Information Diffusion in Social Networks

EE599: Social Network Systems

Keith M. Chugg

Fall 2014

USC University of Southern California

Overview

- Network robustness/resilience and percolation theory
  - Cascades
- Information diffusion and epidemics
  - Network search
- Learning and consensus formation

Easley & Kleinberg Chapter 19

Cascades in Networks

- Thus far we have considered spread of information/diseases, etc. at the macro level
  - Will a disease spread throughout a network or will it die out?
  - What will be the eventual size of the outbreak?
- Cascade analysis considers these phenomena at the micro scale
  - What stops an idea from propagating?
  - What makes or kills a meme?

Information Diffusion

- Social scientists have studied the adoption of new ideas in communities for many years
  - Example: hybrid corn adoption by farmers
  - Adoption trends are lead by a small set of "early adopters" (recall Bass model)
    - Early adopters tend to have higher social-economic status, travel frequently, have access to other communities
- What makes for a successful innovation?
  - Significant **relative advantage** (vs. current)
  - **Low complexity** of adoption
  - **Easy to observe** that others are adopting (social pressure)
  - **Easy to try** out
  - **Compatibility** with other current technologies

Simple Model of Adoption

- Consider two behaviors A and B (e.g., adopt new and keep old)
- A simple benefit model for adoption:

$$
\begin{array}{cc|c|c|}
 & & \multicolumn{2}{c}{w} \\
 & & A & B \\
\cline{3-4}
\multirow{2}{*}{v} & A & a, a & 0, 0 \\
\cline{3-4}
 & B & 0, 0 & b, b \\
\cline{3-4}
\end{array}
$$

Figure 19.1: A-B Coordination Game

Easley & Kleinberg

Captures the factors associated with new adoption discussed on previous slide

Simple Model of Adoption

- What happens to a node involved in the trade-off (game) with all of its neighbors?

Pay-off for v

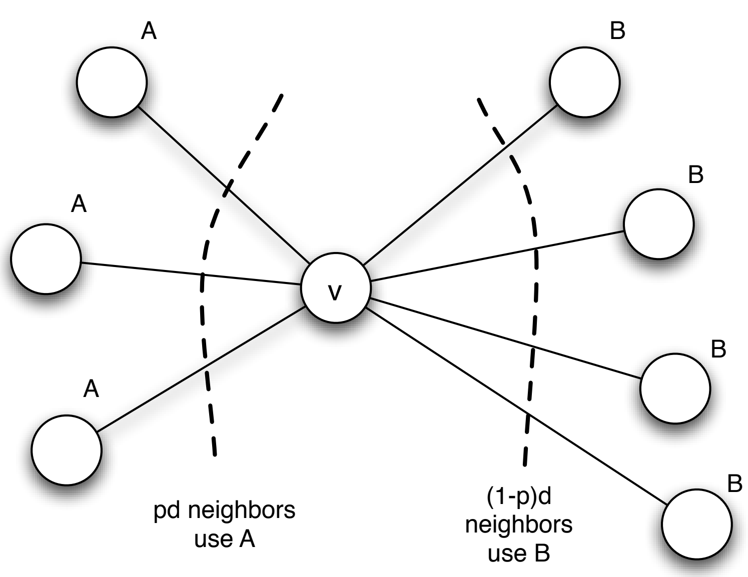

Figure 19.2: v must choose between behavior A and behavior B, based on what its neighbors are doing.

Easley & Kleinberg

Choose A: pda

Choose B: $(1-p)db$

Node v will choose A iff: $pda \geq (1-p)db \iff p \geq \dfrac{b}{a+b} = q$

Simple Model of Adoption

- If the fraction of your neighbors who have adopted A (p) is greater than a threshold q, you will adopt…

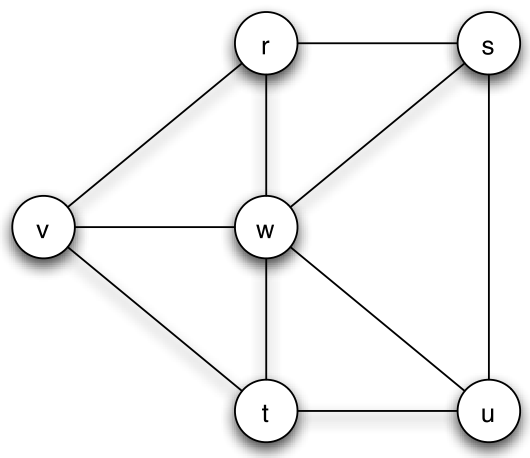

(a) *The underlying network*

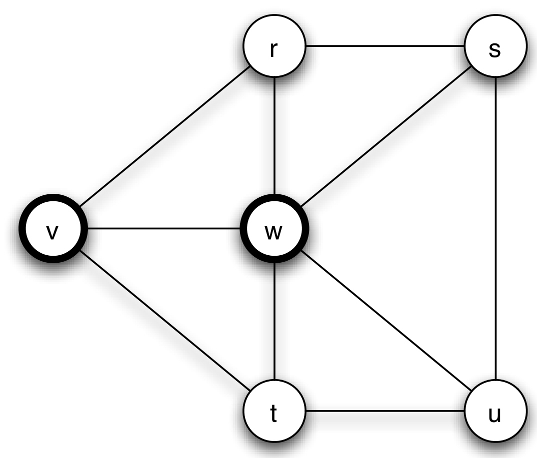

(b) *Two nodes are the initial adopters*

a=3, b=2

q=2/5

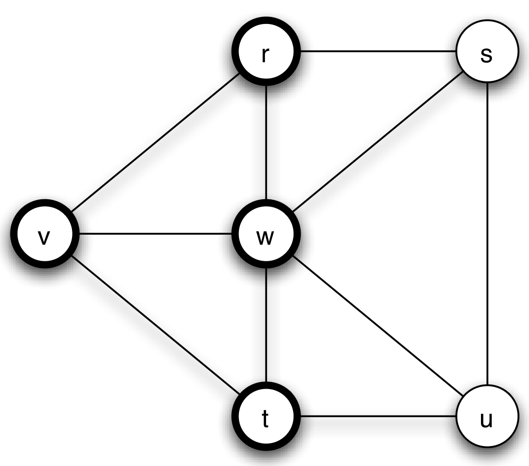

(c) *After one step, two more nodes have adopted*

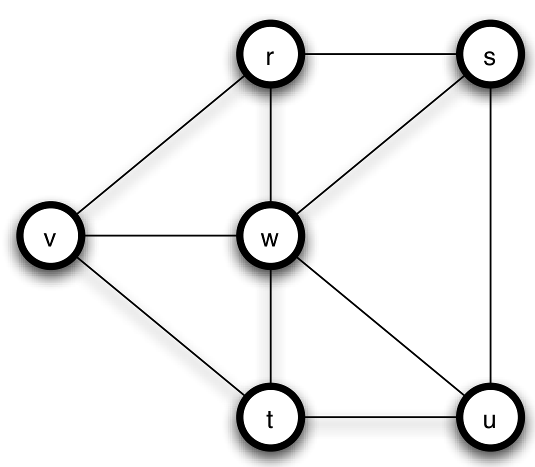

(d) *After a second step, everyone has adopted*

all nodes adopt A after 2 steps with an initial seed of {v,w}

Figure 19.3: Starting with v and w as the initial adopters, and payoffs $a = 3$ and $b = 2$, the new behavior A spreads to all nodes in two steps. Nodes adopting A in a given step are drawn with dark borders; nodes adopting B are drawn with light borders.

Easley & Kleinberg

What Can Result?

- Initial condition
  - All nodes in the network are using B, except for a seed set of early adopters who use A
- Fact: if a node adopts A, it will not switch back to B
  - Only two possible results:
    - Everybody adopts A
    - The spread of A is contained to a finite fraction of network

Consider a set of initial adopters who start with a new behavior A, while every other node starts with behavior B. Nodes then repeatedly evaluate the decision to switch from B to A using a threshold of q. If the resulting cascade of adoptions of A eventually causes every node to switch from B to A, then we say that the set of initial adopters causes a complete cascade at threshold q.

Easley & Kleinberg

A Meme that Dies Out

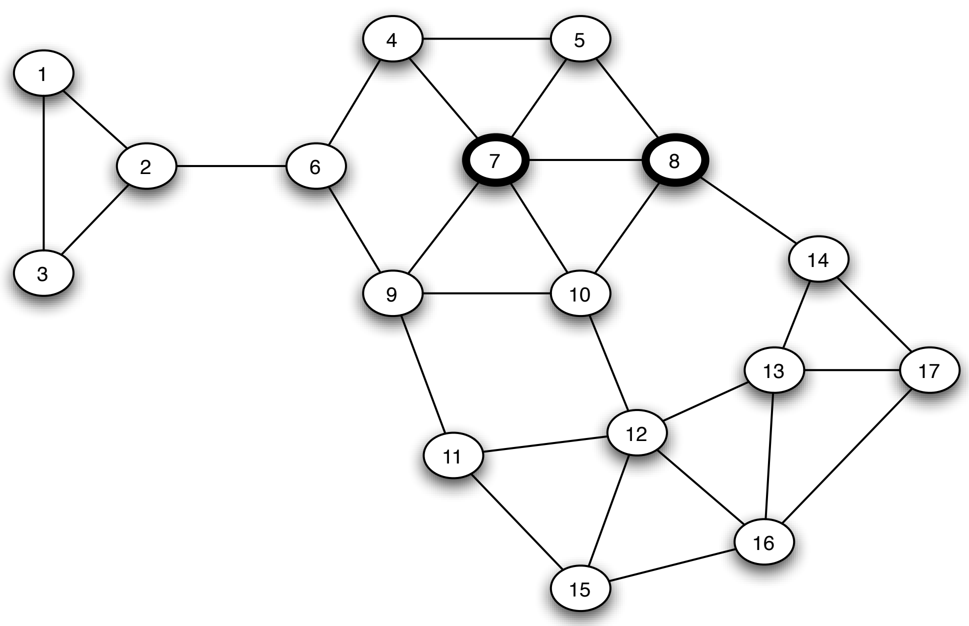

(a) *Two nodes are the initial adopters*

a=3, b=2

q=2/5

spread of A is contained to 7 nodes

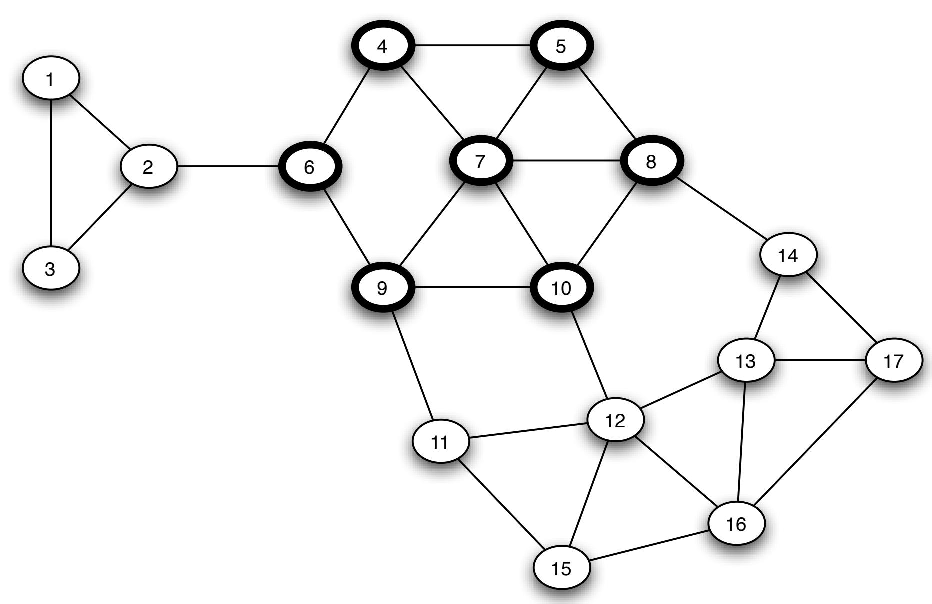

(b) *The process ends after three steps*

Figure 19.5: Starting with nodes 7 and 8 as the initial adopters, the new behavior A spreads to some but not all of the remaining nodes.

Easley & Kleinberg

A Meme that Dies Out

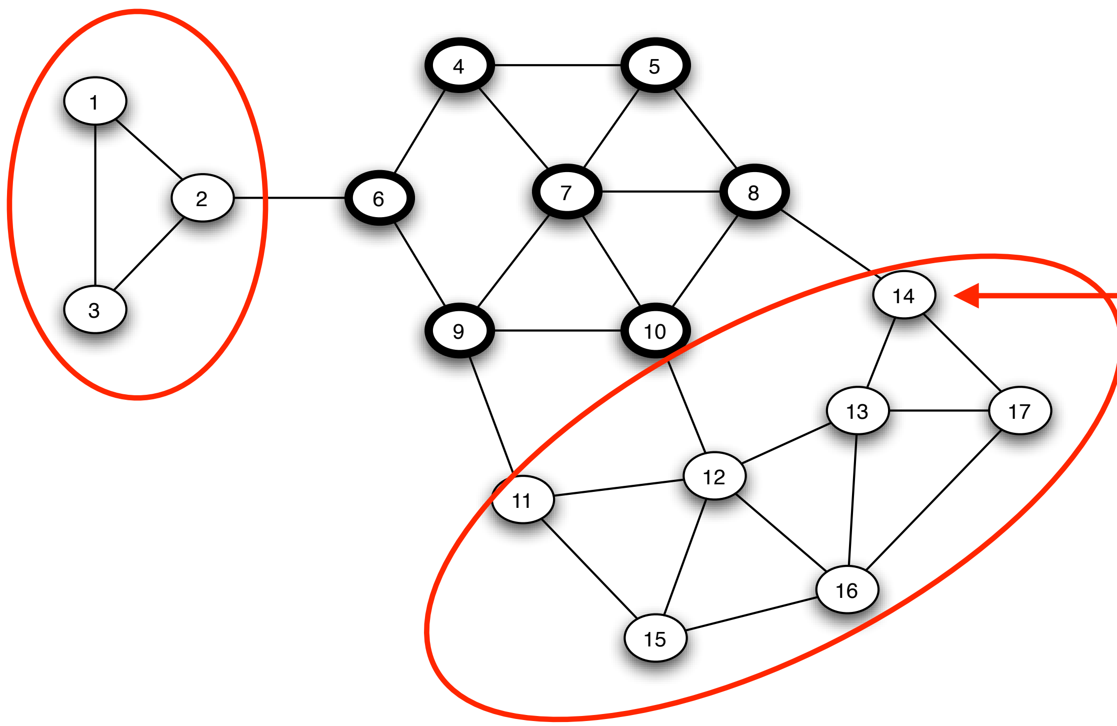

(b) *The process ends after three steps*

Figure 19.5: Starting with nodes 7 and 8 as the initial adopters, the new behavior A spreads to some but not all of the remaining nodes.

Easley & Kleinberg

Node 14 has too many friends in its community that are using B — prevents it from adopting A

Intuitively: tightly knight communities block the spred of new innovations!

A Meme that Dies Out

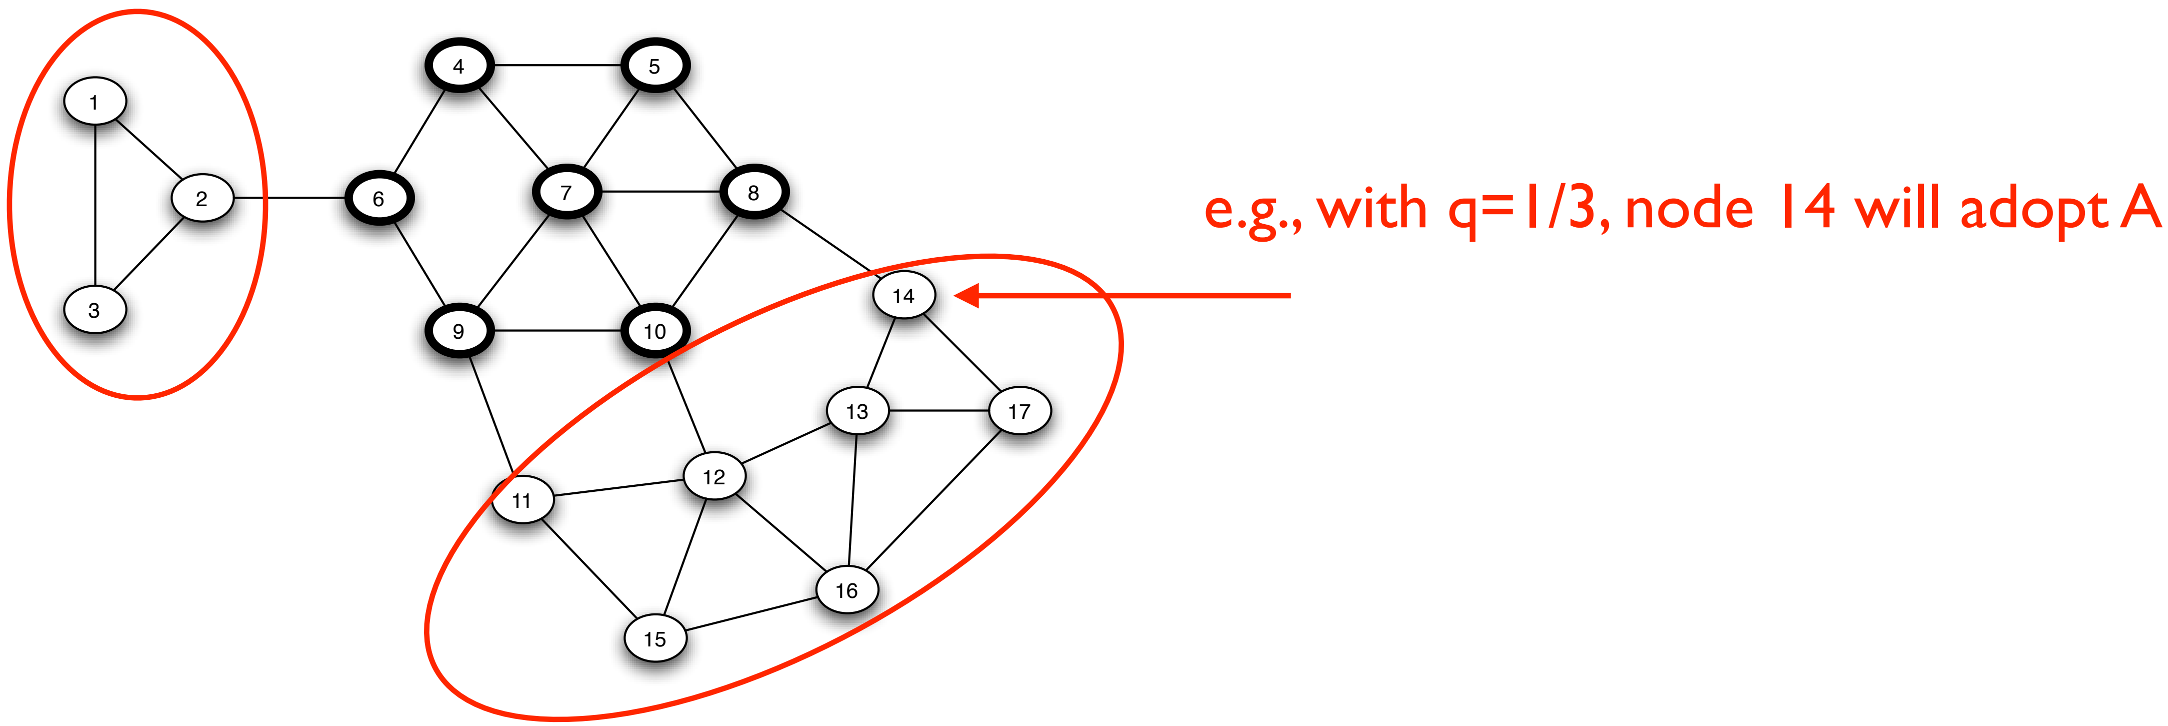

(b) *The process ends after three steps*

Figure 19.5: Starting with nodes 7 and 8 as the initial adopters, the new behavior A spreads to some but not all of the remaining nodes.

Easley & Kleinberg

If a is increased to 4, then q=1/3 and all nodes would adopt A

One direct way to increase adoption is to increase the benefit of your innovation

A Meme that Dies Out

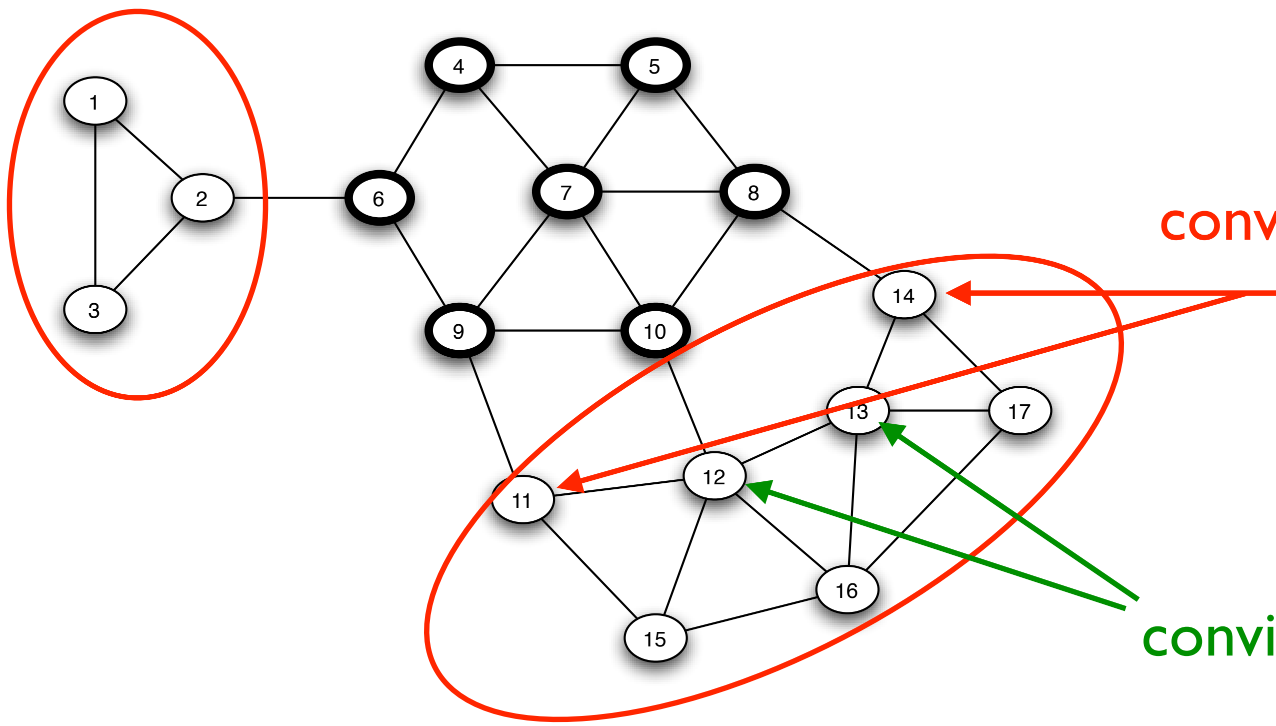

(b) *The process ends after three steps*

convincing 11, 14 to adopt will not allow the diffusion to continue

convincing 12, 13 to adopt will result in all node 11-17 to adopt

Figure 19.5: Starting with nodes 7 and 8 as the initial adopters, the new behavior A spreads to some but not all of the remaining nodes.

Easley & Kleinberg

One could target specific nodes in a community to adopt

Another way to increase adoption is to target key nodes to adopt, thus allowing a cascade to continue into a community

Cascades and Clusters

- Cluster of density p is a collection of nodes for which each member has at least a fraction p of its connections to other members of the cluster

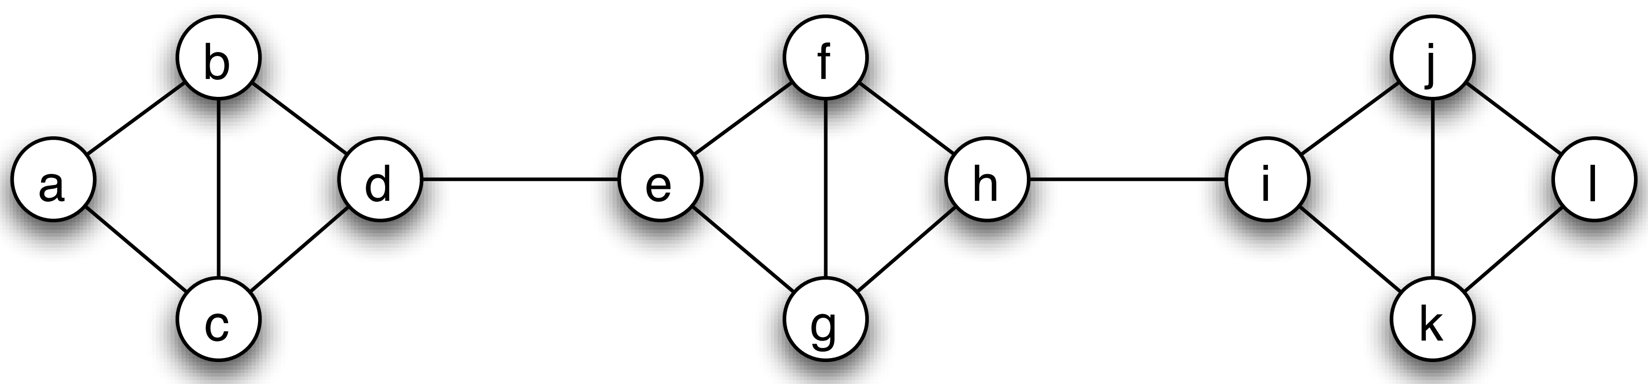

Figure 19.6: A collection of four-node clusters, each of density 2/3.

Easley & Kleinberg

Cascades and Clusters

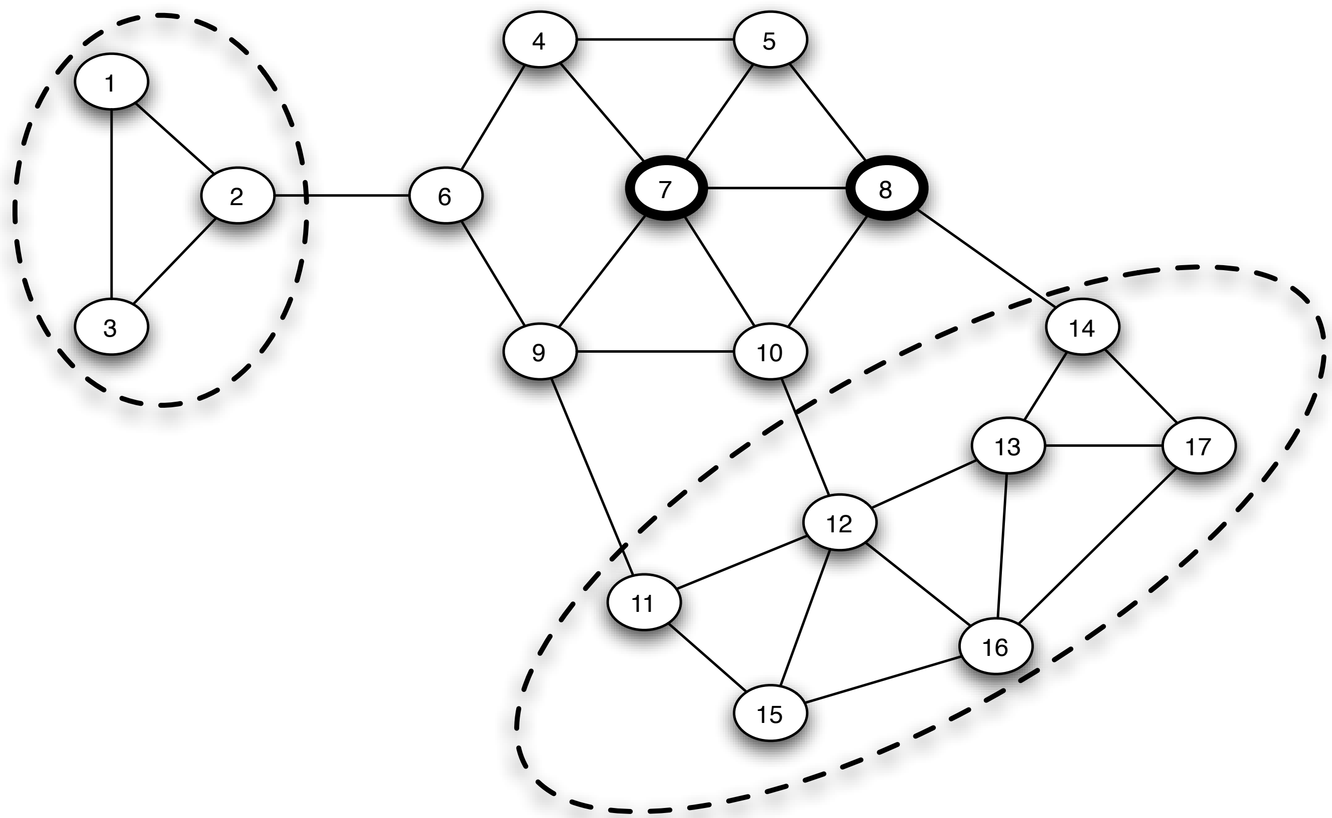

Figure 19.7: Two clusters of density 2/3 in the network from Figure 19.4.

Easley & Kleinberg

These two clusters blocked the cascade

Cascades and Clusters

- Clusters block cascades….

Given an initial set of adopters of A, the behavior A will be adapted by the entire network

The network, excluding initial adopters, does not contain a cluster of density 1-q or greater

Cascades and Clusters

- Clusters block cascades….

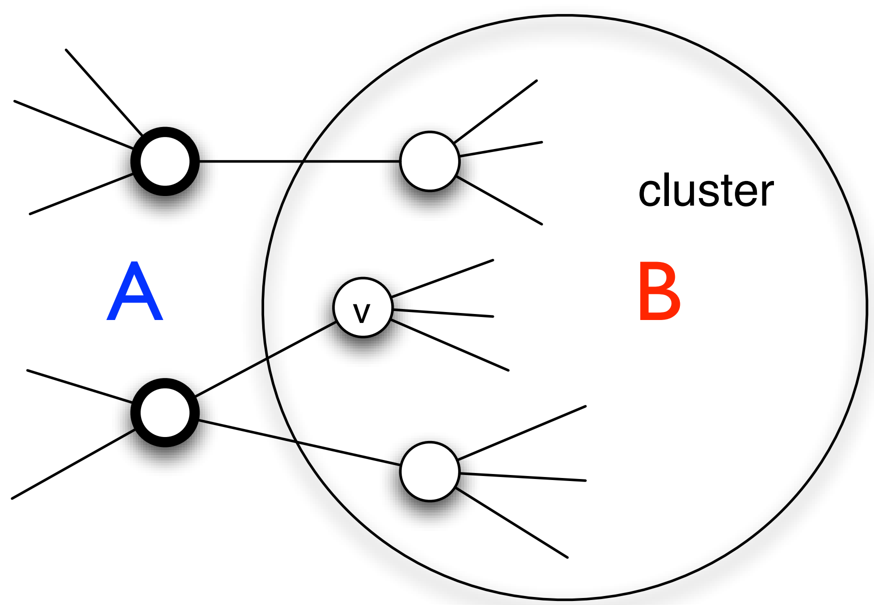

To adopt A, v requires at least q of its neighbors to have adopted A

Figure 19.8: The spread of a new behavior, when nodes have threshold q, stops when it reaches a cluster of density greater than $(1 - q)$.

Easley & Kleinberg

Cascades and Clusters

- Clusters block cascades….

To adopt A, v requires at least q of its neighbors to have adopted A

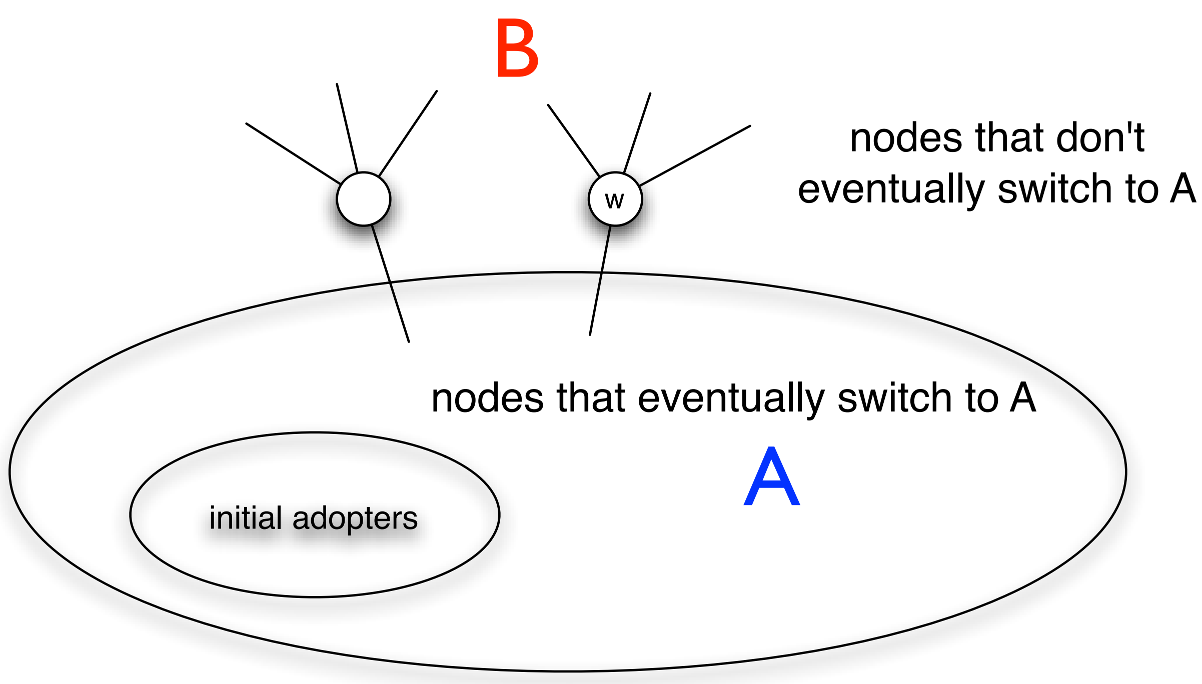

Figure 19.9: If the spread of A stops before filling out the whole network, the set of nodes that remain with B form a cluster of density greater than $1 - q$.

Easley & Kleinberg

Cascades and Weak Ties

- Recall that weak ties are local bridges to new communities
  - Play a key role in introducing those with access to new ideas

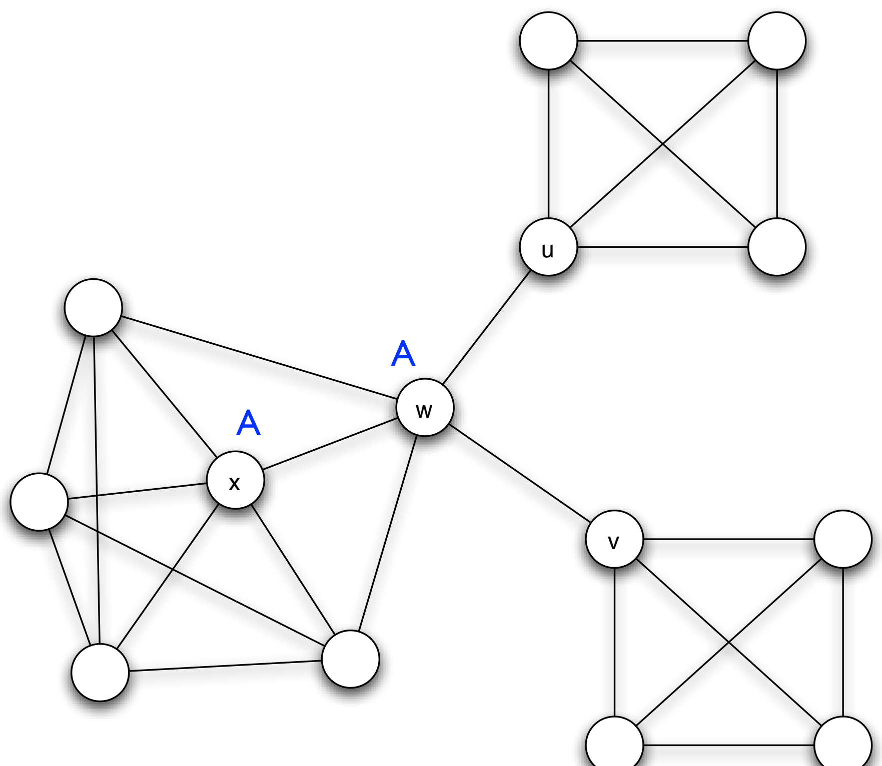

q=1/2 — what will happen?

Figure 19.11: The u-w and v-w edges are more likely to act as conduits for information than for high-threshold innovations.

Easley & Kleinberg

Cascades and Weak Ties

- Recall that weak ties are local bridges to new communities
  - Play a key role in introducing those with access to new ideas

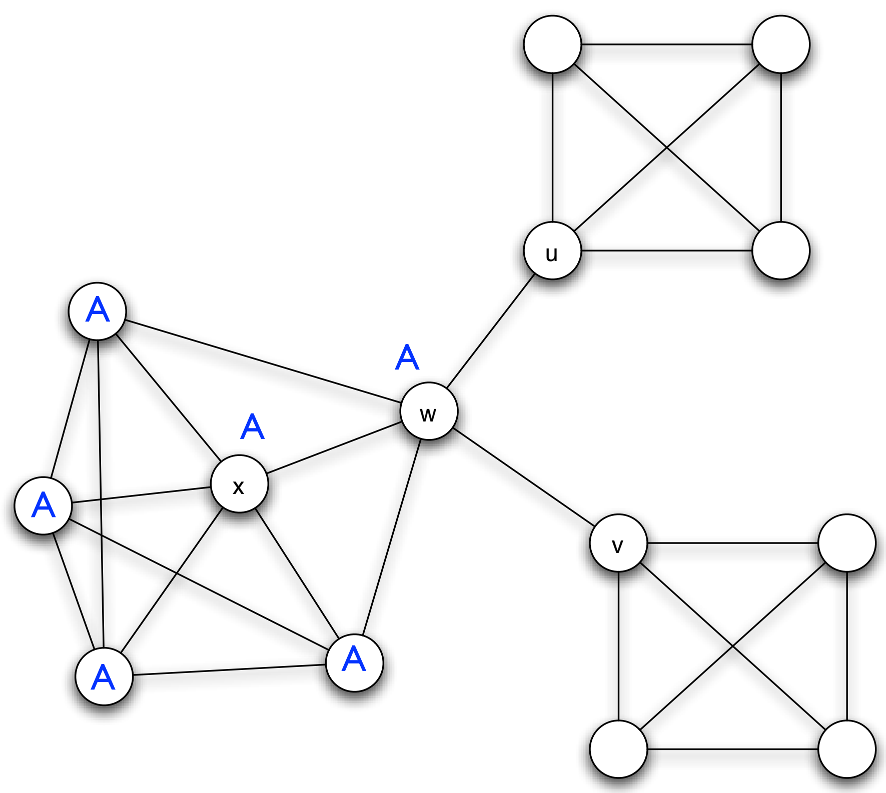

Figure 19.11: The u-w and v-w edges are more likely to act as conduits for information than for high-threshold innovations.

Easley & Kleinberg

q=1/2 — what will happen?

Note that now, u and v have an advantage over others in their community because they have knowledge of A

Cascades in Networks

- Many variations on this simple model
  - node-varying threshold for adoption
  - “bilingual” model with cost
- Cascade capacity
  - largest adoption threshold that will still allow a complete cascade
- Other models for “tipping”
- Finding seed sets for staring a meme (or how to block a meme)

Cascades with Heterogeneous Thresholds

			w	
			A	B
v	A		a_v, a_w	$0, 0$
	B		$0, 0$	b_v, b_w

Figure 19.12: A-B Coordination Game

Easley & Kleinberg

Node v adopts iff:

$$p_v \geq \frac{b_v}{a_v + b_v} = q_v$$

- Each node has its own adoption threshold *q[v]*

Cascades with Heterogeneous Thresholds

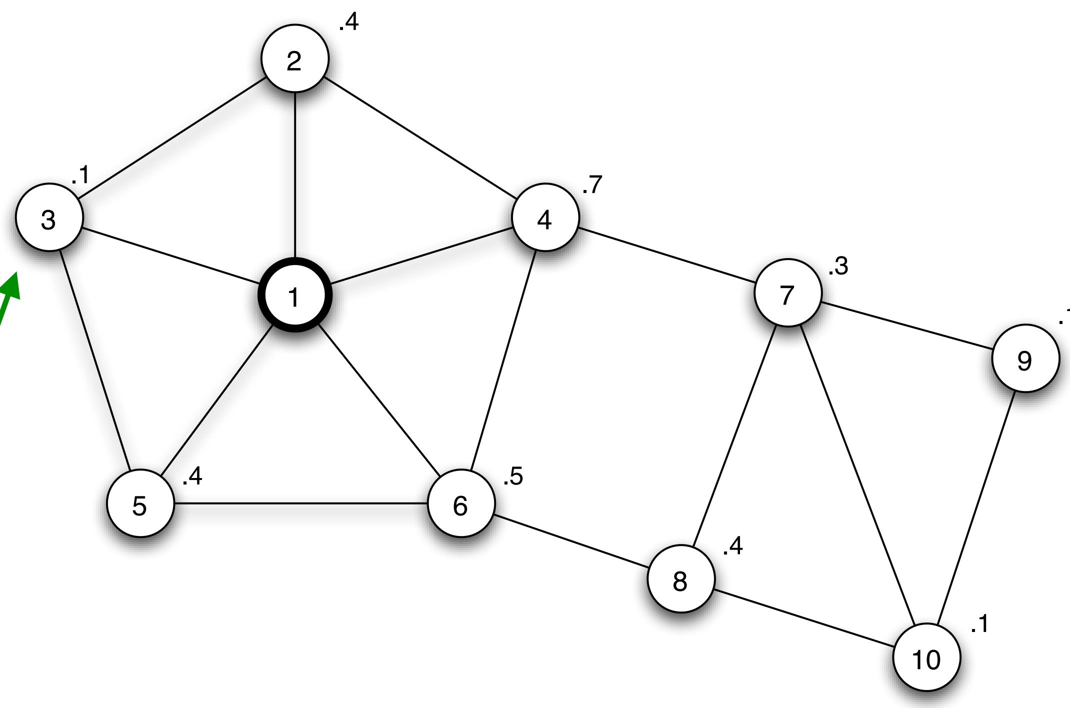

(a) *One node is the initial adopter*

Cluster of some form still stops the cascade

cascade only begins because there is an easily influenced near node 1

(access to both influential people and easily influenced people is important in creating a meme!)

blocking cluster

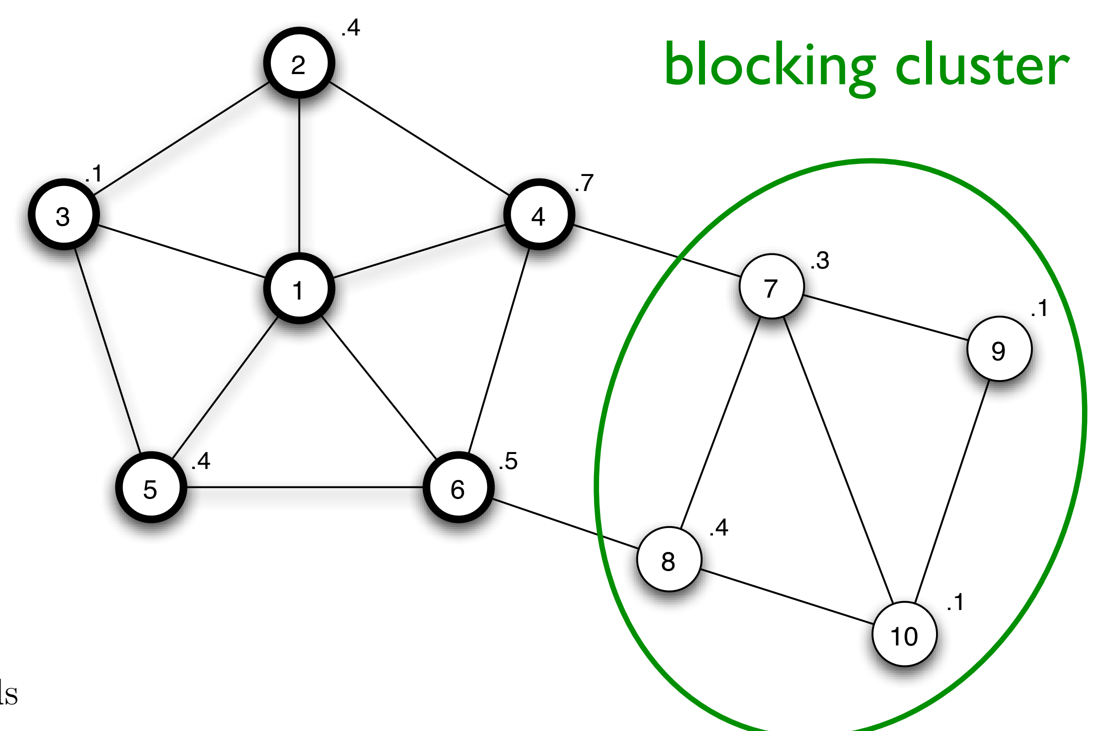

(b) *The process ends after four steps*

all members v have at least 1-q[v] neighbors also in blocking cluster

Figure 19.13: Starting with node 1 as the unique initial adopter, the new behavior A spreads to some but not all of the remaining nodes.

Easley & Kleinberg

Cascade Capacity

- What is the maximum adoption threshold that can be overcome by a small number of early adopters to create a complete cascade?
  - This is the cascade capacity of a network — function of the network topology
- More formally, consider an infinite network and consider the max. threshold for which some finite set of early adopters can cause a complete cascade

Cascade Capacity

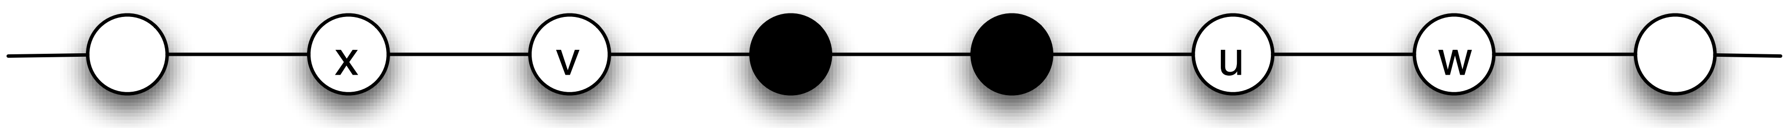

Figure 19.15: An infinite path with a set of early adopters of behavior A (shaded).

Easley & Kleinberg

- Cascade capacity is 1/2 — why?

Cascade Capacity

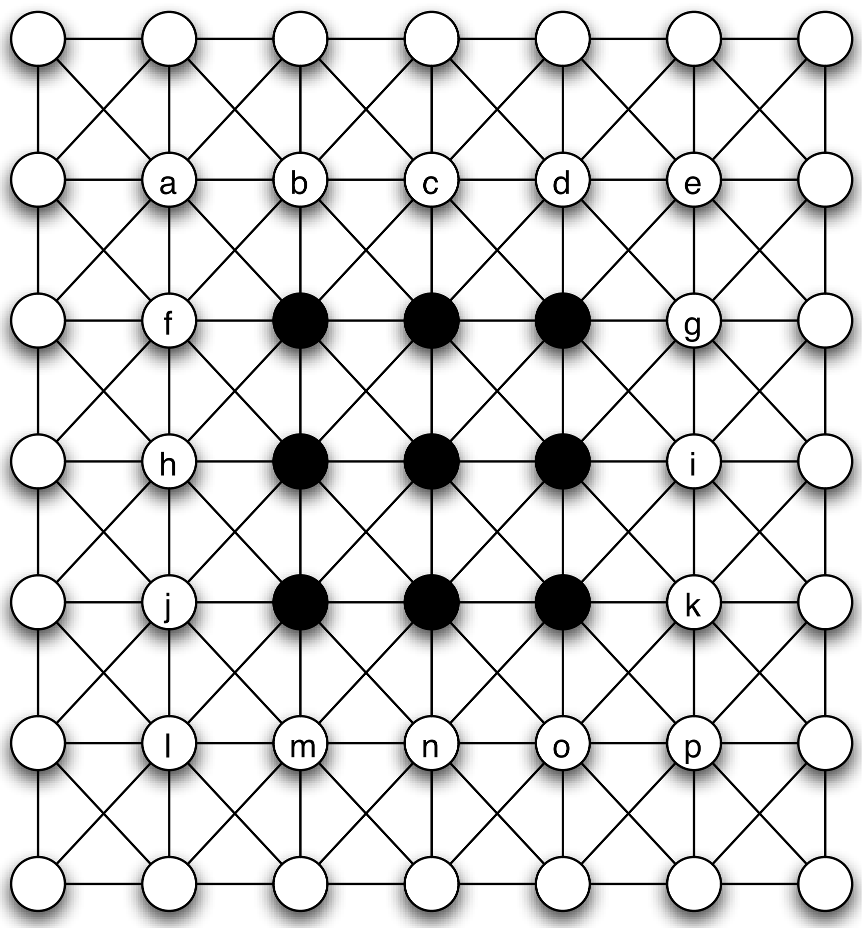

Figure 19.16: An infinite grid with a set of early adopters of behavior A (shaded).

Easley & Kleinberg

- Cascade capacity is 3/8 — why?

Maximum Cascade Capacity

- The maximum cascade capacity of any network is 1/2
  - Makes sense intuitively since $q>1/2$ means that B is favorable to A or that the old way of doing business is better
    - Cannot expect an inferior technology to displace a superior, entrenched technology (in this simple model)

Maximum Cascade Capacity

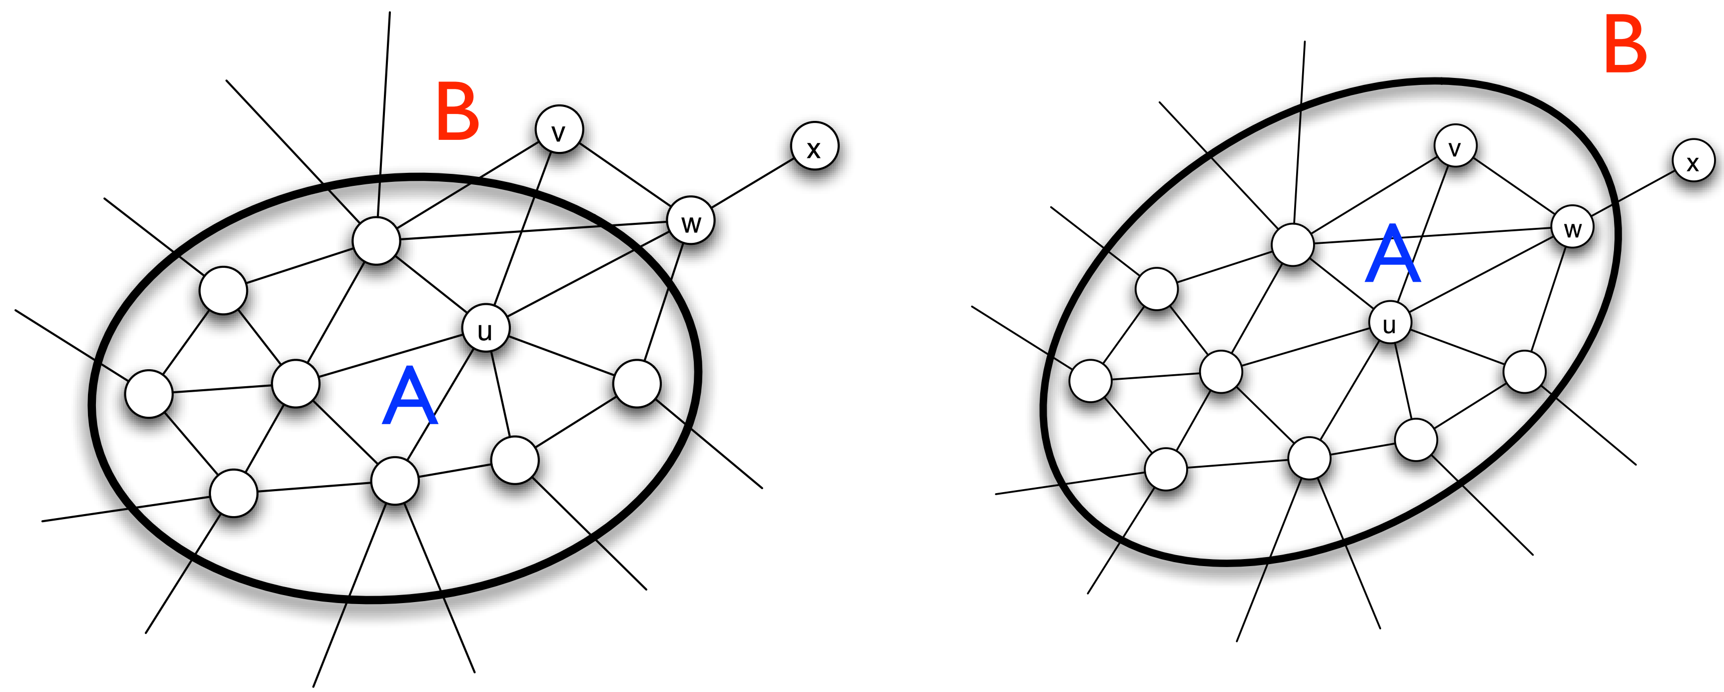

(a) *Before v and w adopt A*   (b) *After v and w adopt A*

Figure 19.17: Let the nodes inside the dark oval be the adopters of A. One step of the process is shown, in which v and w adopt A: after they adopt, the size of the interface has strictly decreased. In general, the size of the interface strictly decreases with each step of the process when $q > \frac{1}{2}$.

Easley & Kleinberg

- Look at the boundary of edges with A-B connections
  - Does this contract? If so, the meme will die out — q>1/2 always dies out

“Bilingual” Model of Adoption

- In many cases, instead of abandoning B for A in a wholesale manner, people will use both A and B — becoming “bilingual”
  - Windows and OS X (virtual machines, multiple machines)
  - Being bilingual should incur an extra cost — c
- This new state “AB” can be transient or stable…
- How does this change cascades and cascade capacity
  - Makes things much more complicated!

“Bilingual” Model of Adoption

			w	
		A	B	AB
	A	a, a	$0, 0$	a, a
v	B	$0, 0$	b, b	b, b
	AB	a, a	b, b	$(a, b)^+, (a, b)^+$

AB costs the adopting node c

Figure 19.18: A Coordination Game with a bilingual option. Here the notation $(a, b)^+$ denotes the larger of a and b.

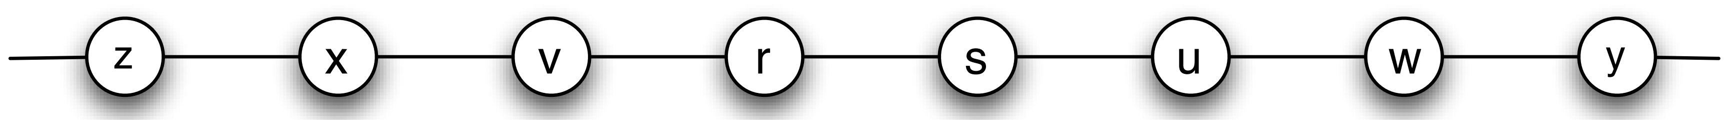

Figure 19.19: An infinite path, with nodes r and s as initial adopters of A.

Easley & Kleinberg

- Consider only the 1-d line network for this more complex adoption rule model…

“Bilingual” Model of Adoption

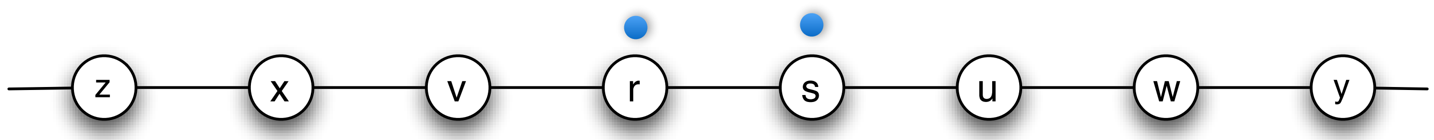

Start	B	B	B	A	A	B	B	B
Step 1	B	B	AB	A	A	AB	B	B
Step 2	B	AB	AB	A	A	AB	AB	B
Step 3	AB	AB	A	A	A	A	AB	AB
Step 4	AB	A	A	A	A	A	A	AB

Figure 19.20: With payoffs $a = 5$ and $b = 3$ for interaction using A and B respectively, and a cost $c = 1$ for being bilingual, the strategy A spreads outward from the initial adopters r and s through a two-phase structure. First, the strategy AB spreads, and then behind it, nodes switch permanently from AB to A.

Easley & Kleinberg

What happens if a=2, b=3, c=1?

"Bilingual" Model of Adoption

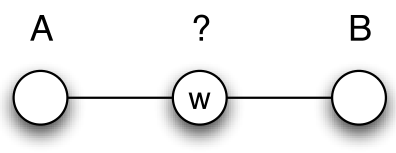

payoff from choosing A: a
payoff from choosing B: 1
payoff from choosing AB: a + 1 - c

b is normalized to 1

Figure 19.21: The payoffs to a node on the infinite path with two neighbors using A and B.

Easley & Kleinberg

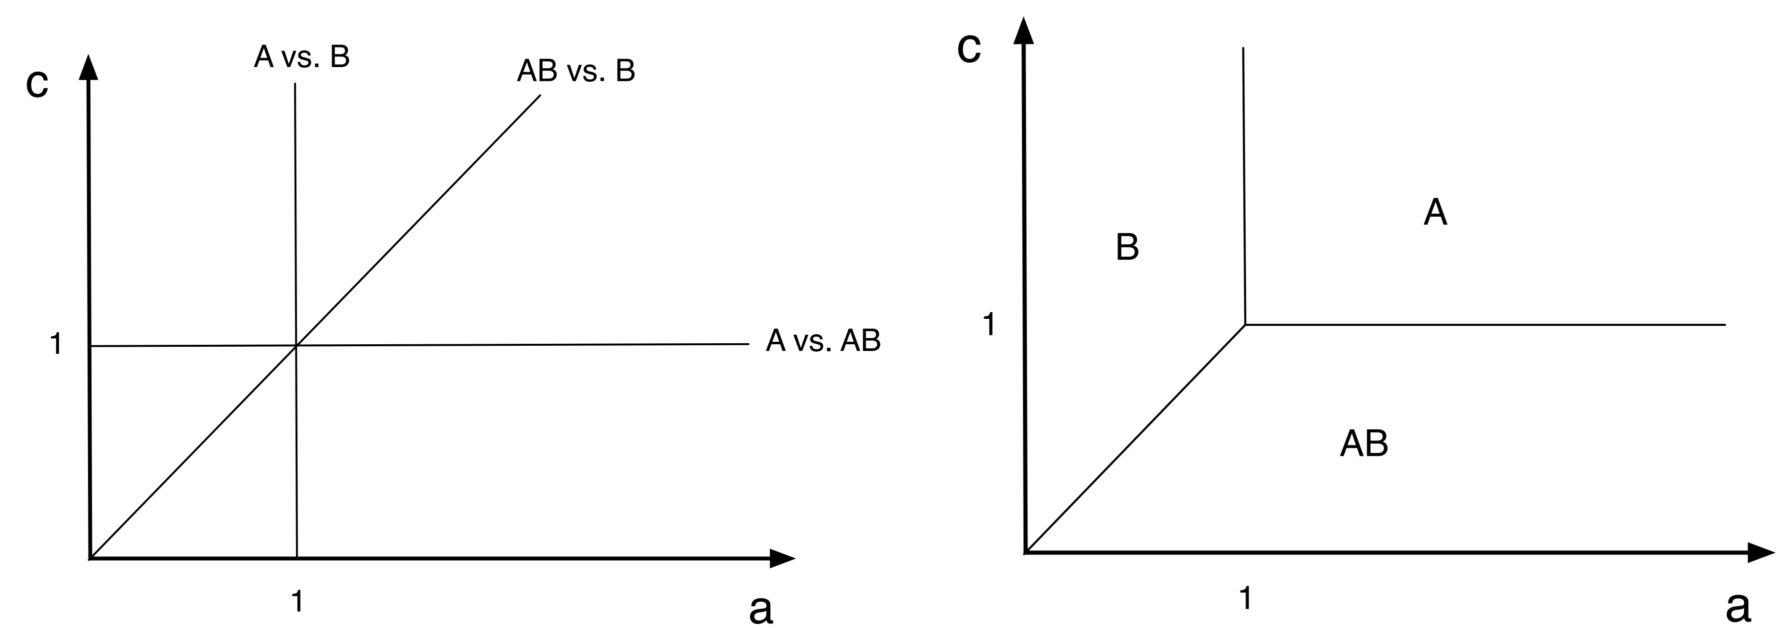

(a) *Lines showing break-even points between strategies.* (b) *Regions defining the best choice of strategy.*

Figure 19.22: Given a node with neighbors using A and B, the values of a and c determine which of the strategies A, B, or AB it will choose. (Here, by re-scaling, we can assume $b = 1$.) We can represent the choice of strategy as a function of a and c by dividing up the (a, c)-plane into regions corresponding to different choices.

“Bilingual” Model of Adoption

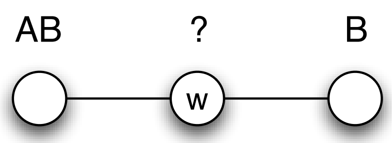

b is normalized to 1

payoff from choosing A: a
payoff from choosing B: 2
payoff from choosing AB: a + 1 - c (if A is better)

Figure 19.23: The payoffs to a node on the infinite path with two neighbors using AB and B.

Easley & Kleinberg

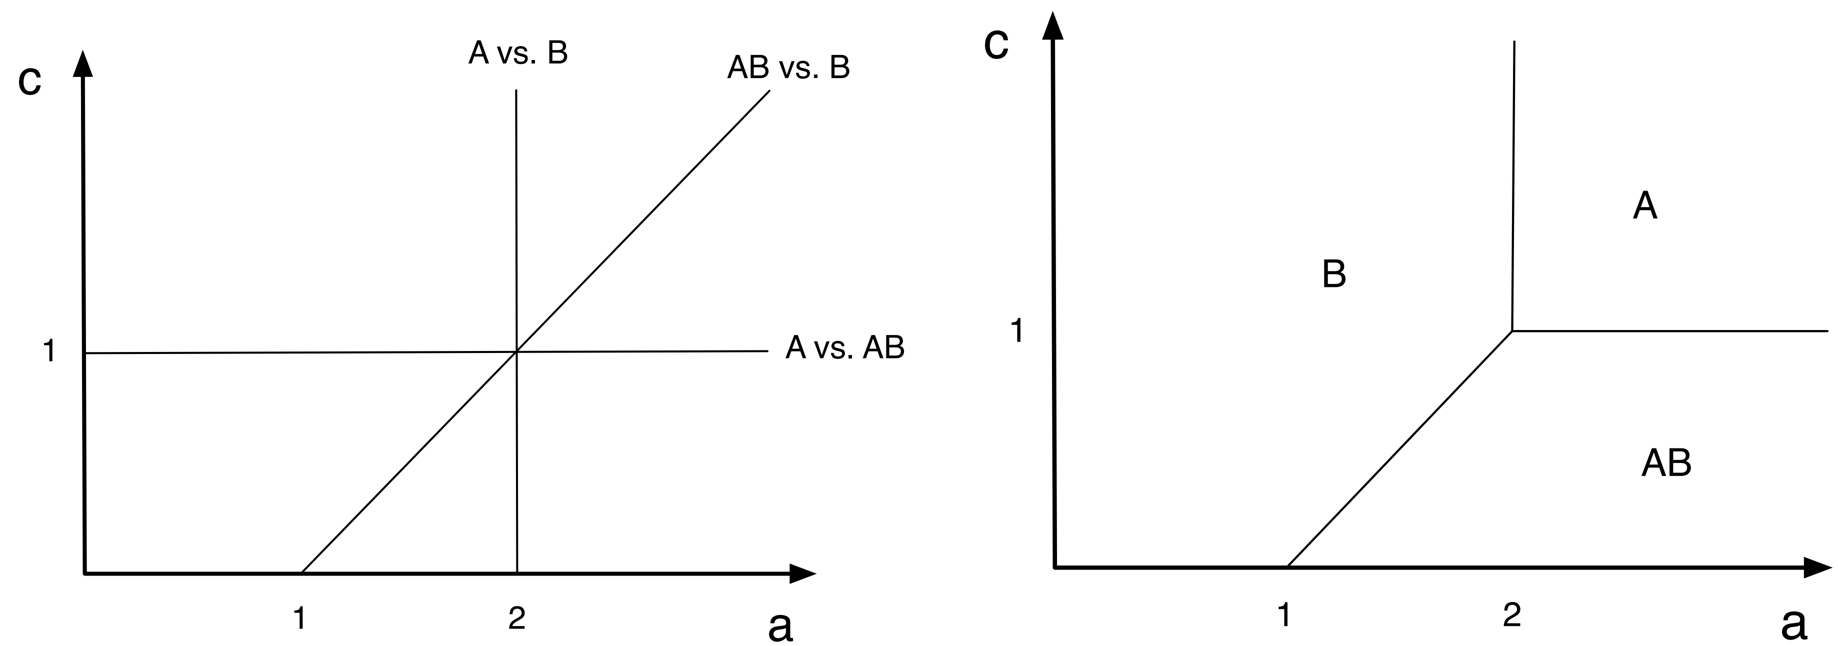

(a) *Lines showing break-even points between strategies.* (b) *Regions defining the best choice of strategy.*

Figure 19.24: Given a node with neighbors using AB and B, the values of a and c determine which of the strategies A, B, or AB it will choose, as shown by this division of the (a, c)-plane into regions.

"Bilingual" Model of Adoption

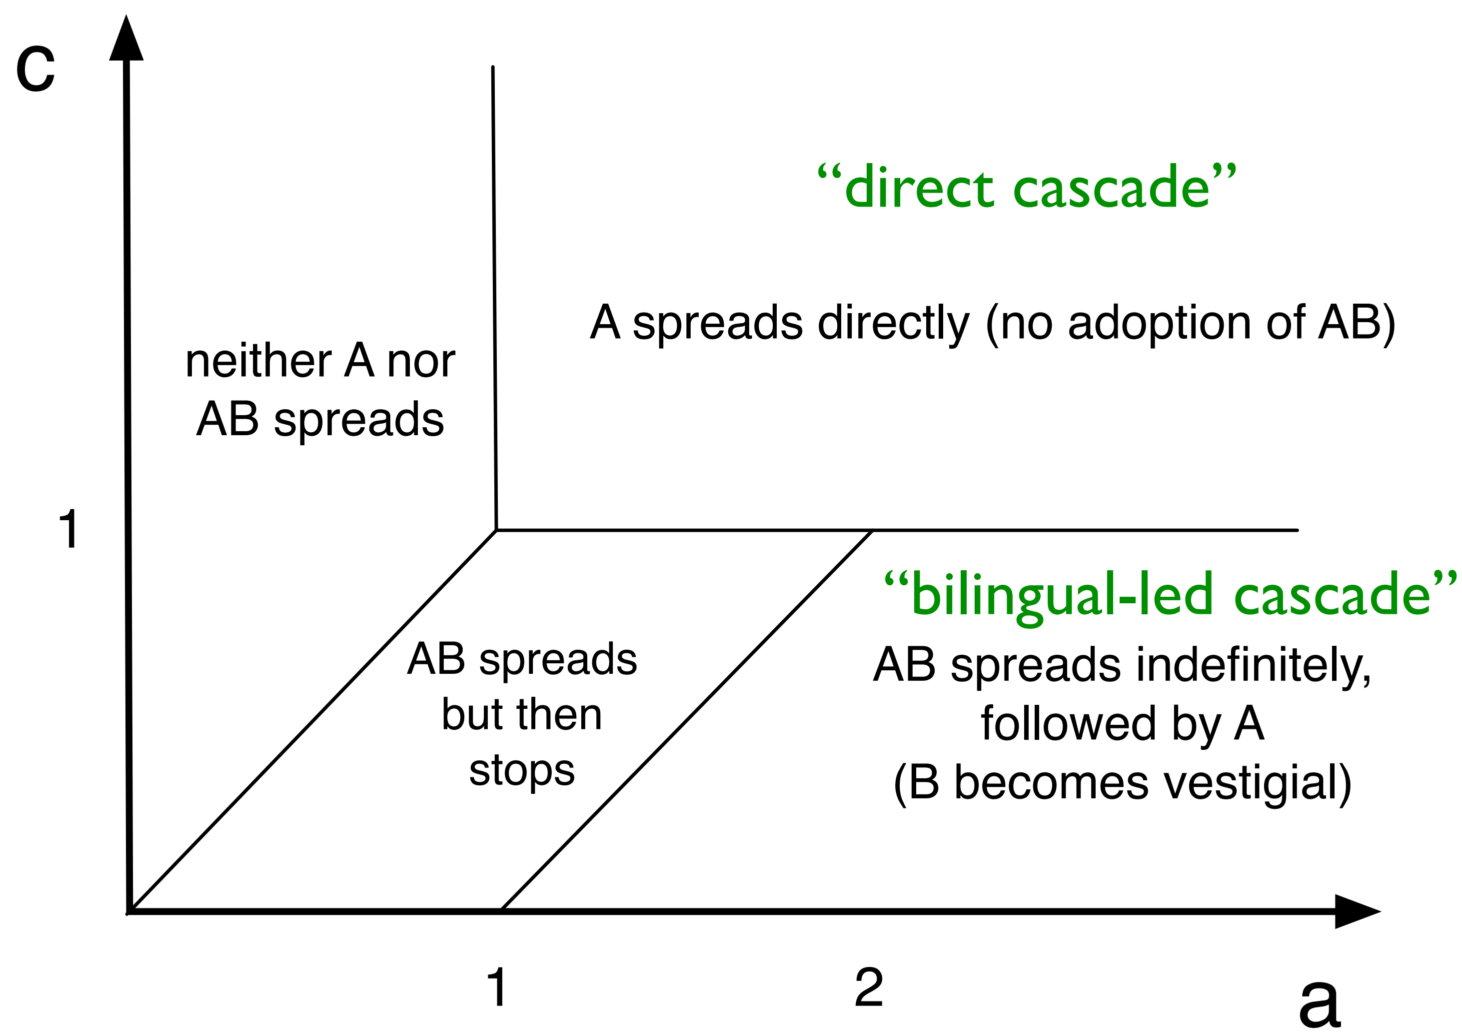

Figure 19.25: There are four possible outcomes for how A spreads or fails to spread on the infinite path, indicated by this division of the (a, c)-plane into four regions.

Easley & Kleinberg

“Bilingual” Model of Adoption

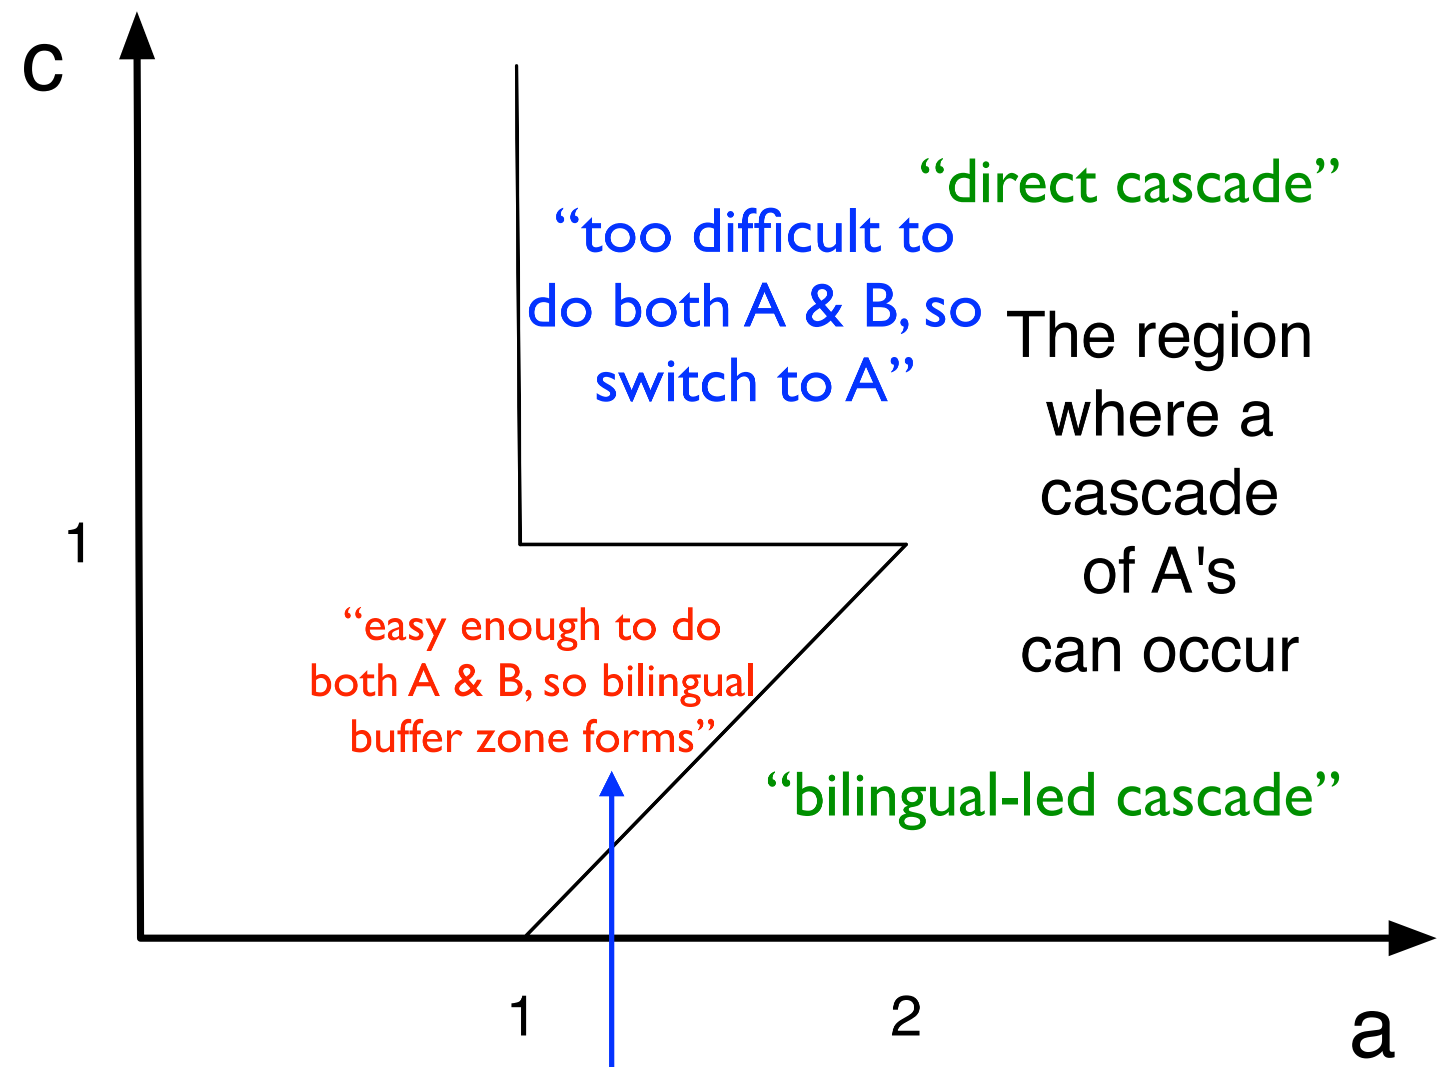

moderate compatibility inhibits full adoption

Figure 19.26: The set of values for which a cascade of A’s can occur defines a region in the (a, c)-plane consisting of a vertical line with a triangular “cut-out.”

Easley & Kleinberg

A is still superior to B here

Thresholds for Collective Action

pluralistic ignorance

Nodes know neighbors threshold

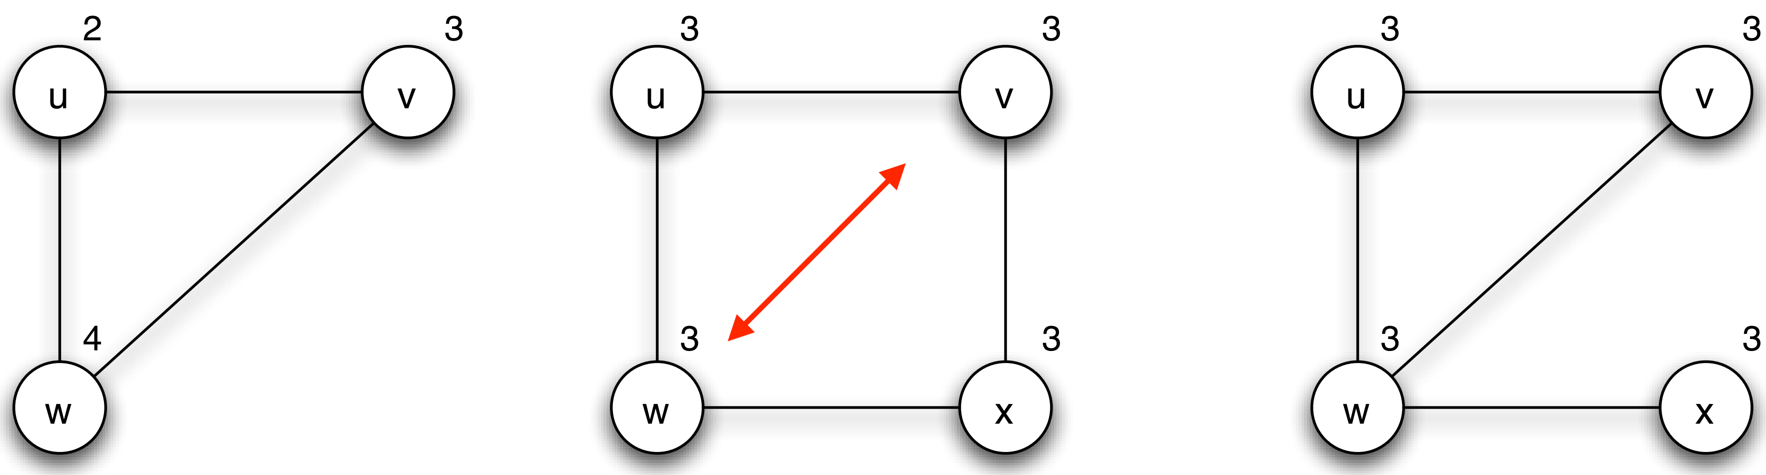

(a) *An uprising will not occur* (b) *An uprising will not occur* (c) *An uprising can occur*

Figure 19.14: Each node in the network has a threshold for participation, but only knows the threshold of itself and its neighbors.

Easley & Kleinberg

- In order to join the "rebellion" a node needs to exceed its threshold of required participants AND know that others others exceed theirs as well

Need tightly coupled/informed communities for risky initiatives

http://www.youtube.com/watch?v=axSnW-ygU5g

Thresholds for Collective Action

- Apple Macintosh commercial during 1984 Super Bowl
  - “Best commercial of all time”
- Mac was a high risk adoption — requiring collective action
  - Mac was expensive and incompatible
  - Fewer applications
- Big part of advertisment’s effectiveness was ensuring collective knowledge
  - Potential adopters now know that other potential adopters know!

http://www.youtube.com/watch?v=axSnW-ygU5g

Maximum Influence Problem

- What is the smallest size set of nodes in a network that if initialized as early adopters will cause a cascade?
  - NP-hard (need to find heuristics for large networks)

West Point Network Science Center
(Pre-Print Manuscript)

A Scalable Heuristic for Viral Marketing Under the Tipping Model

Paulo Shakarian · Sean Eyre · Damon Paulo

recent approach that seems simple and effective

http://arxiv.org/abs/1309.2963

MIP Heuristic (TIP_DECOMP)

- Consider a specific example

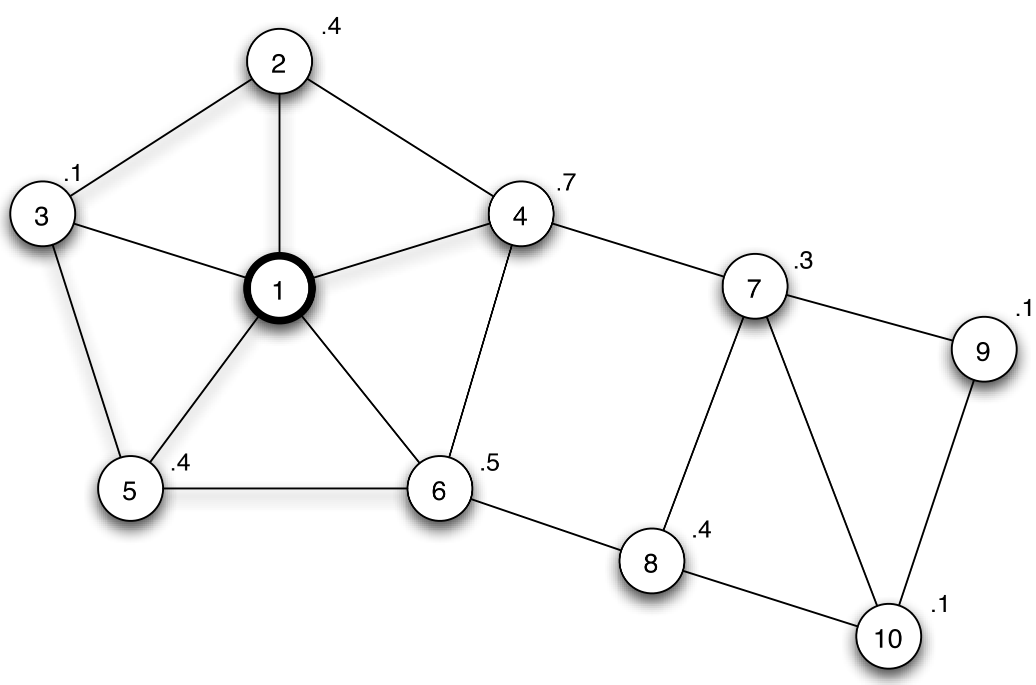

(a) *One node is the initial adopter*

can we find a seed set?

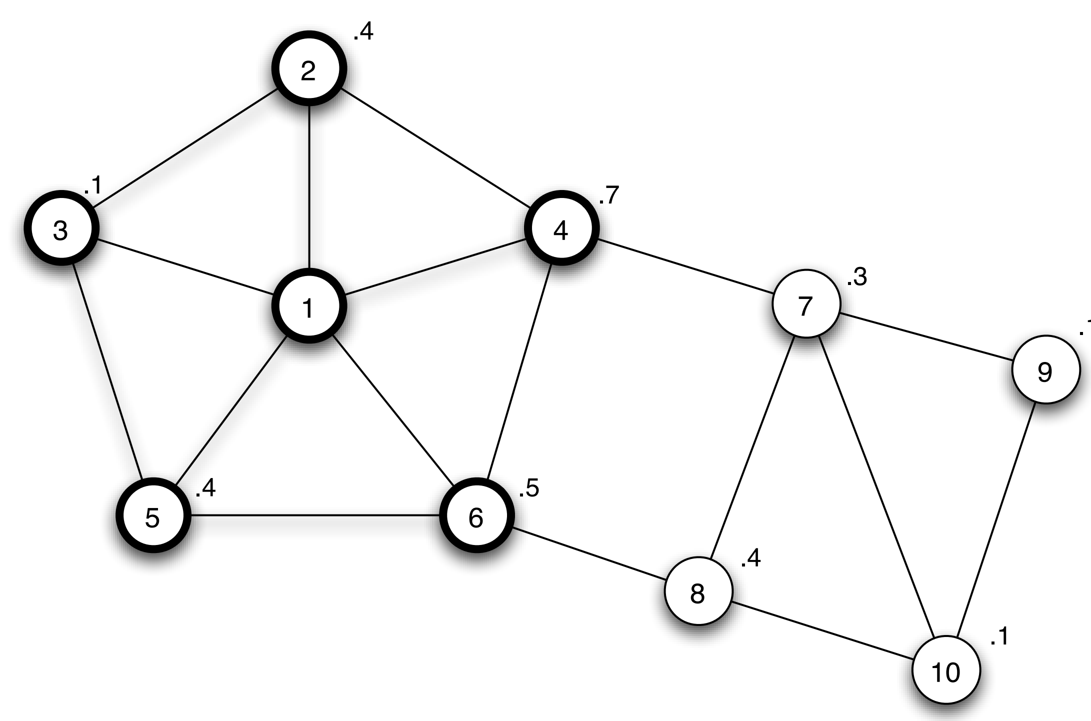

(b) *The process ends after four steps*

Easley & Kleinberg

MIP Heuristic (TIP_DECOMP)

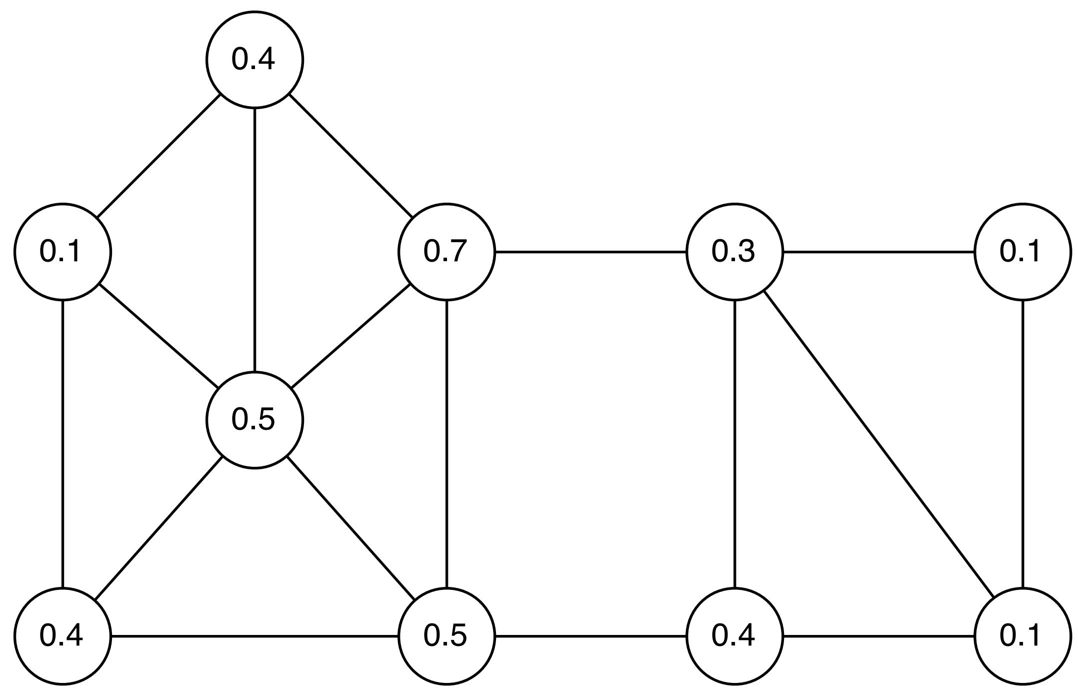

start with thresholds — q[i]

MIP Heuristic (TIP_DECOMP)

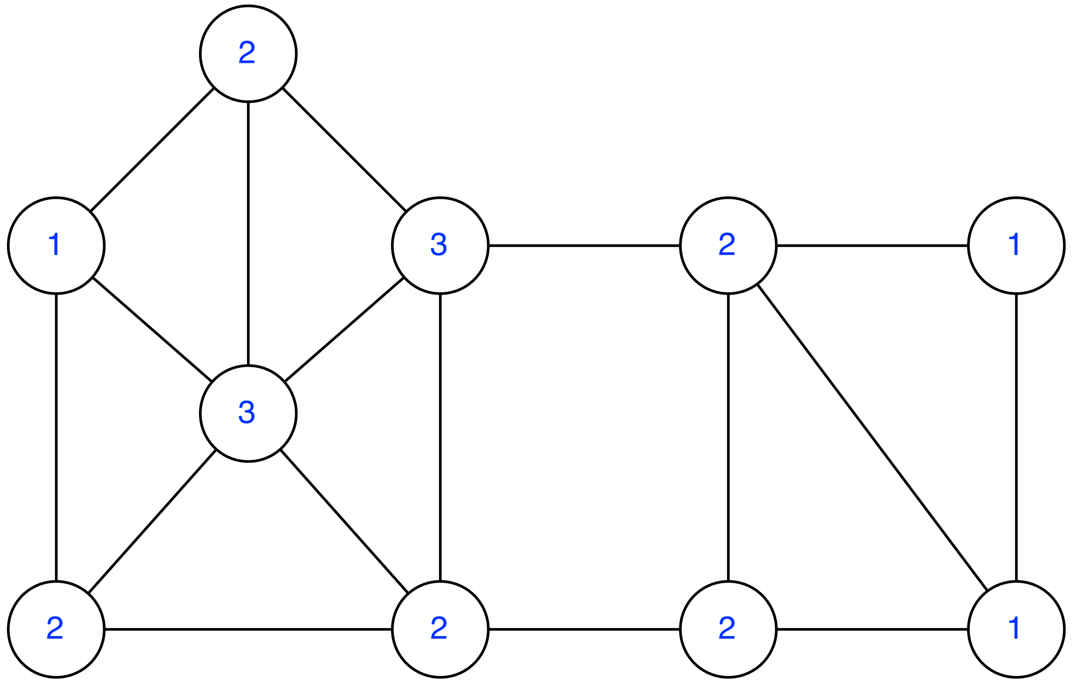

replace thresholds by "edge threshold"

$$m_i = \lceil q_i k_i \rceil$$

MIP Heuristic (TIP_DECOMP)

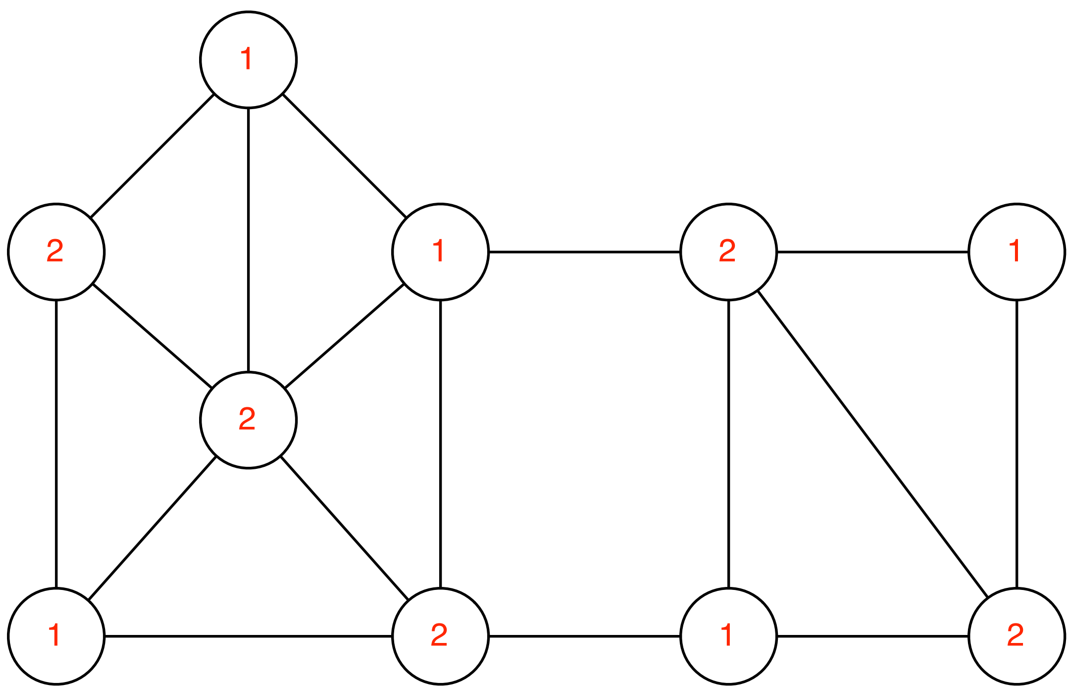

label nodes by "free edges"

$$\Delta_i = k_i - m_i = k_i - \lceil q_i k_i \rceil$$

MIP Heuristic (TIP_DECOMP)

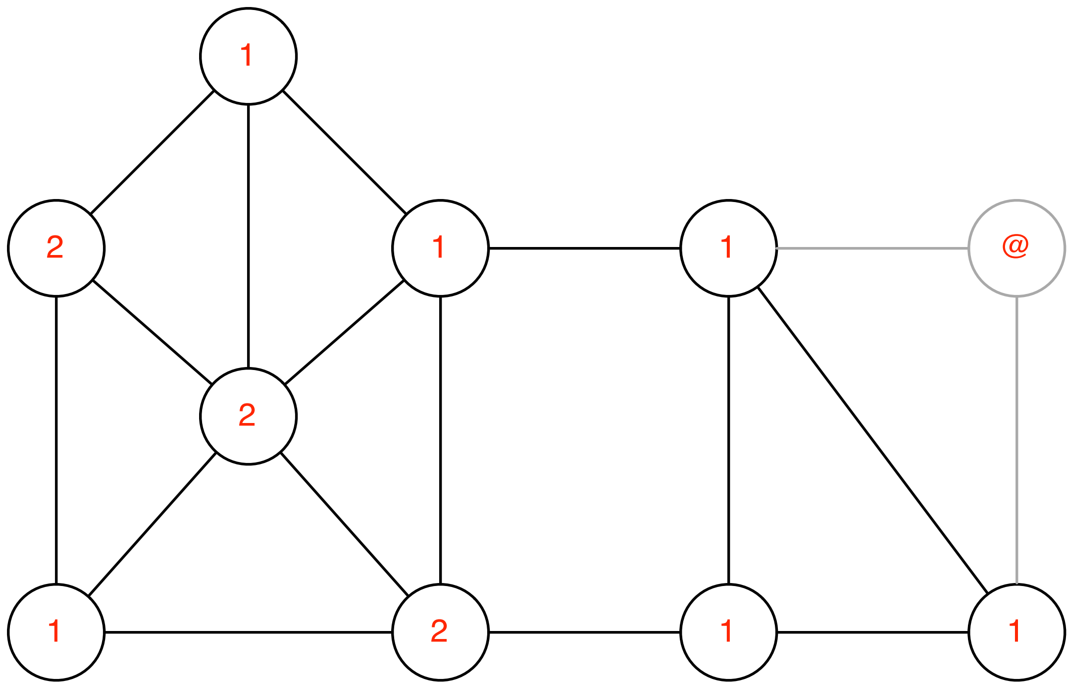

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node

MIP Heuristic (TIP_DECOMP)

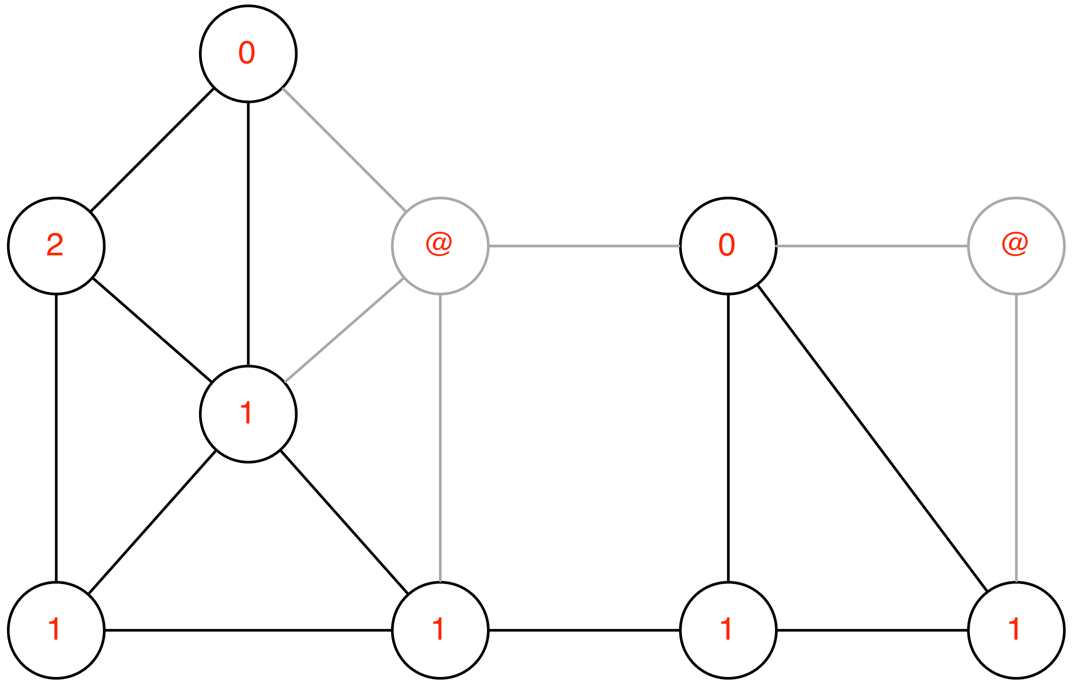

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node

MIP Heuristic (TIP_DECOMP)

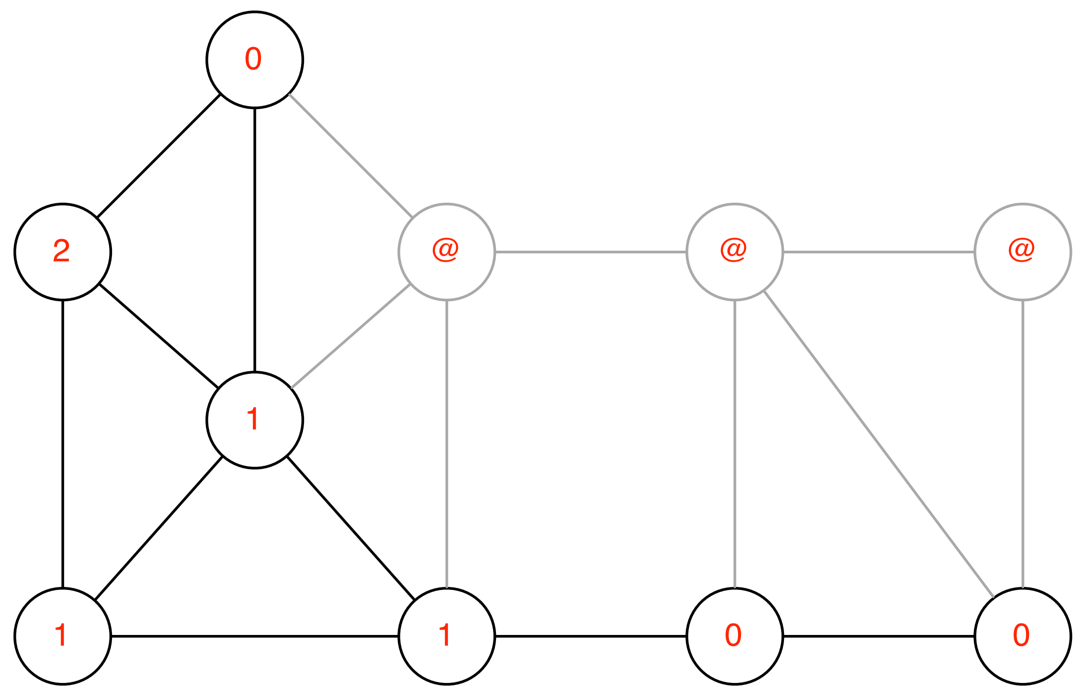

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node

MIP Heuristic (TIP_DECOMP)

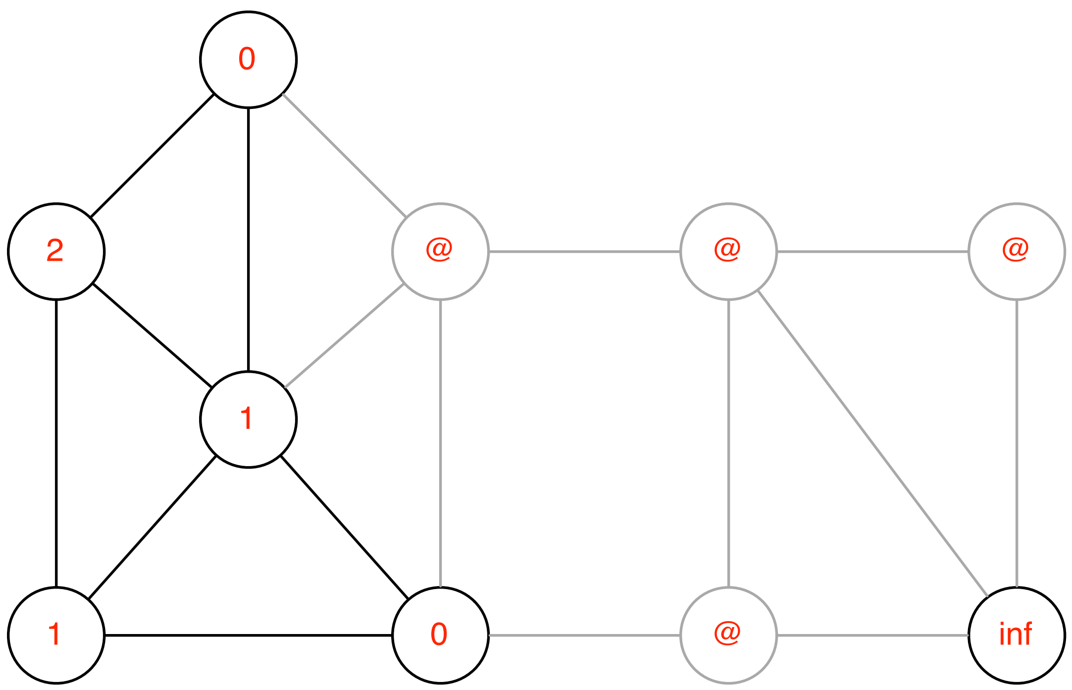

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node
3. Label nodes decremented from delta=0 by infinity

MIP Heuristic (TIP_DECOMP)

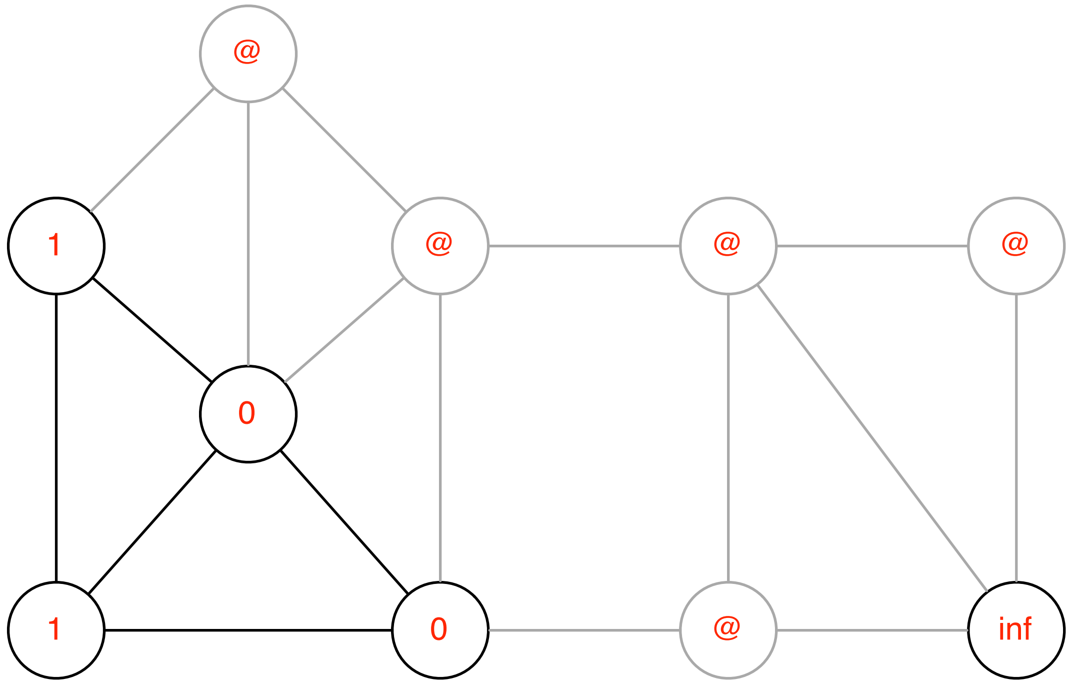

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node
3. Label nodes decremented from delta=0 by infinity

MIP Heuristic (TIP_DECOMP)

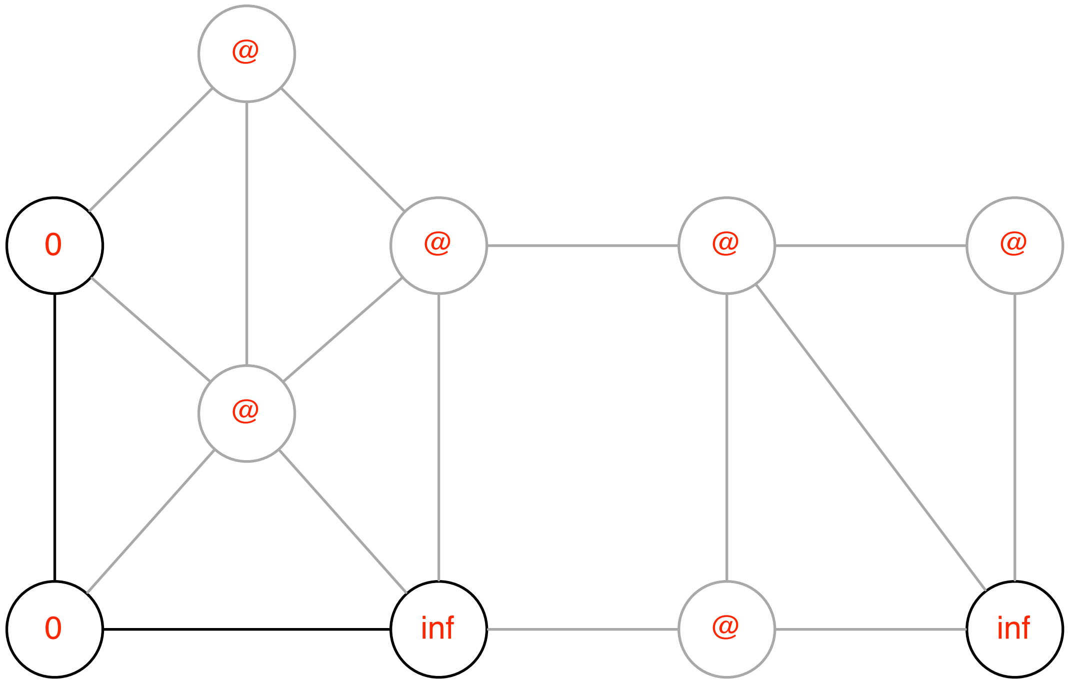

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node
3. Label nodes decremented from delta=0 by infinity

MIP Heuristic (TIP_DECOMP)

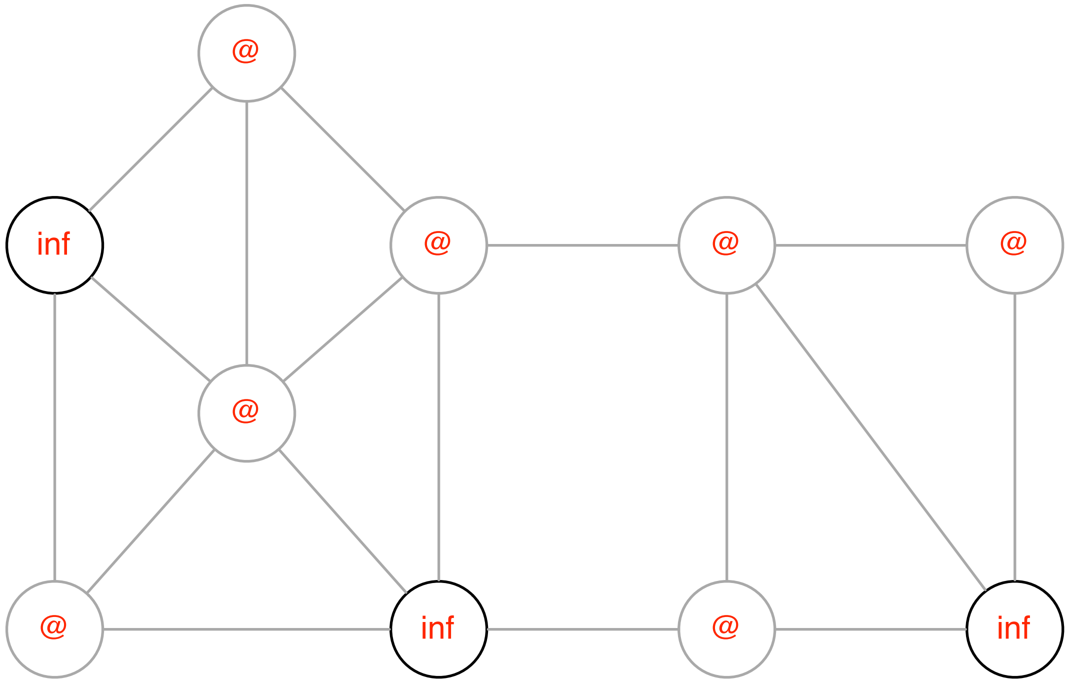

1. Remove node with smallest delta
2. Decrement the delta of all nodes connected to removed node
3. Label nodes decremented from delta=0 by infinity
4. Stop when only "inf" nodes left — these are the seed set

MIP Heuristic (TIP_DECOMP)

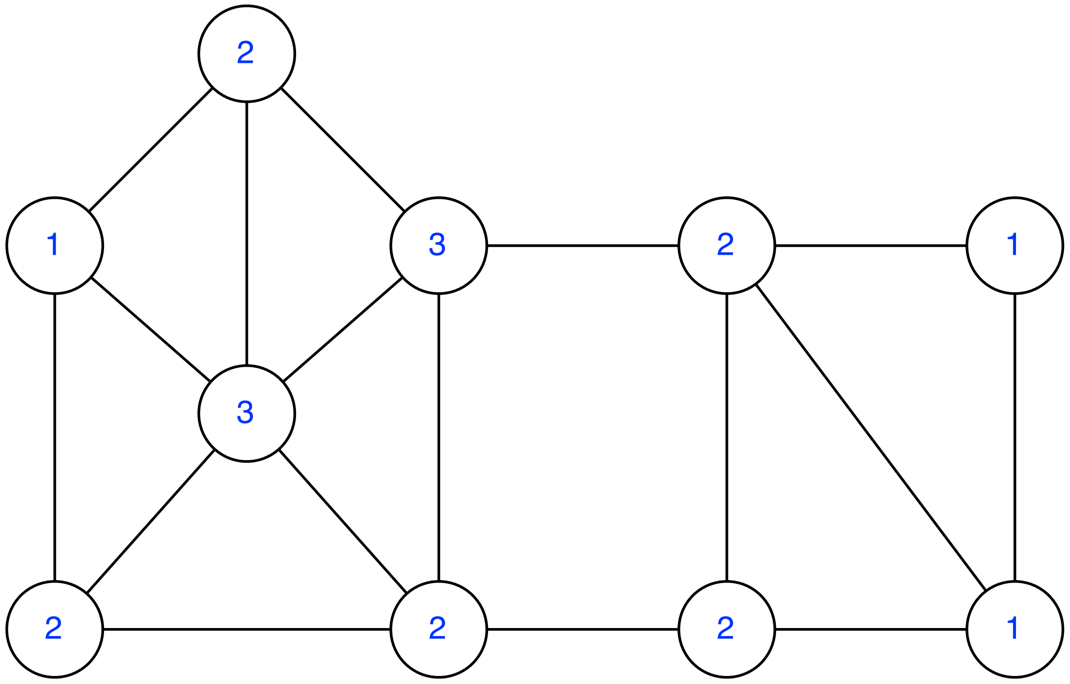

MIP Heuristic (TIP_DECOMP)

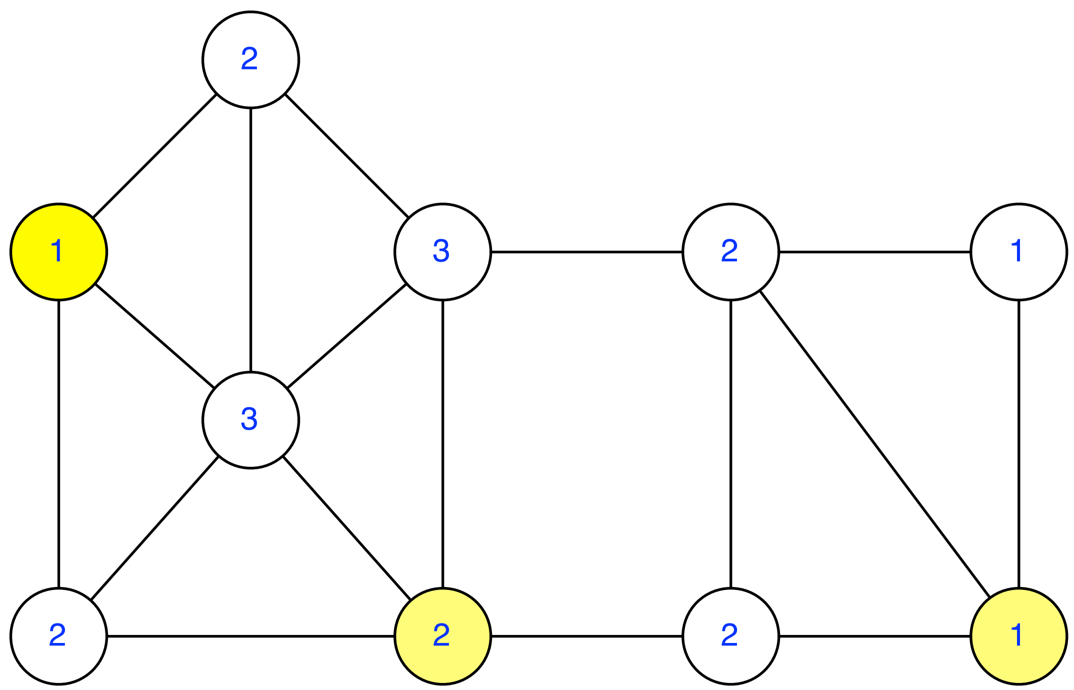

MIP Heuristic (TIP_DECOMP)

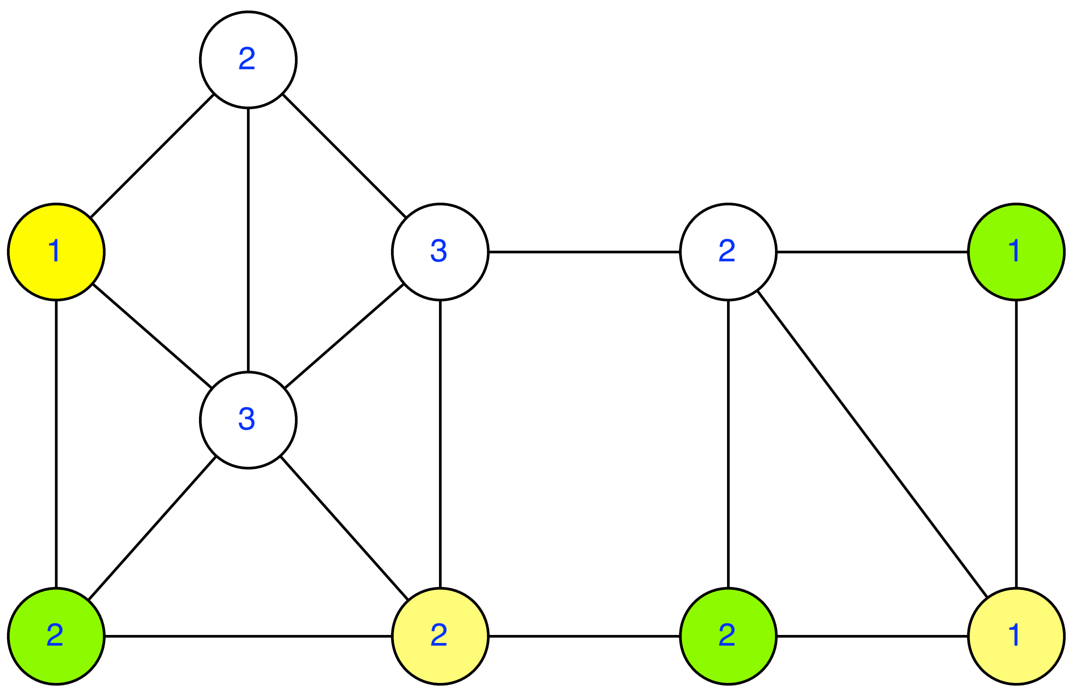

MIP Heuristic (TIP_DECOMP)

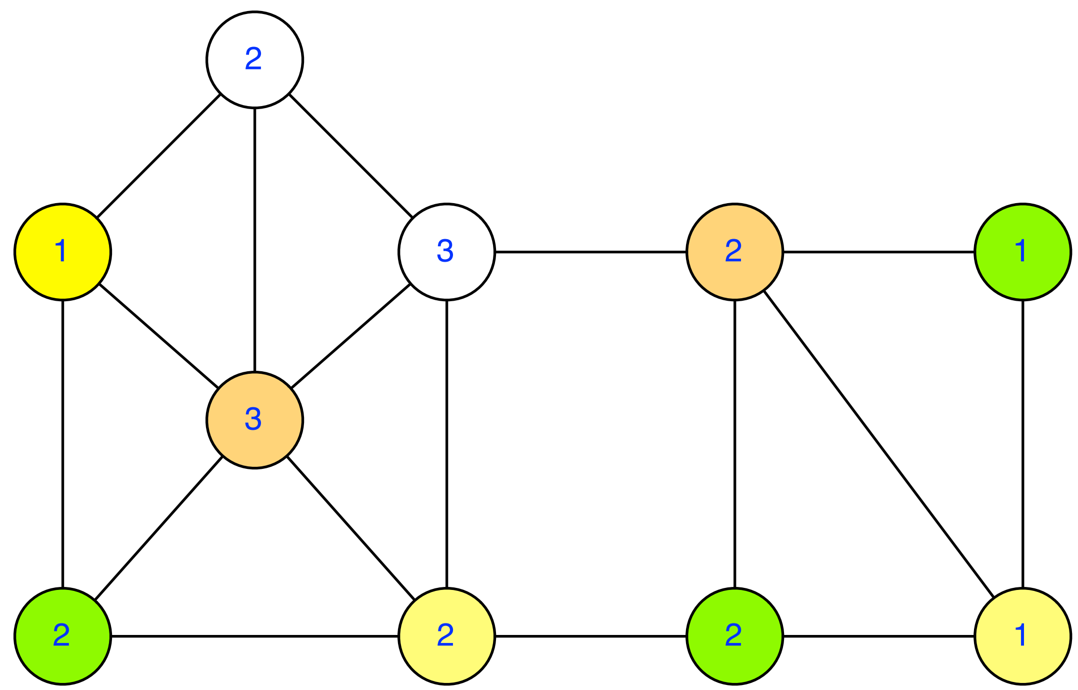

MIP Heuristic (TIP_DECOMP)

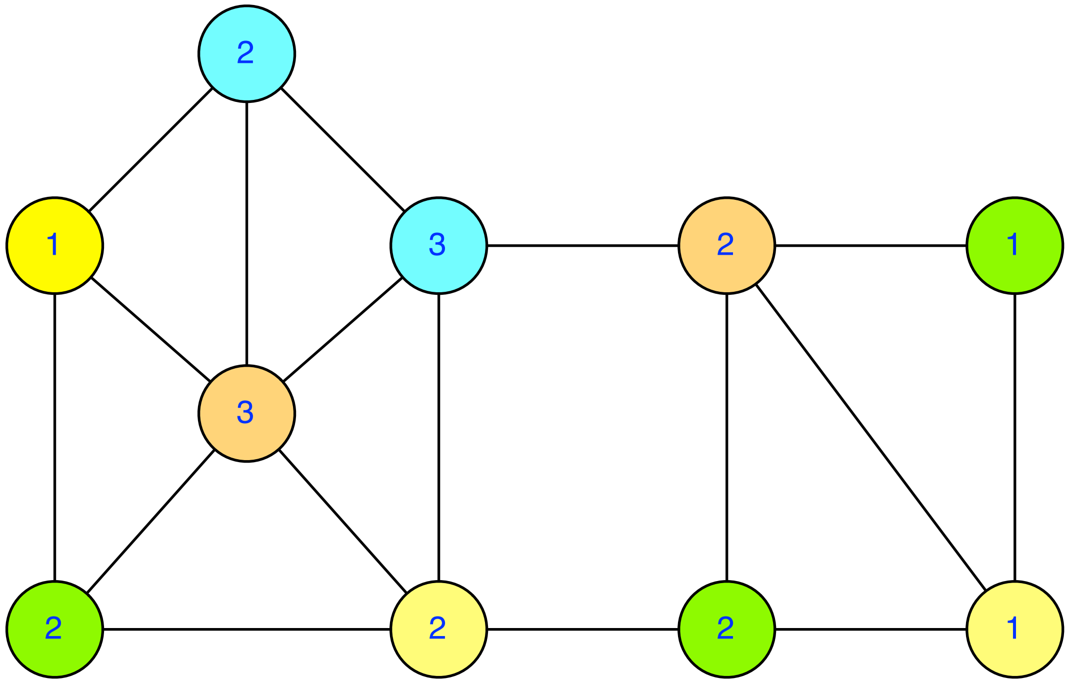

MIP Heuristic (TIP_DECOMP)

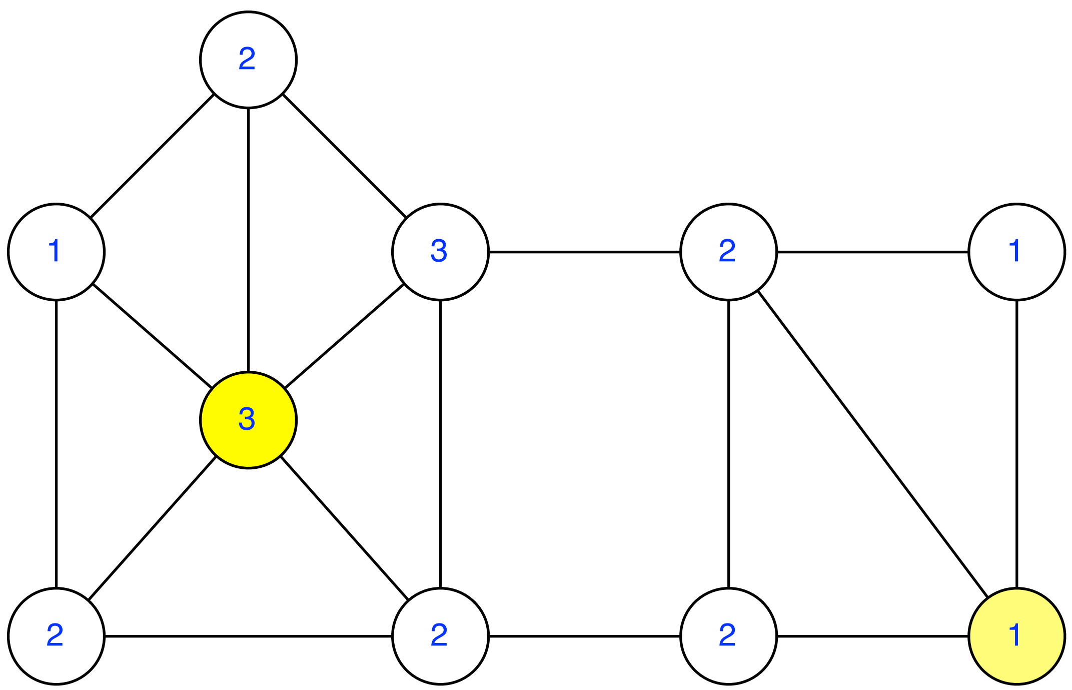

The previous seed set was not minimal

MIP Heuristic (TIP_DECOMP)

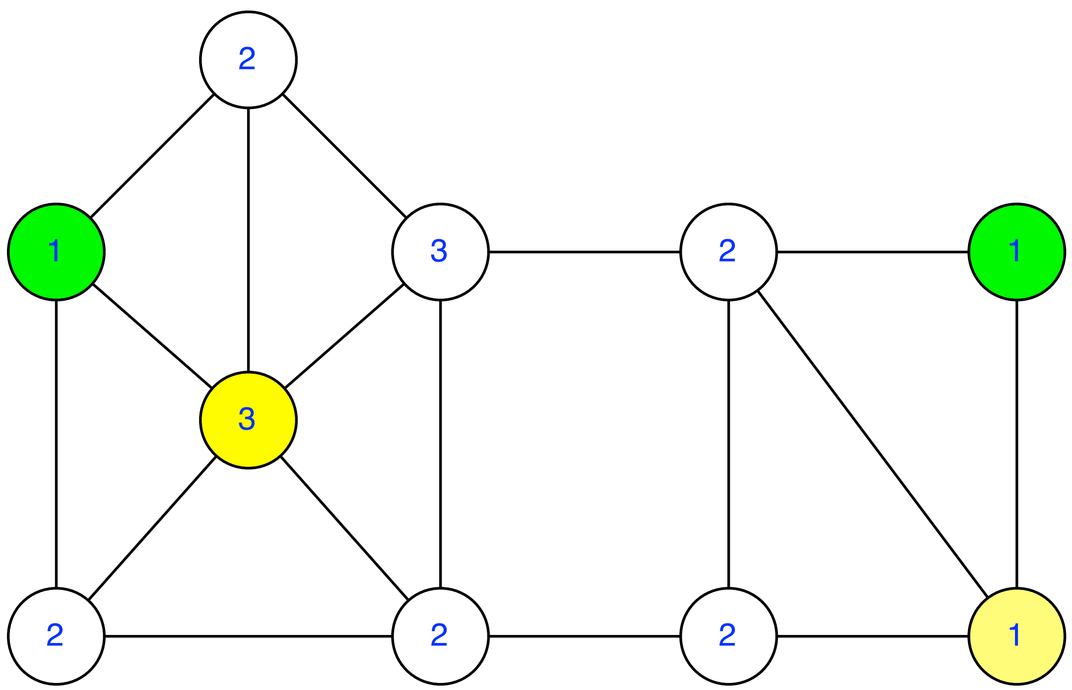

The previous seed set was not minimal

MIP Heuristic (TIP_DECOMP)

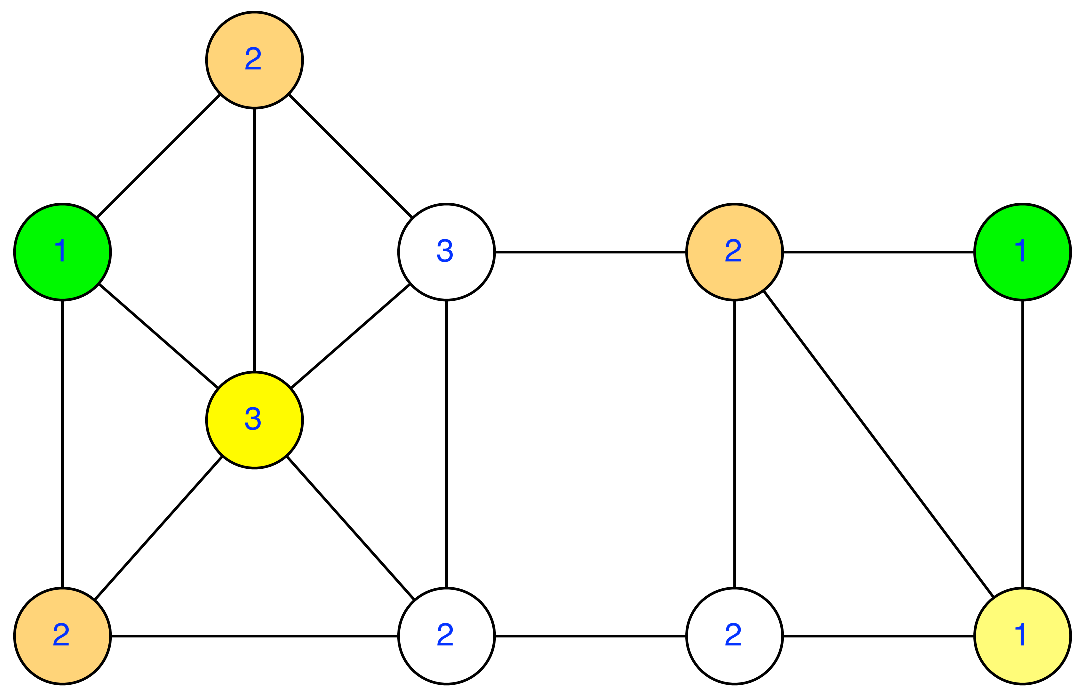

The previous seed set was not minimal

MIP Heuristic (TIP_DECOMP)

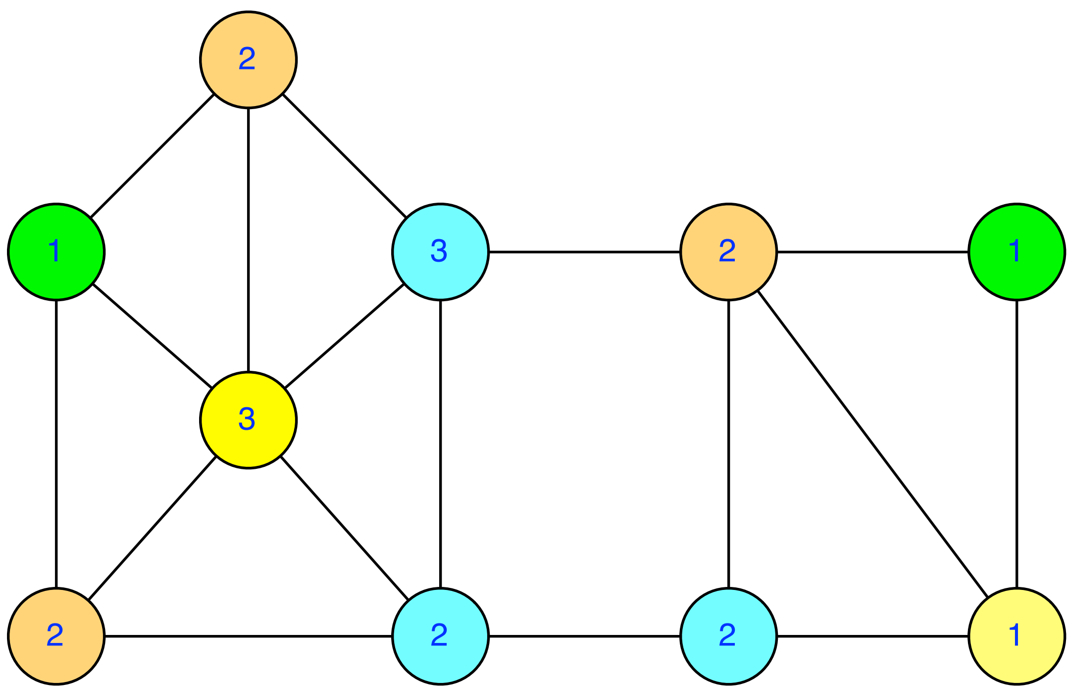

The previous seed set was not minimal

MIP Heuristic (TIP_DECOMP)

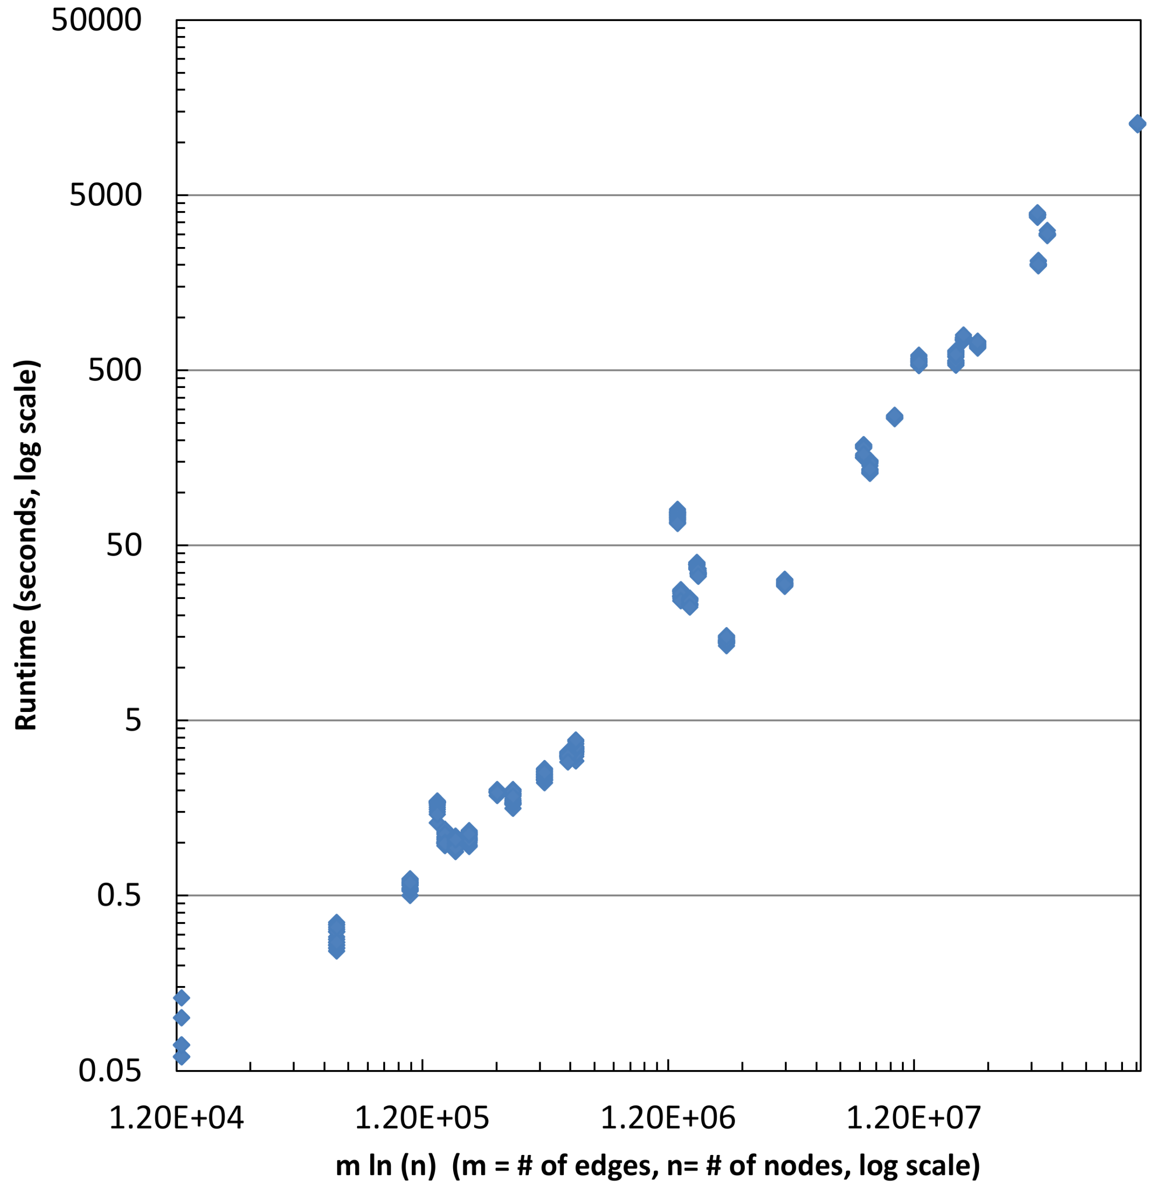

Run time is O(L*log(N))

Fig. 2 $m \ln n$ vs. runtime in seconds (log scale, m is number of edges, n is number of nodes). The relationship is linear with $R^2 = 0.9015$, $p = 2.2 \cdot 10^{-16}$.

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Seed size is <4% for fixed threshold up to 10

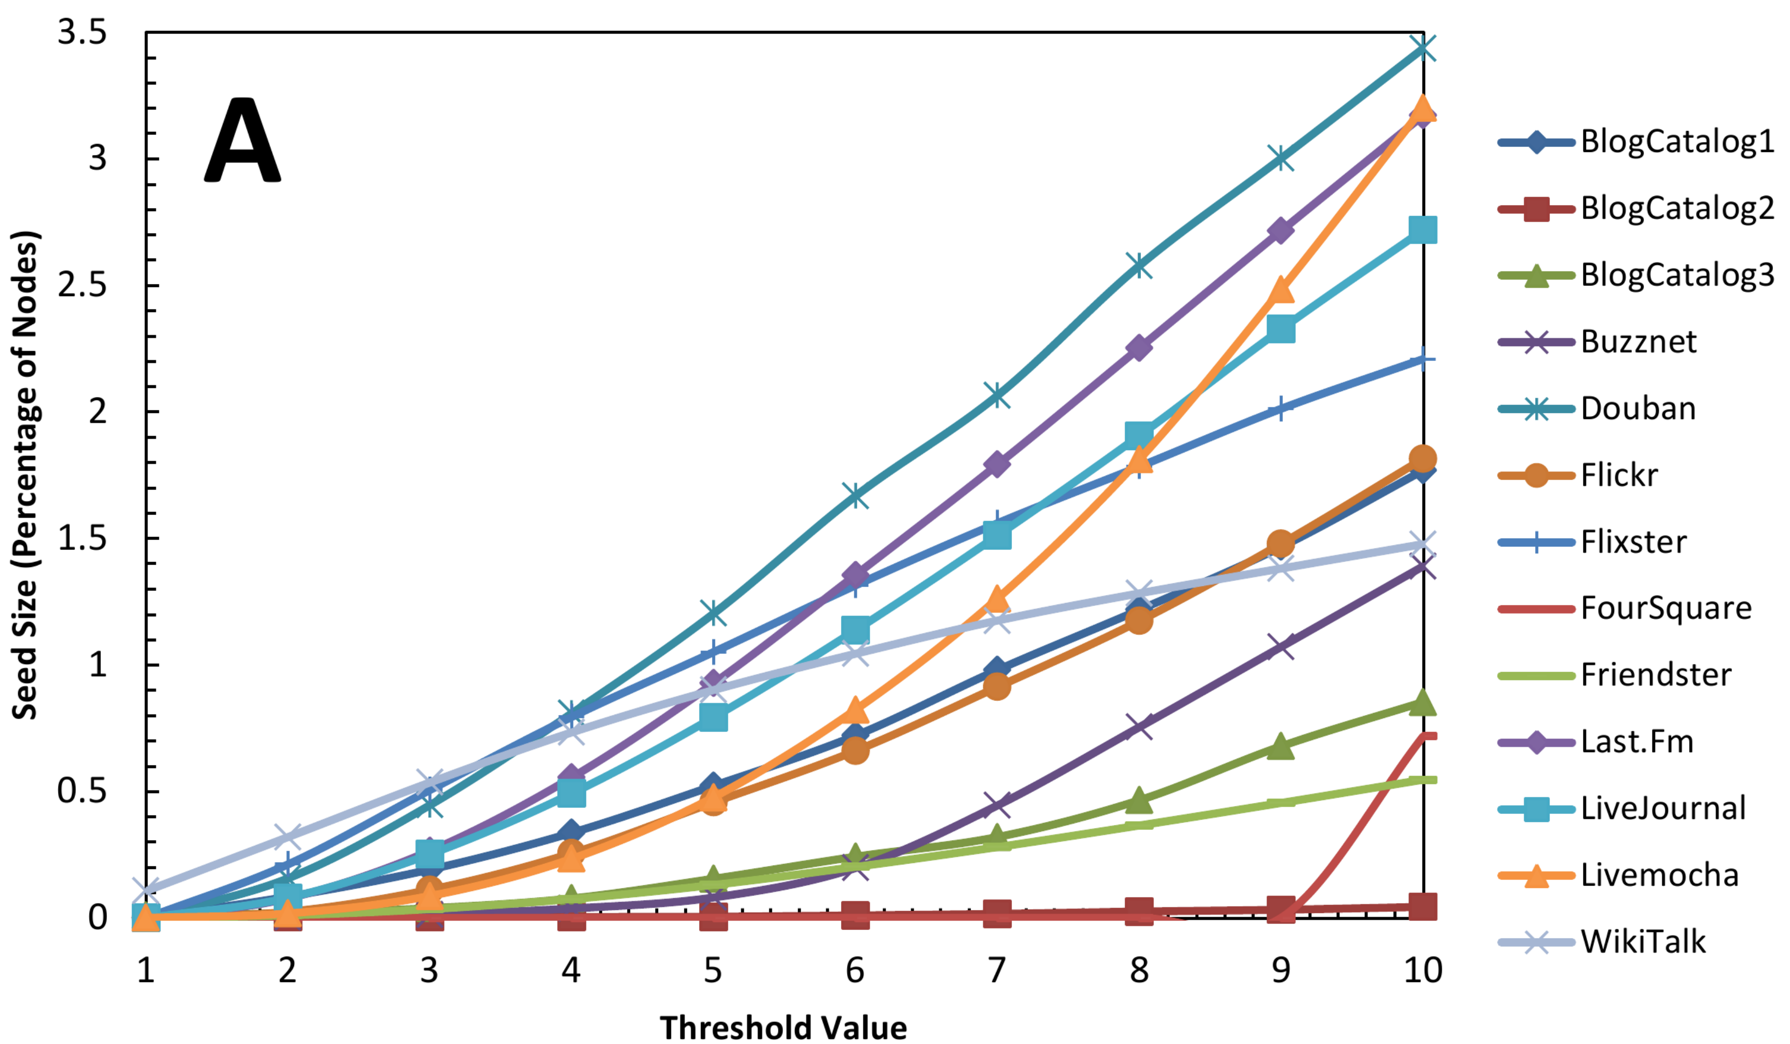

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Seed size vs. q[i]

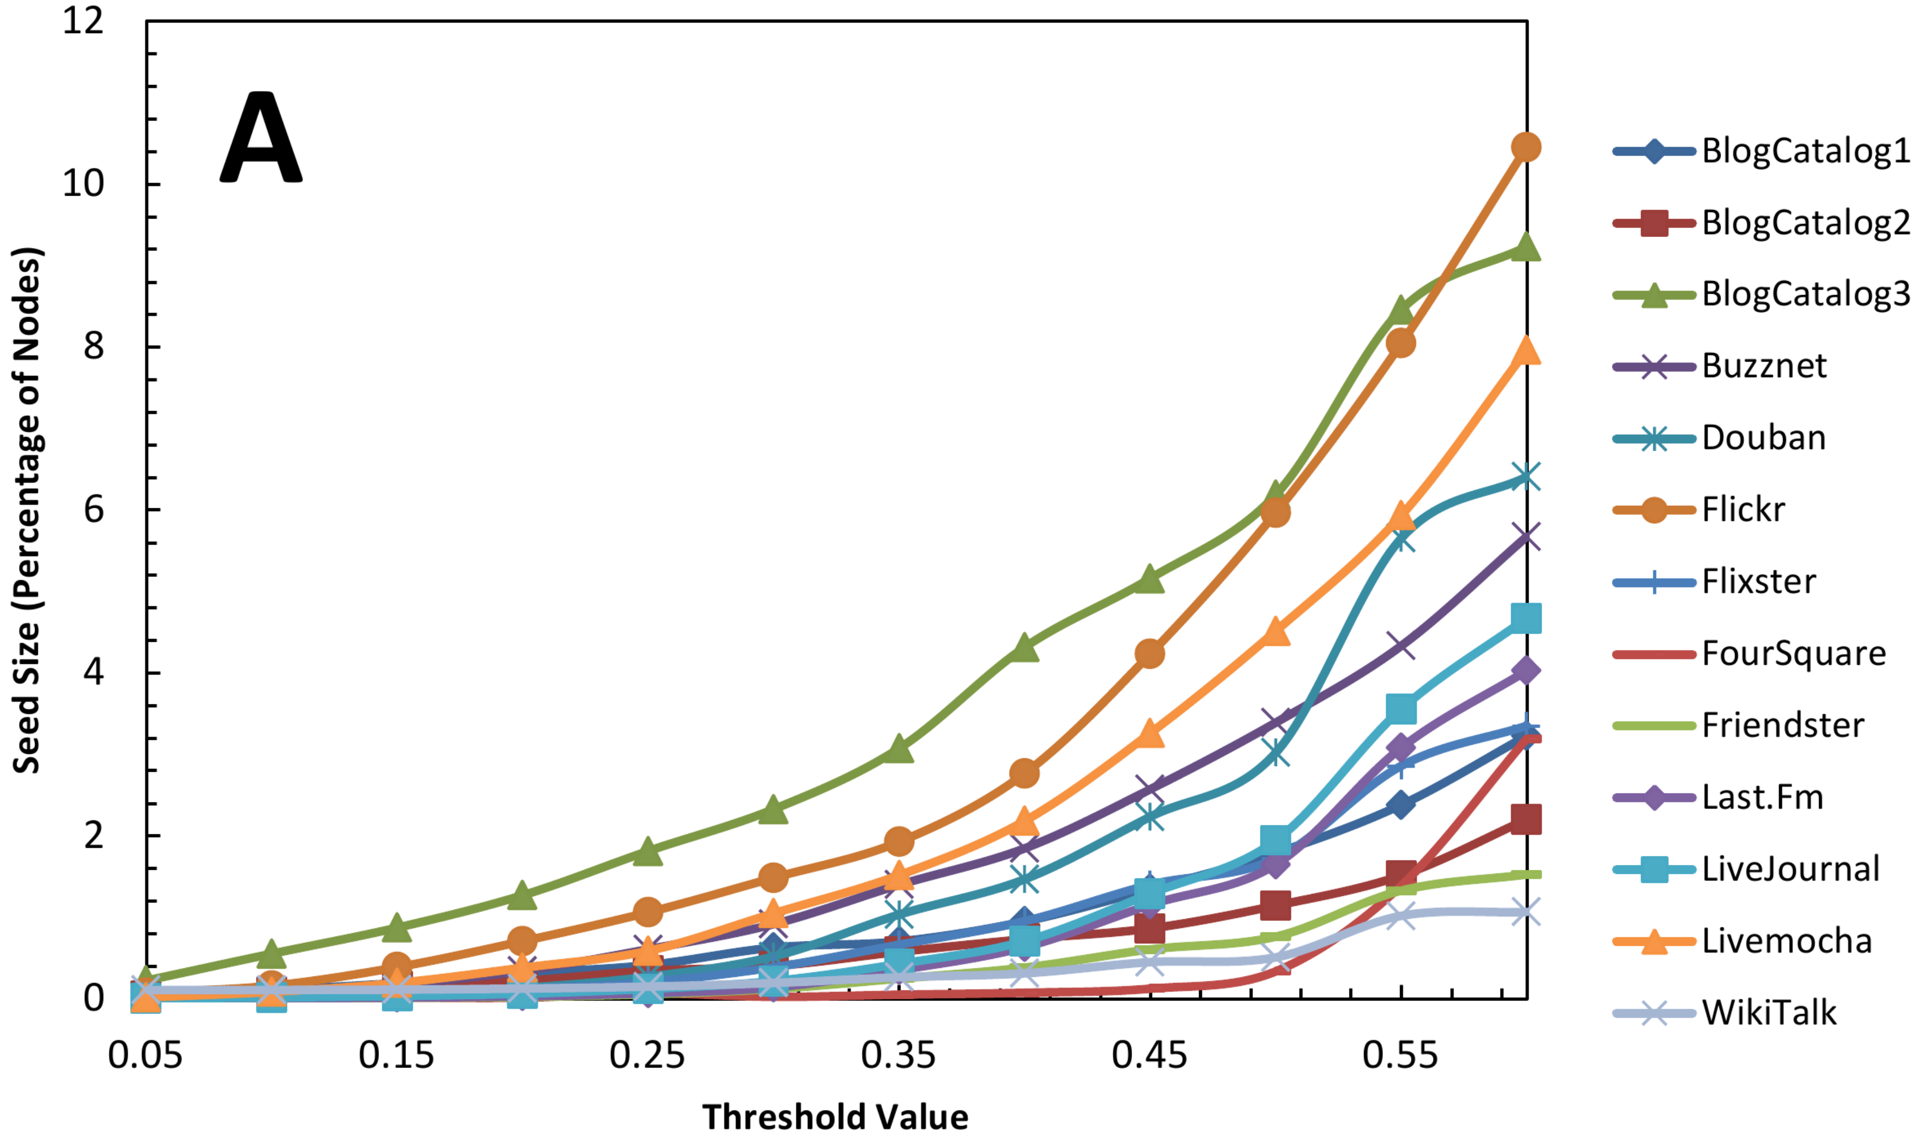

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Significantly larger seed sets for citation networks

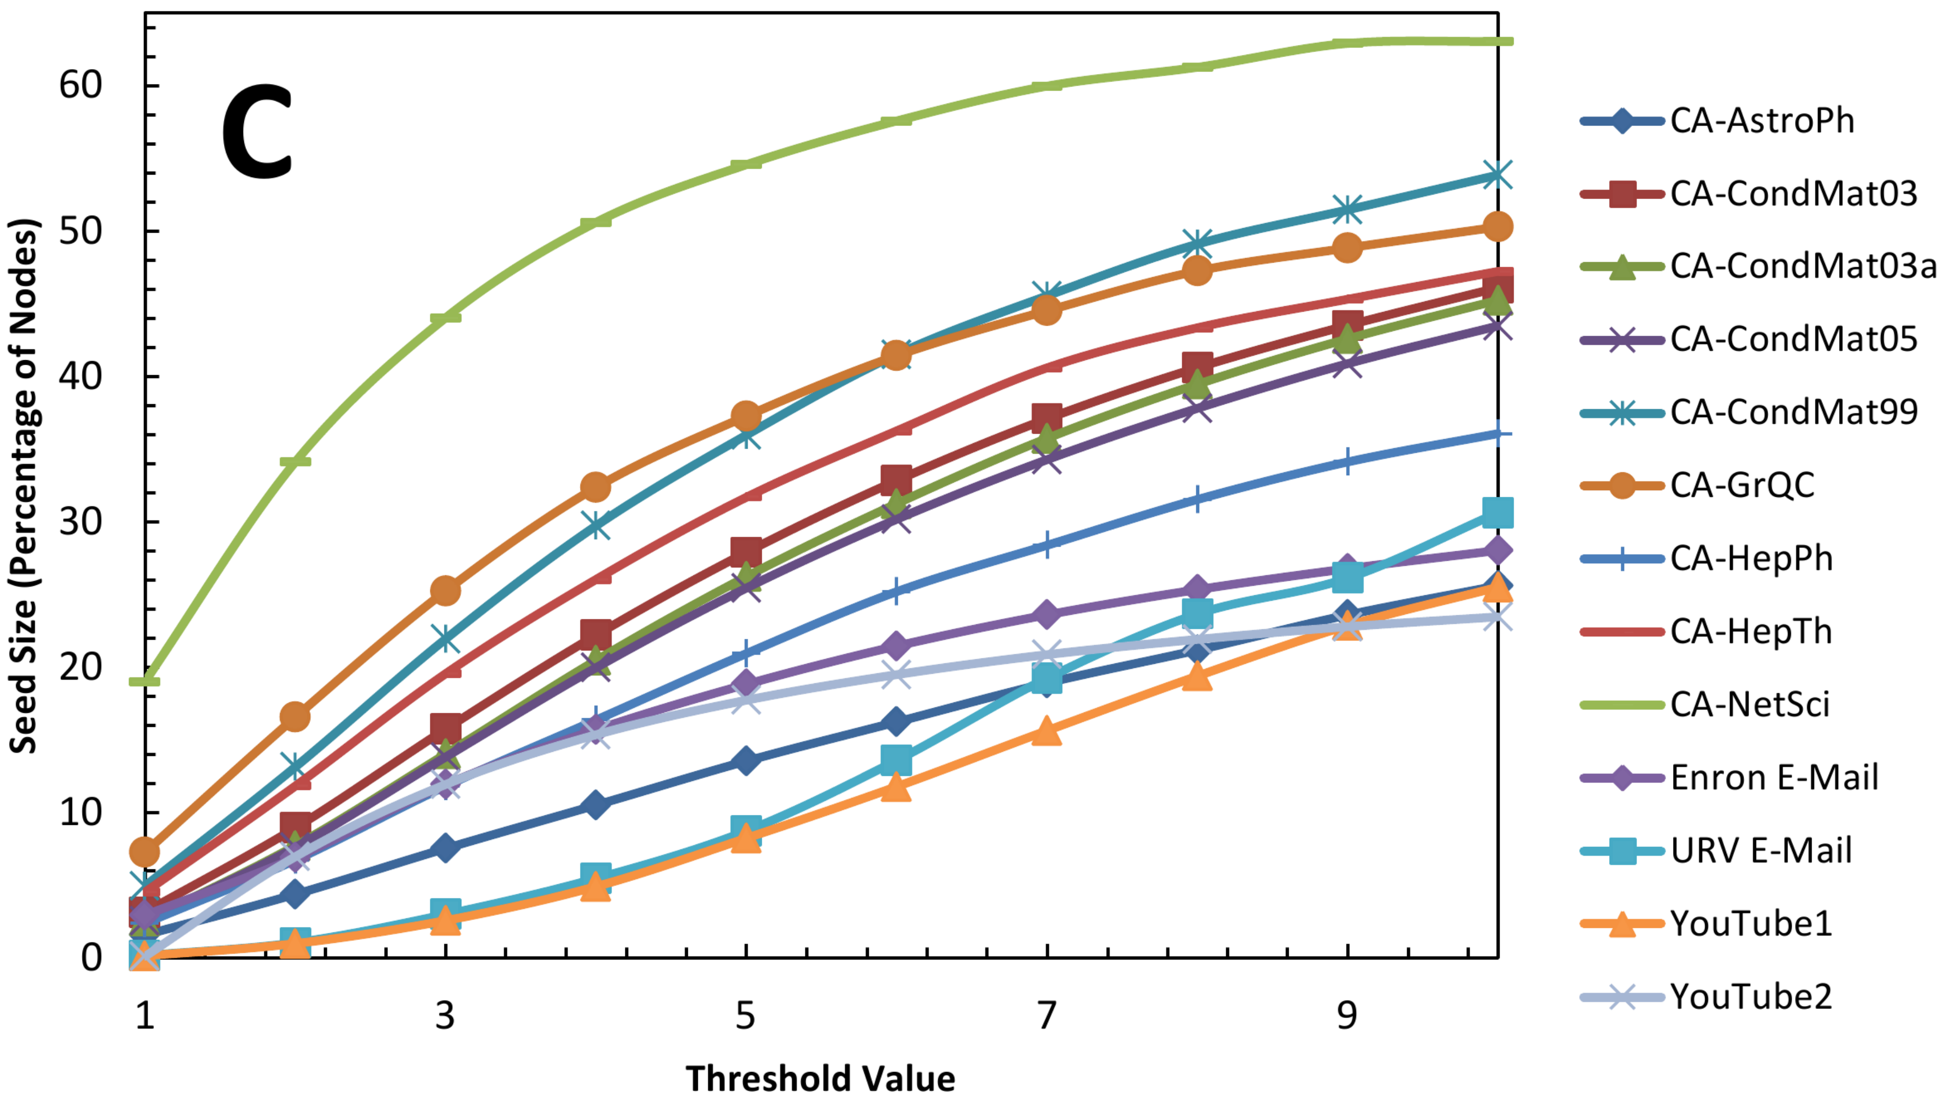

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Significantly larger seed sets for citation networks

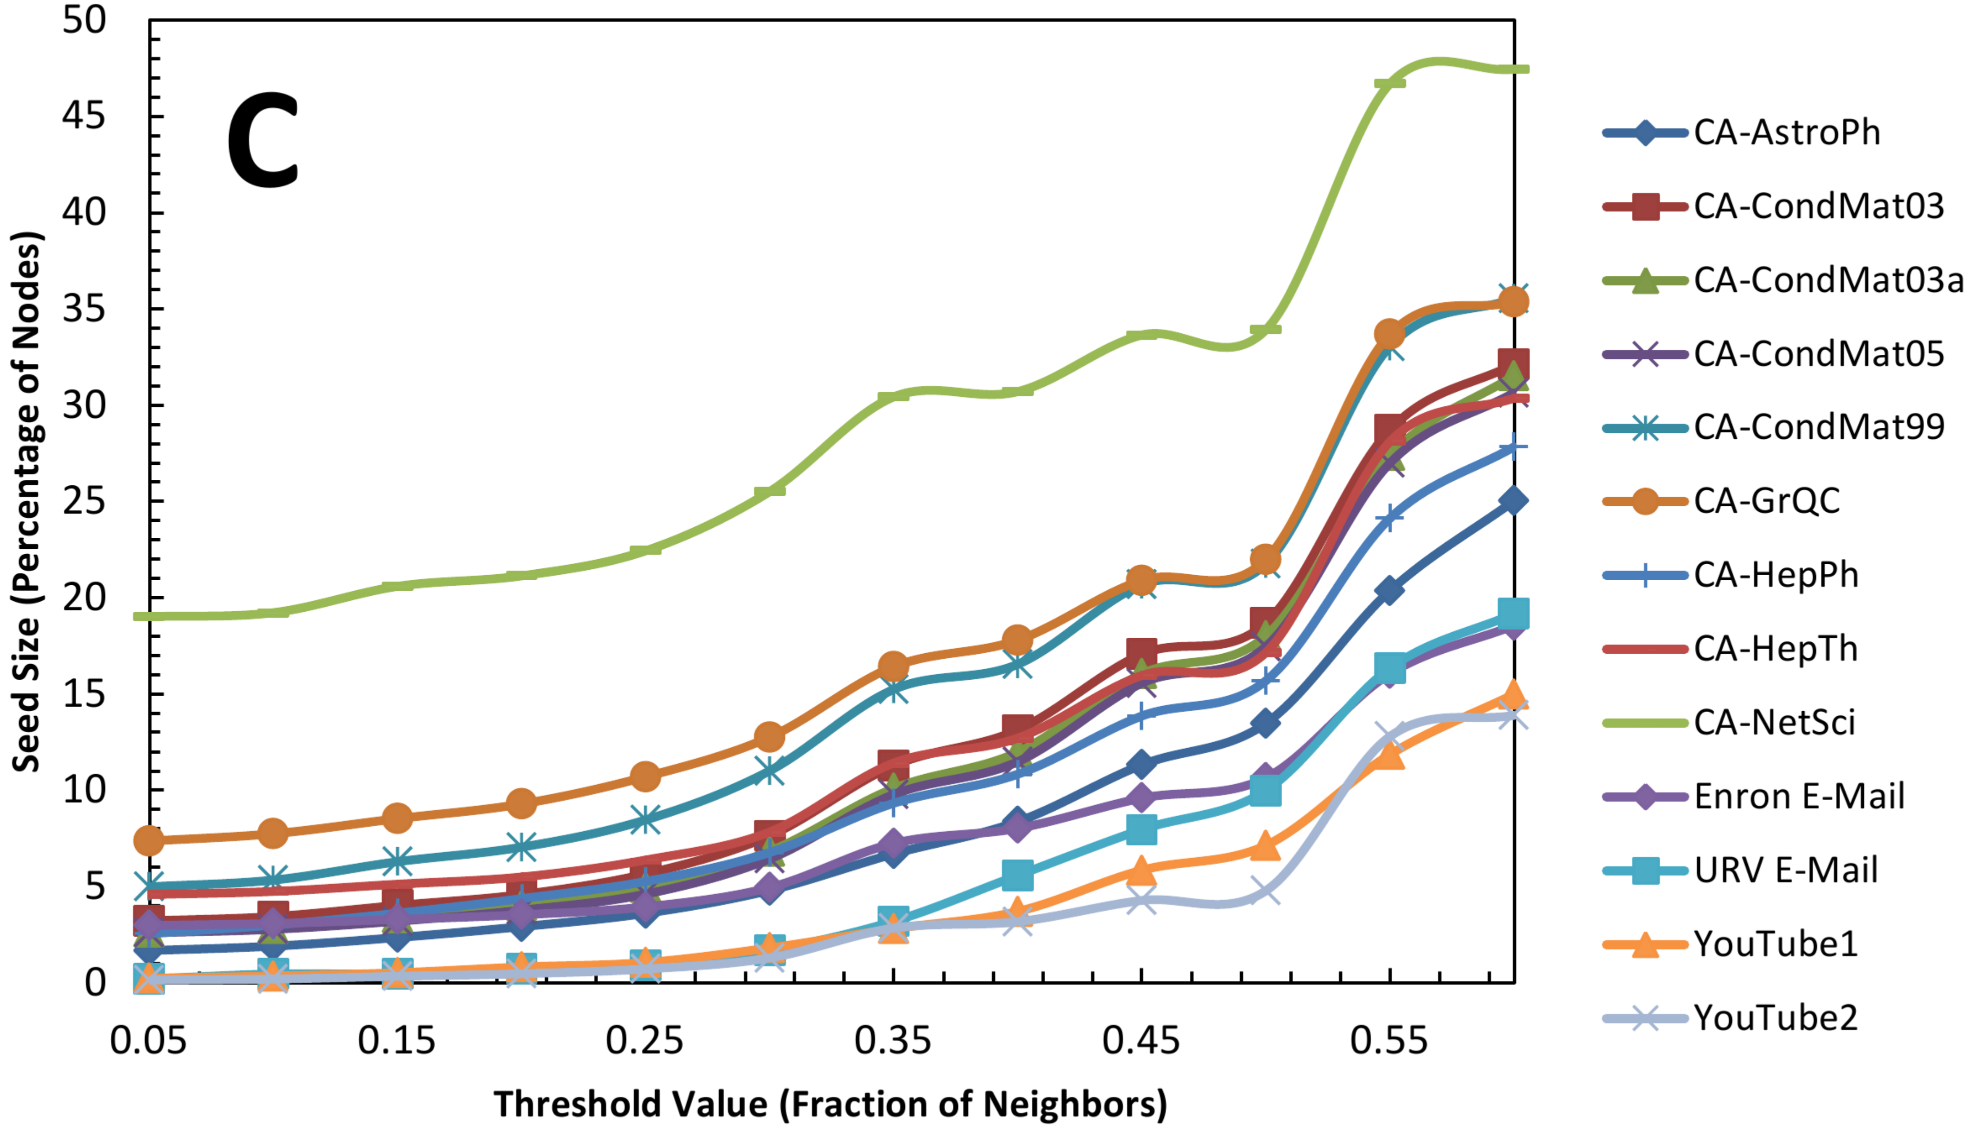

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Example Adoption Rate Dynamics

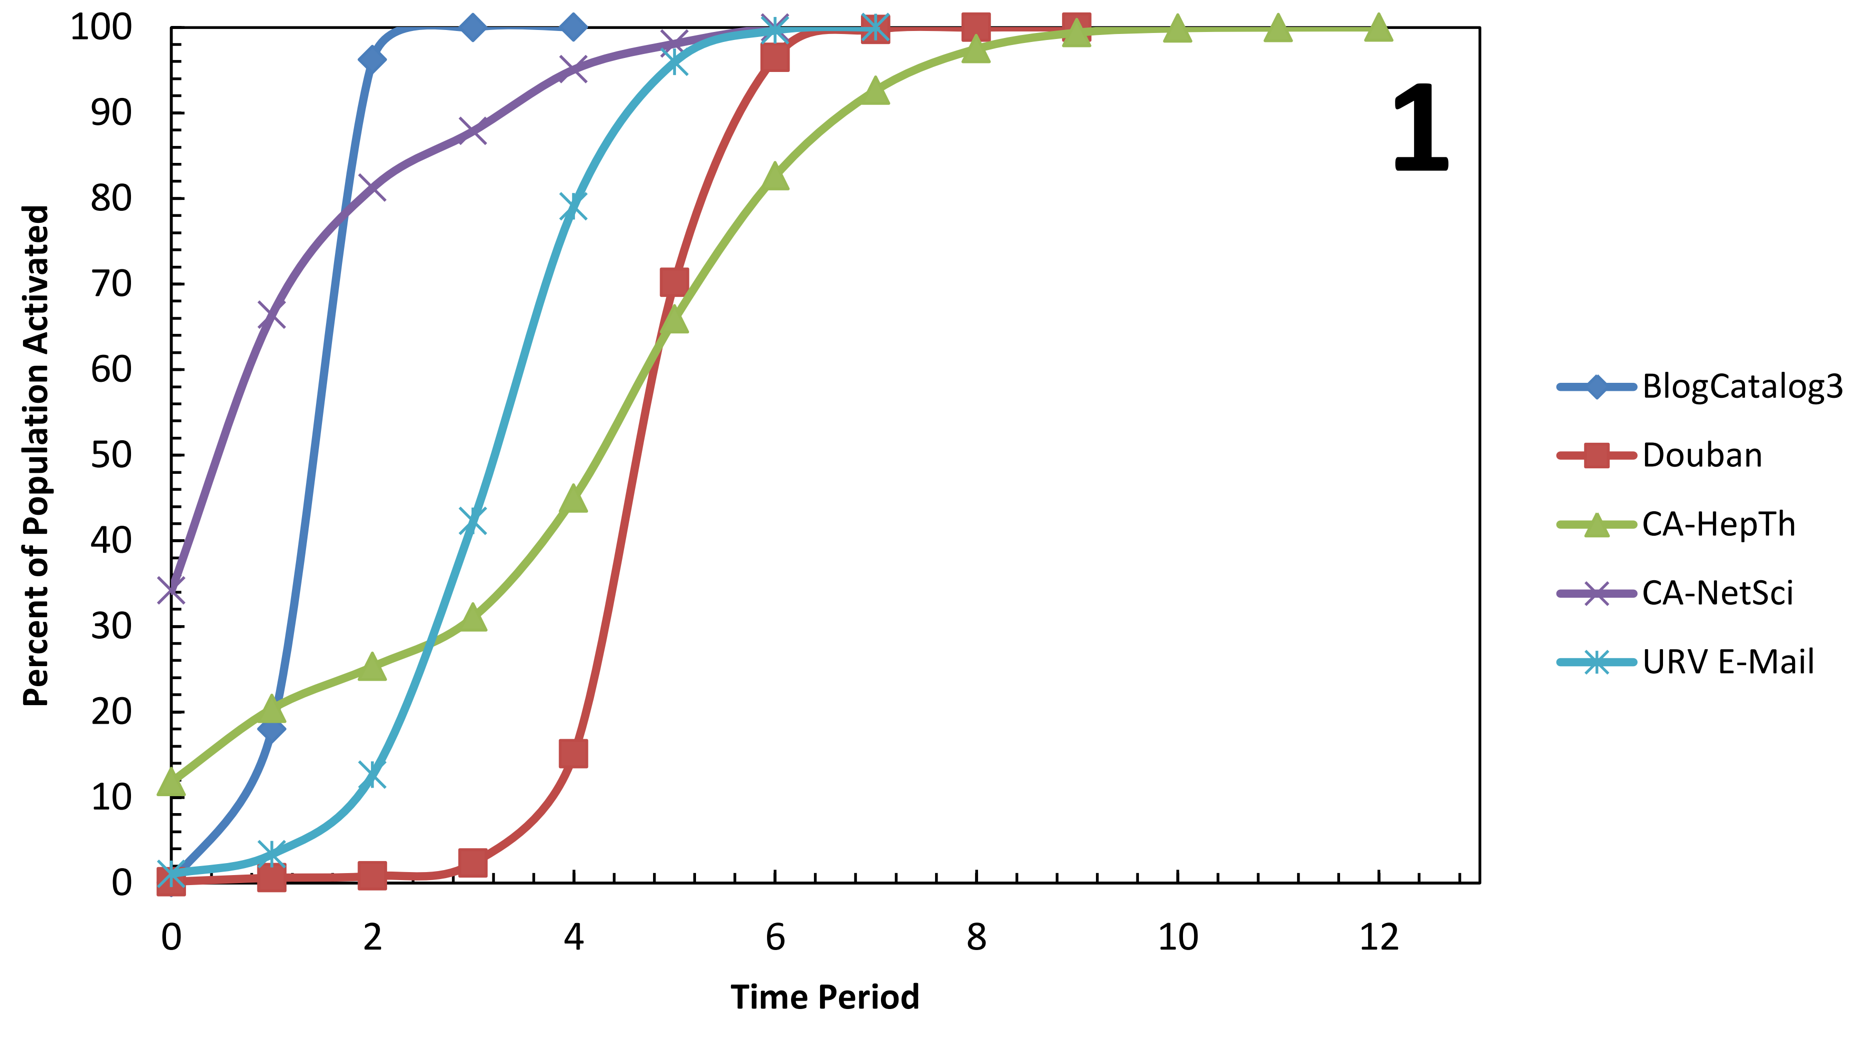

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

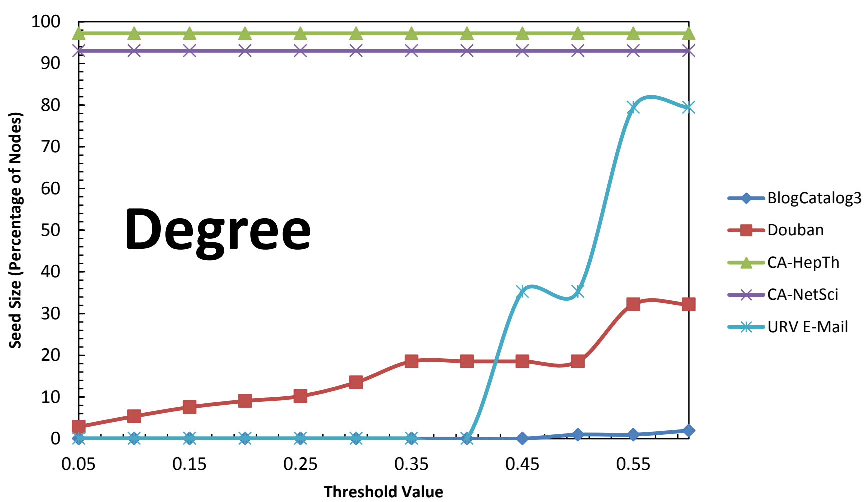

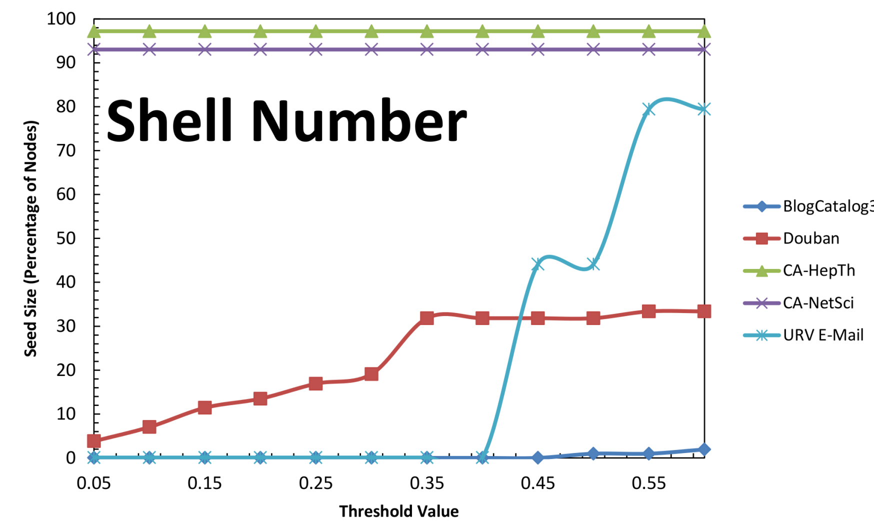

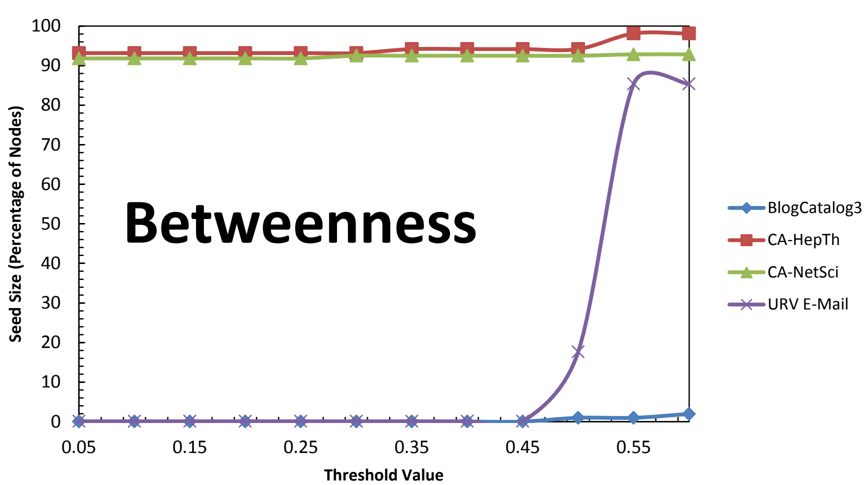

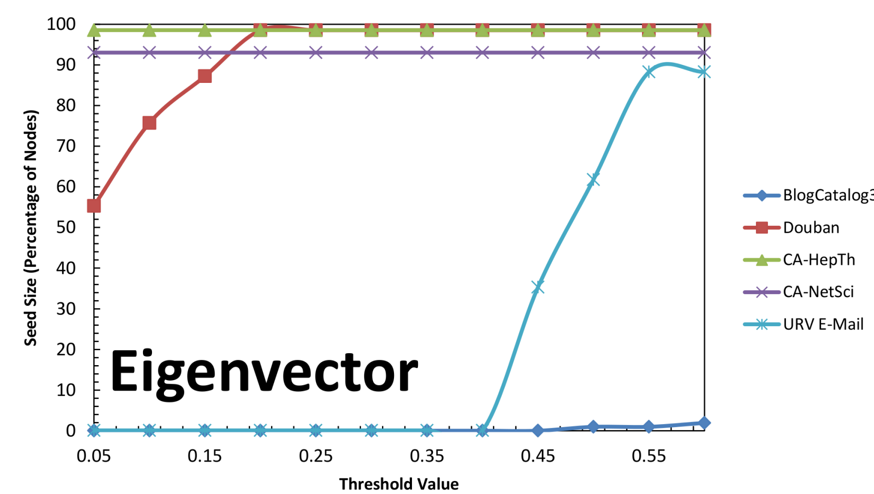

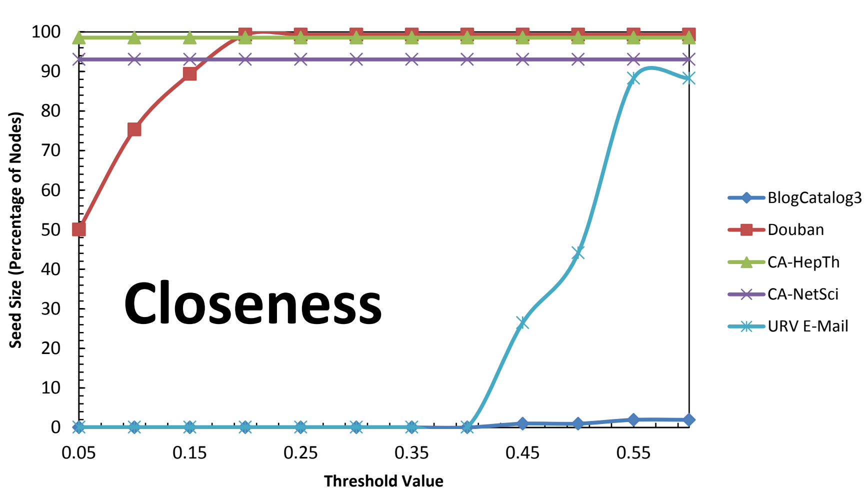

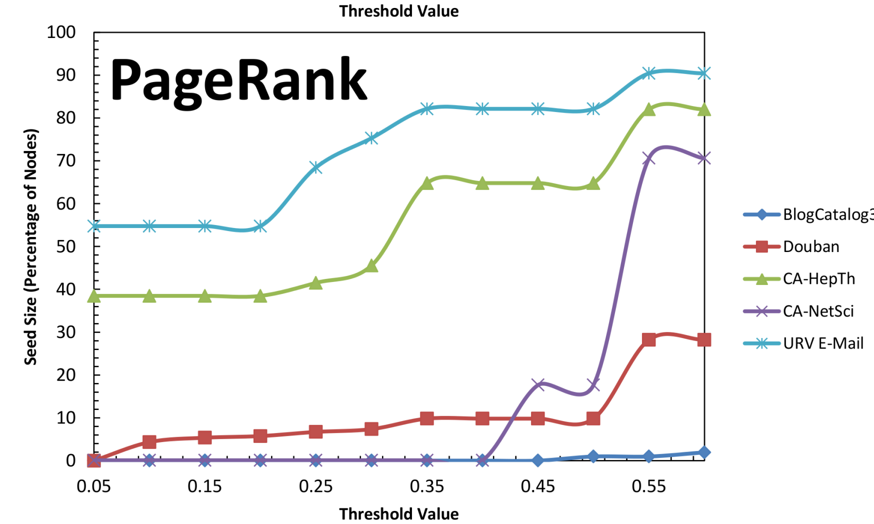

Seed set size when seeded by centrality measures

Shakarian, Eyre, Damon

MIP Heuristic (TIP_DECOMP)

Again, community structure hinders cascades

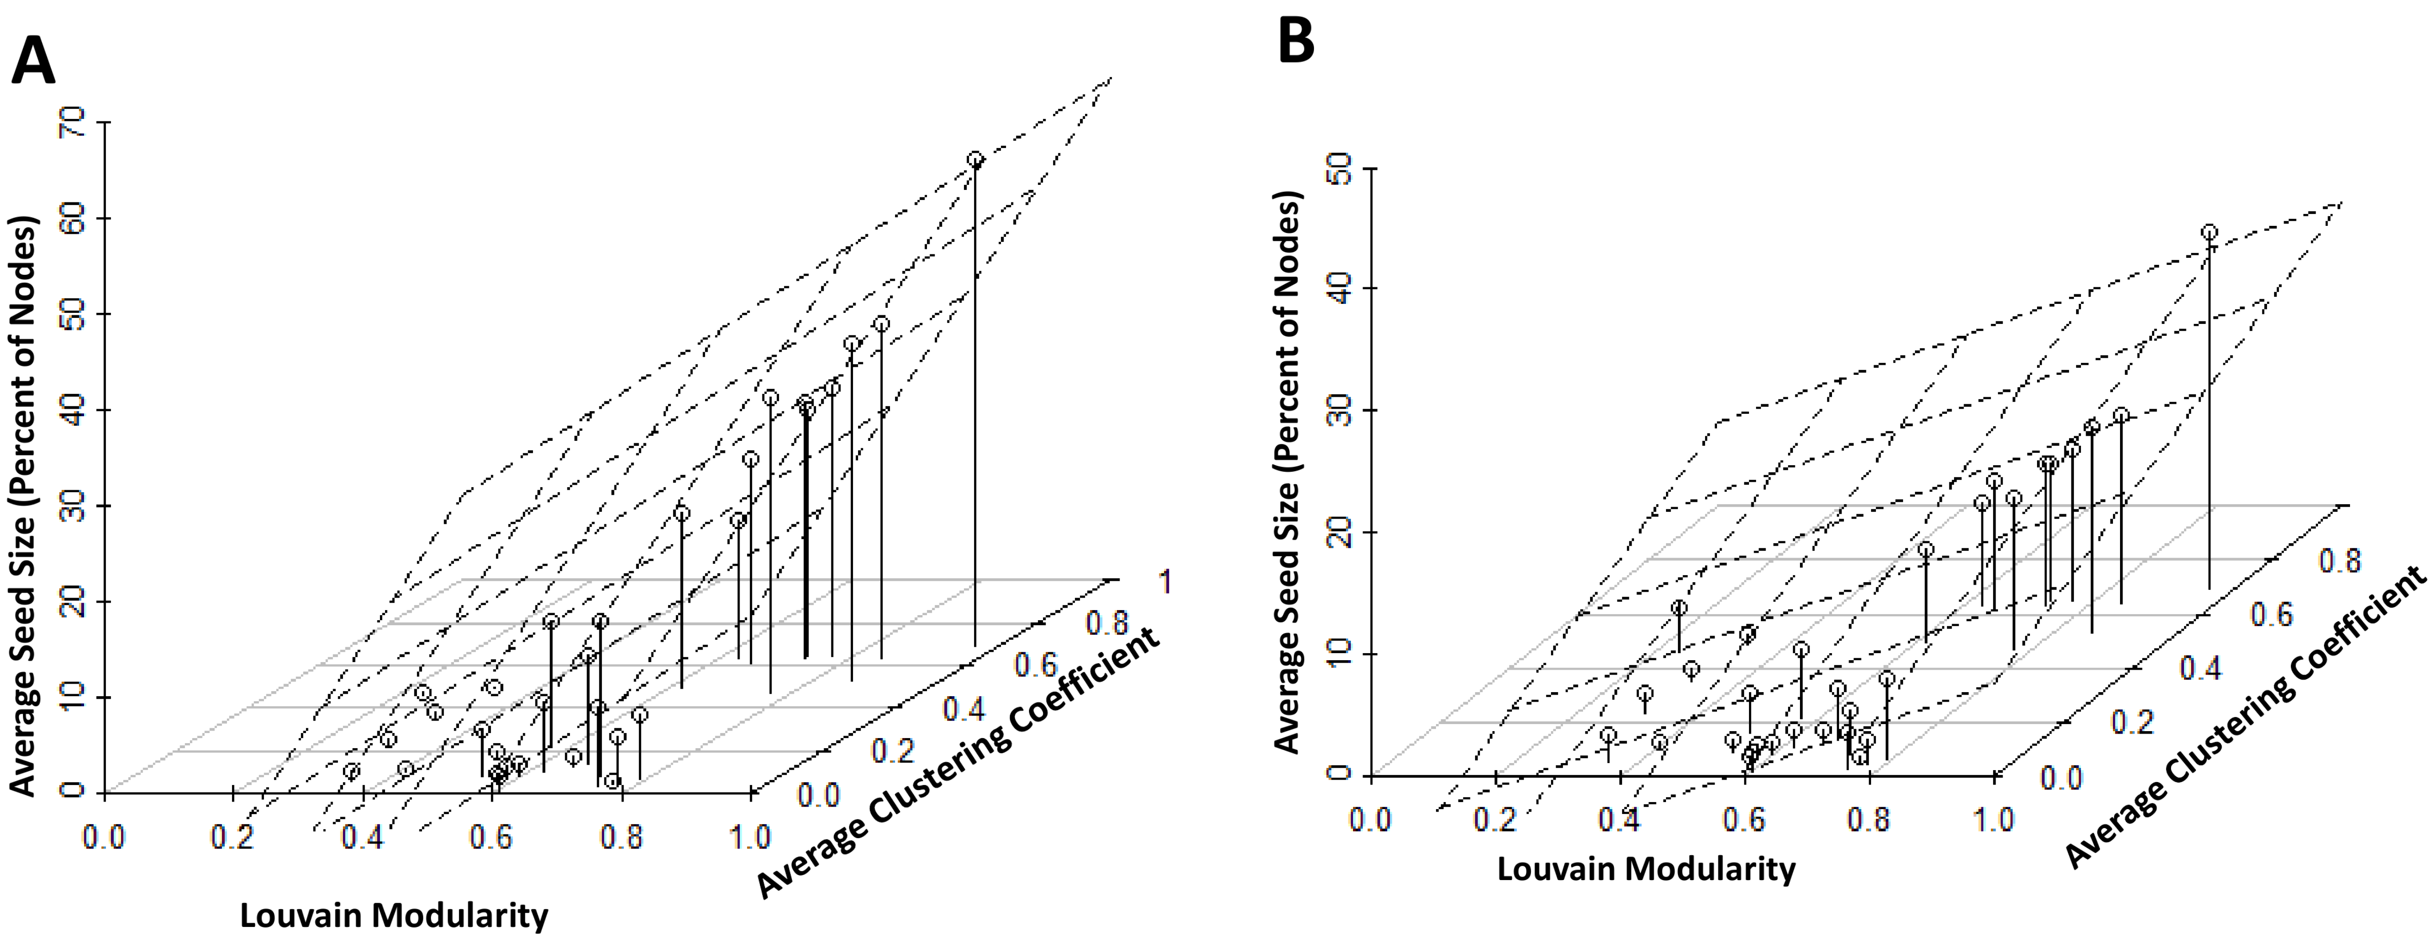

Fig. 14 **(A)** Louvain modularity (M) and average clustering coefficient (C) vs. the average seed size (S). The planar fit depicted is $S = 43.374 \cdot M + 33.794 \cdot C - 24.940$ with $R^2 = 0.8666$, $p = 5.666 \cdot 10^{-13}$. **(B)** Same plot at (A) except the averages are over the 12 percentage-based threshold values. The planar fit depicted is $S = 18.105 \cdot M + 17.257 \cdot C - 10.388$ with $R^2 = 0.816$, $p = 5.117 \cdot 10^{-11}$.

Shakarian, Eyre, Damon

Questions…

- Is there a "centrality" measure inherent in this algorithm for finding seed sets?
- Can you be more greedy — i.e., remove all nodes that tie for minimum delta?
    - (check example given to see that you cannot simply remove all "tie" nodes in parallel)
- Can this algorithm be adopted to a message-passing algorithm to find the seed set in a distributed fashion
- High speed implementation
- Is the seed set size more correlated to the number of communities rather the the modularity?
- How to estimate the threshold or diffusion model from temporal social network data?

Probability Review Items

- Some important random variables
  - Bernoulli, Binomial, Poisson, Gaussian
- Bayes Law & Theorem of Total Probability
- Moments and (Moment) Generating Functions
- Linear MMSE estimation
- Statistics
  - Law of Large Numbers
  - Central Limit Theorem
  - Confidence Intervals
  - Linear Regression
- Markov Chains

Reference:

A. Leon-Garcia, Probability, Statistics, and Random Processes for Electrical Engineer- ing, 3rd Edition, Addison Wesley, 2012.