duality of persons and groups.
Social Forces, 53(2):181–190, 1974. pdf (⊞)

[6] P. S. Dodds, R. Muhamad, and D. J. Watts.
An experimental study of search in global social networks.
Science, 301:827–829, 2003. pdf (⊞)

[7] M. Granovetter.
The strength of weak ties.
Am. J. Sociol., 78(6):1360–1380, 1973. pdf (⊞)

[8] J. Kleinberg.
Navigation in a small world.
Nature, 406:845, 2000. pdf (⊞)

References III

[9] P. L. Krapivsky and S. Redner.
Organization of growing random networks.
Phys. Rev. E, 63:066123, 2001. pdf (⊞)

[10] R. Milo, N. Kashtan, S. Itzkovitz, M. E. J. Newman, and U. Alon.
On the uniform generation of random graphs with prescribed degree sequences, 2003. pdf (⊞)

[11] G. Pickard, W. Pan, I. Rahwan, M. Cebrian, R. Crane, A. Madan, and A. Pentland.
Time-critical social mobilization.
Science, 334:509–512, 2011. pdf (⊞)

[12] G. Simmel.
The number of members as determining the sociological form of the group. I.
American Journal of Sociology, 8:1–46, 1902.

References IV

[13] J. Travers and S. Milgram.
An experimental study of the small world problem.
Sociometry, 32:425–443, 1969. pdf (⊞)

[14] D. J. Watts, P. S. Dodds, and M. E. J. Newman.
Identity and search in social networks.
Science, 296:1302–1305, 2002. pdf (⊞)

[15] D. J. Watts and S. J. Strogatz.
Collective dynamics of ‘small-world’ networks.
Nature, 393:440–442, 1998. pdf (⊞)