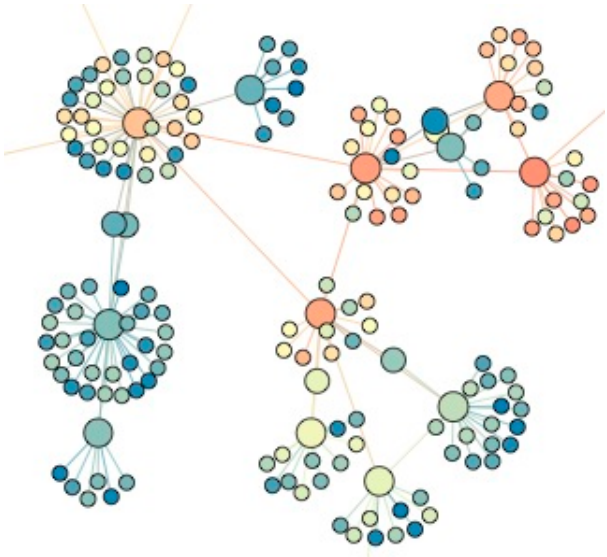

Topic mash I: games, learning & signed edges

CS224W

When networks get interesting/complicated

- Evolution
  - local processes
  - global optimization
  - games (each node connects to maximize utility)
- Learning
  - innovation
  - coordination (graph coloring)
- Signed edges
  - balance
  - status

Rewiring through global optimization

- Assign properties to nodes (e.g. spatial location, group membership)
- Add or rewire links according to some rule
  - optimize for a particular property (simulated annealing)
  - add links with probability depending on property of existing nodes, edges (preferential attachment, link copying)
  - simulate nodes as agents ‘deciding’ whether to rewire or add links

Rewiring through search

- When searching, rewire to an intermediary. Can produce "navigable" networks with $P(d) \sim d^{-\alpha}$

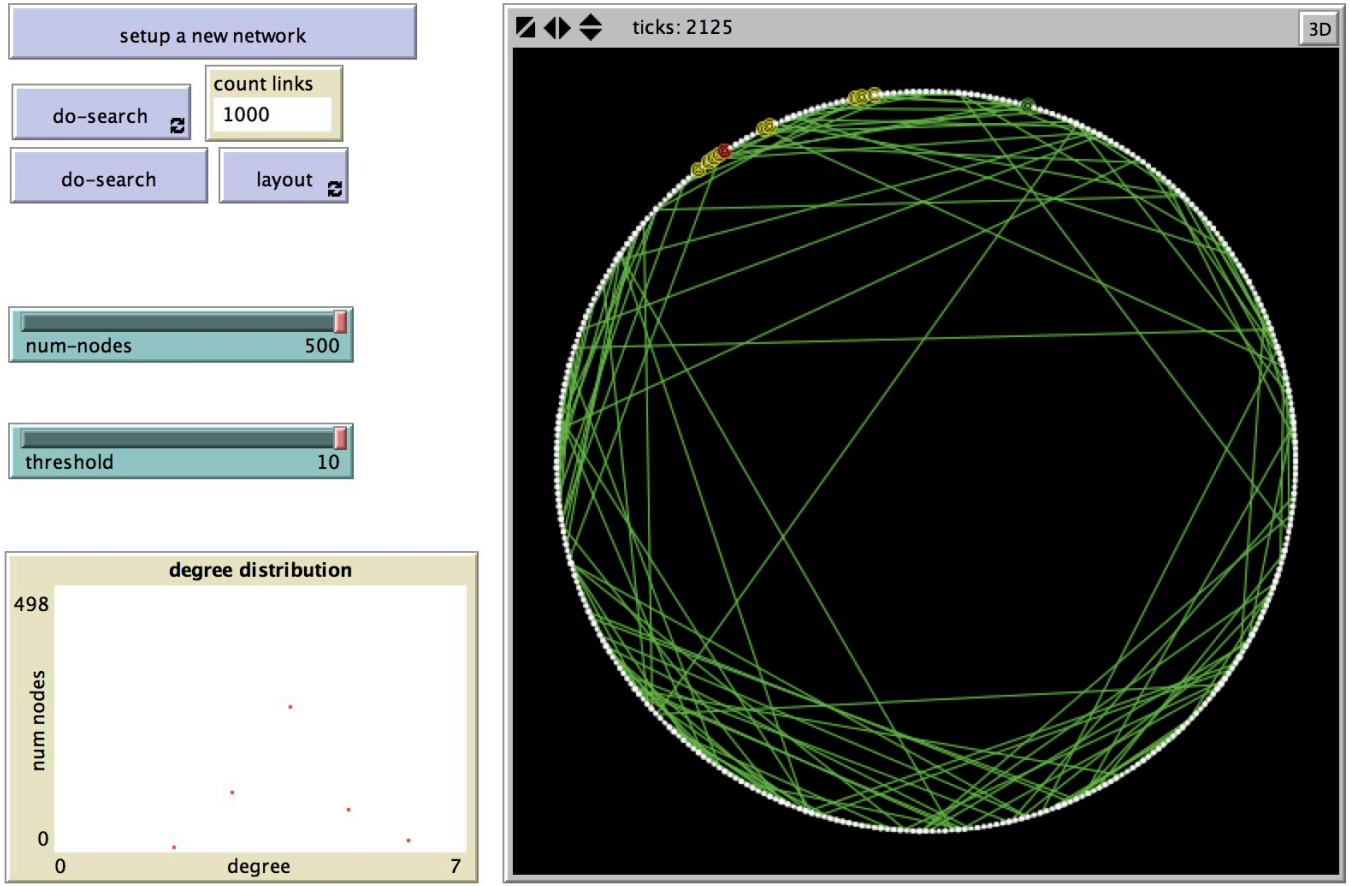

http://web.stanford.edu/class/cs224w/NetLogo/SearchNetwork.nlogo

Clauset & Moore: how to networks become navigable

Rewiring through global optimization

Small worlds: How and Why, Nisha Mathias and Venkatesh Gopal

$$E = \lambda L + (1 - \lambda) W$$

$$L = \frac{1}{n(n-1)} \sum_{i \neq j} d_{ij}$$

$$W = \sum_{e_{ij}} \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$$

- E is the ‘energy’ cost we are trying to minimize
- L is the average shortest path in ‘hops’
- W is the total length of wire used

optimized networks

- rewire using simulated annealing
- sequence is shown in order of increasing λ

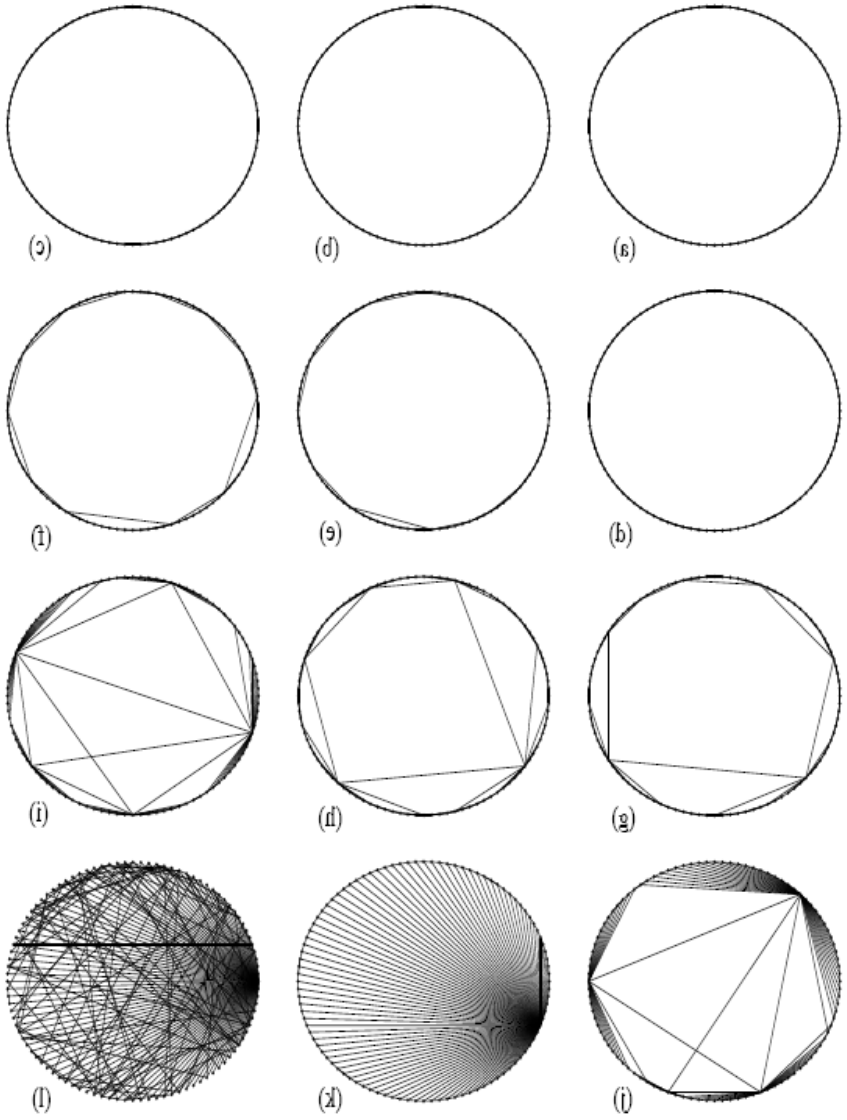

Source: Small worlds: How and Why, Nisha Mathias and Venkatesh Gopal
http://link.aps.org/doi/10.1103/PhysRevE.63.021117 DOI: 10.1103/PhysRevE.63.021117

another view of optimized networks

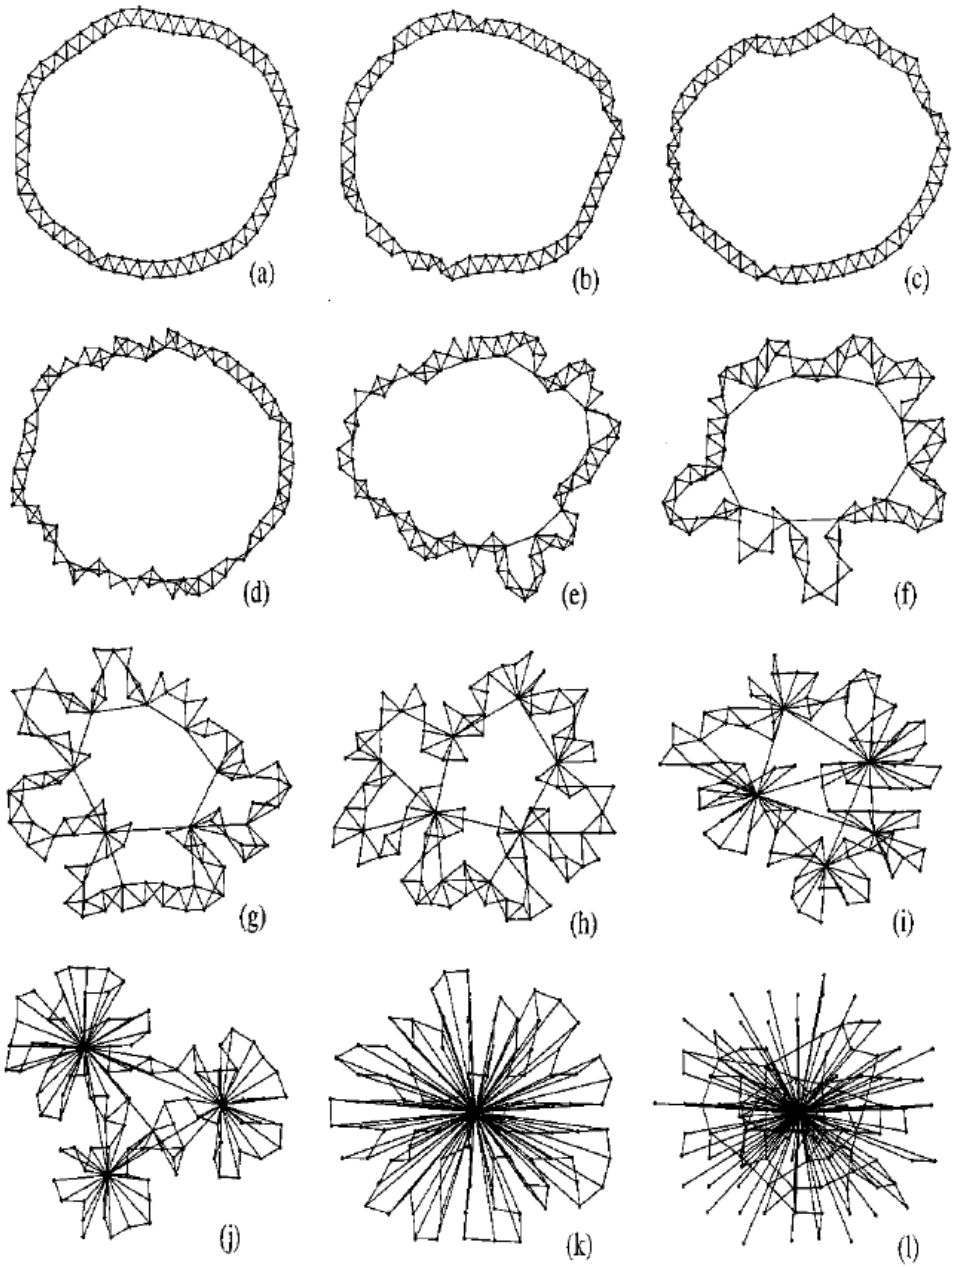

- same networks, but the vertices are allowed to move using a spring layout algorithm
- wiring cost associated with the physical distance between nodes

Source: Small worlds: How and Why, Nisha Mathias and Venkatesh Gopal
http://link.aps.org/doi/10.1103/PhysRevE.63.021117 DOI: 10.1103/PhysRevE.63.021117

optimizing from scratch

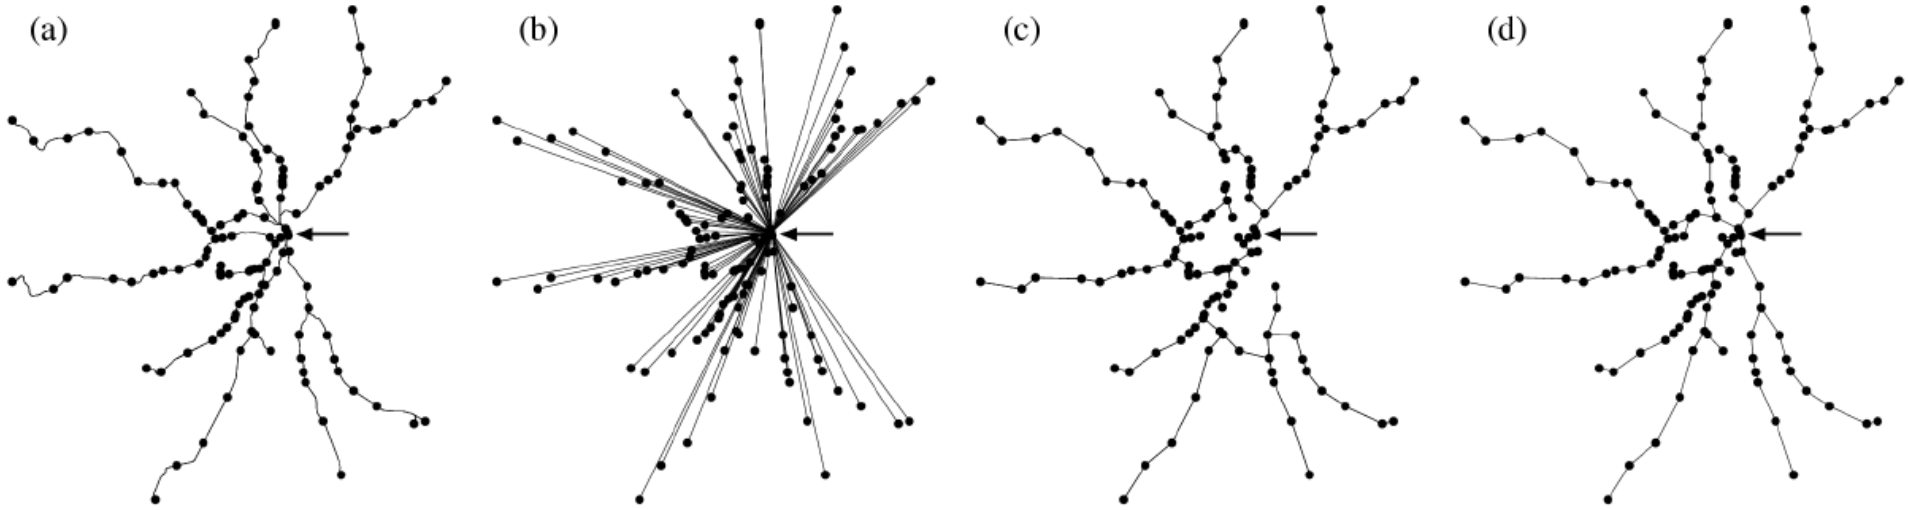

(a) Commuter rail network in the Boston area. The arrow marks the assumed root of the network.

(b) Star graph.

(c) Minimum spanning tree.

(d) The model applied to the same set of stations.

add edge with smallest weight $w'_{ij} = d_{ij} + \beta l_{j0}$

hops to root node

Euclidean distance between i and j

Source: The Spatial Structure of Networks, M. T. Gastner and M. E.J. Newman

http://www.springerlink.com/content/p26t67882668514q DOI: 10.1140/epjb/e2006-00046-8

reminiscent of

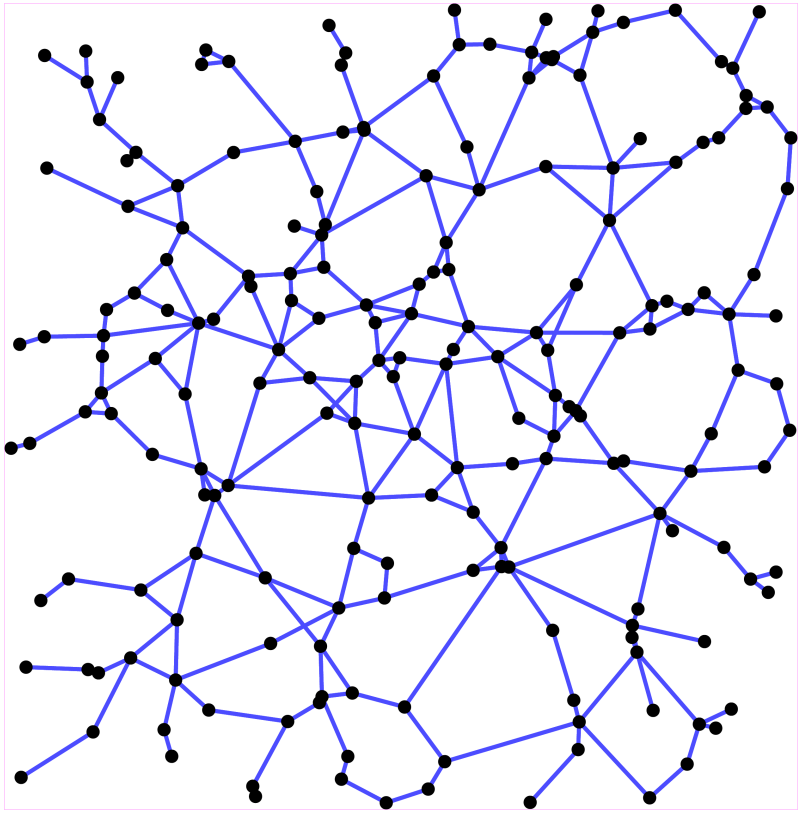

Roads

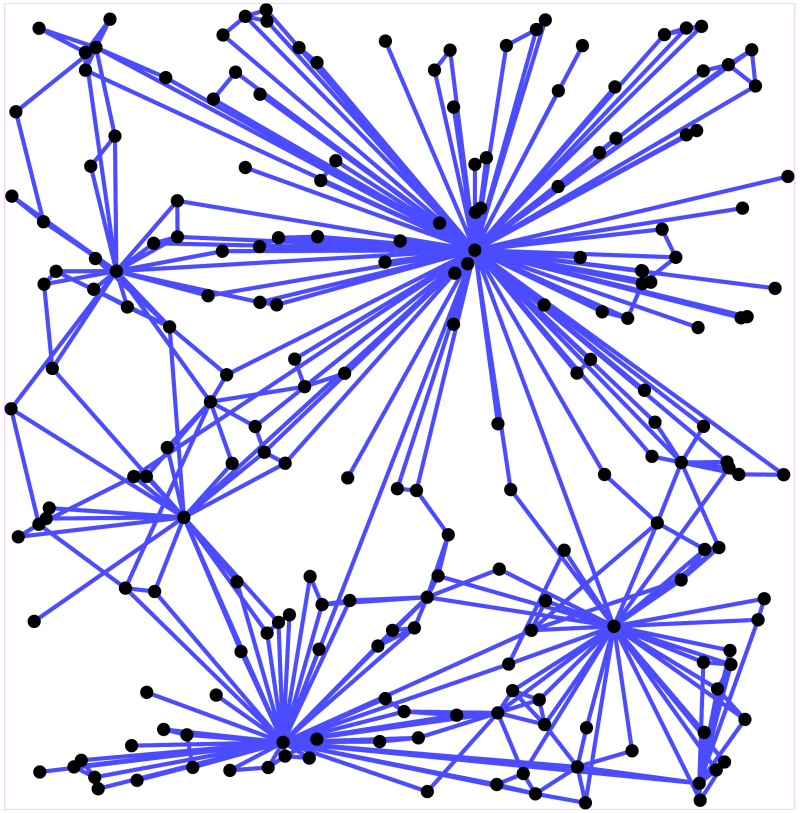

Air routes

Source: The Spatial Structure of Networks, M. T. Gastner and M. E.J. Newman
http://www.springerlink.com/content/p26t67882668514q DOI: 10.1140/epjb/e2006-00046-8

Quiz Q:

- A network that contains many hubs with far reaching edges is indicative of(check all that apply)
  - high cost of distance traveled
  - low cost of distance traveled
  - high cost of making many hops
  - low cost of making many hops

http://web.stanford.edu/class/cs224w/NetLogo/howwhysmallworlds.nlogo

HW 4 games

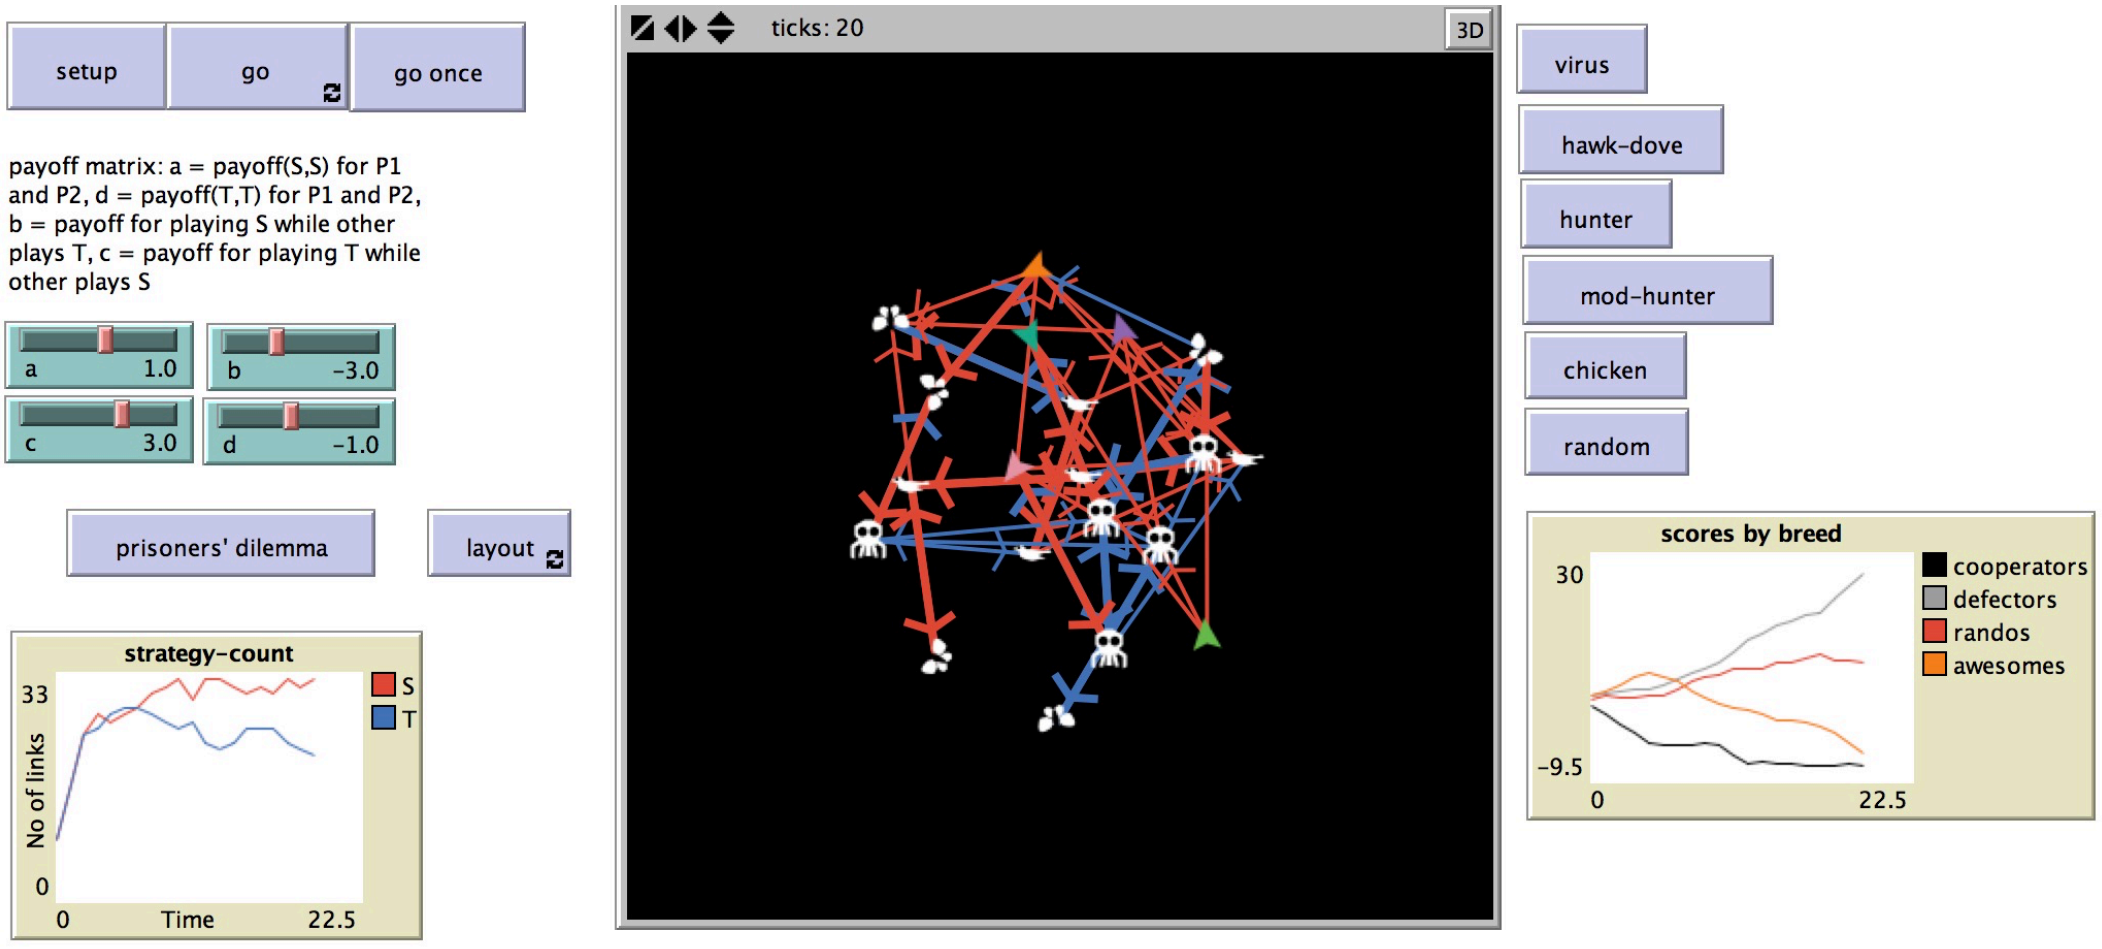

http://web.stanford.edu/class/cs224w/NetLogo/GameCompetitionShell.nlogo

Virus game

- Phage Φ6 has a mutational variant Phage ΦH2.
- On its own Φ6 replicates better than ΦH2.
- ΦH2 is able to take advantage of chemical products produced by Φ6, which gives ΦH2 a fitness advantage when it is in the presence of Φ6.

		Virus 2	
		Φ6	ΦH2
Virus 1	Φ6	1.00, 1.00	0.65, 1.99
	ΦH2	1.99, 0.65	0.83, 0.83

Easley & Kleinberg Ch7: Evolutionary game theory

Stag hunt game

- Hunting the stag gives both players higher reward
- but if one hunts the stag, they fail alone, while the one who hunts the hare gets a small catch
- if both hunt hare, they both get a small catch

		Hunter 2	
		Hunt Stag	*Hunt Hare*
Hunter 1	*Hunt Stag*	$4, 4$	$0, 3$
	Hunt Hare	$3, 0$	$3, 3$

Easley & Kleinberg Ch7: Evolutionary game theory

Modified hunt game

- Same, but if one decides to hunt the stag, the other doesn't have competition in hunting hares, and so gets higher payoff

		Hunter 2	
		Hunt Stag	*Hunt Hare*
Hunter 1	*Hunt Stag*	$4, 4$	$0, 4$
	Hunt Hare	$4, 0$	$3, 3$

Hawk dove game

- The dove strategy means acting meekly sharing food
- The hawk strategy means being aggresive and grabbing most of the food (at the expense of the dove)
- However, two hawk will hurt each other, and destroy the food, reciving no payoff

		Animal 2	
		D	H
Animal 1	D	$3, 3$	$1, 5$
	H	$5, 1$	$0, 0$

Easley & Kleinberg Ch7: Evolutionary game theory

game of chicken

- Two cars drive straight toward each other
- The first to swerve is "chicken" and endures the ridicule of the gloating other
- However, if neither swerves, the cost is high

	go straight	swerve
go straight	-10,-10	1,0
swerve	0,1	0,0

Innovation & coordination

innovation in networks

- network topology influences who talks to whom
- who talks to whom has important implications for innovation and learning

better to innovate or imitate?

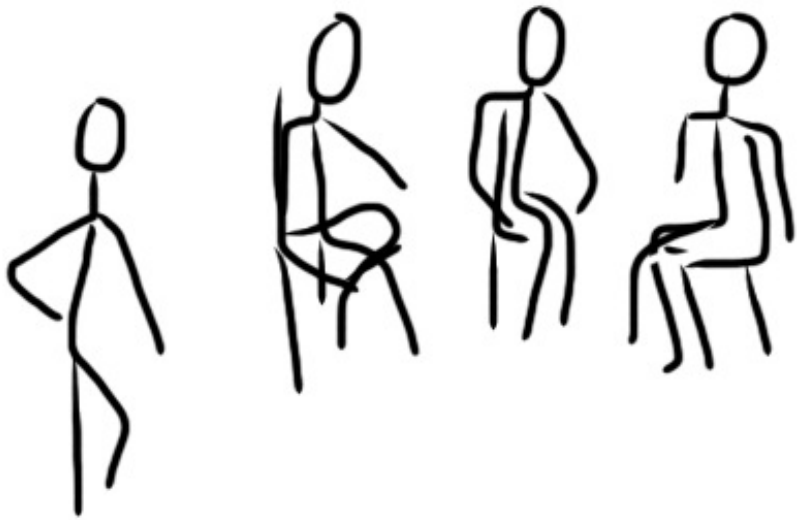

brainstorming:
more minds together,
but also danger of groupthink

working in isolation:
more independence
slower progress

in a network context

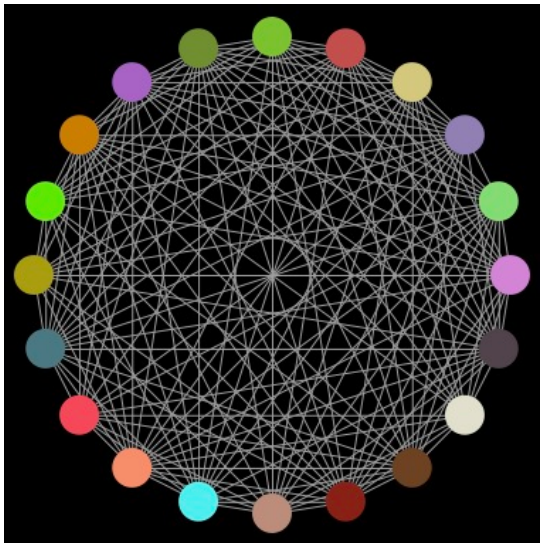

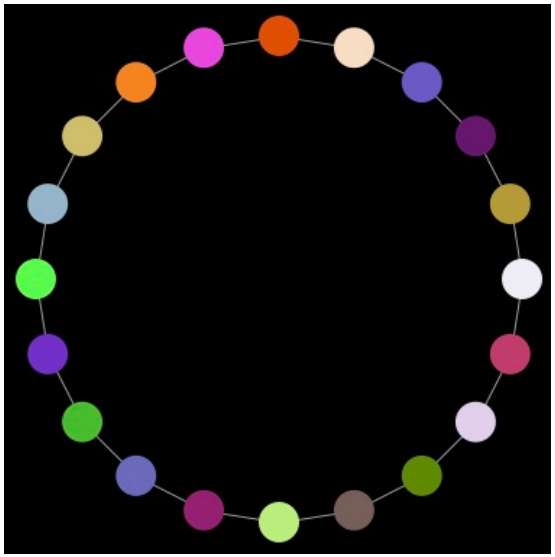

modeling the problem space

- Kauffman's NK model
- N dimensional problem space
  - N bits, each can be 0 or 1
- K describes the smoothness of the fitness landscape
  - how similar is the fitness of sequences with only 1-2 bits flipped (K = 0, no similarity, K large, smooth fitness)

Kauffman's NK model

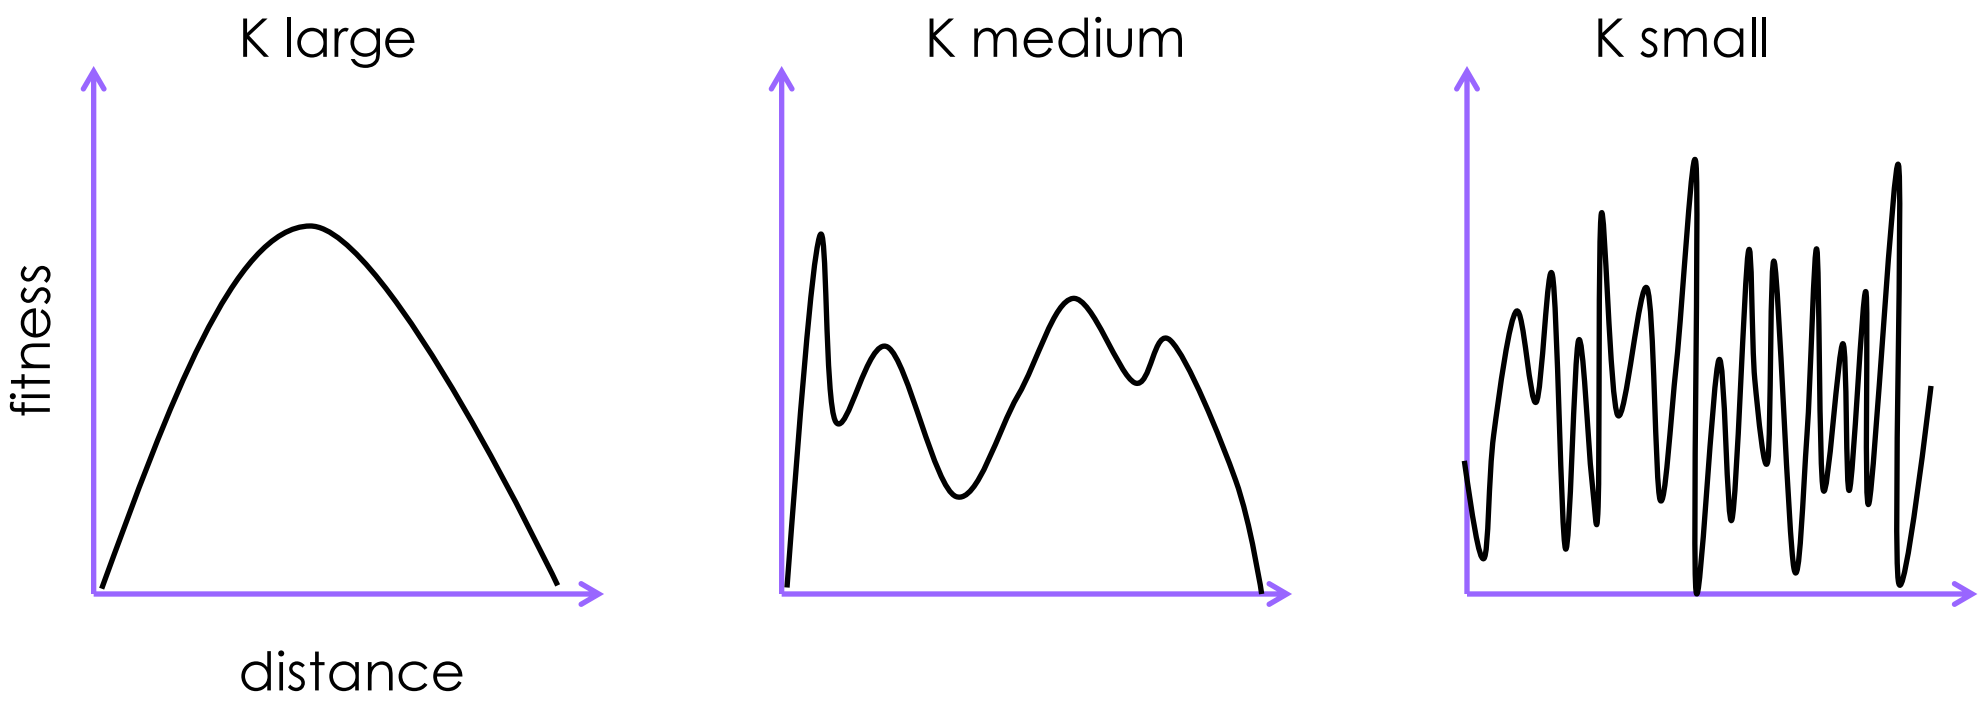

Update rules

- As a node, you start out with a random bit string
- At each iteration
  - If one of your neighbors has a solution that is more fit than yours, imitate (copy their solution)
  - Otherwise innovate by flipping one of your bits

NetLogo model

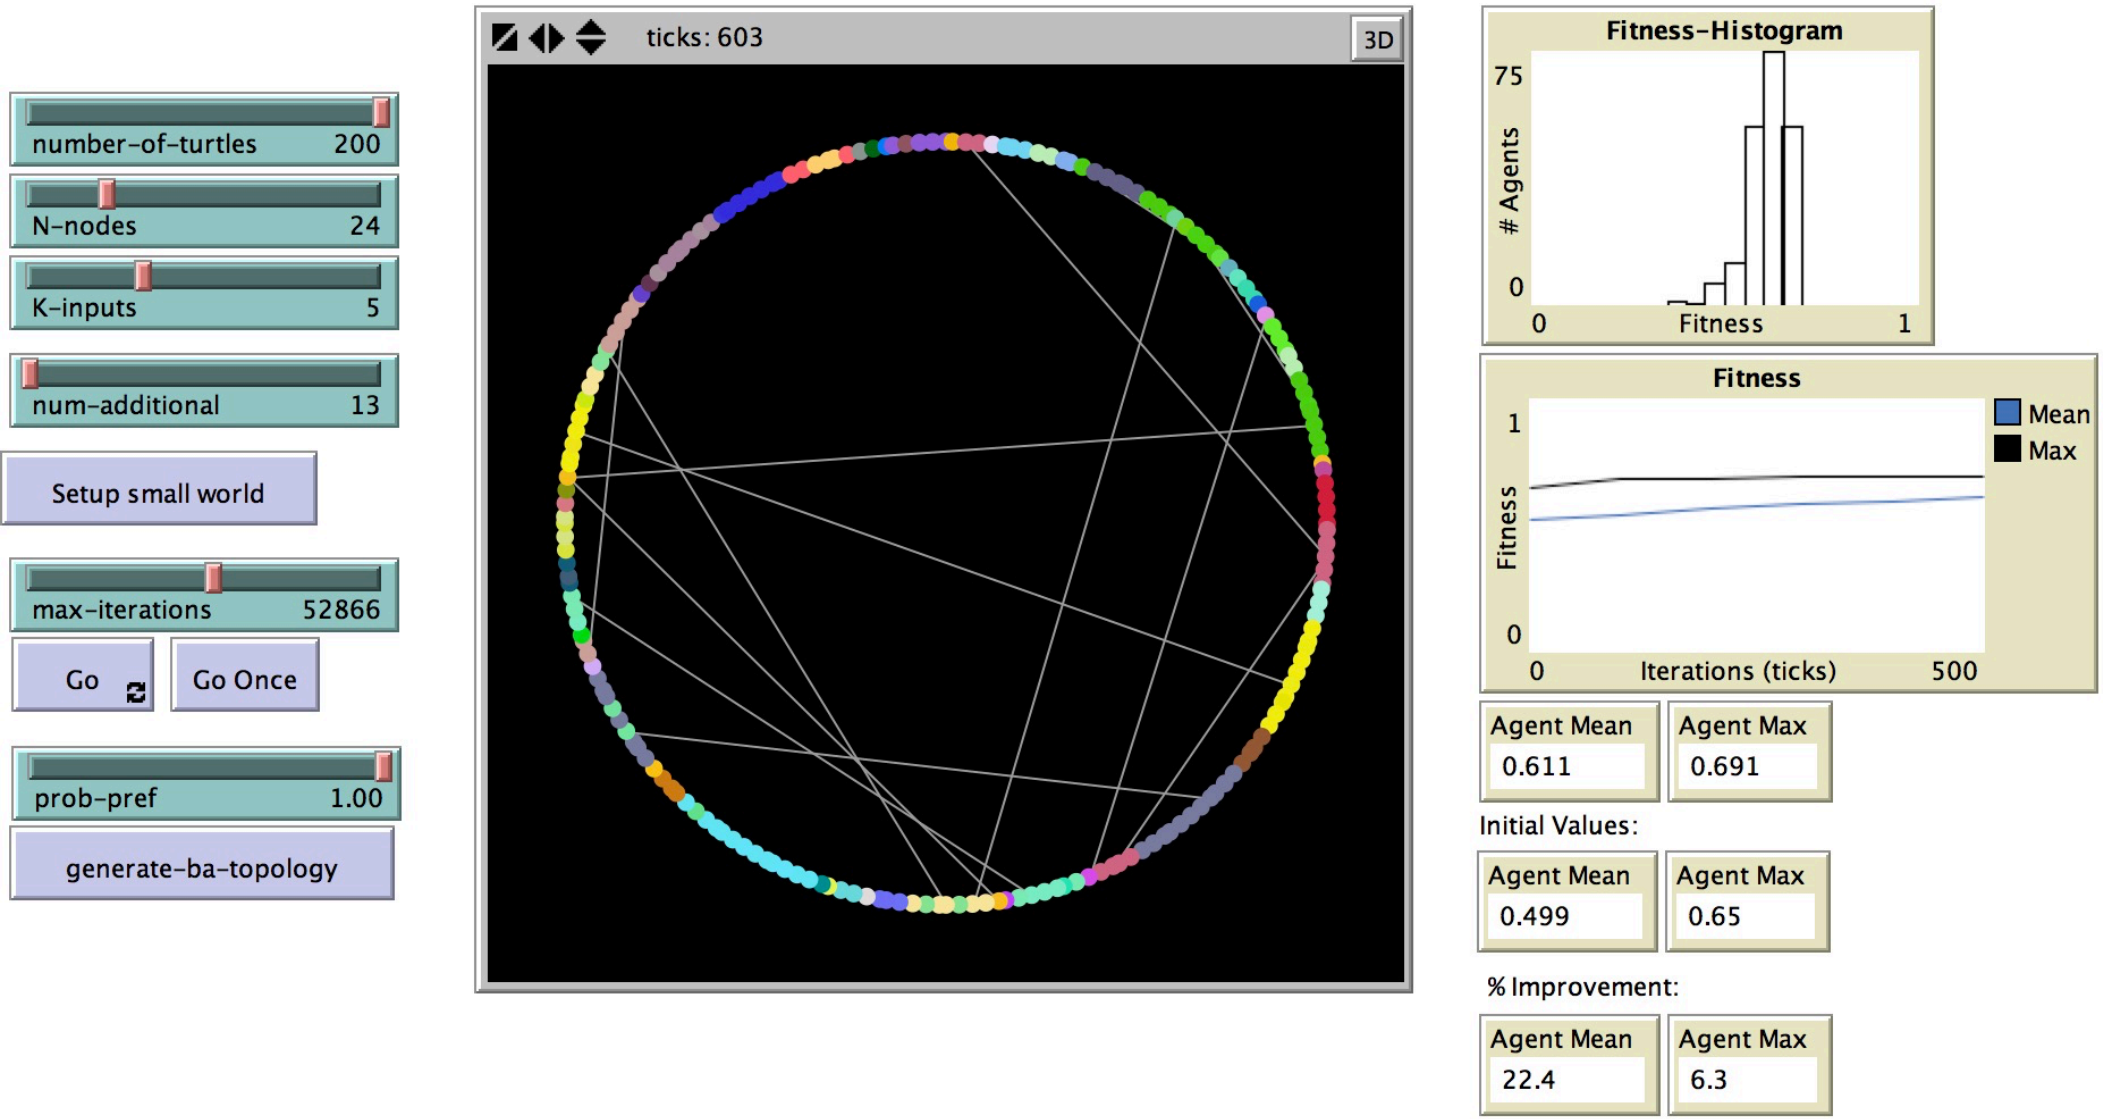

http://web.stanford.edu/class/cs224w/NetLogo/SmallWorldInnovation.nlogo

Quiz Q:

- Relative to the regular lattice, the network with many additional, random connections has on average:
  - slower convergence to a local optimum
  - smaller improvement in the best solution relative to the initial maximum
  - more oscillations between solutions

Coordination: graph coloring

- Application: coloring a map: limited set of colors, no two adjacent countries should have the same color

graph coloring on a network

- Each node is a human subject. Different experimental conditions:
  - knowledge of neighbors' color
  - knowledge of entire network
- Compare:
  - regular ring lattice
  - small-world topology
  - scale-free networks

Kearns et al., 'An Experimental Study of the Coloring Problem on Human Subject Networks', Science, 313(5788), pp. 824-827, 2006

simulation

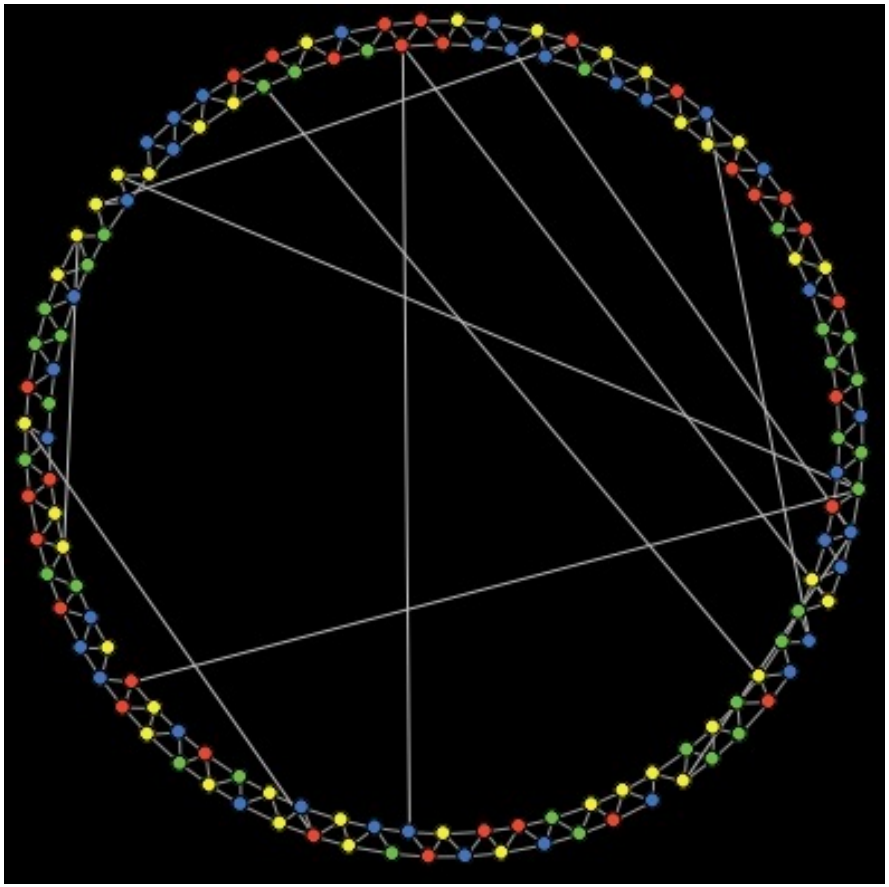

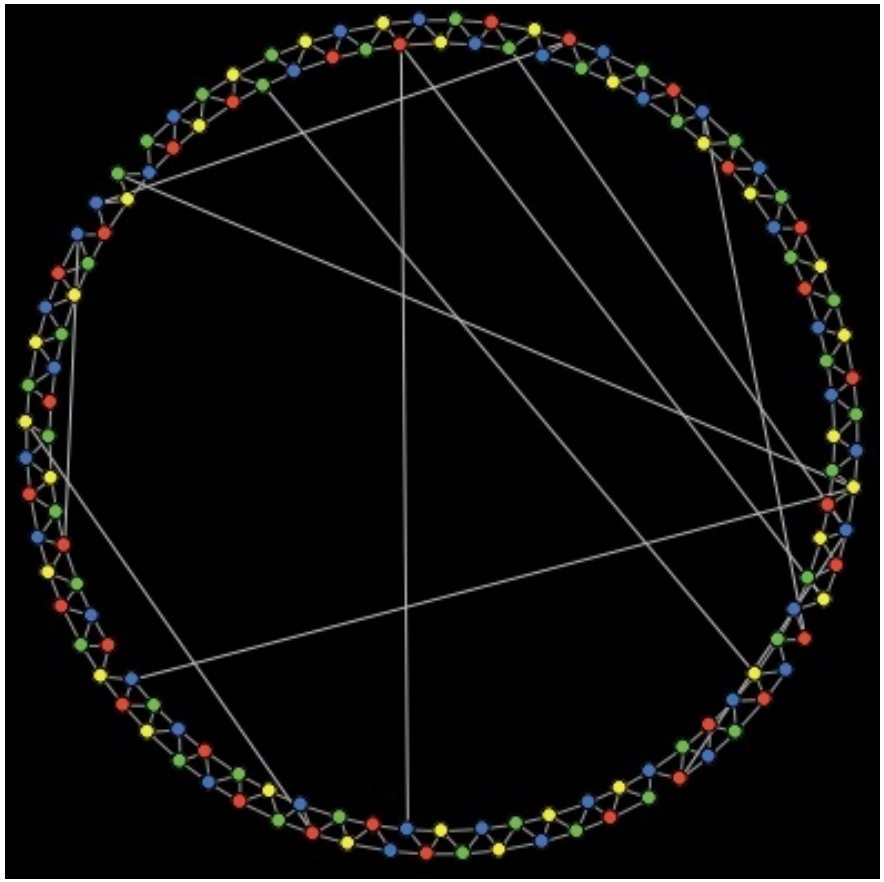

http://web.stanford.edu/class/cs224w/NetLogo/GraphColoring.nlogo

Quiz Q:

- As the rewiring probability is increased from 0 to 1 the following happens:
  - the solution time decreases
  - the solution time increases
  - the solution time initially decreases then increases again

Structural Balance

- **Start with the intuition [Heider '46]:**
  - Friend of my friend is my friend
  - Enemy of enemy is my friend
  - Enemy of friend is my enemy
- Look at connected triples of nodes:

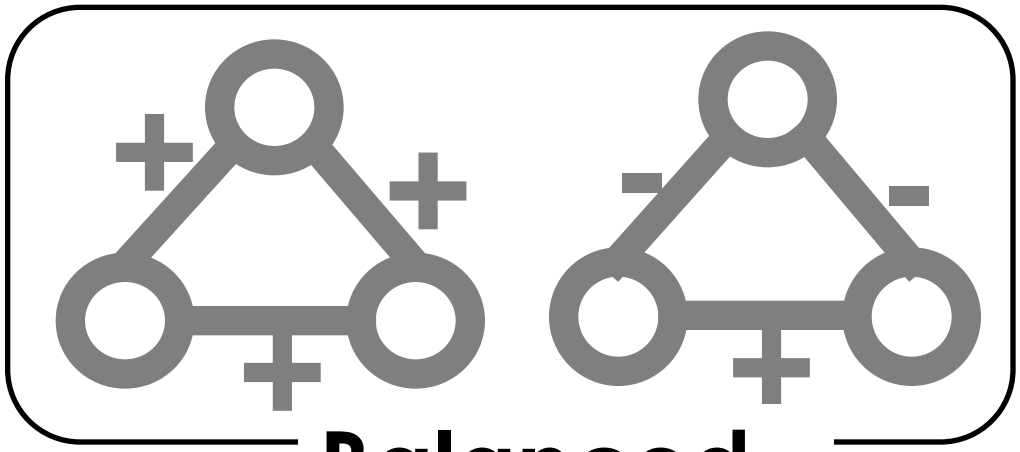

Balanced

Consistent with "friend of a friend" or "enemy of the enemy"

intuition

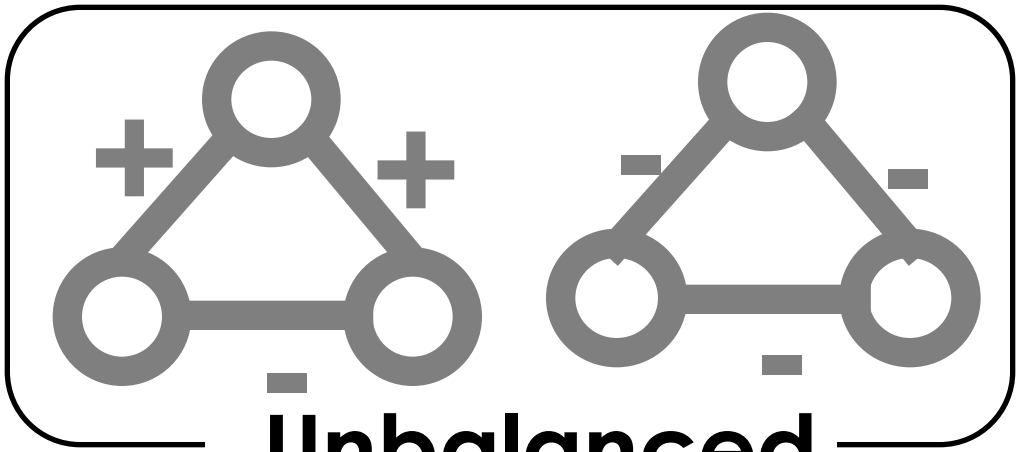

Unbalanced

Inconsistent with the "friend of a friend" or "enemy of the enemy"

intuition

Balance in General Networks

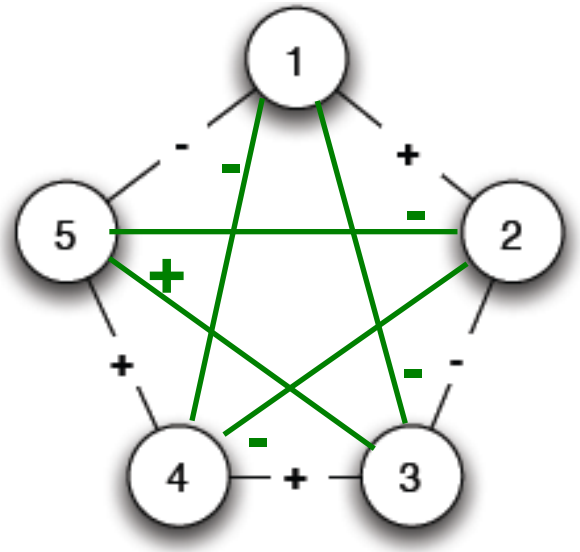

Balanced?

Def 1: Local view

Fill in the missing edges to achieve balance

Def 2: Global view

Divide the graph into two coalitions

The 2 definitions are equivalent!

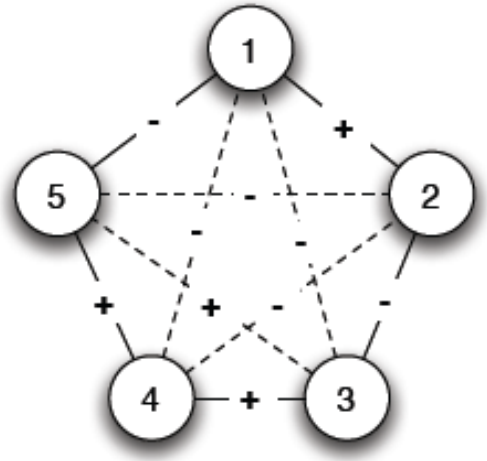

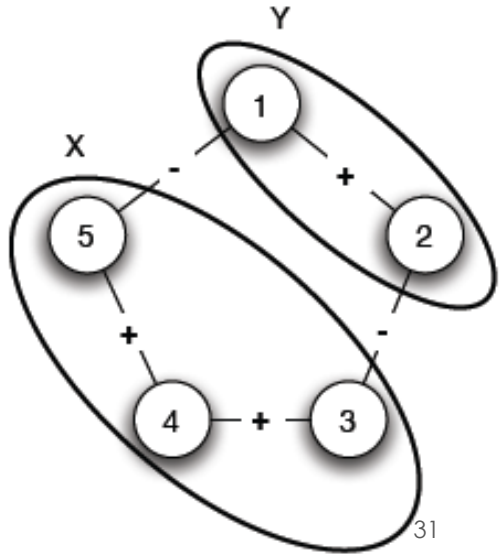

Is a Signed Network Balanced?

- Graph is **balanced** if and only if it contains **no cycle with an odd number of negative** edges

- **How to compute this?**
  - Find connected components on + edges
    - If we find a component of nodes on +edges that contains a – edge ⇒ **Unbalanced**
  - For each component create a super-node
  - Connect components A and B if there is a negative edge between the members
  - Assign super-nodes to sides using BFS

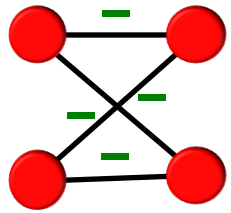

Even length cycle

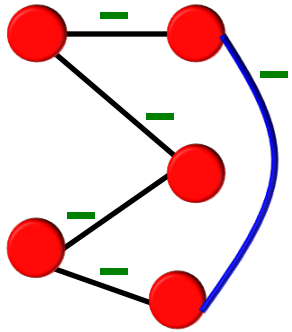

Odd length cycle

Signed Graph: Is it Balanced?

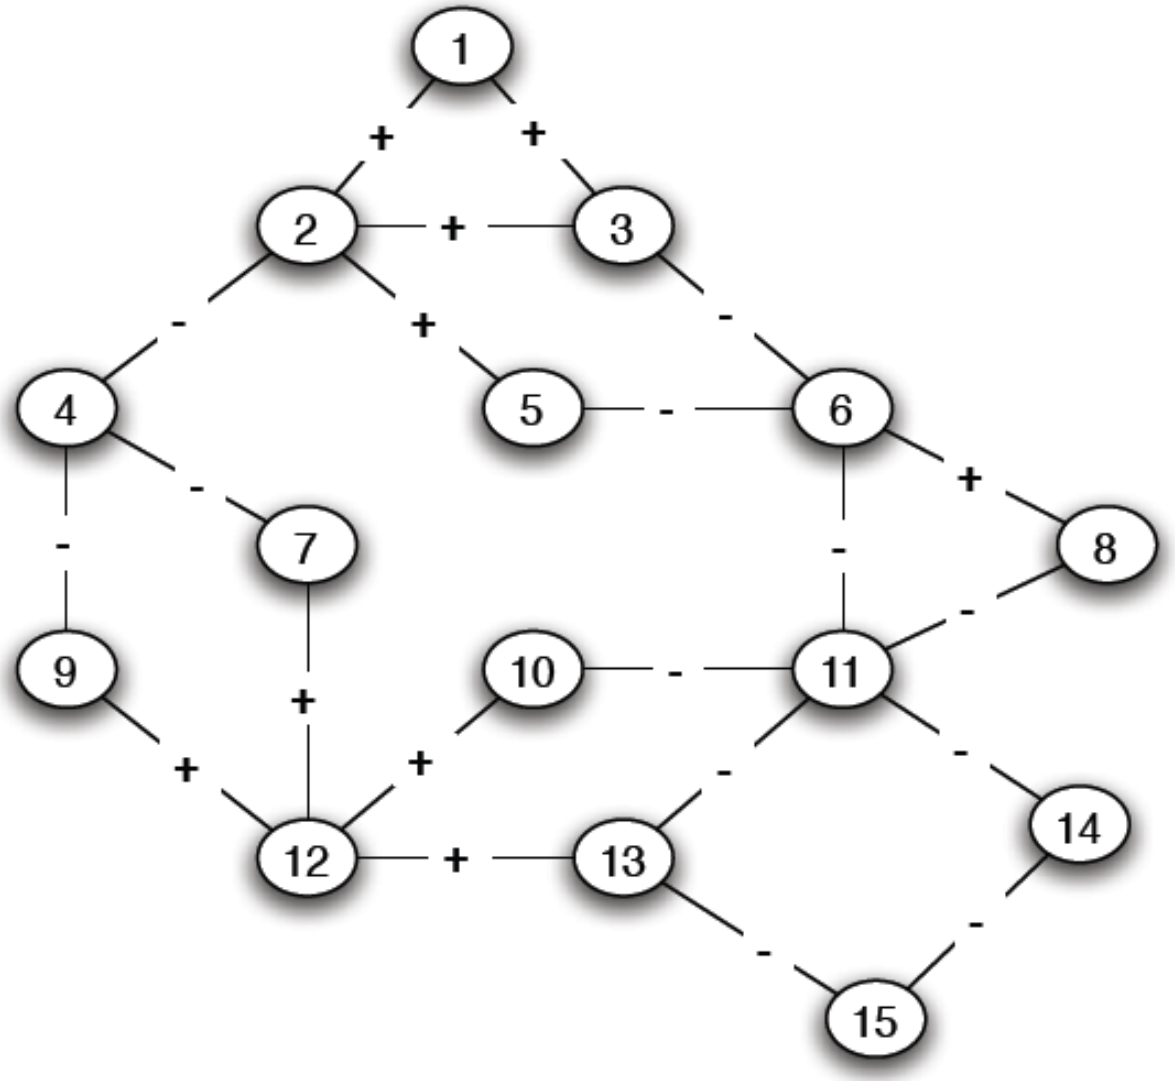

Positive Connected Components

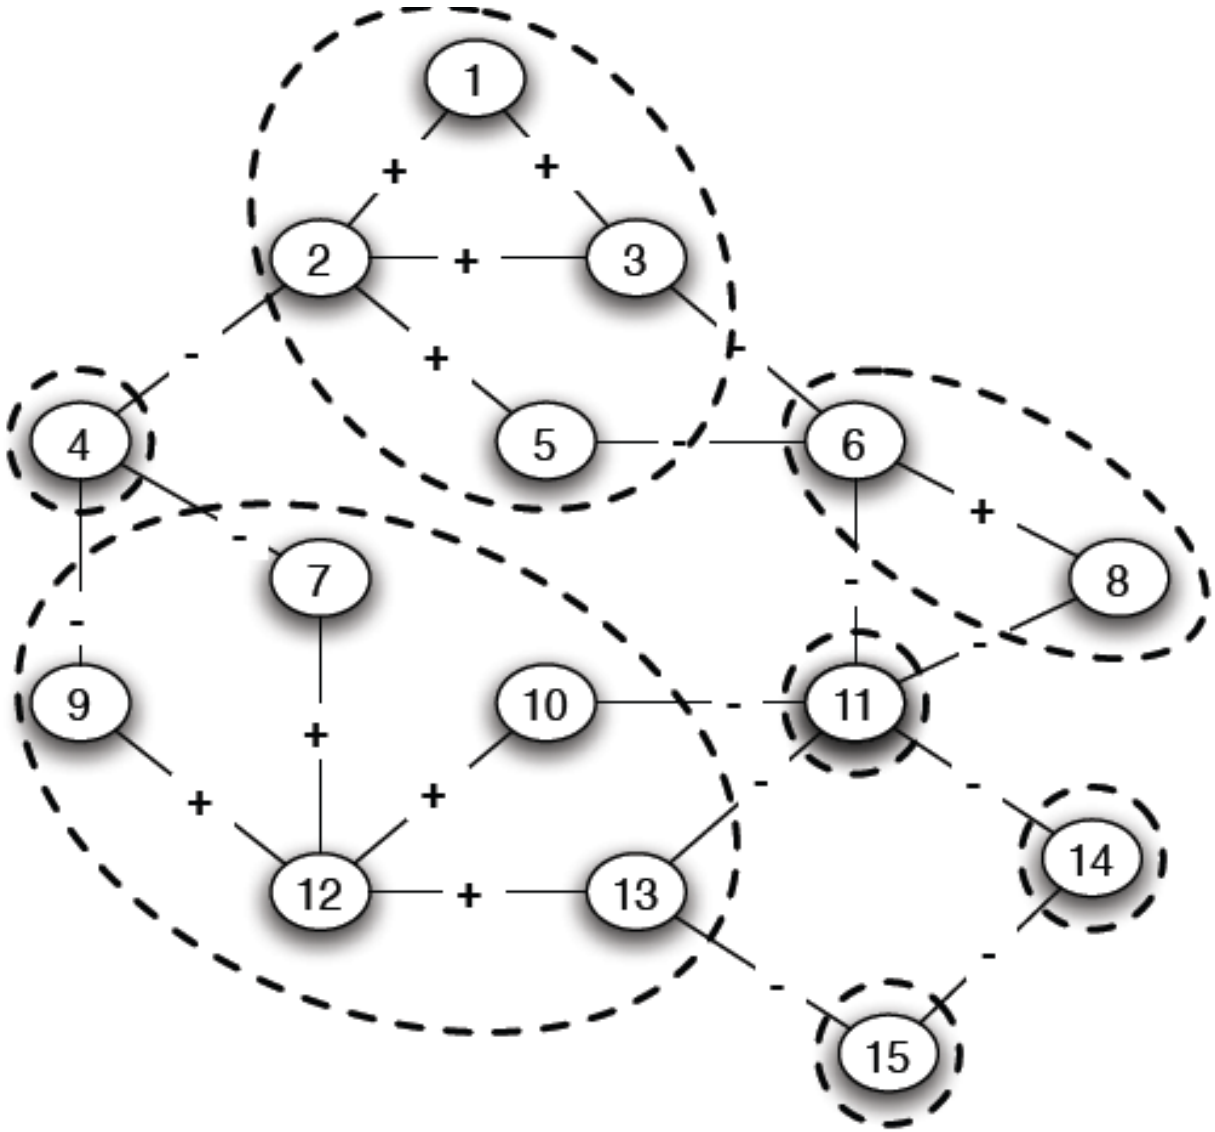

Reduced Graph on Super-Nodes

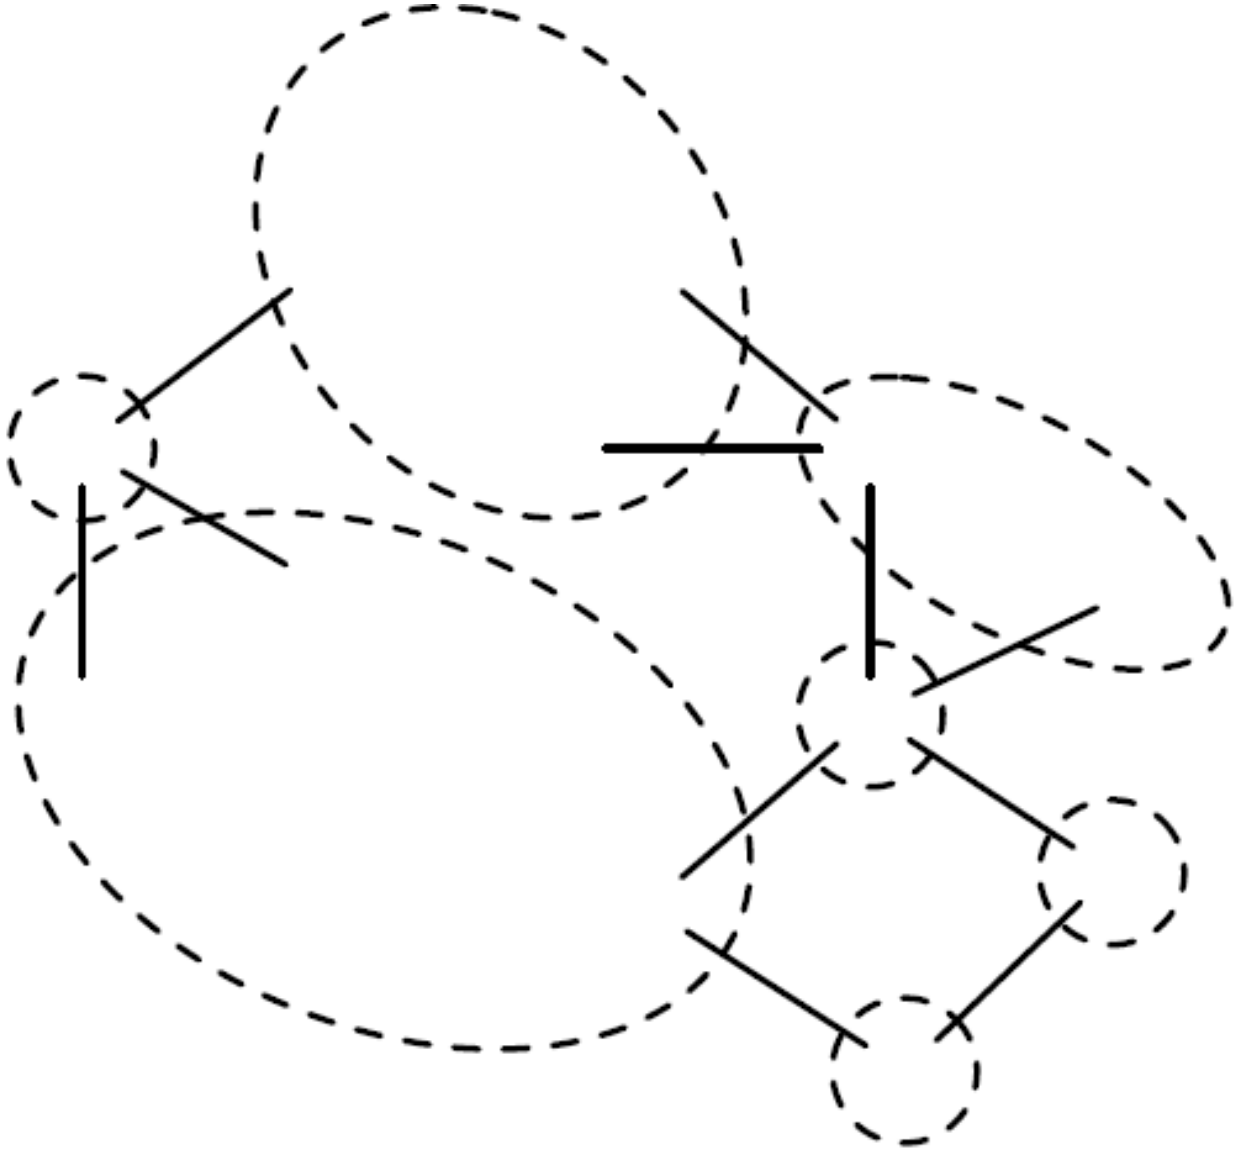

BFS on Reduced Graph

- Using BFS assign each node a **side**
- Graph is **unbalanced** if any two connected super-nodes are assigned the **same side**

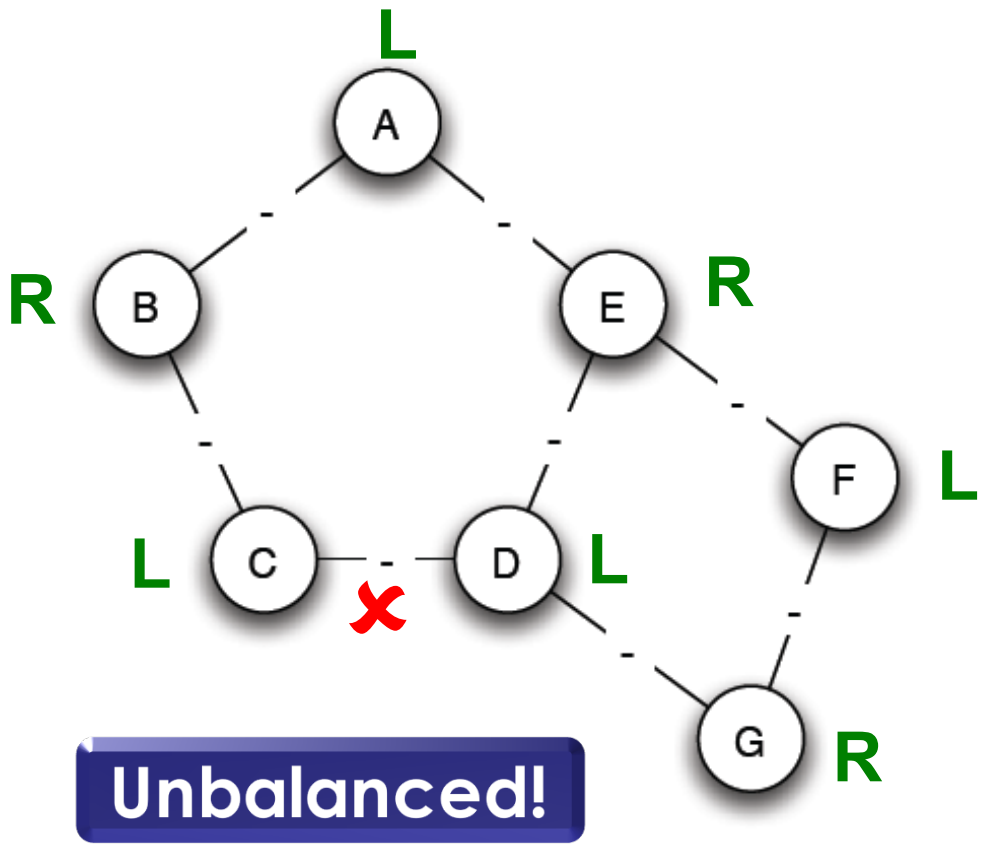

Exploring Real Data

Real Large Signed Networks

- **A→B**
  - **E**pinions: Trust/Distrust
    - Does A trust B's product reviews? (only positive links are visible to users)
  - **W**ikipedia: Support/Oppose
    - Does A support B to become Wikipedia administrator?
  - **S**lashdot: Friend/Foe
    - Does A like B's comments?
  - **Other examples:**
    - Online multiplayer games

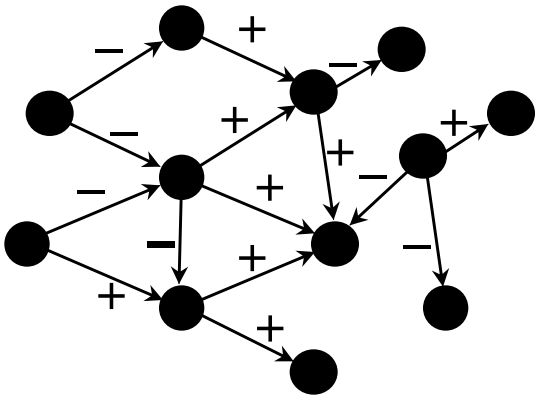

	Epinions	Slashdot	Wikipedia
Nodes	119,217	82,144	7,118
Edges	841,200	549,202	103,747
+ edges	85.0%	77.4%	78.7%
− edges	15.0%	22.6%	21.2%

Balance in Our Network Data

Does structural balance hold?

- Compare frequencies of signed triads in real and "shuffled" signs

	Triad	Epinions		Wikipedia		Consistent with Balance?
		P(T)	$P_0(T)$	P(T)	$P_0(T)$	
Balanced	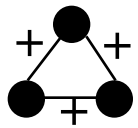	0.87	0.62	0.70	0.49	✓
Unbalanced	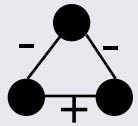	0.07	0.05	0.21	0.10	✓
Balanced	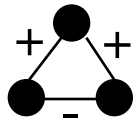	0.05	0.32	0.08	0.49	✓
Unbalanced	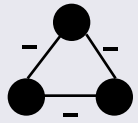	0.007	0.003	0.011	0.010	✗

P(T) … fraction of a triads

$P_0(T)$… triad fraction if the signs would appear at random

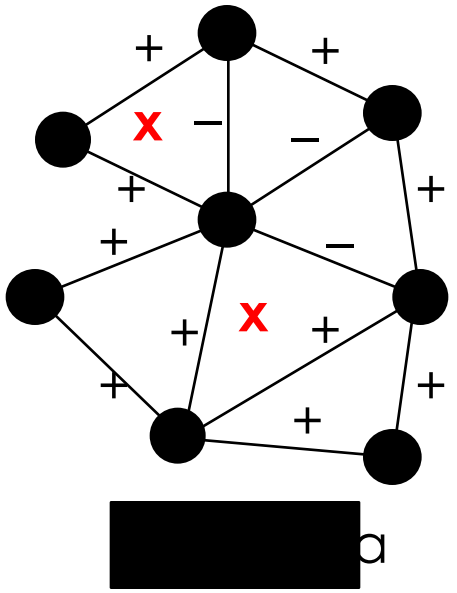

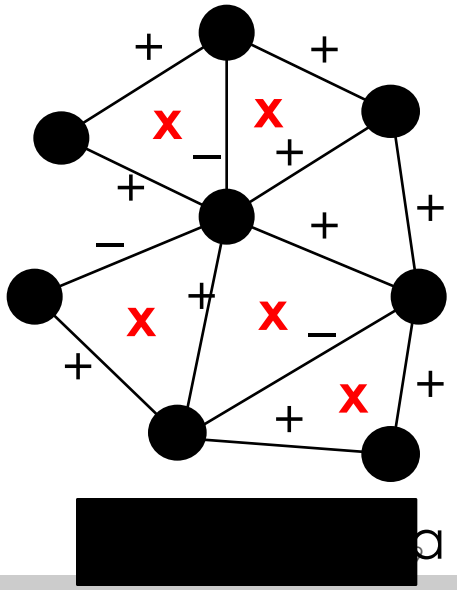

Global Structure of Signed Nets

- Intuitive picture of social network in terms of densely linked clusters

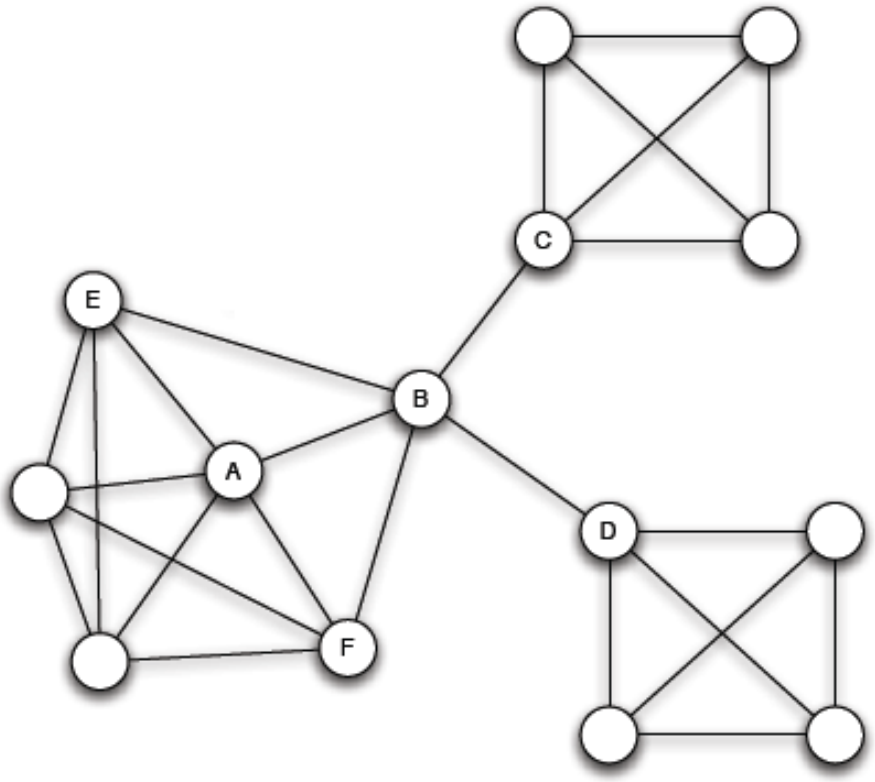

- **How does structure interact with links?**

- **Embeddedness of link (A,B):** Number of shared neighbors

Global Factions: Embeddedness

- **Embeddedness of ties:**
  - Positive ties tend to be **more** embedded

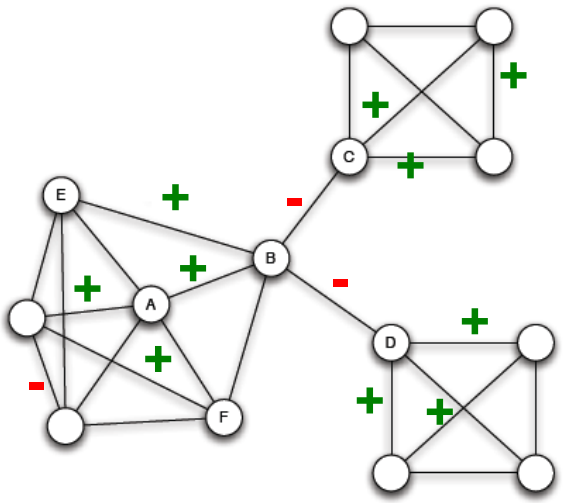

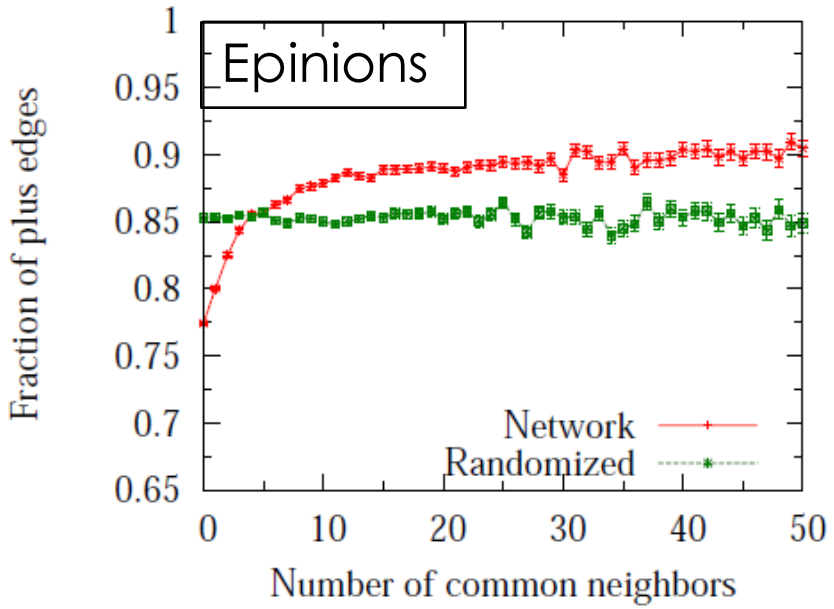

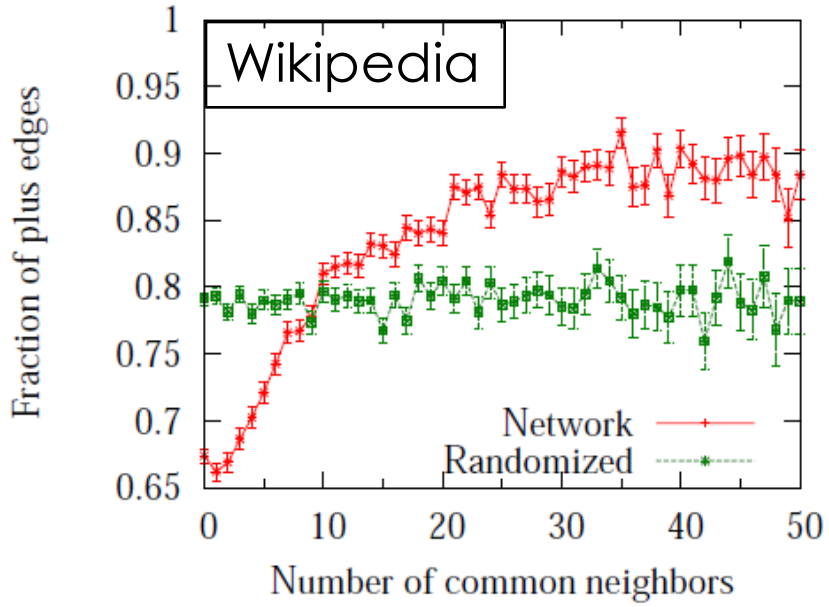

Evolving Directed Networks

- **New setting:** Links are **directed**, created over time
  - Node **A** links to **B**
  - Directions and signs of links from/to X provide context

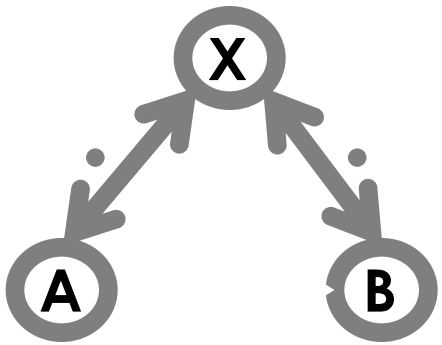

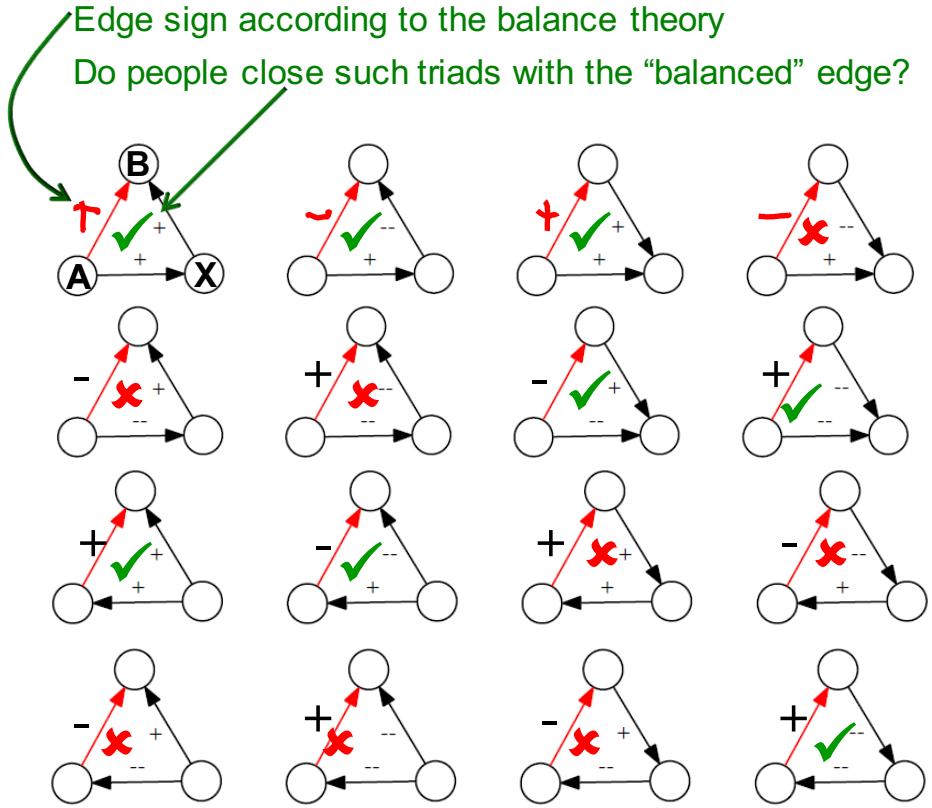

16 signed directed triads

(in directed networks people traditionally applied balance by ignoring edge directions)

- **How many △ are now explained by balance?**
  - (8 out of 16)

Alternate Theory: Status

- **Status** in a network [Davis-Leinhardt '68]
  - **A** $\xrightarrow{+}$ **B** :: **B** has **higher** status than **A**
  - **A** $\xleftarrow{-}$ **B** :: **B** has **lower** status than **A**
    - Note the notion of status is now implicit and governed by the network (rather than the number of edits)
  - **Apply this principle transitively over paths**
    - Can replace each **A** $\xrightarrow{-}$ **B** with **A** $\xleftarrow{+}$ **B**
    - Obtain an all-positive network with same status interpretation

Status vs. Balance

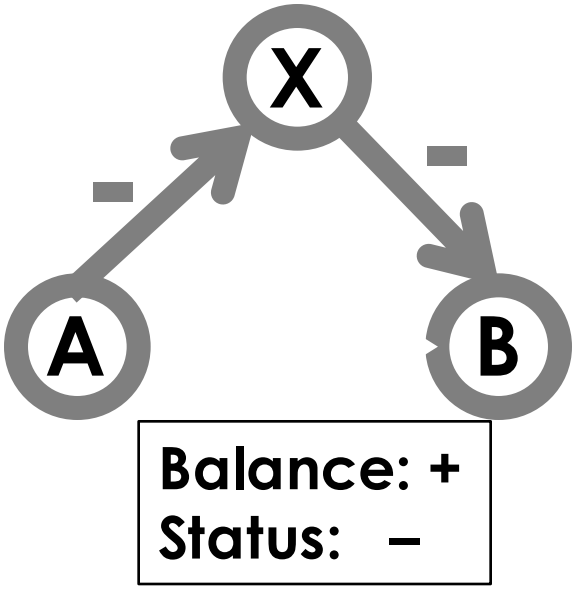

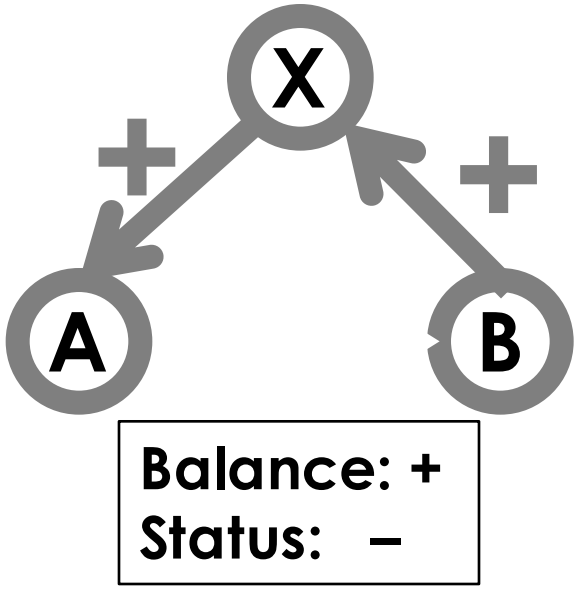

Status and balance give different predictions!

Status vs. Balance

At a global level (in the ideal case):

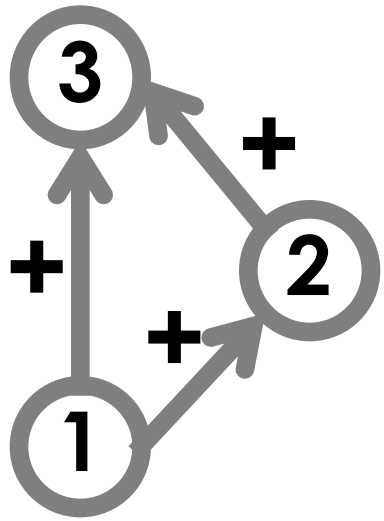

- **Status ⇒ Hierarchy**
  - All-positive directed network should be approximately **acyclic**

- **Balance ⇒ Coalitions**
  - Balance ignores directions and implies that subgraph of negative edges should be approximately **bipartite**

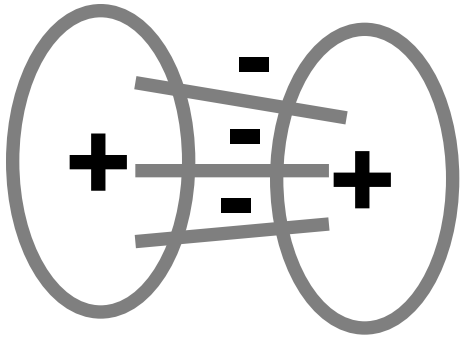

Theory of Status

- **Edges are directed:**
  - X has links to A and B
  - Now, A links to B (triad A-B-X)
  - **How does sign of A->B depend signs from/to X?**
    P(A->B | X) vs. P(A->B)
    + **+**

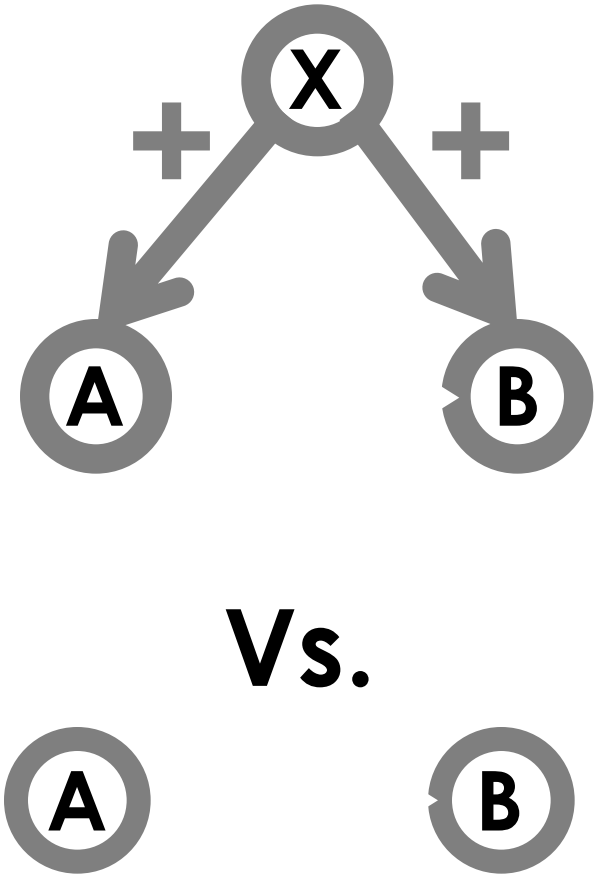

Links are embedded in triads

- Link A->B appears in context X:
  A->B | X

- 16 possible contexts:

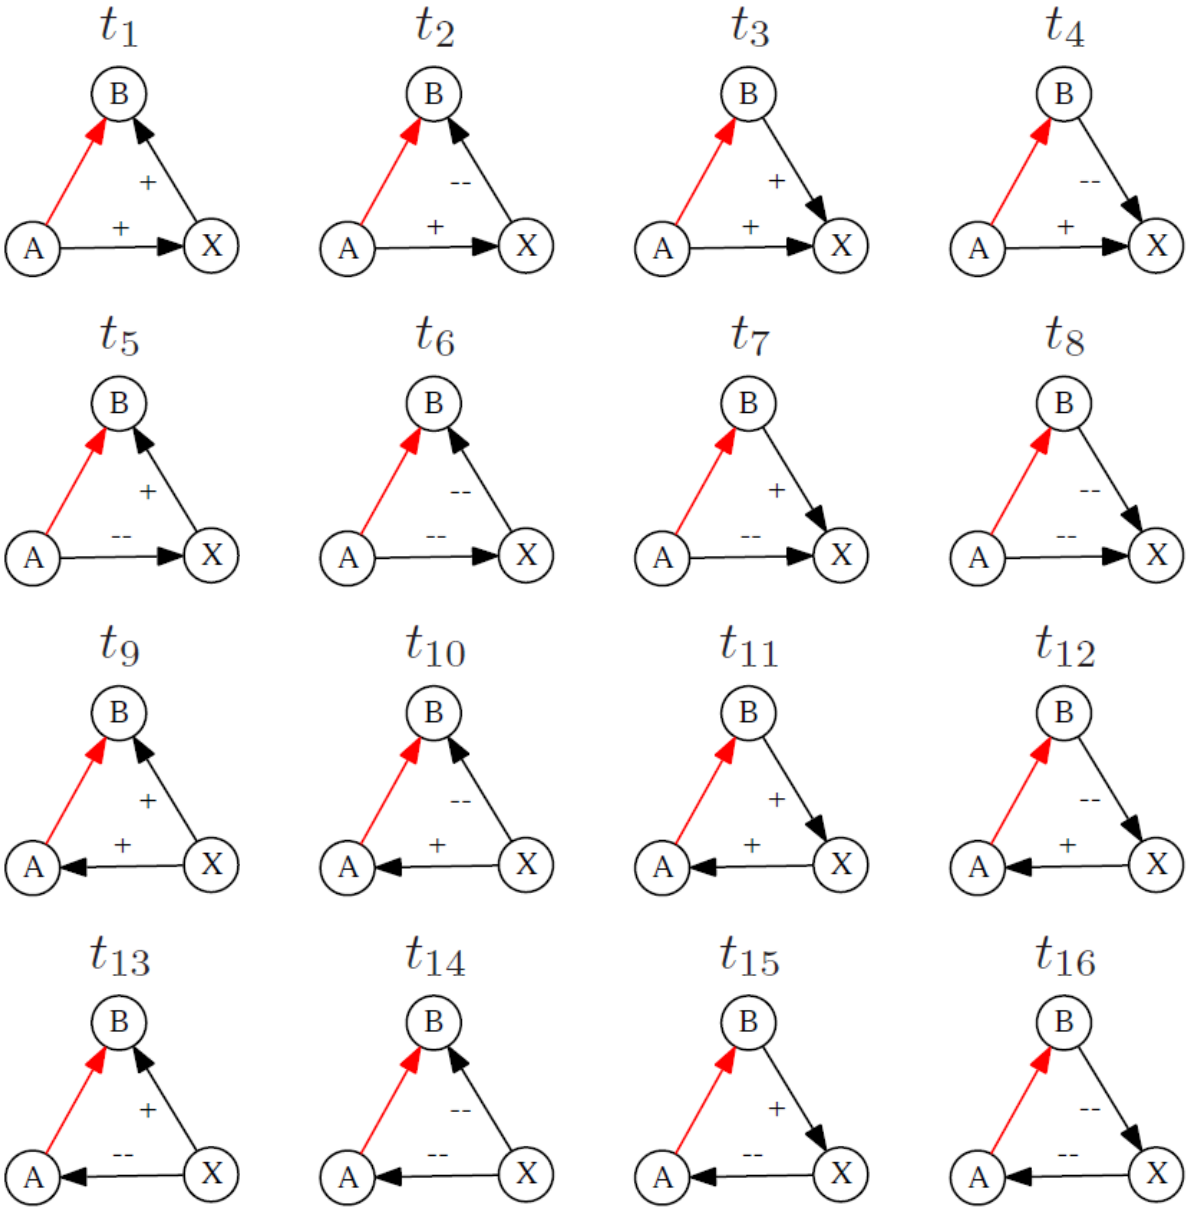

Note: Context of a link is uniquely determined by the directions and signs of links from/to **X**

Predicting Edge Signs

Edge sign prediction problem

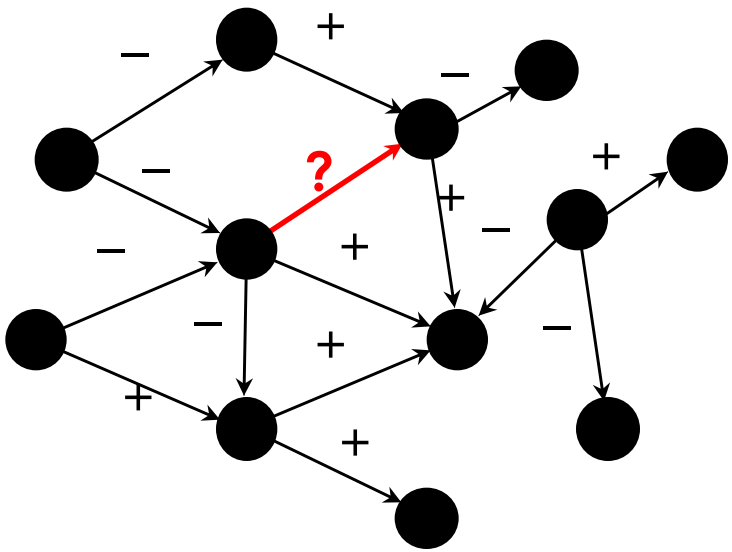

- Given a network and signs on all but one edge, predict the missing sign
- **Friend recommendation:**
  - Predicting whether you know someone vs. Predicting what you think of them
- **Setting:**
  - **Given edge *(A,B)*, predict its sign:**
  - **Let's look at signed triads *(A,B)* belongs to:**

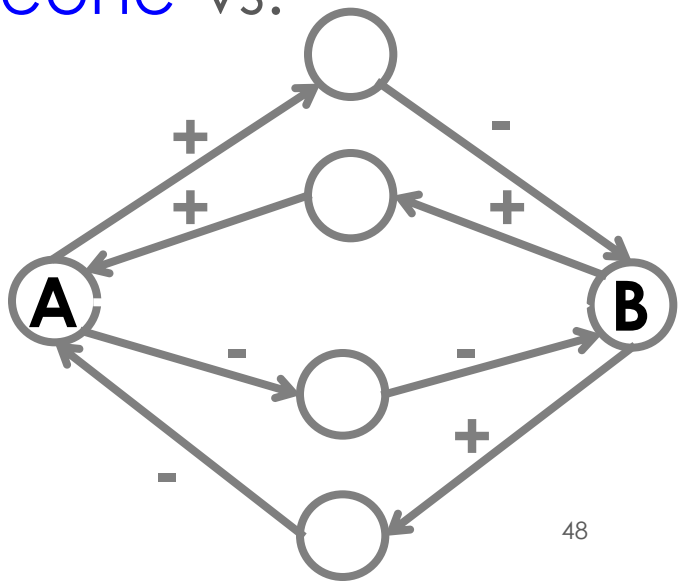

Features for Learning

For the edge (A,B) we examine

Its network context:

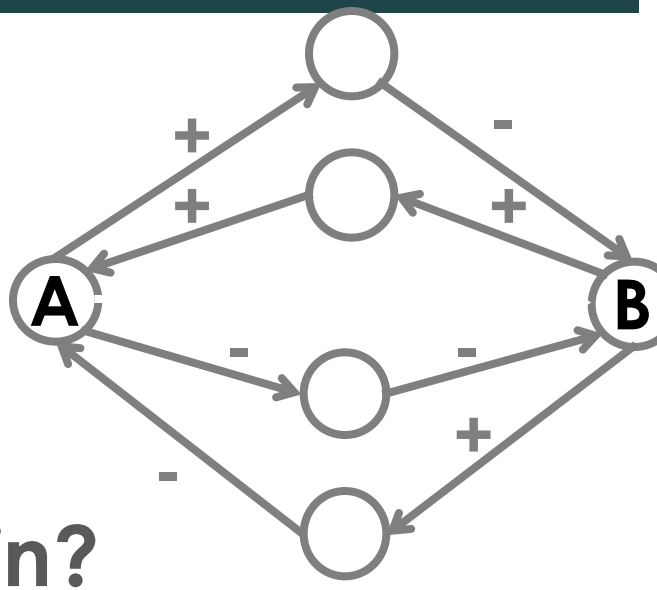

- **In what types of triads does our red-edge participate in?**

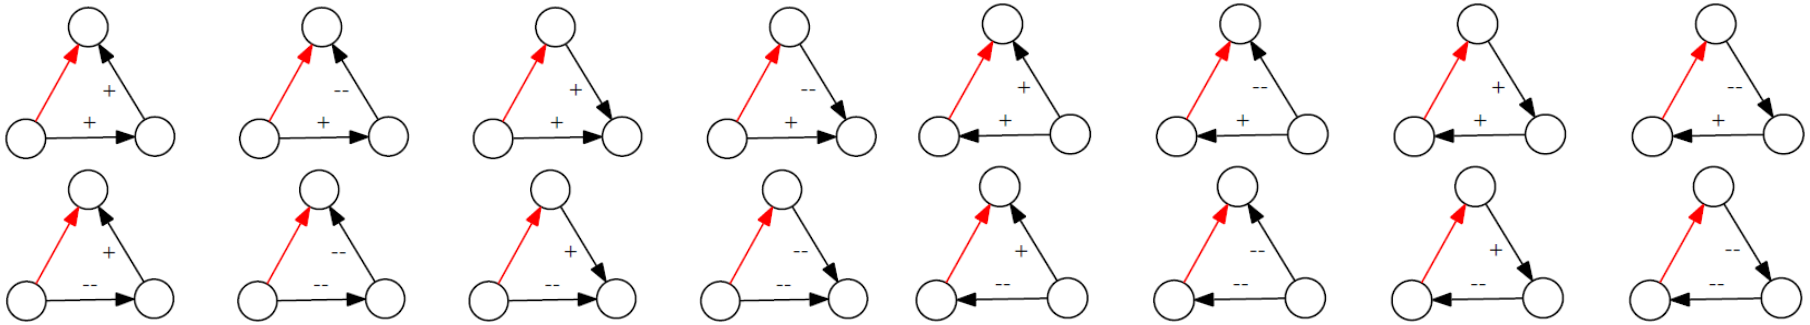

- Each triad then "votes" and we determine the sign

Balance and Status: Complete Model

Triad	Bal	Stat
const		
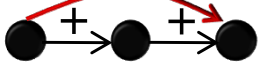	1	1
	-1	0
	-1	0
	1	-1
	1	0
	-1	1
	-1	-1
	1	0
	1	0
	-1	-1
	-1	1
	1	0
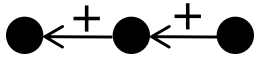	1	-1
	-1	0
	-1	0
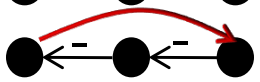	1	1

Balance and Status: Complete Model

Triad	Bal	Stat	**Epin**	**Slashd**	**Wikip**
const			-0.2	0.02	-0.2
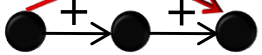	1	1	**0.5**	**0.9**	**0.3**
	-1	0	**-0.5**	**-0.9**	**-0.4**
	-1	0	**-0.4**	**-1.1**	**-0.3**
	1	-1	**-0.7**	**-0.6**	**-0.8**
	1	0	0.3	0.4	0.05
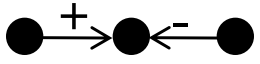	-1	1	-0.01	-0.1	-0.01
	-1	-1	**-0.9**	**-1.2**	**-0.2**
	1	0	0.04	-0.07	-0.03
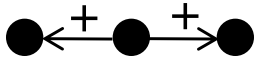	1	0	0.08	0.4	0.1
	-1	-1	**-1.3**	**-1.1**	**-0.4**
	-1	1	-0.1	-0.2	0.05
	1	0	0.08	-0.02	-0.1
	1	-1	-0.09	-0.09	-0.01
	-1	0	-0.05	-0.3	-0.02
	-1	0	-0.04	-0.3	0.05
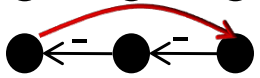	1	1	-0.02	0.2	-0.2

Edge Sign Prediction

Prediction accuracy:

	Balance	Status	Triads
Epinions	80%	82%	93.5%
Slashdot	84%	72%	94.4%
Wikipedia	64%	70%	81%

Observations:

- Signs can be modeled from local network structure alone!
- Triad counts perform less well for less embedded edges (E_m)
  - Wikipedia is harder to model:
    - Votes are publicly visible

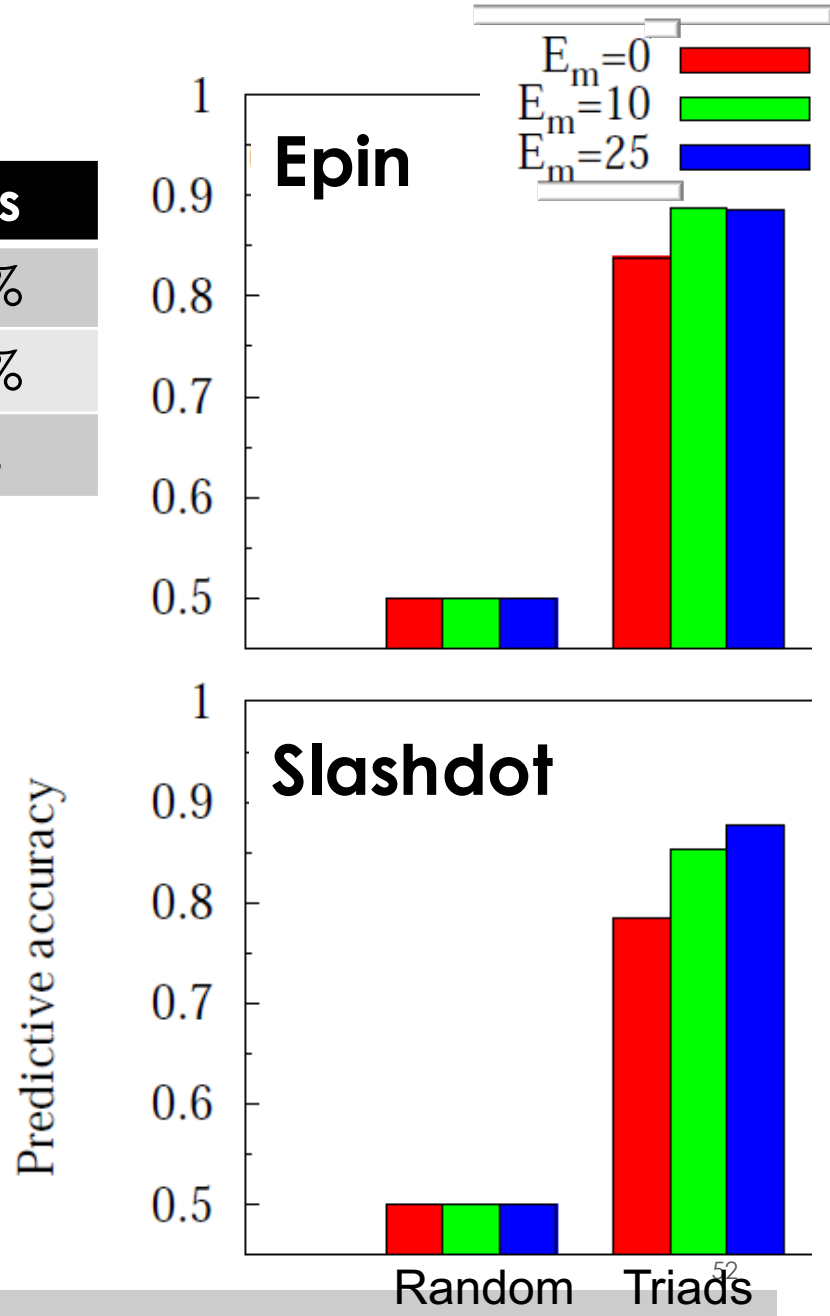

Generalization

- **Do people use these very different linking systems by obeying the same principles?**
  - How generalizable are the results across the datasets?
    - Train on row "dataset", predict on "column"

Train on row, test on column	Epinions	Slashdot	Wikipedia
Epinions	0.9342	0.9289	0.7722
Slashdot	0.9249	0.9351	0.7717
Wikipedia	0.9272	0.9260	0.8021

- Nearly **perfect generalization** of the models even though networks come from very different applications!

Summary: Signed Networks

- **Signed networks provide insight into how social computing systems are used:**
  - Status vs. Balance
  - Role of embeddedness and public display
  - More evidence that networks are globally organized based on status
- **Sign of relationship can be reliably predicted from the local network context**
  - ~90% accuracy sign of the edge
  - People use signed edges consistently regardless of particular application
    - Near perfect generalization of models across datasets

Why online reputation systems matter

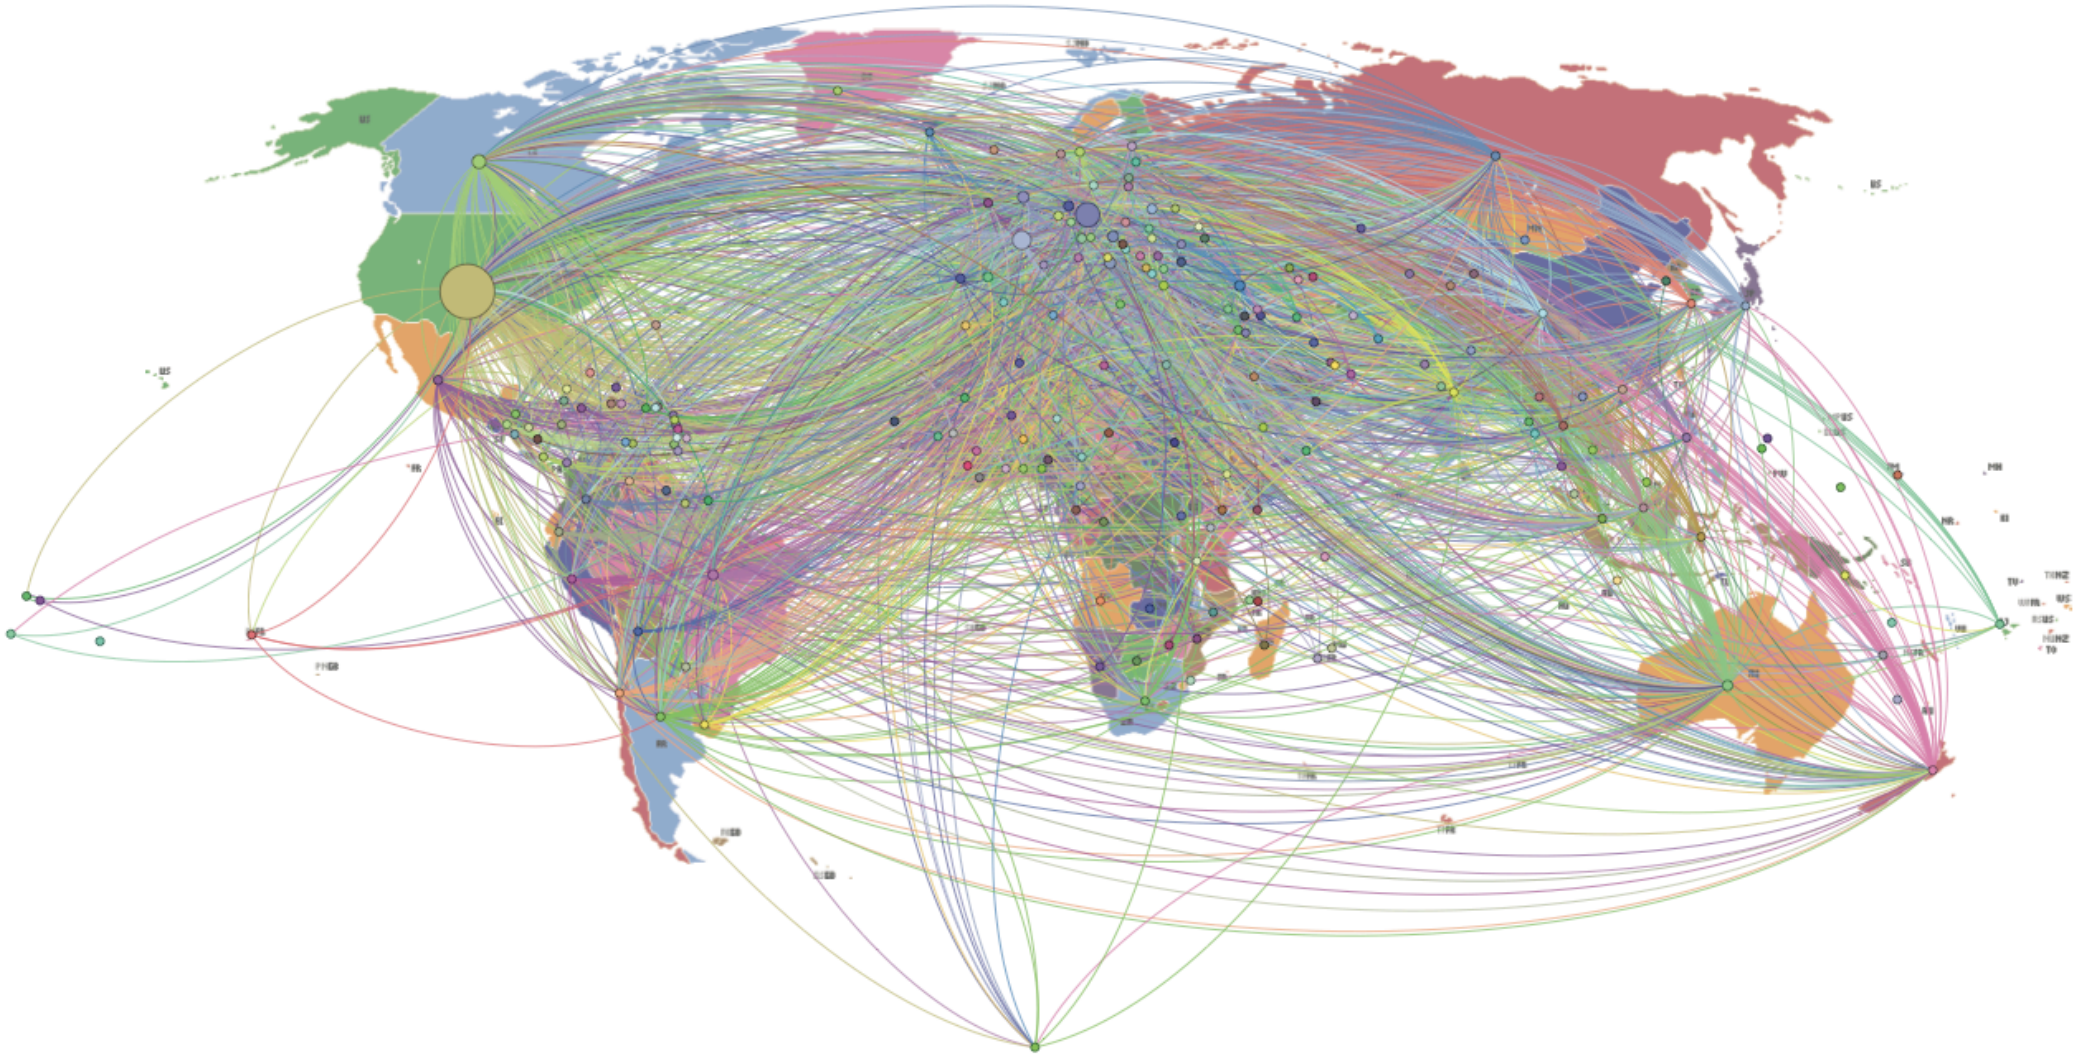

IEEE SocialCom-09/SIN-09

hospitality exchange on CouchSurfing.org

Instant trust: host or surf from the start

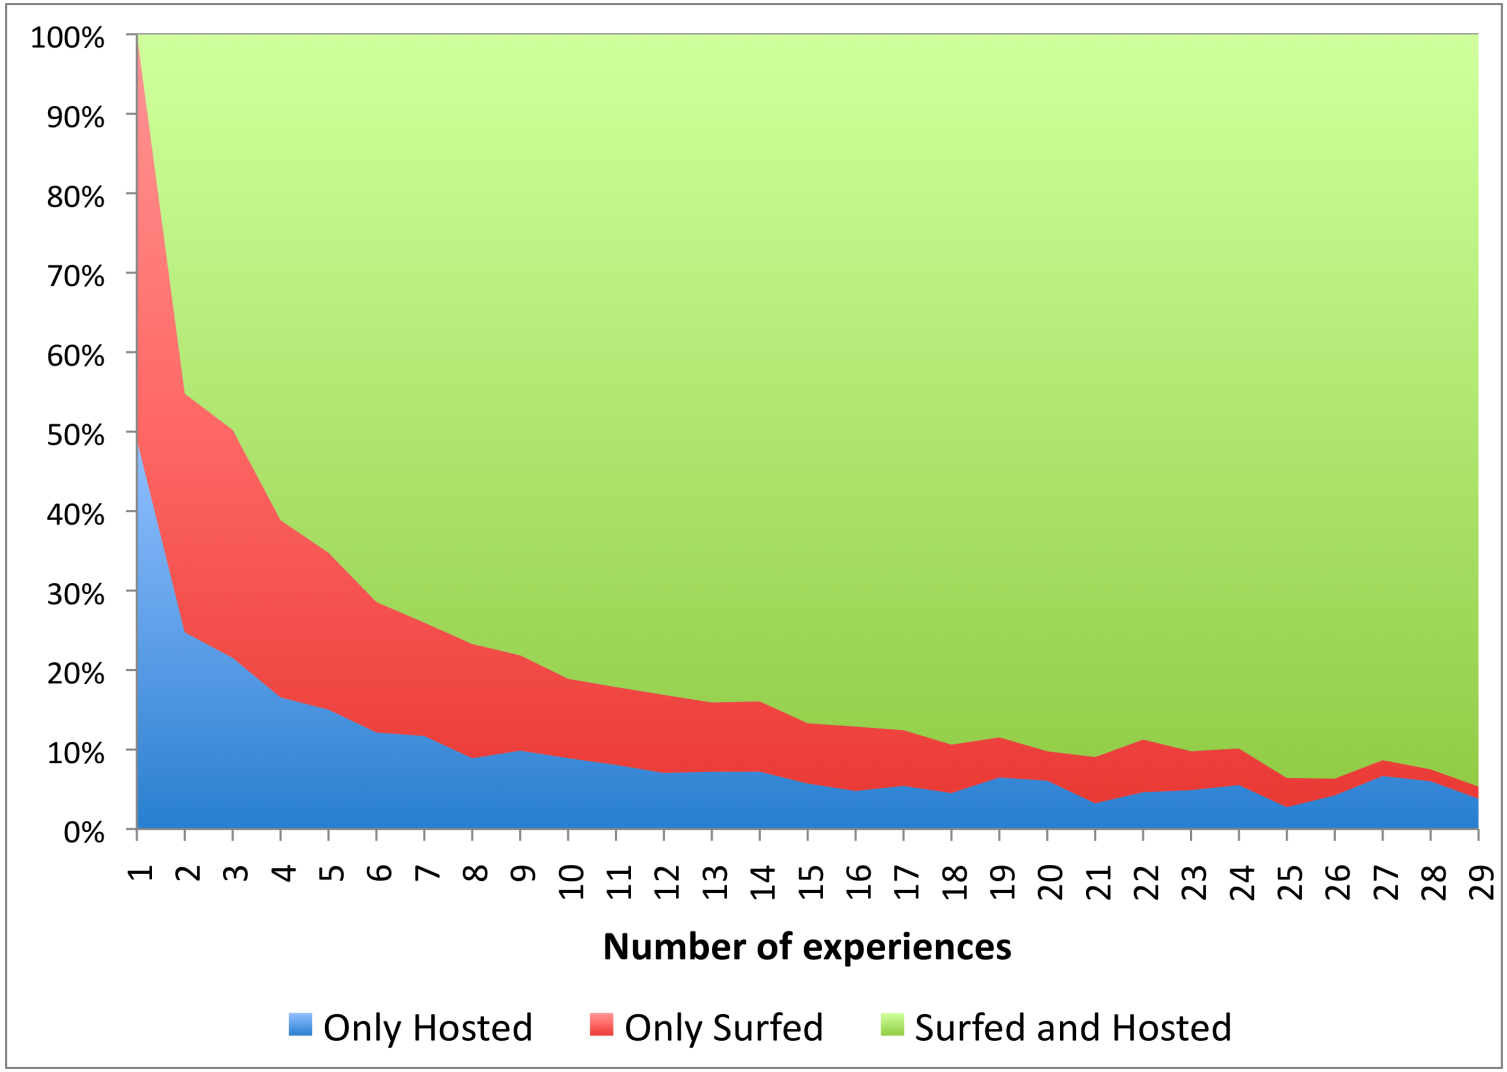

How can this be?

- *It's the exact same [] as if you had just met, casually met somebody while you were traveling and they said, "You know, you could stay on my couch if you want".* [P9]
- *We arrived in Brussels at 9 o'clock in the morning and the guy that we were staying with, our host, had to come to his work and he just handed his house keys right over to us and said, ``This is where I live. Go here, go here and my wife and I won't be home until after 5."*

Direct reciprocity

- 12-18% of stays are directly reciprocated

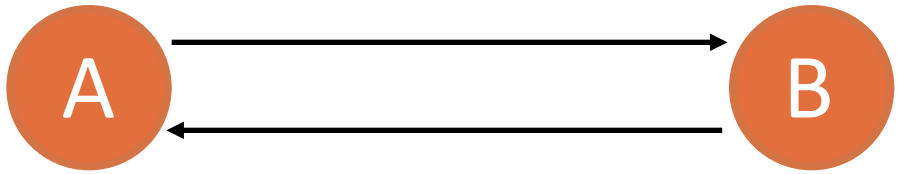

Generalized reciprocity

- Who hosted whom:
  - Largest strongly connected component (1/3 of active users)

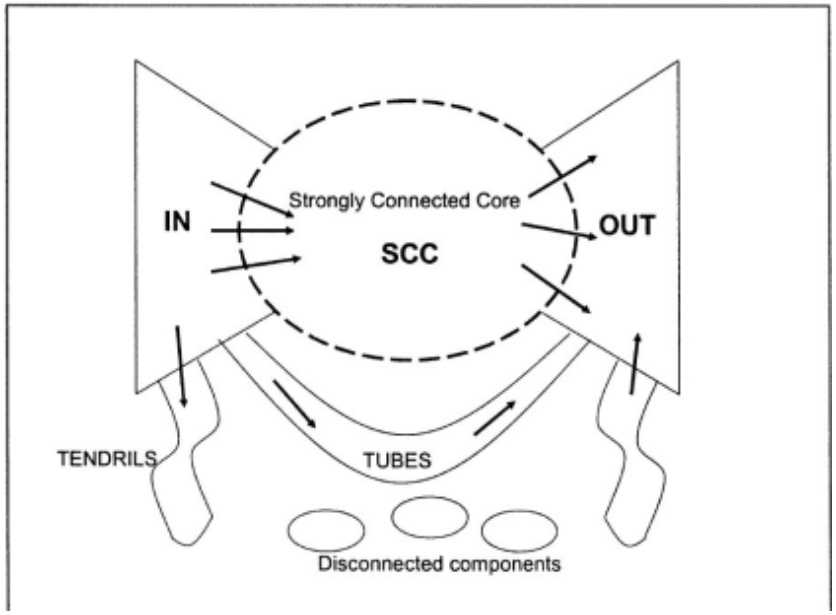

But how are ratings influenced by design choices?

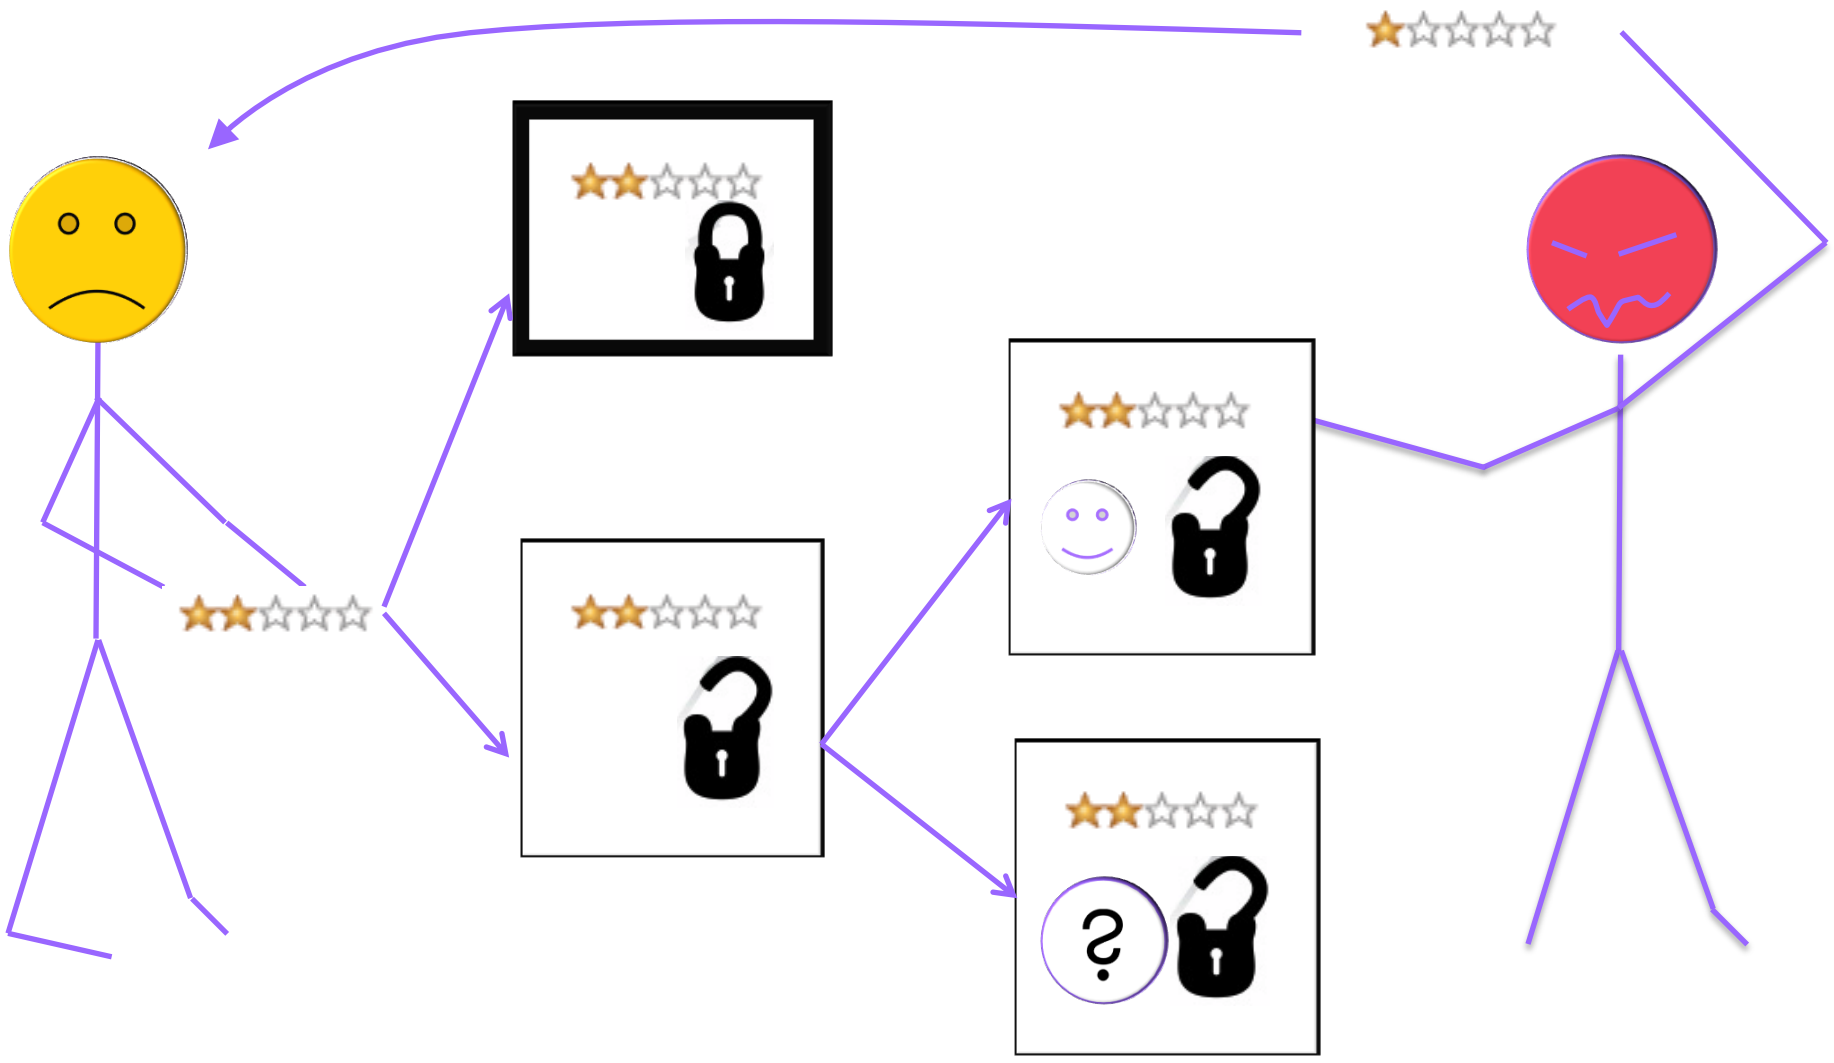

Research Questions

- How do design choices in online social networking & recommendation sites influence ratings?
- Are there other factors affecting ratings?
- Can friendship serve as a proxy for trust?
- What are trust and friendship anyway?

Datasets: Who is rating what/whom?

Amazon

- Amazon.com provides a platform which allows users to review products
- Users can decide to use a pen name or real name to review products
- 15 thousand reviews from top 1500 reviewers (about one half using pen names)

Effect of anonymity in absence of reciprocity

amazon.com

attribute	pen name	REAL NAME™	statistically significant
product rating # stars	4.19	4.21	no
# reviews	498	551	yes
length of review (words)	364	377	yes
# of fan voters	28.6	37.1	yes

PUBLIC | **ANONYMOUS** | ~~**RECIPROCAL**~~

Epinions

- Epinions.com allows users to share product reviews.
- Users can write reviews, rate other users' reviews, and specify which users they "trust" or "distrust"
  - ~800K user-to-user ratings (trust or not)
  - ~100K users and 3 million articles

Anonymity when identity invites reciprocity

Epinions.com

- Anonymous ratings are lower (3.84) on average than identified ratings (4.71)
- For the same user, anonymous ratings still average lower (4.01) than identified ones (4.76)

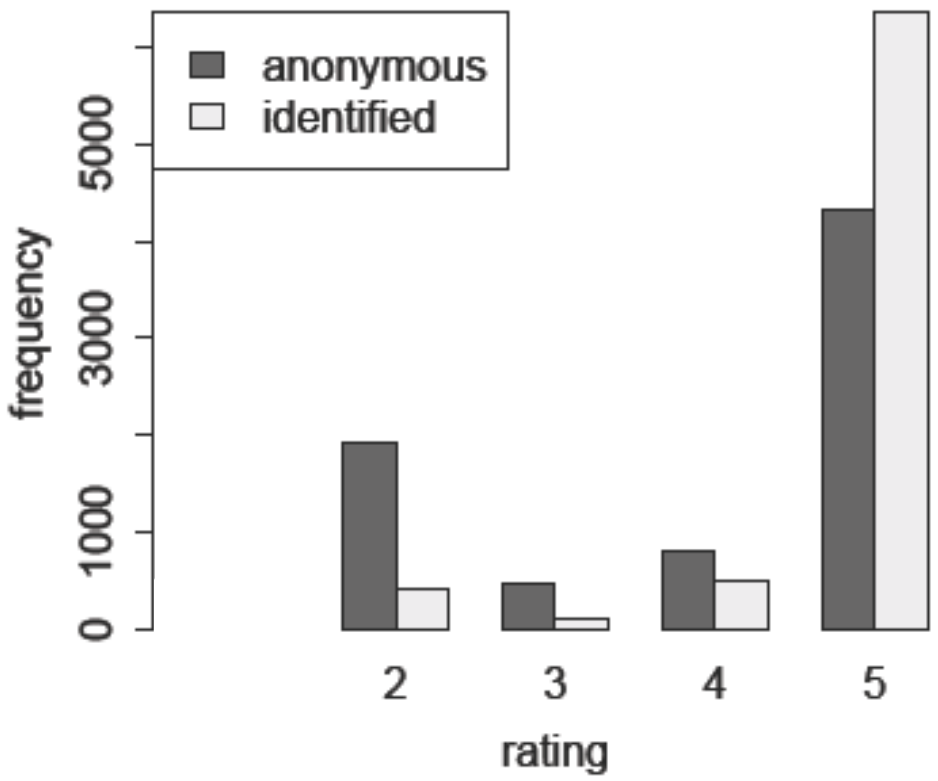

PUBLIC	ANONYMOUS	RECIPROCAL

Evidence of reciprocity?

- We average multiple user-to-article ratings into user-to-user ratings.
- Rating from A to B is correlated with rating from B to A ($\rho = 0.48$)
- # of ratings from A to B and B to A also correlated ($\rho = 0.49$)
- ***anonymously*** given ratings between users have much ***lower*** correlation ($\rho = 0.14$)

PUBLIC	ANONYMOUS	RECIPROCAL

Privacy enables negative ratings

- Epinions allows users to "trust" others publicly, but "distrust" privately
- non-trivial fraction (14.7%) are "distrust" ratings.
- For pairs of users who mutually rated one another (35% of public trust, and 6% of private distrust were reciprocated):

A->B, B->A	observed	expected
trust, trust	97.1%	72.8%
trust,distrust	1.1%	25.1%
distrust,distrust	1.8%	2.2%

~~PUBLIC~~ ~~ANONYMOUS~~ RECIPROCAL

CouchSurfing

- data: 600K users, 3 million edges
- ~ 500 survey respondents
- 18 interviews

- Users can do the following for other users:
  - specify friendship level (e.g. acquaintance, friend, best friend)
  - specify how much they trust them (e.g. "somewhat", "highly")
  - vouch for them
  - leave positive, neutral, or negative references

PUBLIC

PRIVATE

PUBLIC

PUBLIC

scarcity of public, negative ratings when identified

- Users leave a positive reference for 87.7 % of those they host and for 90.1% of those who host them
- Neutral/missing references are confounded in data
- The ratio of positive to negative references is 2500:1!

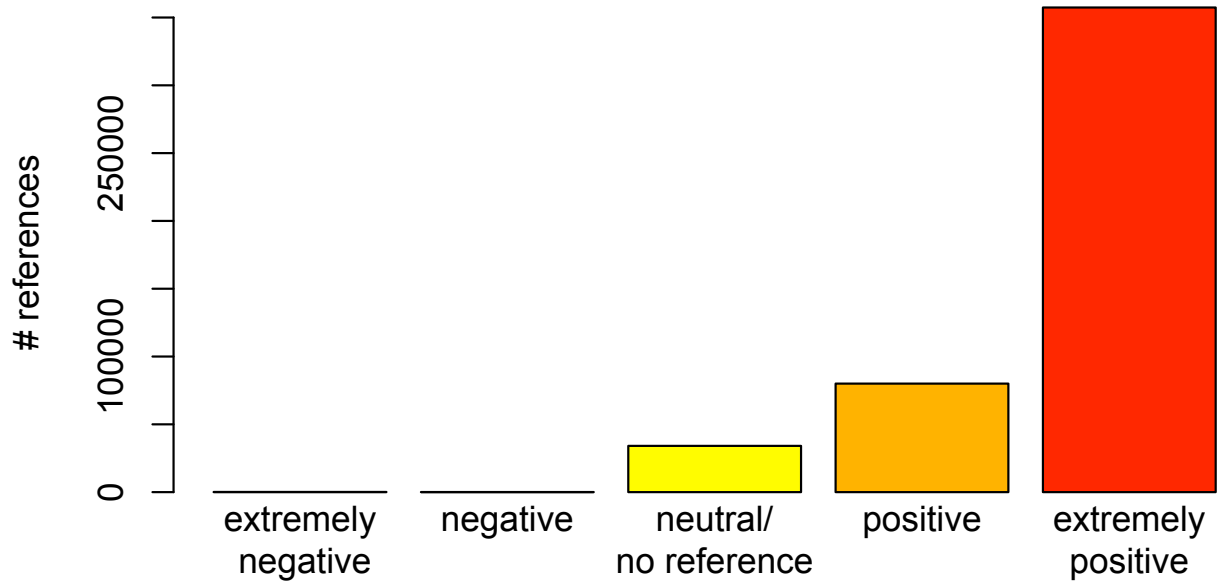

PUBLIC **~~ANONYMOUS~~** **RECIPROCAL**

Why only positive references?

- 55 % say they always leave references

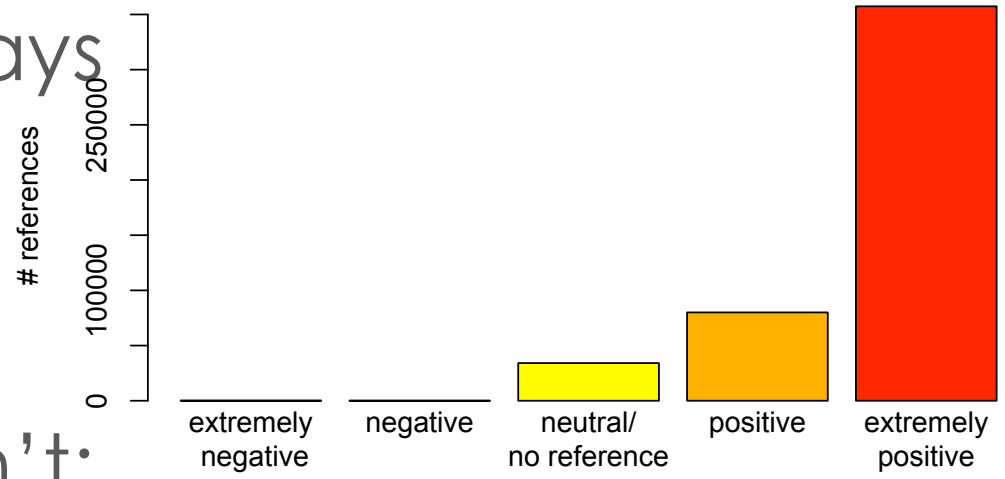

- For those who don't:

percentage of respondents	reason for not leaving a reference
51.3%	too busy
31.7%	neutral experience and didn't want to state it
12.1%	negative experience and didn't want to state it

Is there reciprocity in references?

- *generally, I prefer the host or the surfer to leave a reference first, so I can kind of... And I do gauge. I mean I gauge on the way that their reference was. If their reference was very detailed and inclusive of our expense then I will meter that, and if it was a little more scoped. [laughter] then I will narrow mine down too. [P9]*
- *I usually don't write references to those I hosted and didn't leave a reference on my profile [S399]*

Lack of negative references

- fear of reciprocal action
  - *But the big problem is that if you leave a bad reference, what happens then. What will that person say about you. You leave a bad reference and he can do the same. And its not true. [S37]*
- results in lack of information about negative experiences
  - *I chose not to leave a reference because I just felt uncomfortable [] then I actually ended up speaking to the country ambassador [] and she told me that several other surfers have had the same experience with this particular host, and I just chose on a personal level not to leave a reference for him. Today, I regret that. I wish I had left a negative reference so no one else would have been put in that situation that we were.* [P9]

Lack of negative references

- Concern for others' reputation:
  - *the few times when I had a neutral experience, I believe it was because of personal character differences, and not because I had complaints against the person in question. Somebody else might have a positive experience, why write them a negative/ neutral one and prevent people from considering the person? [S83]*

Reciprocity in CouchSurfing

- *Public* friendship ratings are more highly correlated ($\rho = 0.73$) than *private* trust ratings ($\rho = 0.39$)

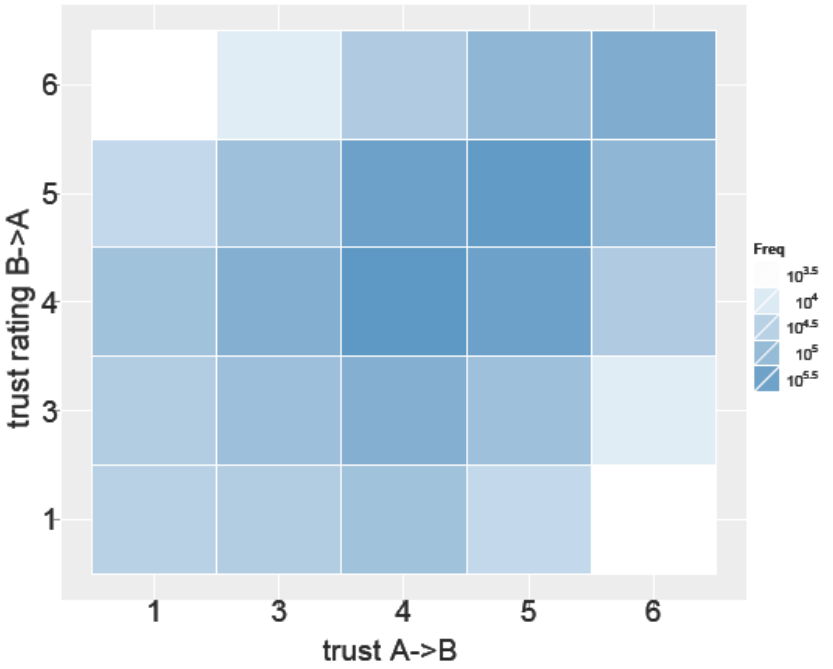

(a) alignment of trust ratings between user pairs

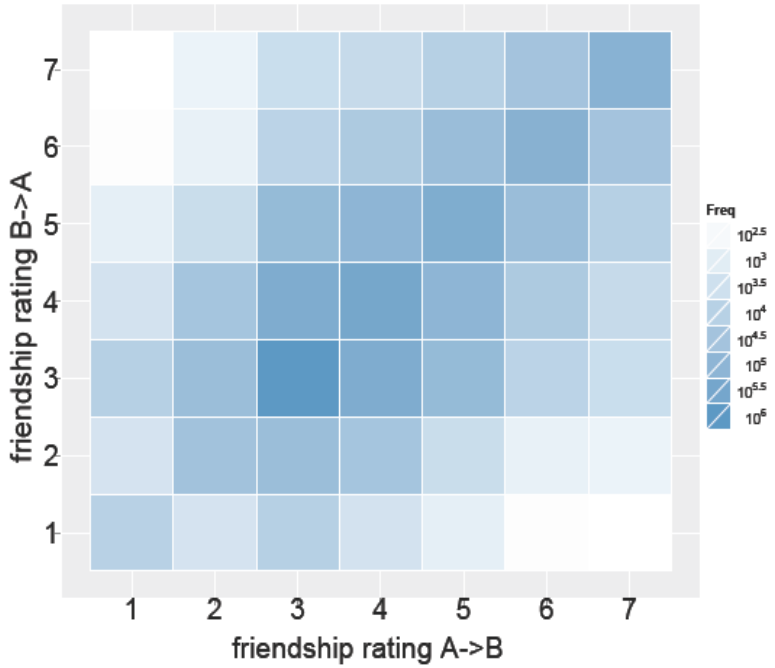

(b) alignment of friendship ratings between user pairs

We omit trust rating of 2 (I don't know the person)

~~PUBLIC~~ ~~ANONYMOUS~~ ~~RECIPROCAL~~ | PUBLIC ~~ANONYMOUS~~ RECIPROCAL

Reciprocity friendship & trust

- reminder: friendship = public, trust = private
  - *It can be difficult to select a friendship level if I am unsure of how the other person may react or if I think they may see our friendship as being at a different level. [S114]*
  - *Cause sometimes you don't want to be unpolite[] ... you want to have the person the same friendship level. [S175]*
  - *the trust level is anonymous, and I tend to trust people more easily. that's why friendship level is more difficult: everyone can see it [S276]*

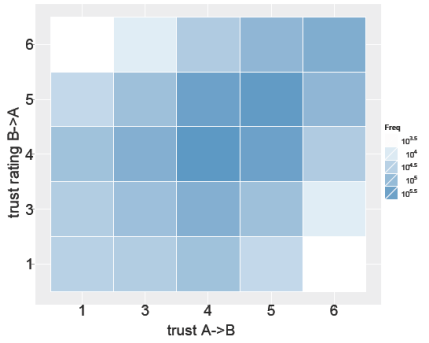

(a) alignment of trust ratings between user pairs

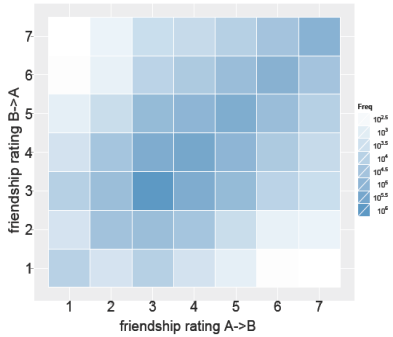

(b) alignment of friendship ratings between user pairs

reactions to misaligned ratings

- Only 41% of users even recalled noticing a misaligned friendship rating. Those who did typically did not attach much importance. But for some:
  - *I once said one girl was a "good friend" - however, she added me as an acquaintance. It actually made me feel quite bad to hear that she didn't even consider me as her friend [S491]*
  - *Not a big deal but yes it feels not great. Because you see that the feelings about the friendships is not really mutual. [S31]*

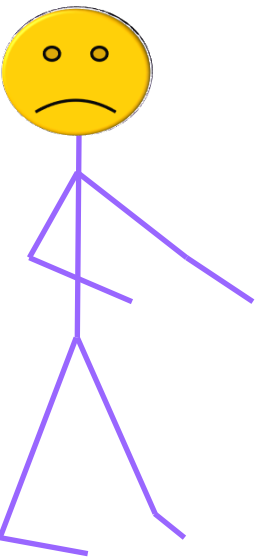

What is a "CouchSurfing friend"?

- They are not really a friend:
  - *We're not friends outside of CouchSurfing. Like Couch- Surfing is the only thing we really have in common.[P09]*
  - *I use it sometimes, when I really didn't have many things to do with the other one. [P18]*

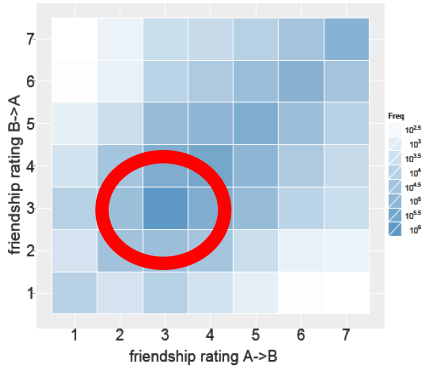

- When it designated how the acquaintance was made:
  - *Friendship level - easy, those who I hosted I mark as CS friends.*

Can trust and friendship be quantified?

- How easy is it to quantify friendship and trust?
  - *One "level" is never enough to point to the correct tone of a human relation. [S291]*
- Can friendship be interpreted as trust?
  - *I think close friends you trust, but I don't think everyone you trust is a close friend. [P12]*
  - *Friendship includes trust. You can trust someone, but still without that person being a friend. I guess friendship is a more elusive concept and therefore more difficult to judge. [S312*]

Is the friendship/trust asymmetry reflected in the data?

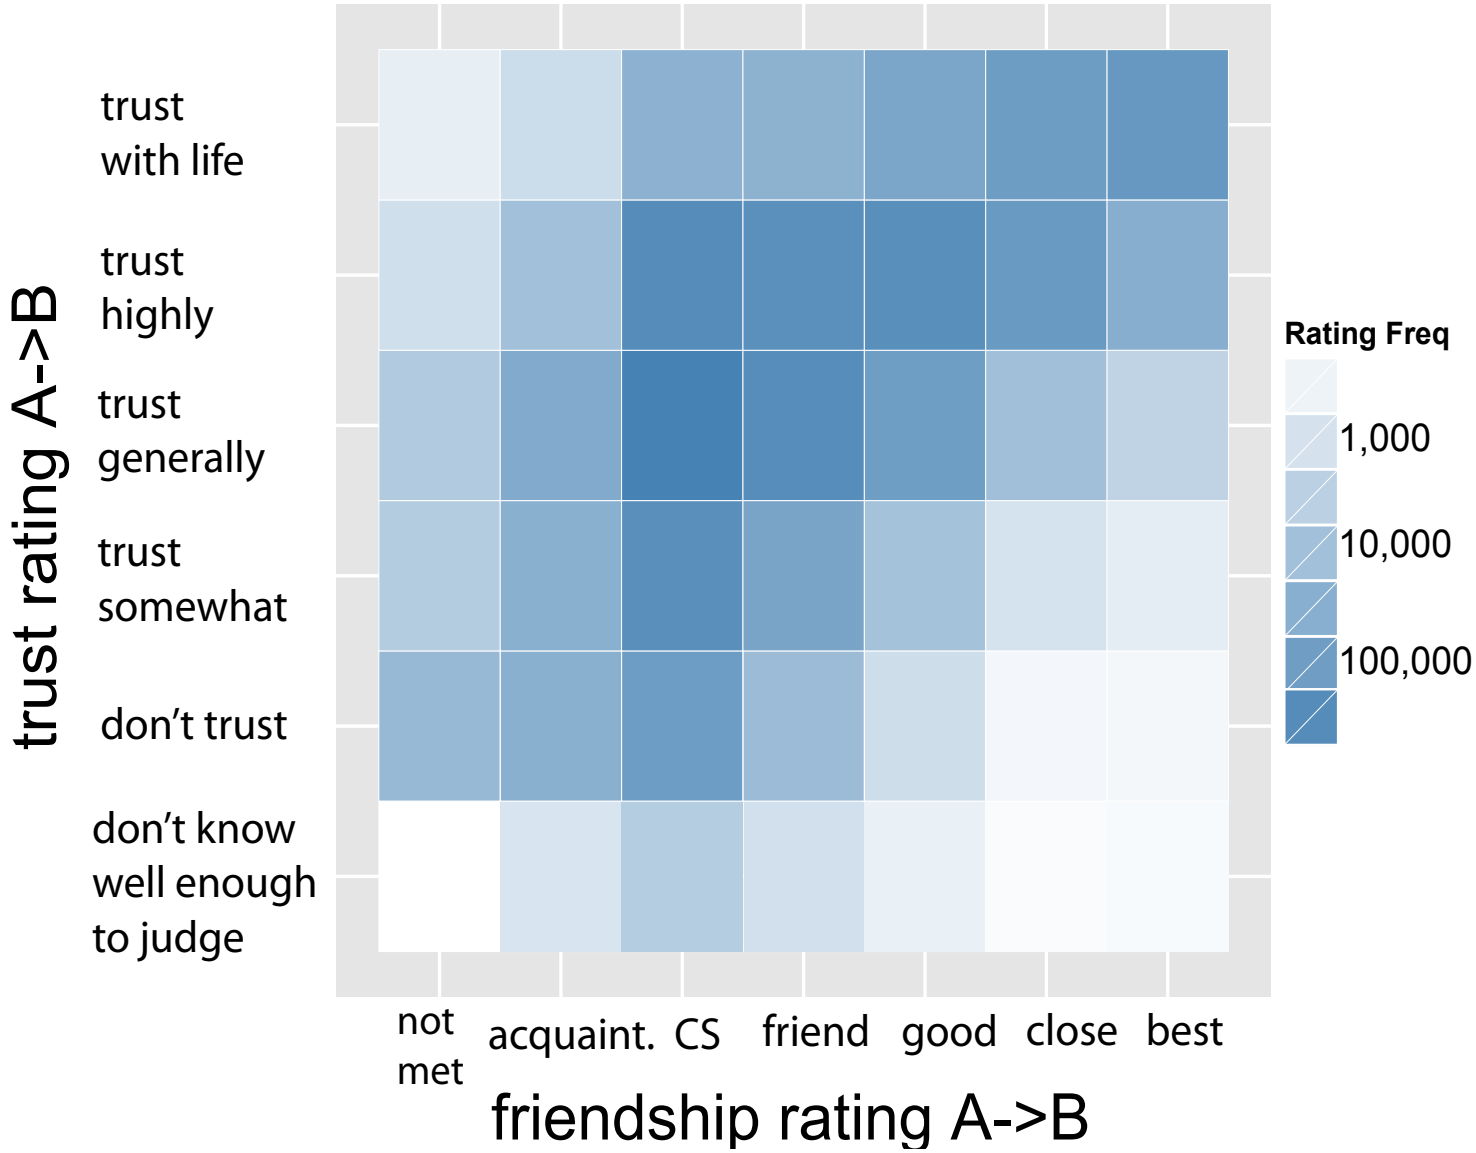

Time and friendship/trust

- *trust takes, you know, in a lot of cases years to build, whereas friendship and that sort of thing can sort of happen instantly [P8]*
- *I have a gut feeling about who I could trust, but not so much about who actually counts as my friend. [S10]*
- *In general I only *really* trust my closest, real life friends. That takes years to earn. [S256]*

Which one takes time?

- both do, but trust plateaus earlier

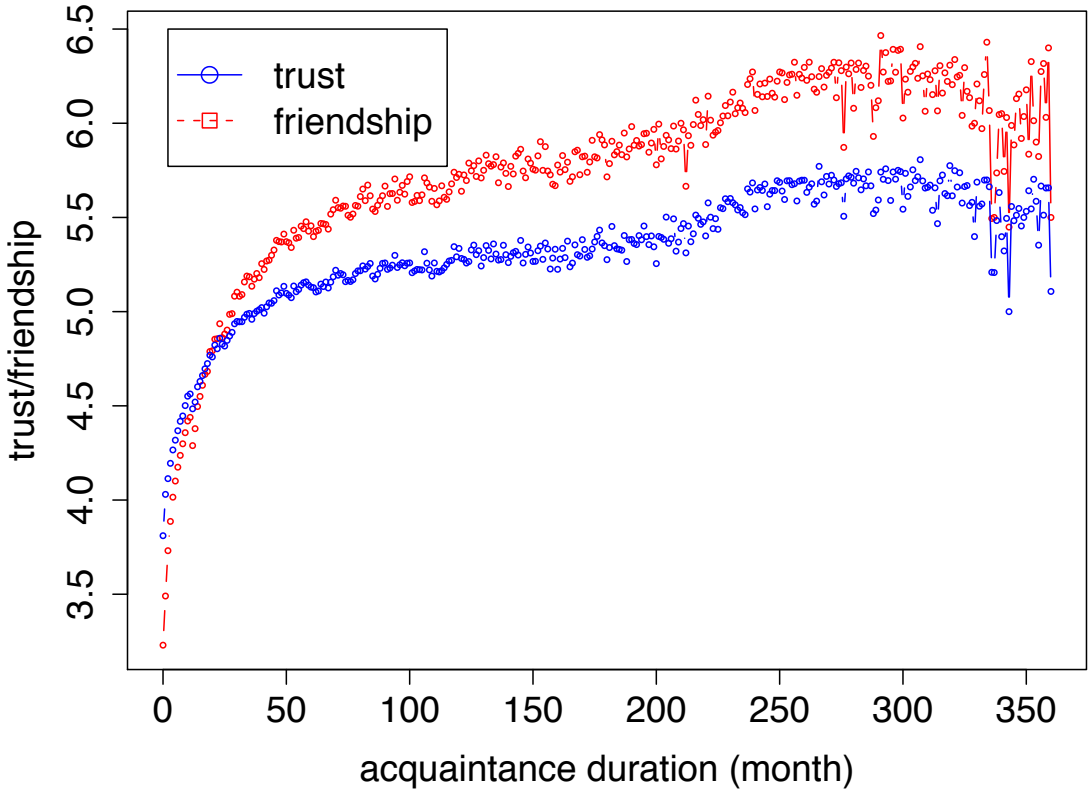

In part, it's a question of hours that we stay together. If we still stay together the level of trust increases... Not always, but increases. And sometimes you understand some limitations so maybe your level of trust can arrive to a certain level and that's it, and doesn't increase over that level. [P10]

What else does it take?

how well you know other	0.719	0.670
log (how long…)	0.592	0.365
days traveled	0.404	0.250
same country	0.212	0.075
abs. age difference	-0.106	-0.068
days surfed	0.268	0.203
days hosted	0.279	0.193
same gender	0.078	0.045

Vouching: less pressure?

- Vouching means you believe that friend to be trustworthy
- You can only vouch for others if you have at least 3 vouches yourself
- Vouching forms a small "web of trust" in the network
  - 6.8% of users have been vouched at least once
  - 1.8% can vouch for others

"Respecting the significance of vouching is essential to the integrity of the network... It is very important that you ONLY vouch for people that you … know well enough to believe that he or she is trustworthy."

Can one predict vouches?

- Logistic regression model (10-fold cross-validation)
- 71% accuracy in predicting whether a random edge is vouched
- Most predictive attributes were friendship degree, rating of experience, how they met

Predicting vouches - network ranking algorithms

- Two-step indirect measure for propagating vouches:

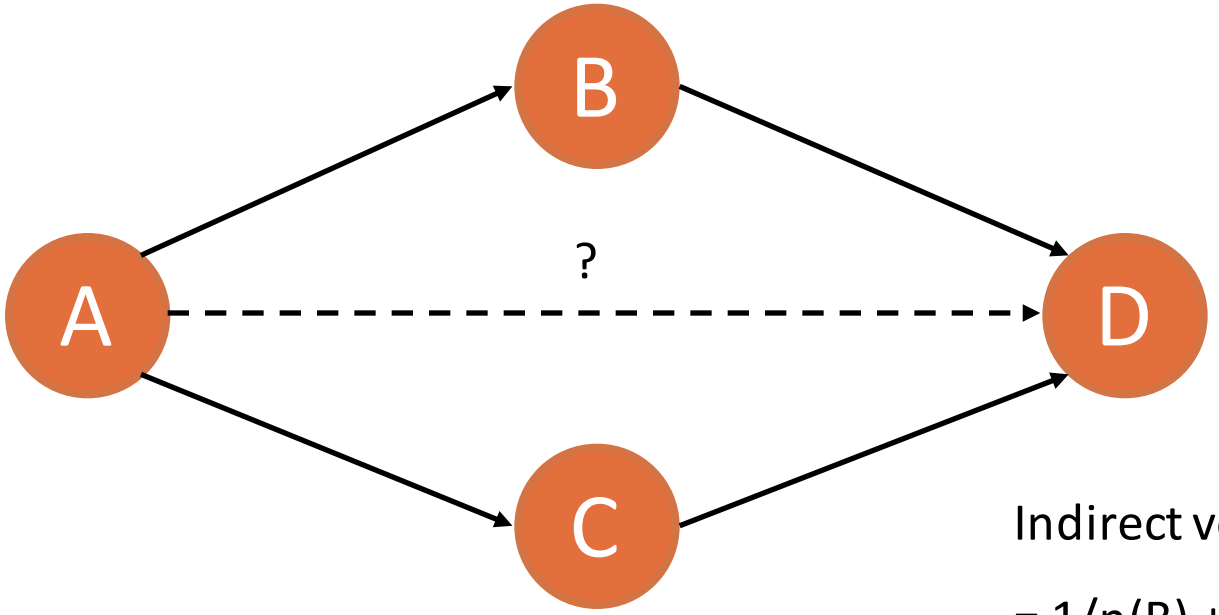

Indirect vouch score for A->D:

= 1/n(B) + 1/n(C)

Predicting vouches - global measures

- Results from logistic regression for each variable alone:

Variable	**Predictive accuracy:**
Friendship degree	67.7%
Jaccard coefficient	55.8%
2-step vouch propagation	54.2%
PageRank	50.6%

- Global measures are poor predictors of whether an edge is vouched

Whom are vouches applied to?

- A high number of vouches are between "CouchSurfing friends"

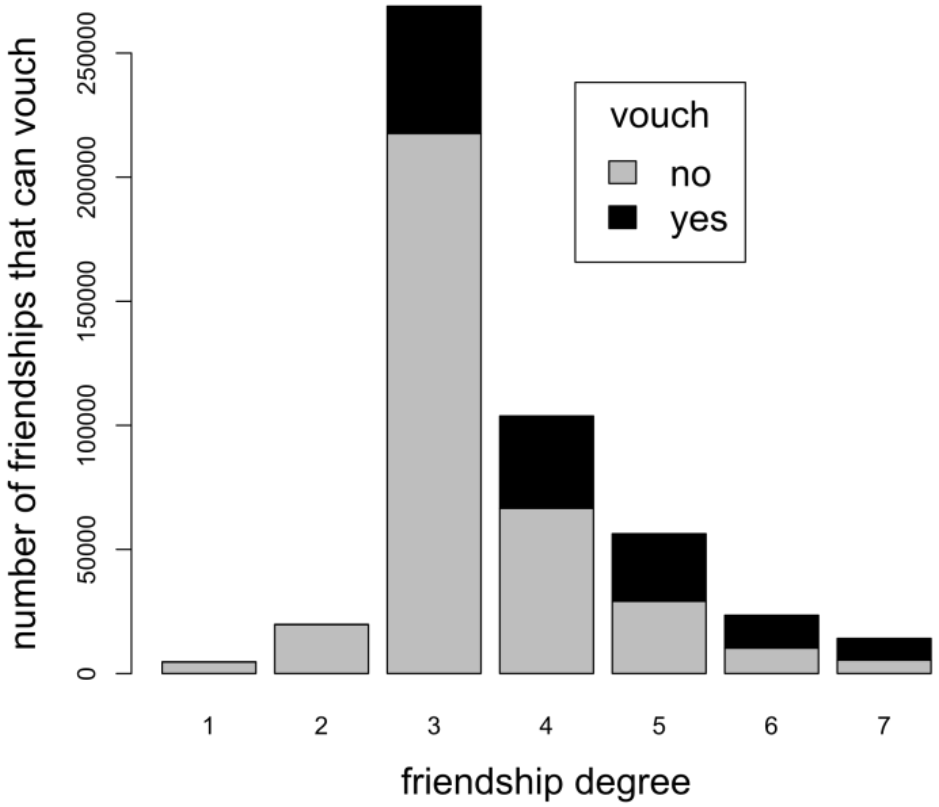

Friendship degree:

1= Haven't met yet

2= Acquaintance

3= CouchSurfing friend

4= Friend

5= Good friend

6= Close friend

7= Best friend

Why do users vouch others?

- Tight web of trust….or vouching too freely?
- Mutual trust….or social pressure to reciprocate?
  - 95% of users with > 10 friends have been vouched
  - 25% of friendships that can be vouched are
  - High rate of reciprocity

Attitudes towards soliciting and reciprocating vouches

- *I know some people are engaged in some sort of vouching competition and they are guys with 1700 vouching. I'm not very much into this sort of thing. [P11]*
- *I didn't vouch back to a guy that vouched me just be- cause he wanted to get some popularity and vouch(es) back - I think [S65]*
- And no one ever asked me and I've never asked anyone to vouch for me. It's kind of like a...[] taboo thing. You hope that they do [P09]

Reciprocity in vouches on CouchSurfing

- If A vouched for B, 70% of the time B also vouched for A
- Mean **private** trust score for *reciprocated* **public** vouches was higher (4.47) than unreciprocated ones (4.19)
  ➔ lack of rating could signal lower trust

Are even truthful ratings reliable?

- Even if one were able to elicit truthful ratings, would there still be biases?
- To answer this we used demographic information from CouchSurfing.com

Gender effects for trust & friendship

- Men rate both men and women about equally on trust and friendship
- Women rate other women more highly on both

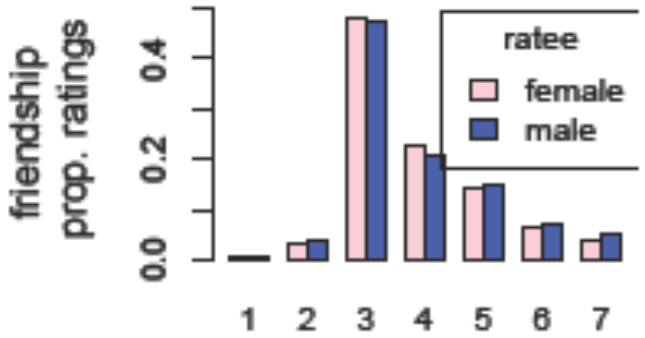

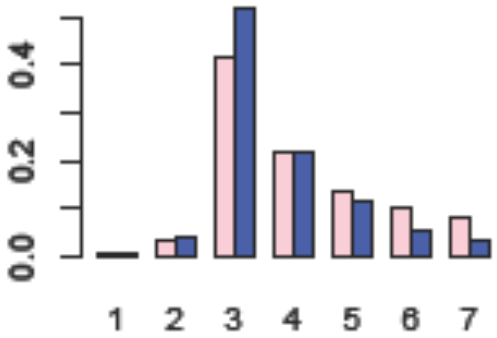

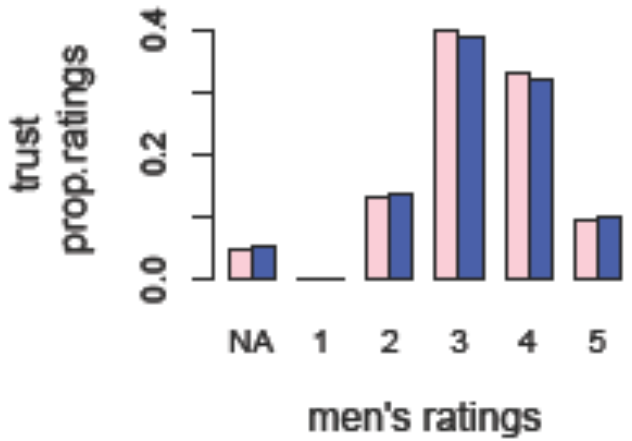

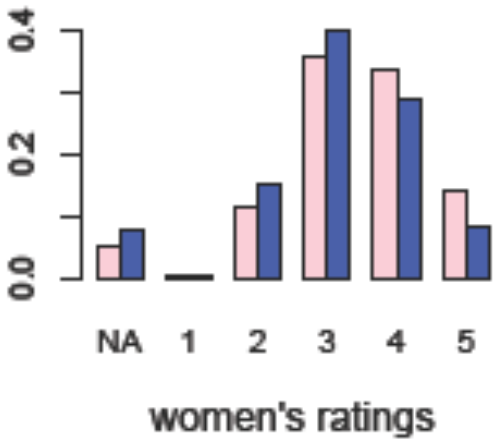

Geography

- Closer friends tend to be geographically proximate
  - Friendship for one's countrymen (4.19) is higher than foreigners (3.65)
  - Trust for one's countrymen is higher than for foreigners (4.33 vs 4.16)

Geography

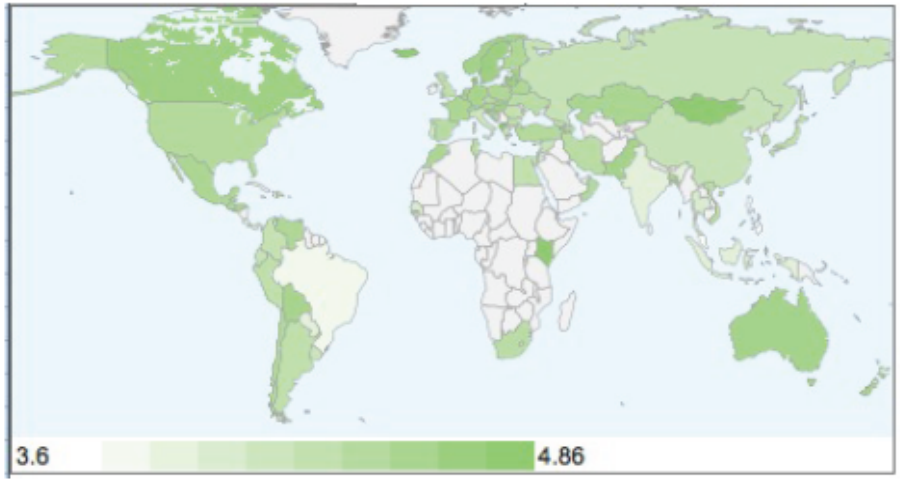

(a) average within-country ratings

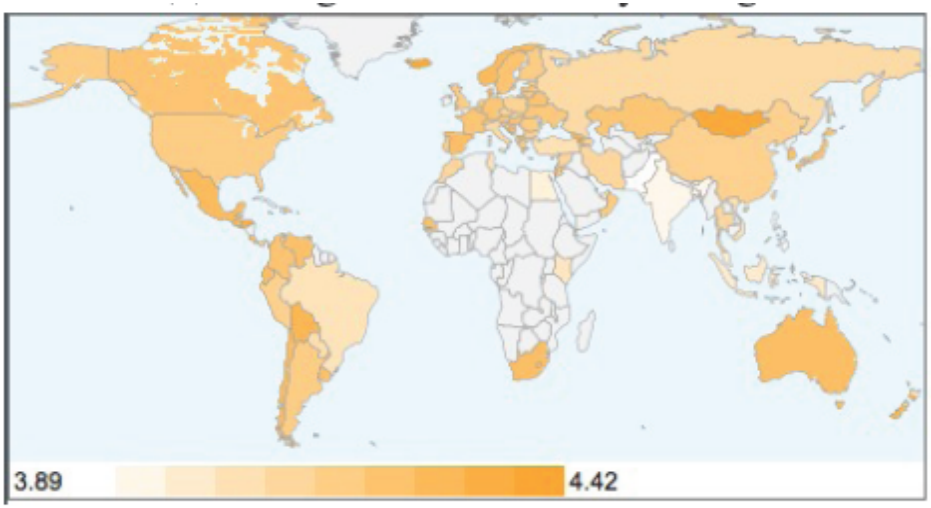

(b) average ratings received by country residents

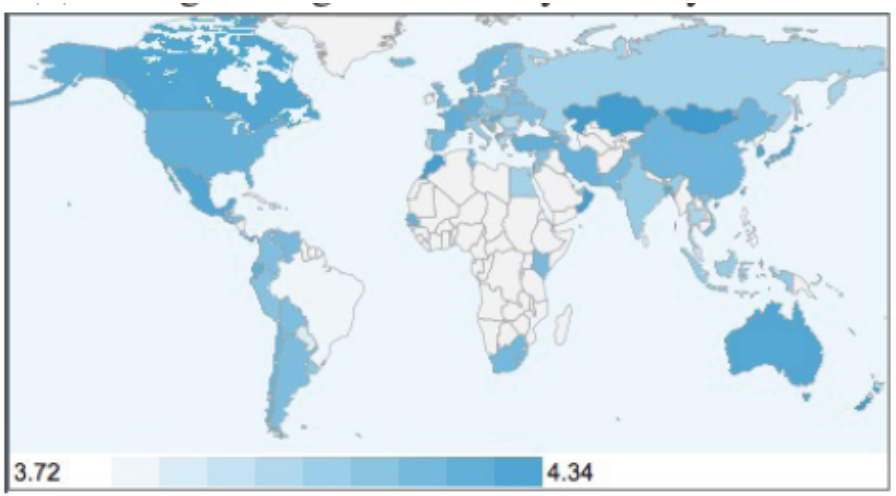

(c) average ratings given by country residents

Are ratings biased by culture?

- From non-American interviewees:
  - *Americans have an interesting way of putting things. they can write "awesome, great, super" but don't mean it [P19]*
  - *[] American people, [laughter] they tend to be very open and very, you know, they help you. Yeah, you're my friend even though you meet them for the first time. [P11]*

	US	Western Europe
av. friend rating	3.91	3.86
% best friend	5.90	4.86

Geography

- Countries with similar cultural background tend to be trusting of one another (e.g. Austria and Germany)
- Sharing a border does not always correspond to greater trust (e.g. Canadians did not rate US contacts more highly)

Language

- Expressions don't always translate:
  - *Let's say that my level of English is the one that you can hear [from] me. Best friend, I understand what that means but between good friend and close friend, I don't know exactly what is the higher level, I'm not sure. [P10]*
  - *Some of these things are, they're with the web site, or directly translated from English to our... Whatever language. To say I trust this person with my life, I think that's such a hyperbole.*

Age

- Trust is very slightly higher the smaller the age difference between rater and ratee (ρ= -0.06)
- Trust depends on age of ratee – typical CouchSurfing demographic preferred?

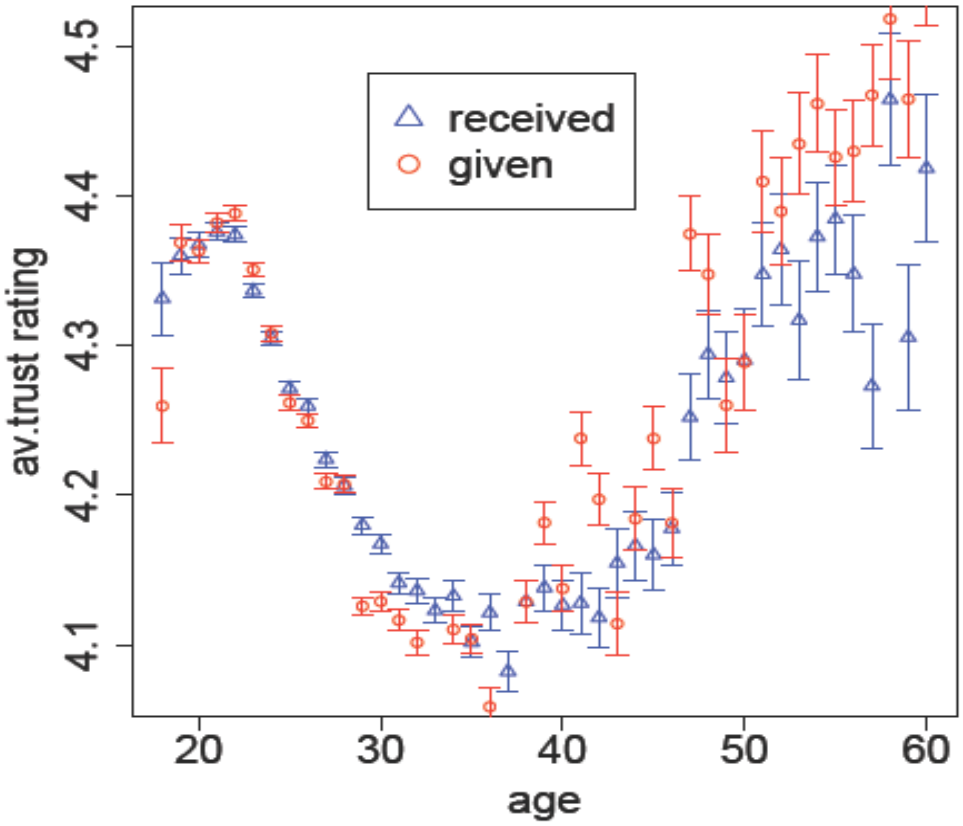

Lest you think it is just CouchSurfing

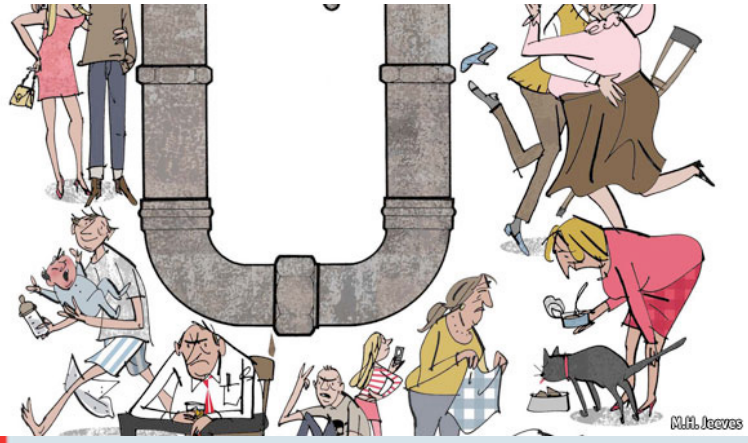

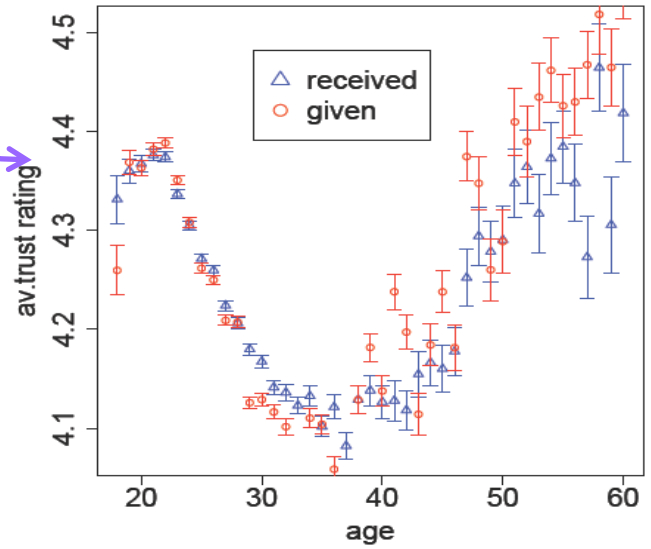

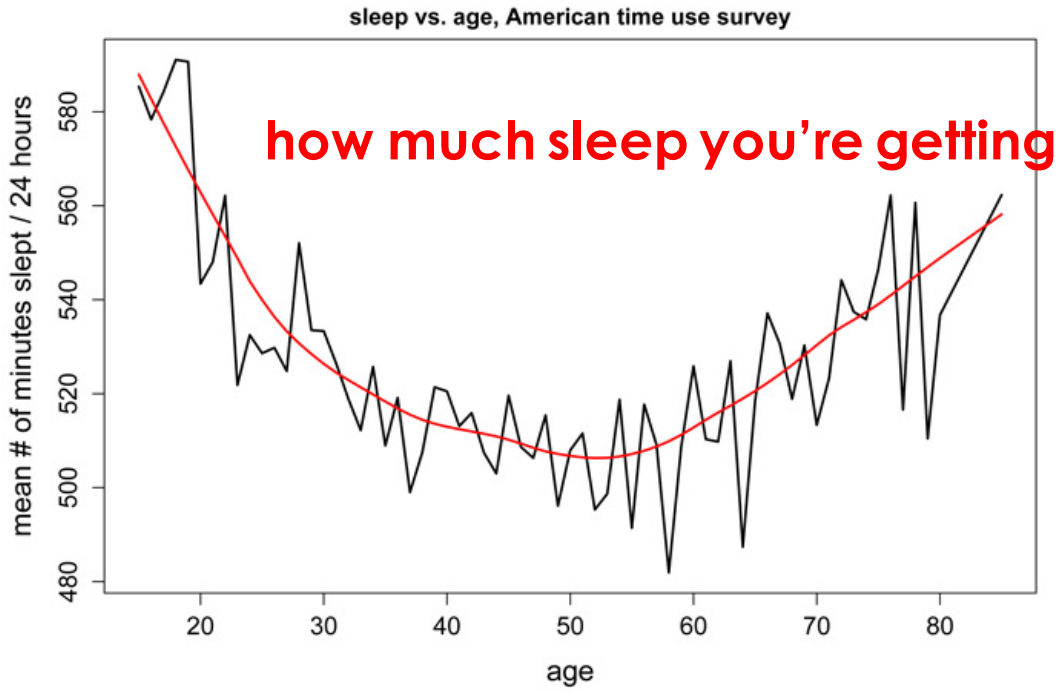

How useful are numerical ratings of human relationships?

- How important are the following when choosing whether to host or be hosted by someone?

	Very Important	Important	Neutral	Unimportant	Very Unimportant
how many vouches they have	8.6%	29.3%	37.1%	15.8%	6.9%
whether they are verified	9.1%	23.0%	36.4%	18.9%	11.6%
number of references received	23.0%	57.0%	13.7%	4.4%	1.1%
text of references received	47.6%	40.8%	8.2%	2.1%	0.8%
number of friends	4.4%	23.6%	37.7%	25.1%	8.0%
friends' friendship level	3.2%	19.0%	35.6%	25.0%	13.5%

Why are textual references more useful?

- Many include information about the individual that signal to others the person's personality and interest
- It is possible to leave a ***neutral*** reference while using a seemingly ***positive*** tone.
  - *I've gone pretty keen on what certain references mean, and you can tell when a reference is just like a simpatico- nice; you-were-a-nice-person-reference: "[She] was great. She was very hospitable. She's a great host." That can mean in a sense you might be kind of boring. [P9]*

Conclusion

- Ratings and online social relationships should not be taken at face ("friendship" or "trust") value
  - Public, identified ratings tend to be positive when there is potential for reciprocity
  - Even truthful ratings can be biased by various factors
  - The framing of the ratings can help improve reliability

Summary: networks can get *really* interesting

- Add signs/ratings to edges
  - remember all the social conventions
- Optimize
- Learn and coordinate