Decentralized Overlay Network Design

Laura Poplawski Ma
PhD Candidate at Northeastern University

4/28/09

joint work with: Nikolaos Laoutaris (Telefonica), Rajmohan Rajaraman (Northeastern), Ravi Sundaram (Northeastern), Shang-Hua Teng (Boston University and Microsoft Research NE)

Talk Overview

- The problem
- Techniques – game theory definitions
- Simple model. (stable)
- More realistic model. (will oscillate)
- Allowing fractional edges (stable)
- Can we use this to fix BGP oscillations?
- Conclusions and open questions

Decentralizing Overlay Networks

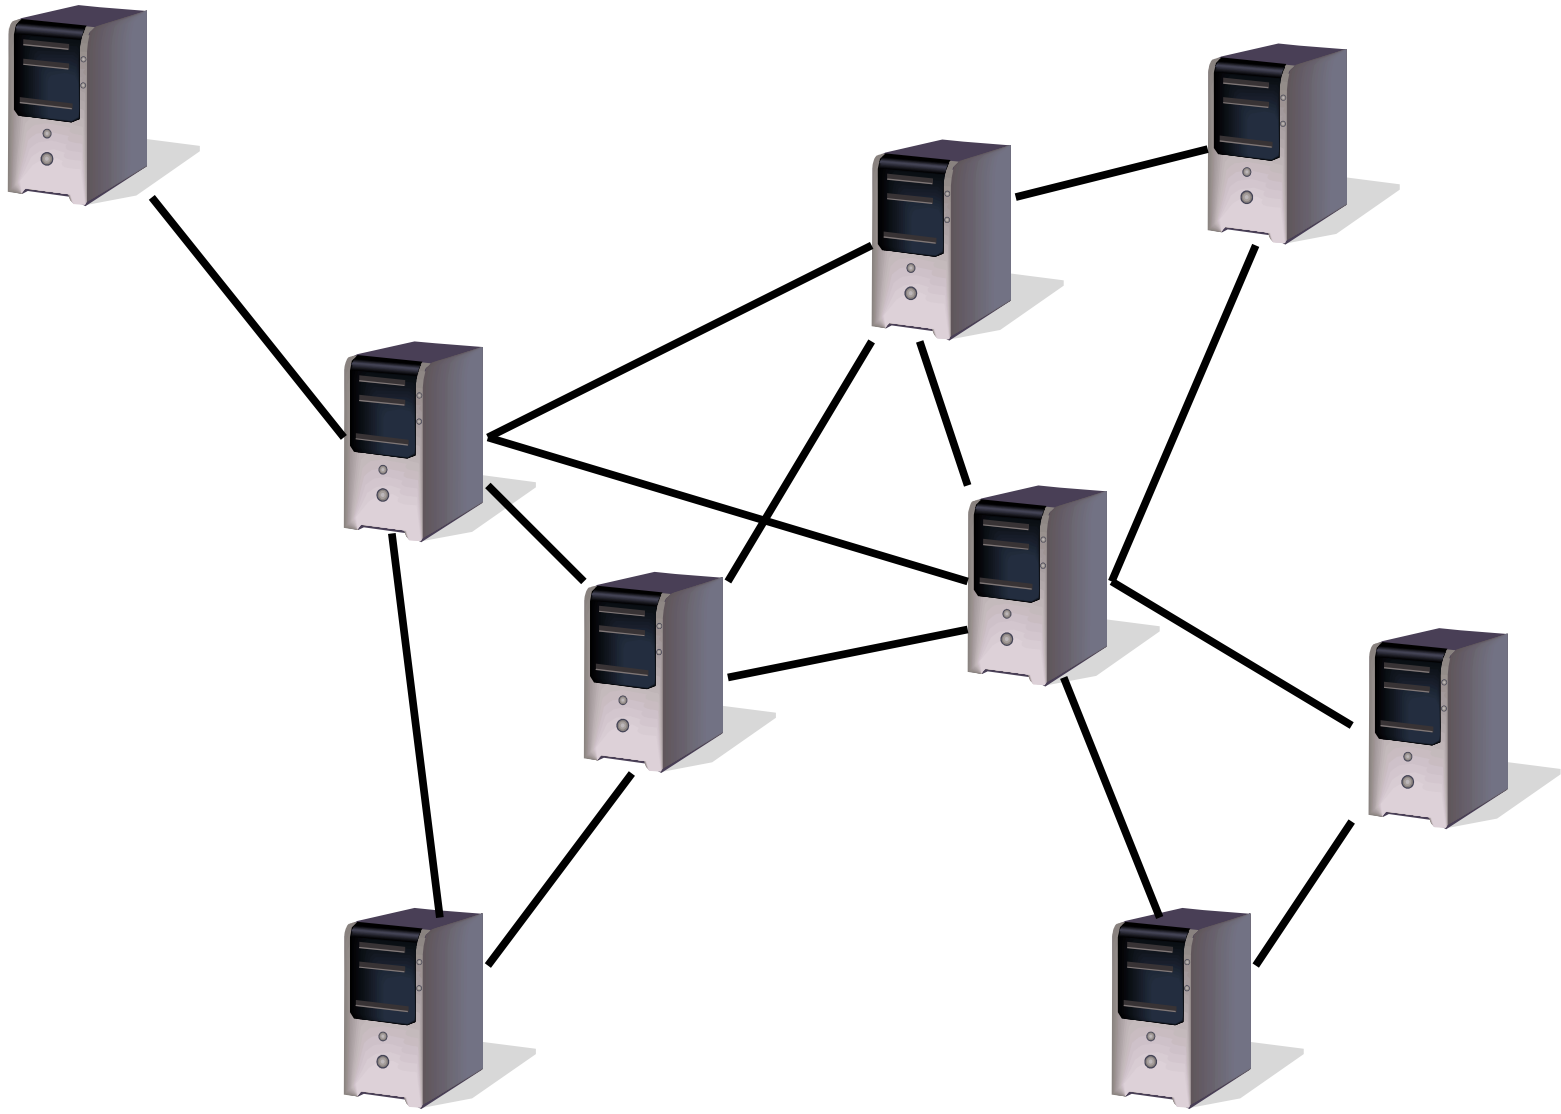

Decentralizing Overlay Networks

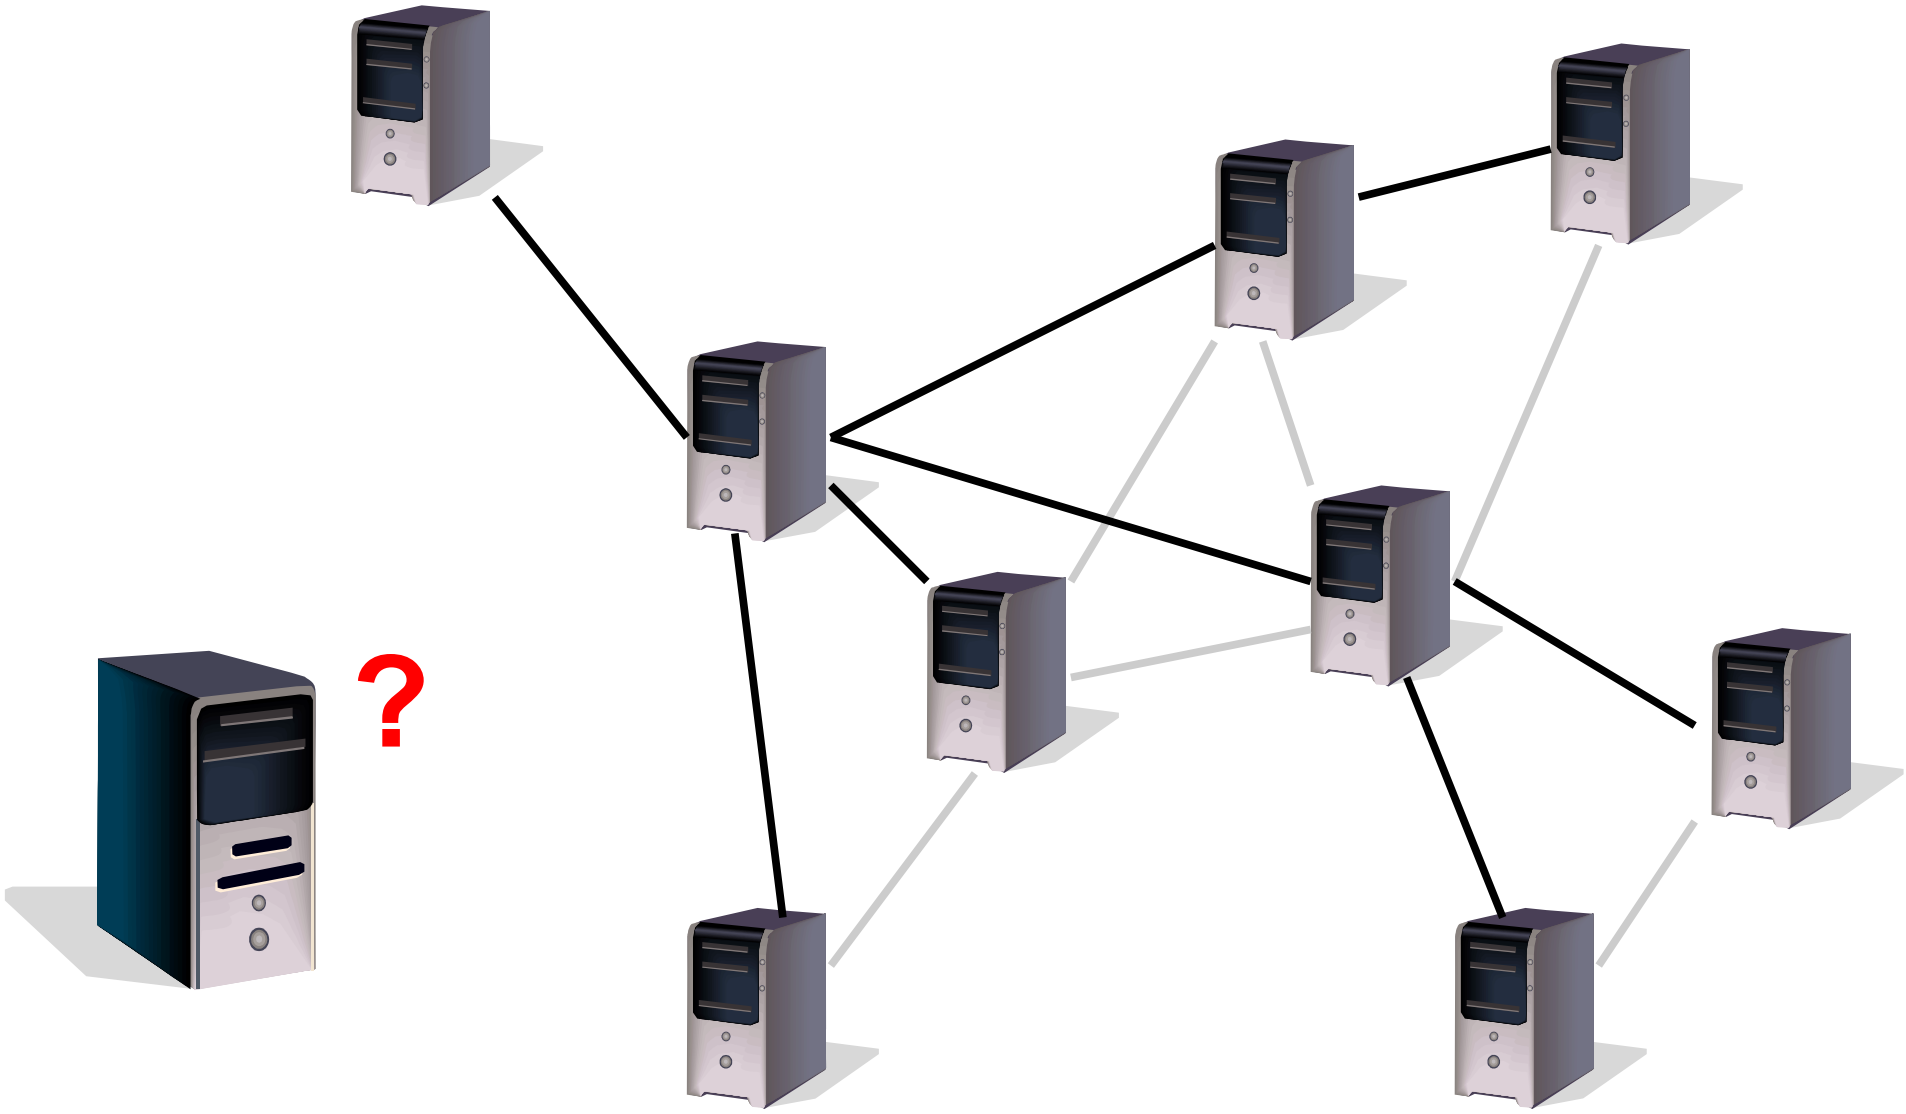

Decentralizing Overlay Networks

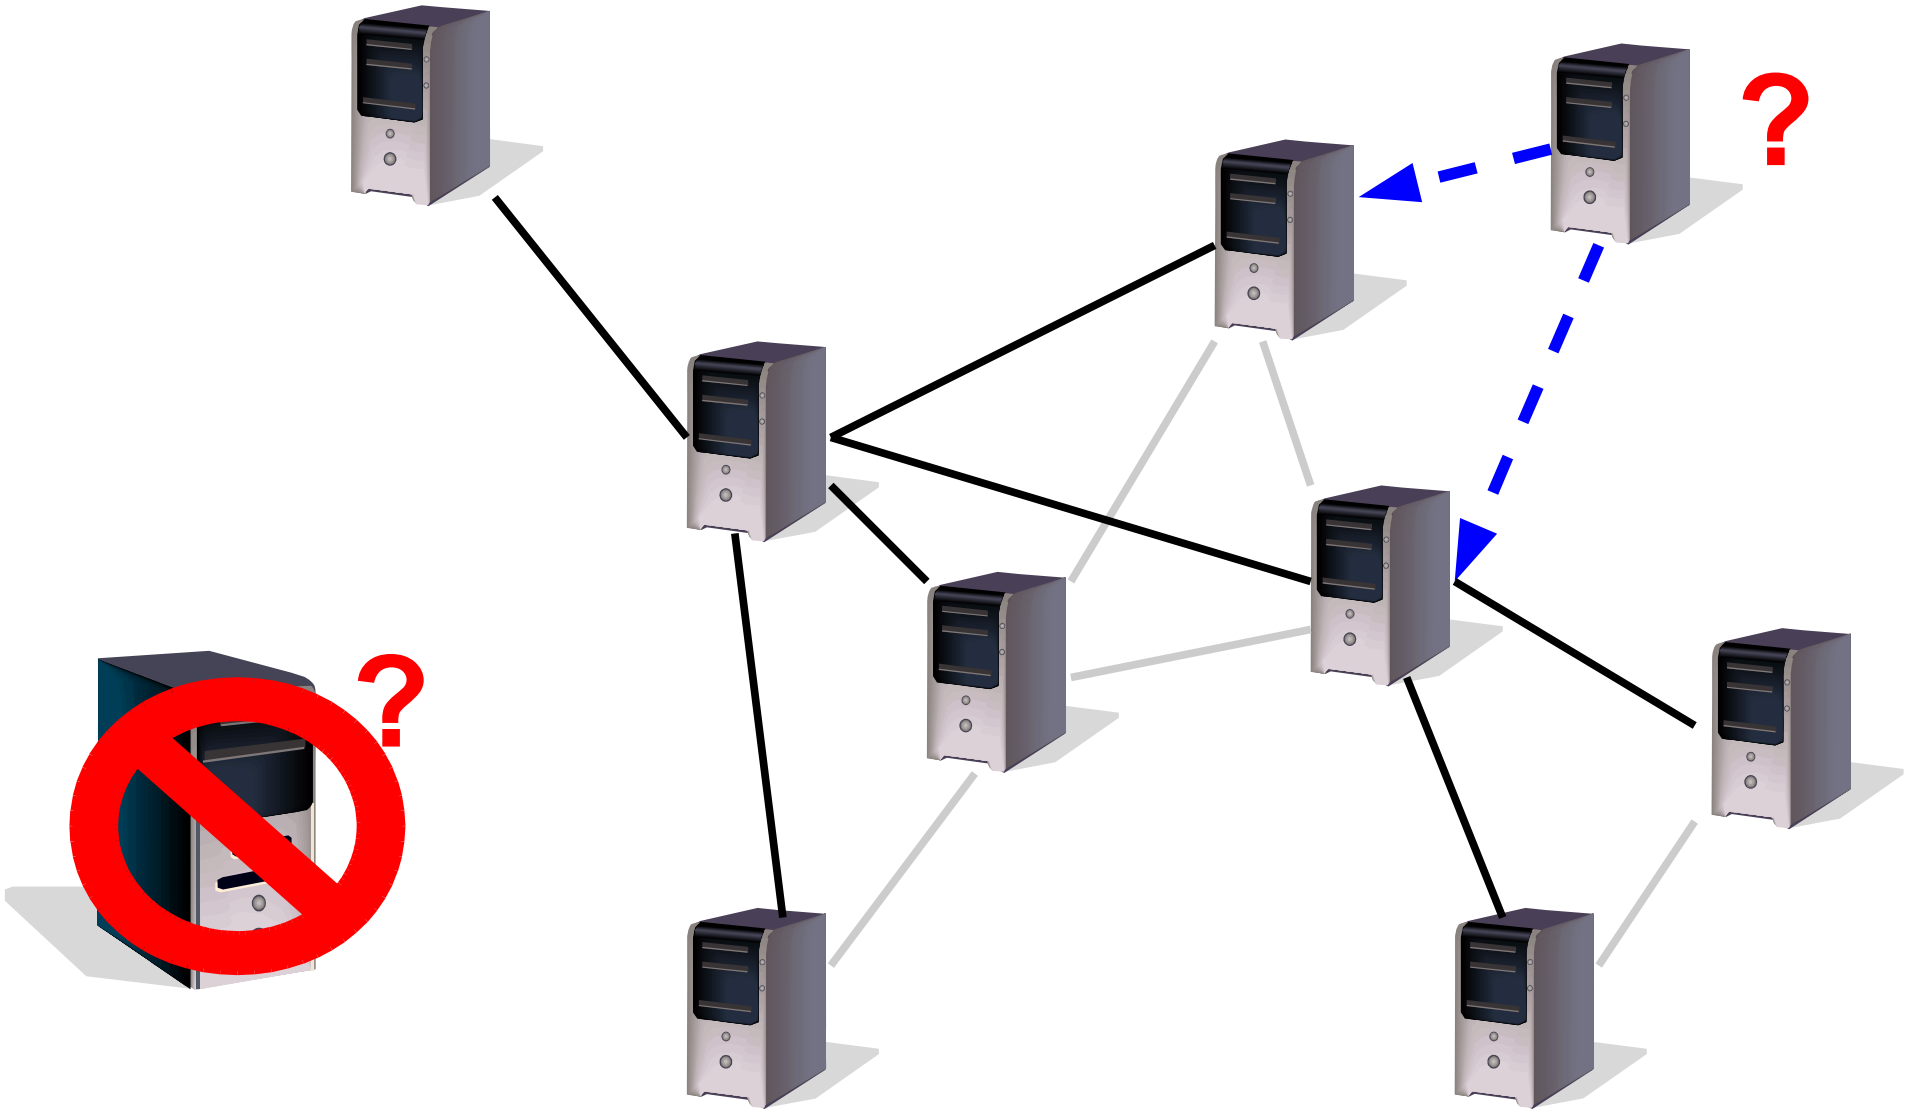

Decentralizing Overlay Networks

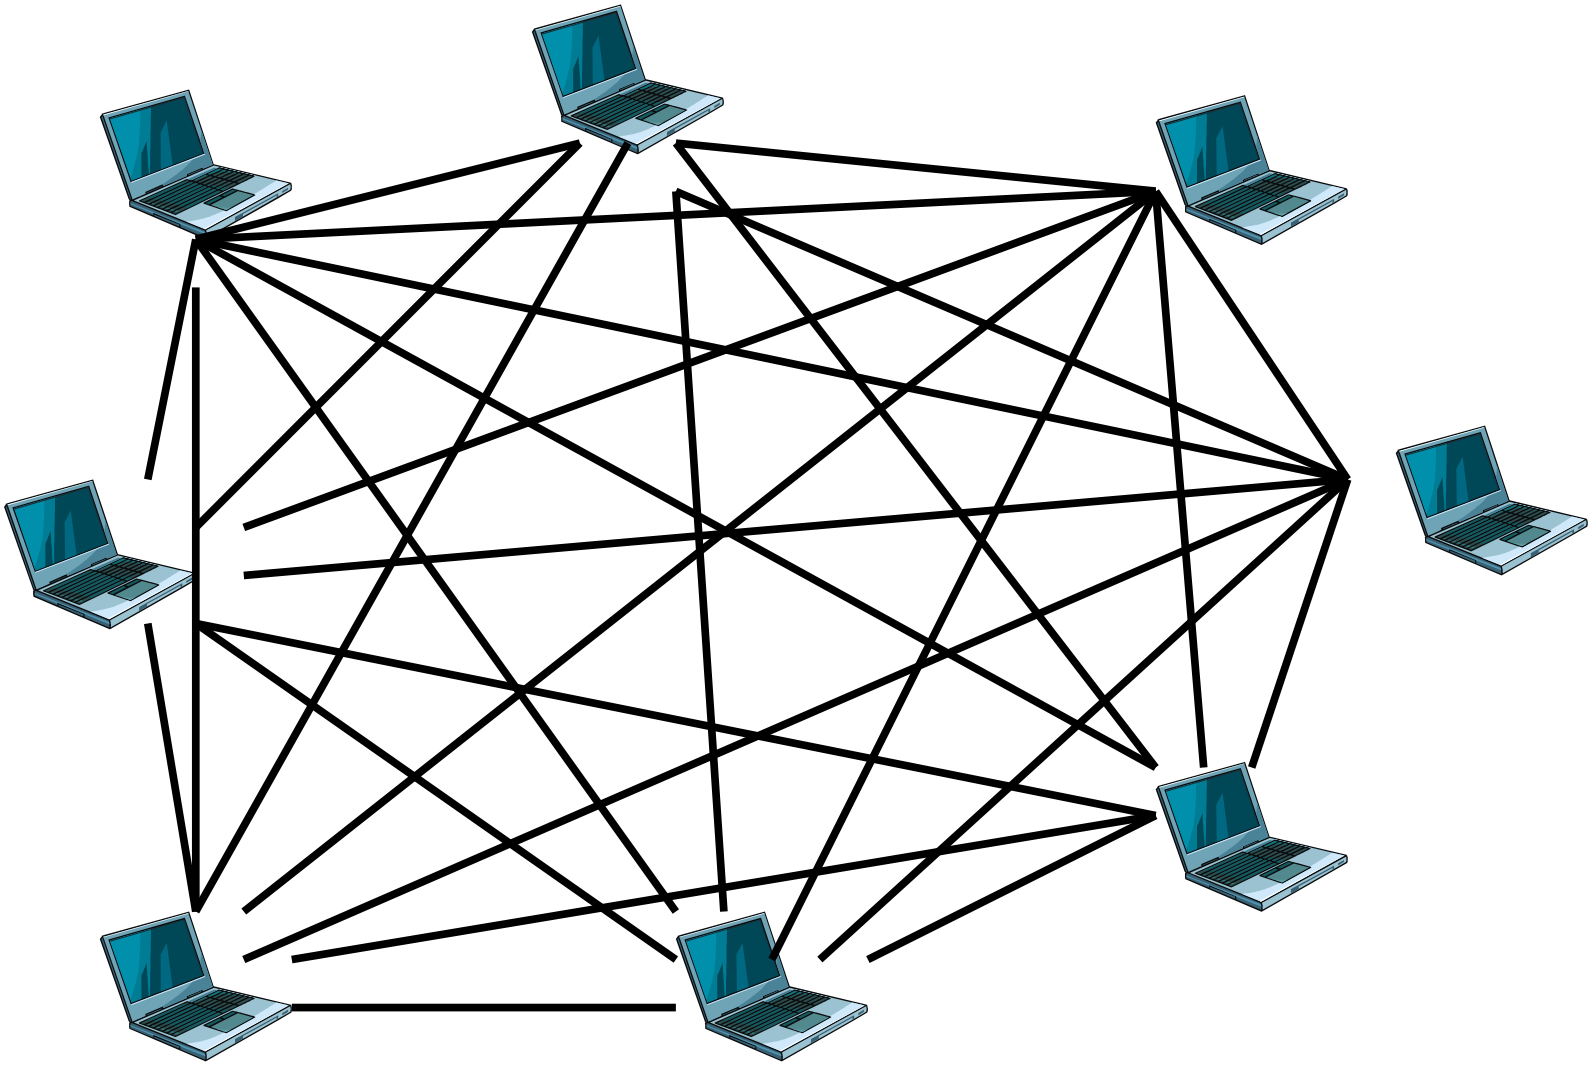

Decentralizing Overlay Networks

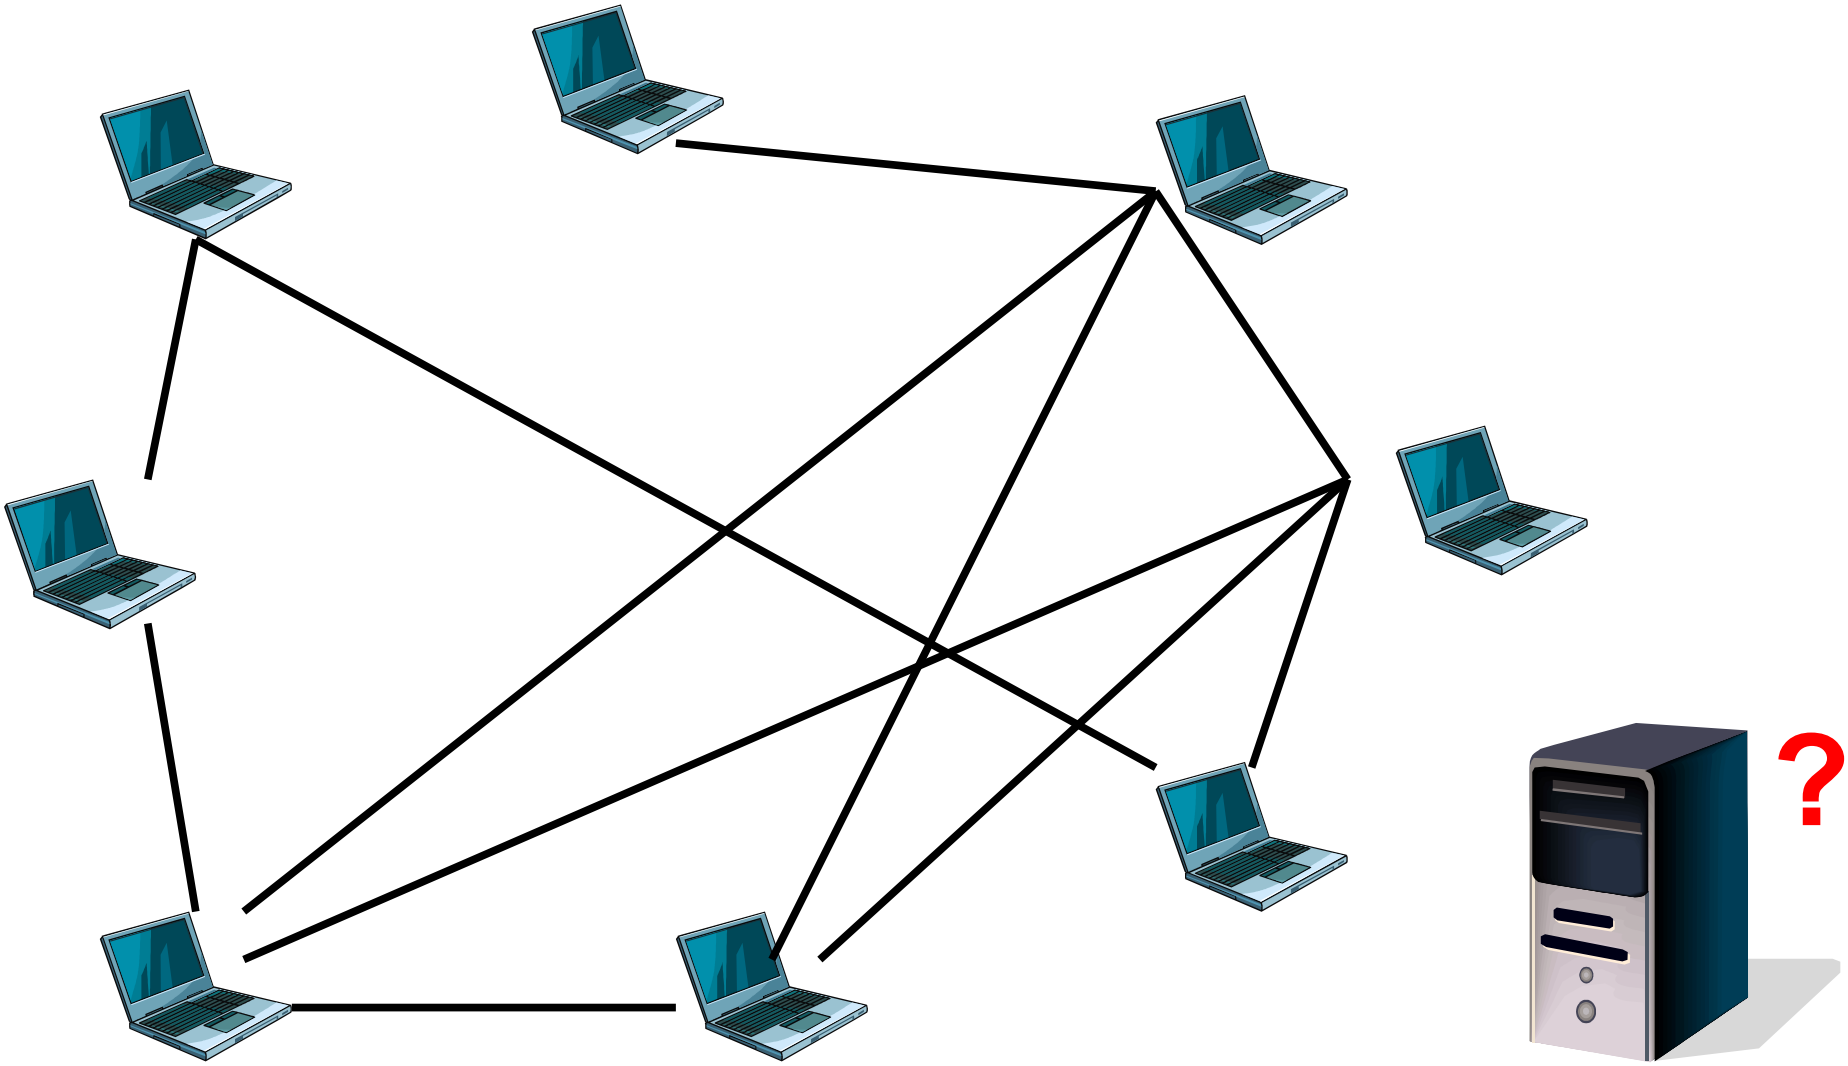

Decentralizing Overlay Networks

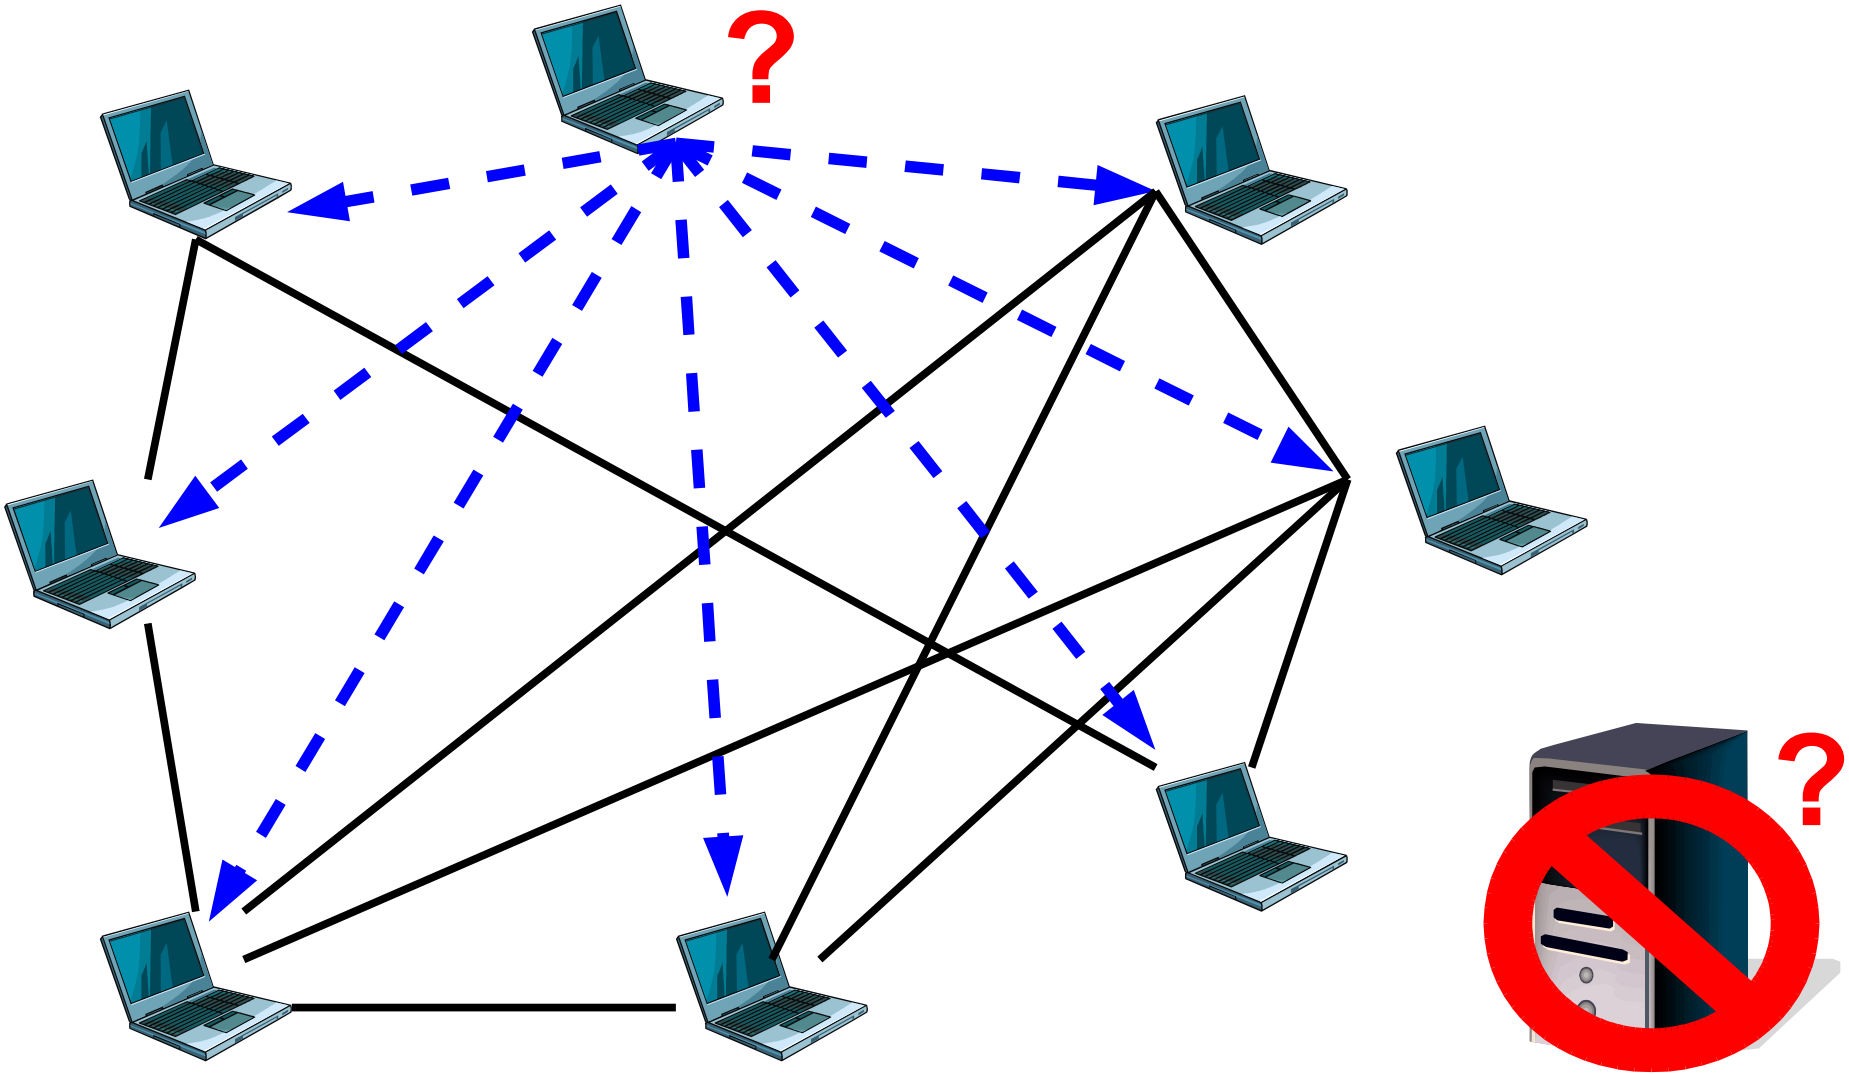

Assumptions

- Full knowledge
- Similar goals/utilities for all nodes
- Oscillation is bad
- Using deterministic routing

Questions

- Is any network resulting from decentralized design:
  - “Good” or “Bad”?

Questions

- Is any network resulting from decentralized design:
  - “Good” or “Bad”?
  - Stable or oscillating?

Questions

- Is any network resulting from decentralized design:
  - “Good” or “Bad”?
  - Stable or oscillating?
  - Converging or cycling?

Questions

- Is any network resulting from decentralized design:
  - “Good” or “Bad”?
  - Stable or oscillating?
  - Converging or cycling?
  - Causing cooperation or competition?

Techniques

- Experimentation
- Mathematical prediction

Techniques

- Experimentation
- Mathematical prediction

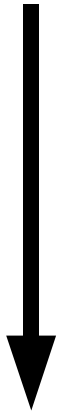

Game theory

Algorithmic Game Theory

- Define a game:
  - Strategies for each player
  - Utility/Cost: numerical value for each [player x combination of strategies]

	Heads	Tails
Heads	-1, -1	1, 1
Tails	1, 1	-1, -1

Algorithmic Game Theory

- Best Response
  - Given what others are doing, you get lowest utility
- Nash Equilibrium (Stable solution)
  - All players are simultaneously playing a best response

	Heads	Tails
Heads	-1, -1	1, 1
Tails	1, 1	-1, -1

Algorithmic Game Theory

- Best Response
  - Given what others are doing, you get lowest utility
- Nash Equilibrium (Stable solution)
  - All players are simultaneously playing a best response

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

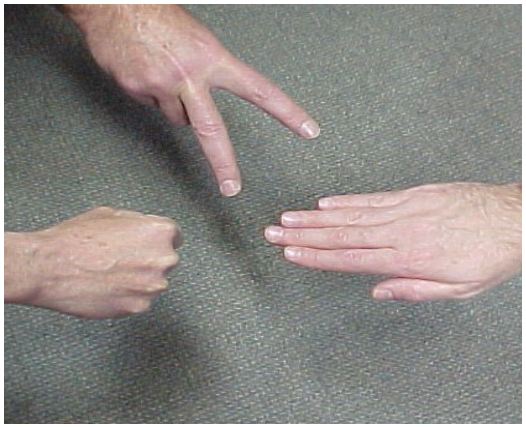

Algorithmic Game Theory

- Pure vs. Mixed Nash Equilibrium
  - Pure ~ deterministic, Mixed ~ probabilistic

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

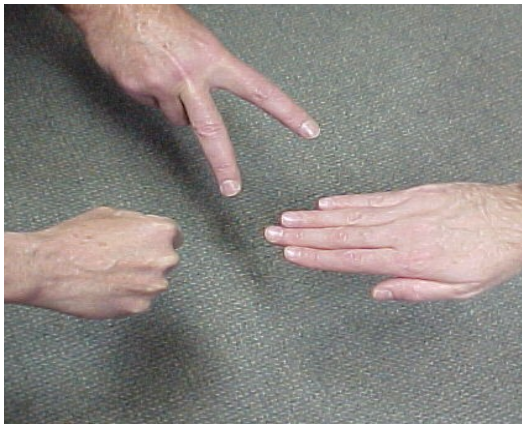

Our Game Definition

- Players = nodes in the network
- Strategies = possible edges they may build
- Utility = Length or quality of routes

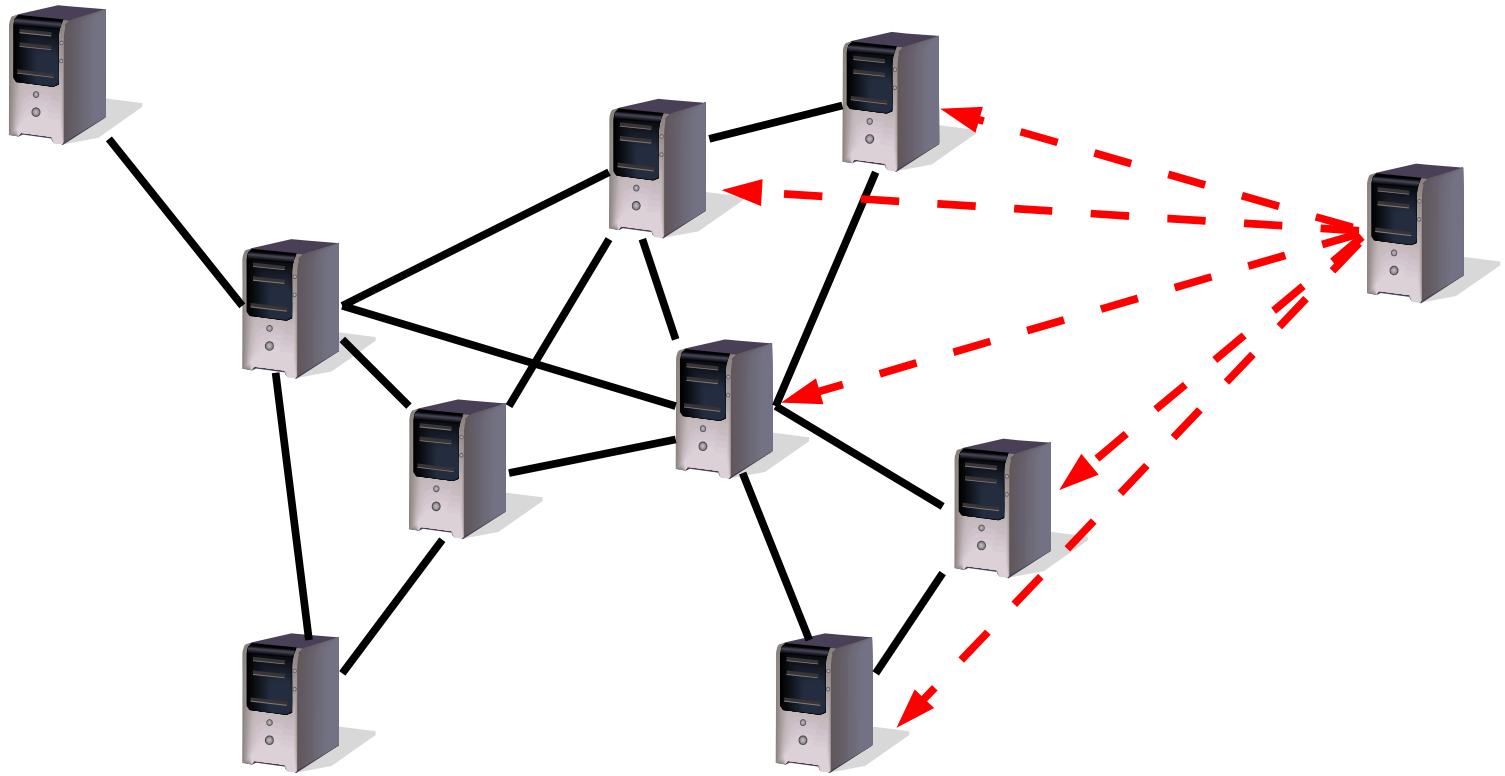

Questions

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?
  - Causing cooperation or competition?

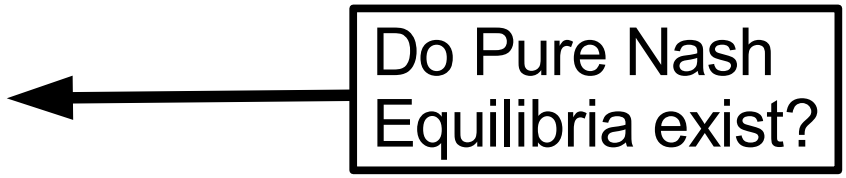

Questions

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?
  - Causing cooperation or competition?

Do Pure Nash Equilibria exist?

Will best response trials converge to an equilibrium?

Questions

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?
  - Causing cooperation or competition?

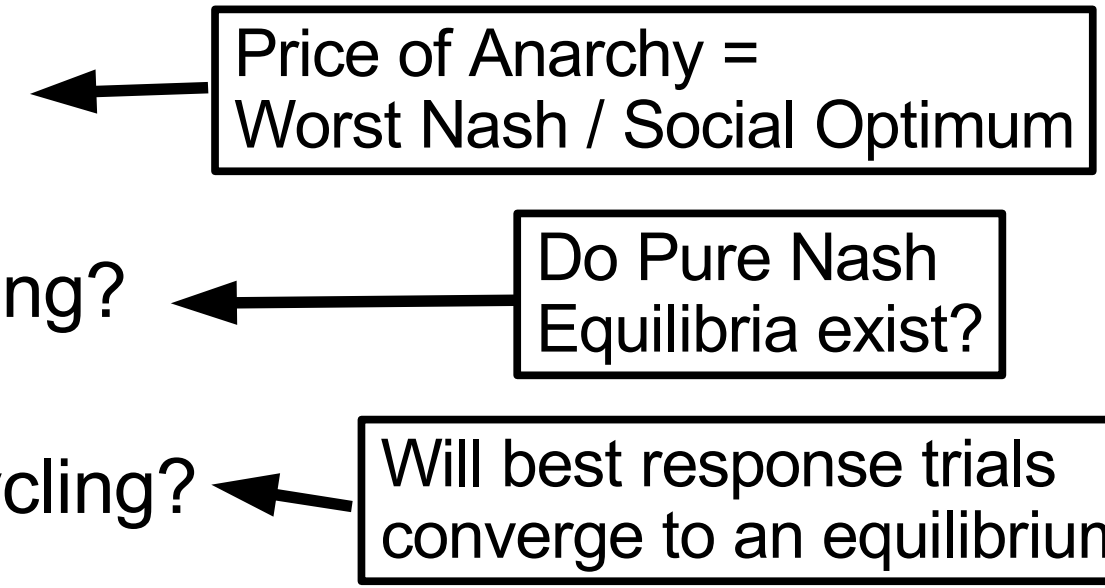

Questions

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?
  - Causing cooperation or competition?

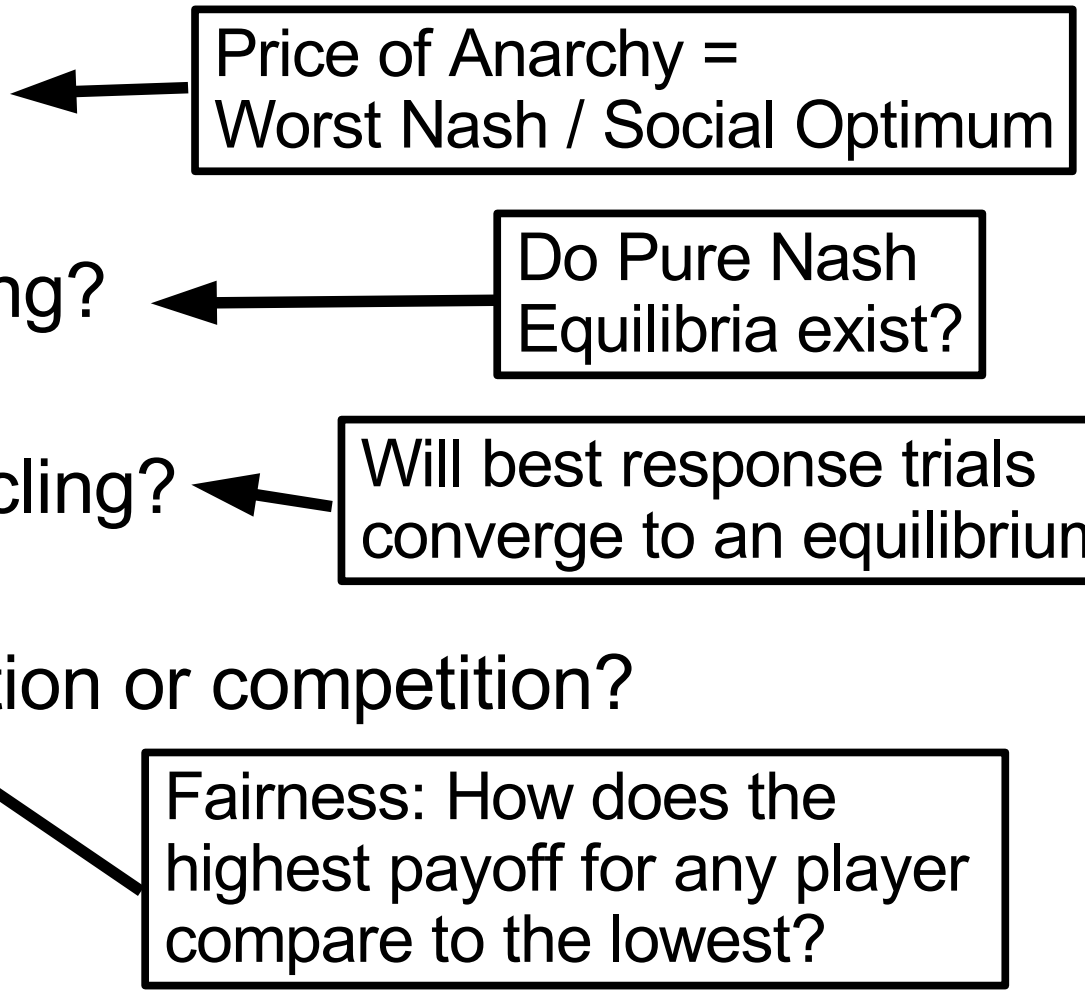

Simplest Game

- Players = Nodes in the network
- Strategies = Connections they can make
  - One edge to any other node
- Utilities = How close am I to the other nodes?
  - Average hop-count distance to all other nodes

Simplest Game

- Players = Nodes in the network
- Strategies = Connections they can make
  - One edge to any other node
- Utilities = How close am I to the other nodes?
  - Average hop-count distance to all other nodes

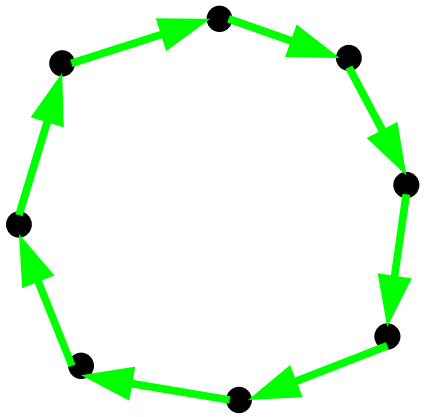

2 Connection Game

- Players = Nodes in the network
- Strategies = Connections they can make
  - Two edges to any other node
- Utilities = How close am I to the other nodes?
  - Average hop-count distance to all other nodes

2 Connection Game

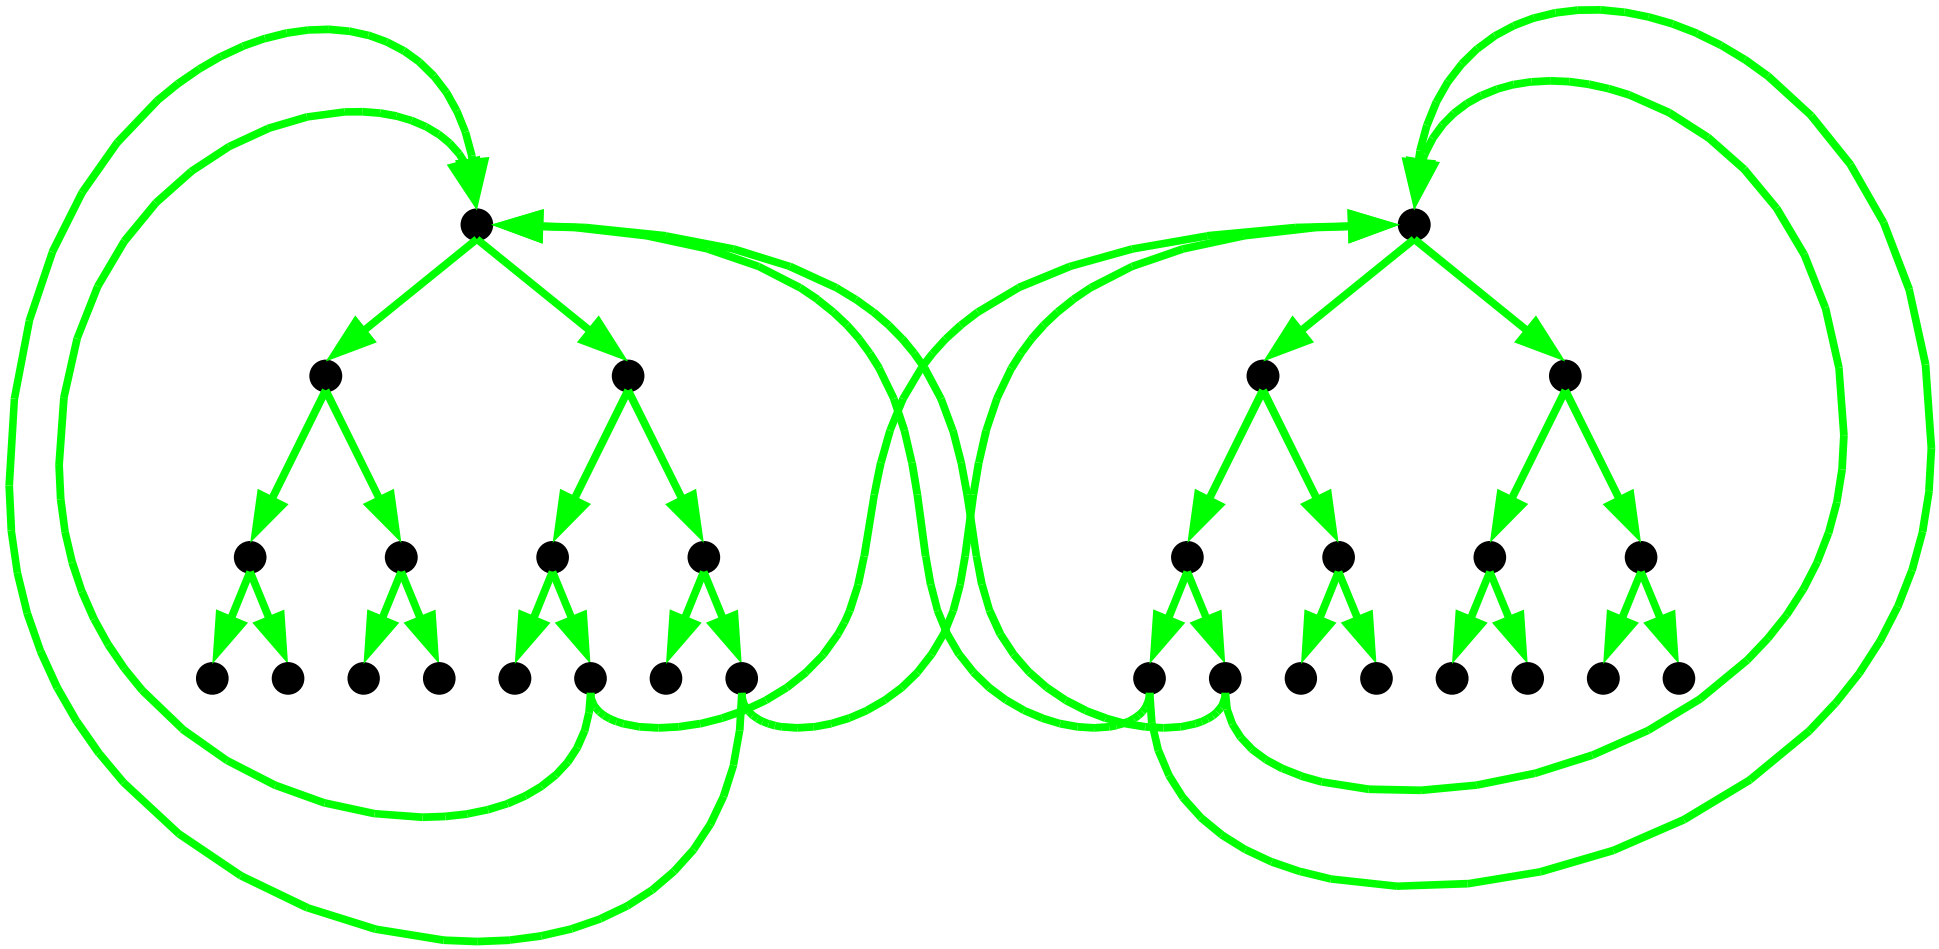

Average utility = $O(\log_2 n)$

2 Connection Game

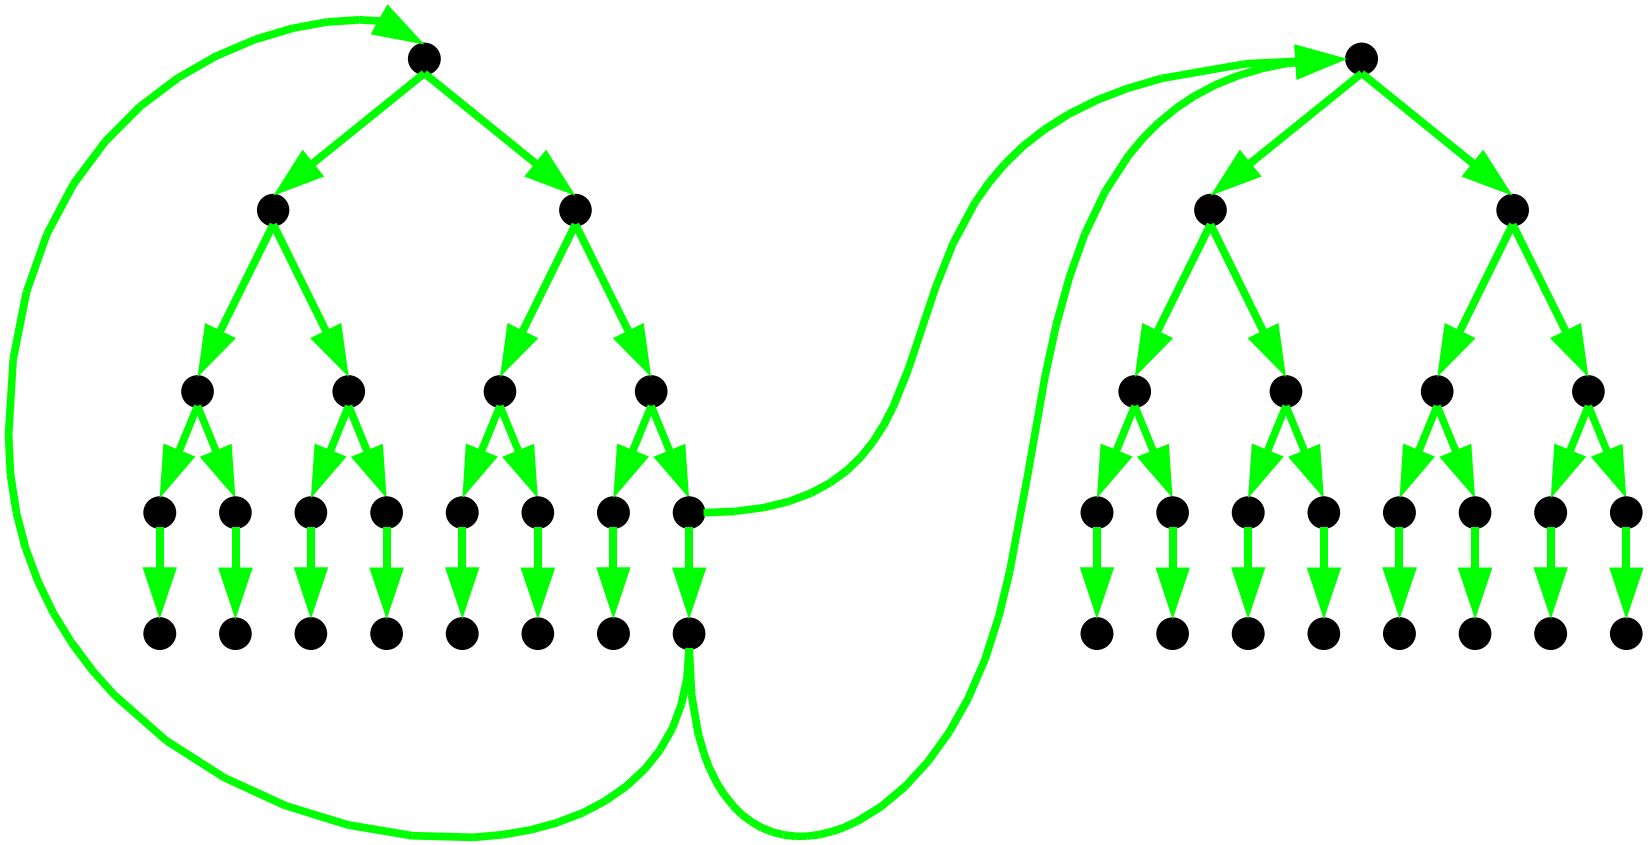

2 Connection Game

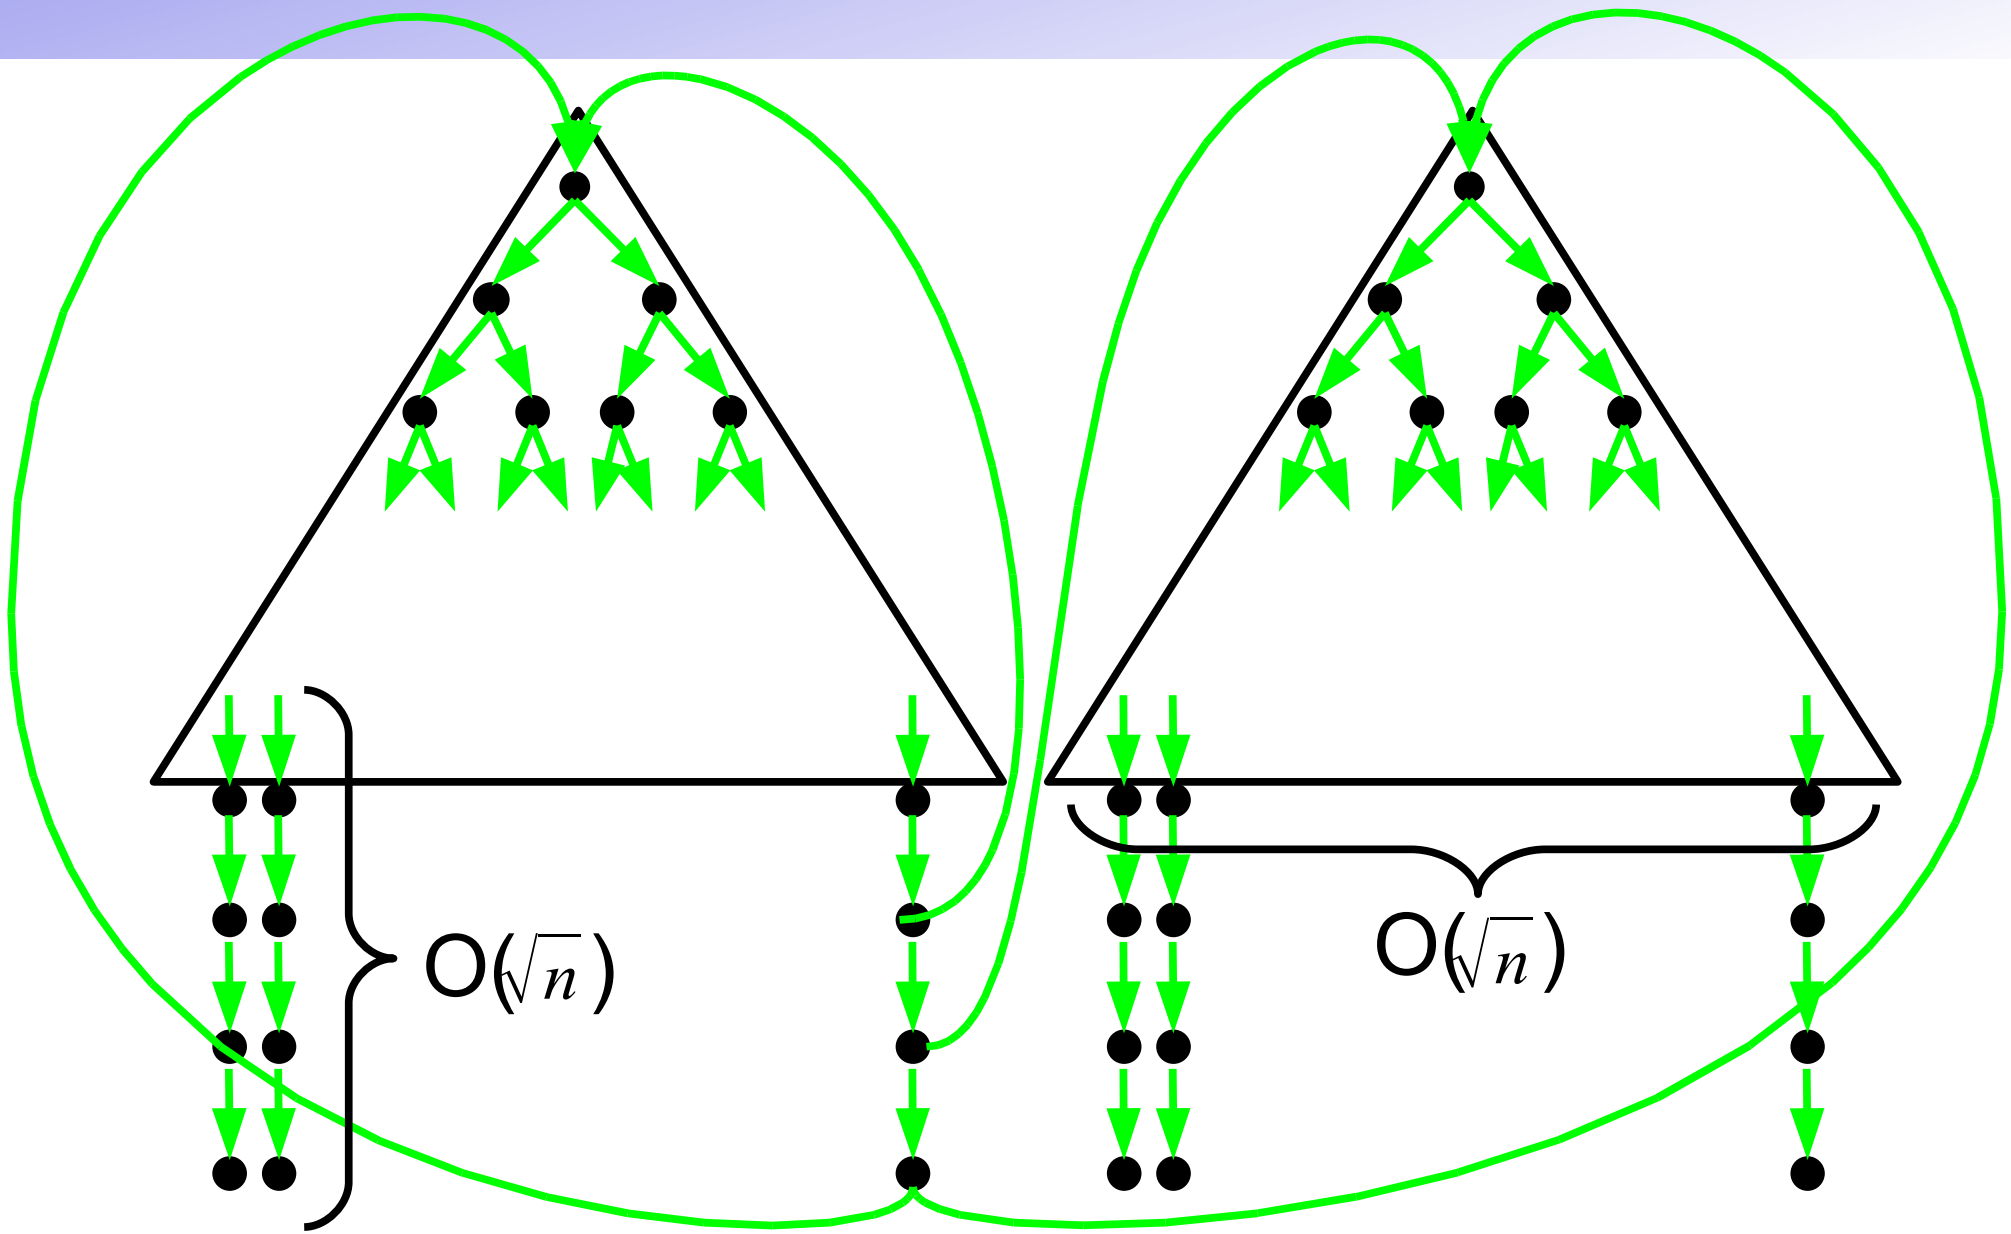

Average utility = $O(\sqrt{n})$

2 Connection Game

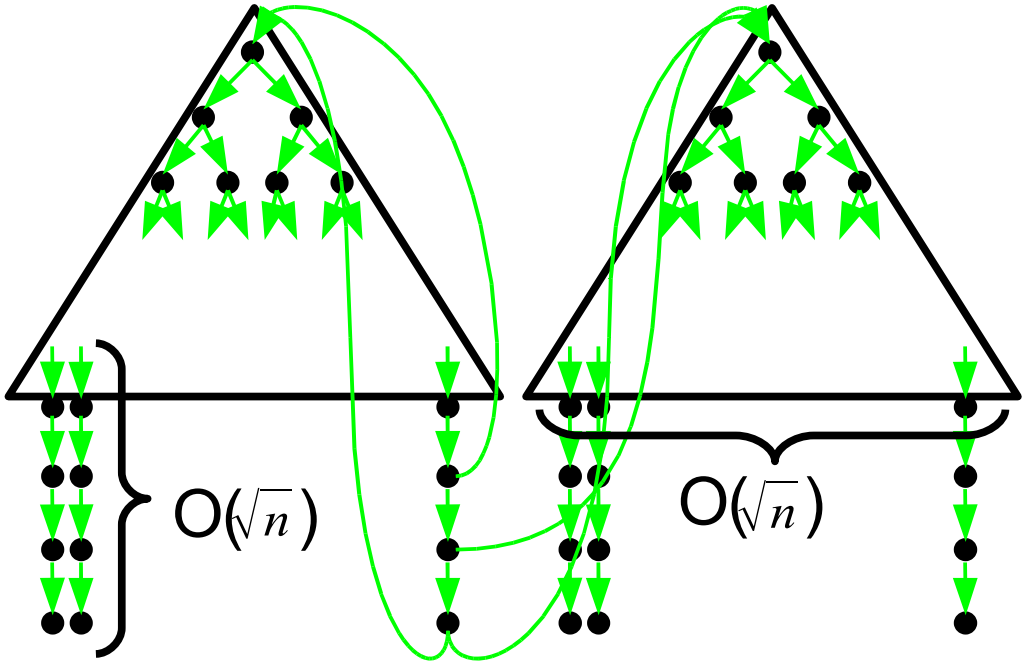

Average utility = $O(\sqrt{n})$

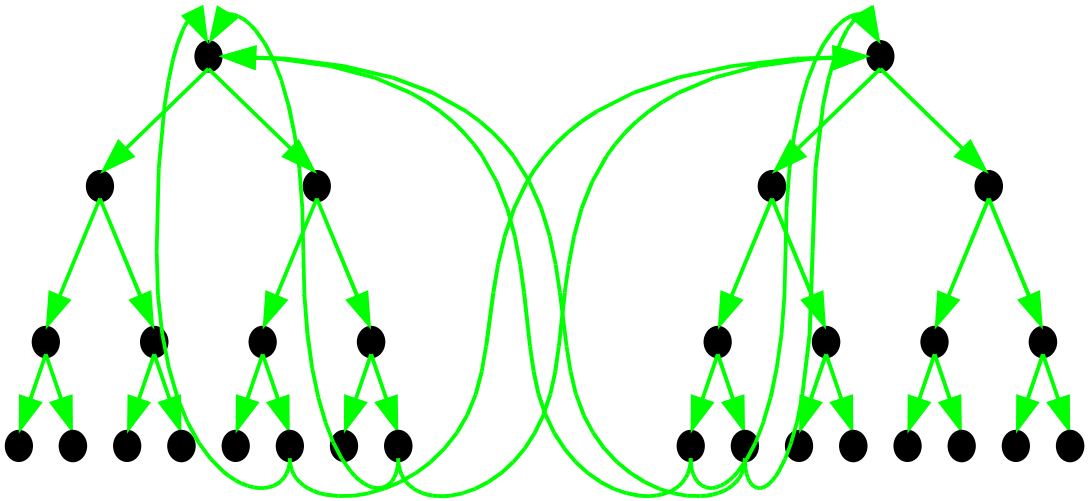

Average utility = $O(\log_2 n)$

Price of Anarchy = $\Omega(\sqrt{n} / \log_2 n)$

k-Connection Game

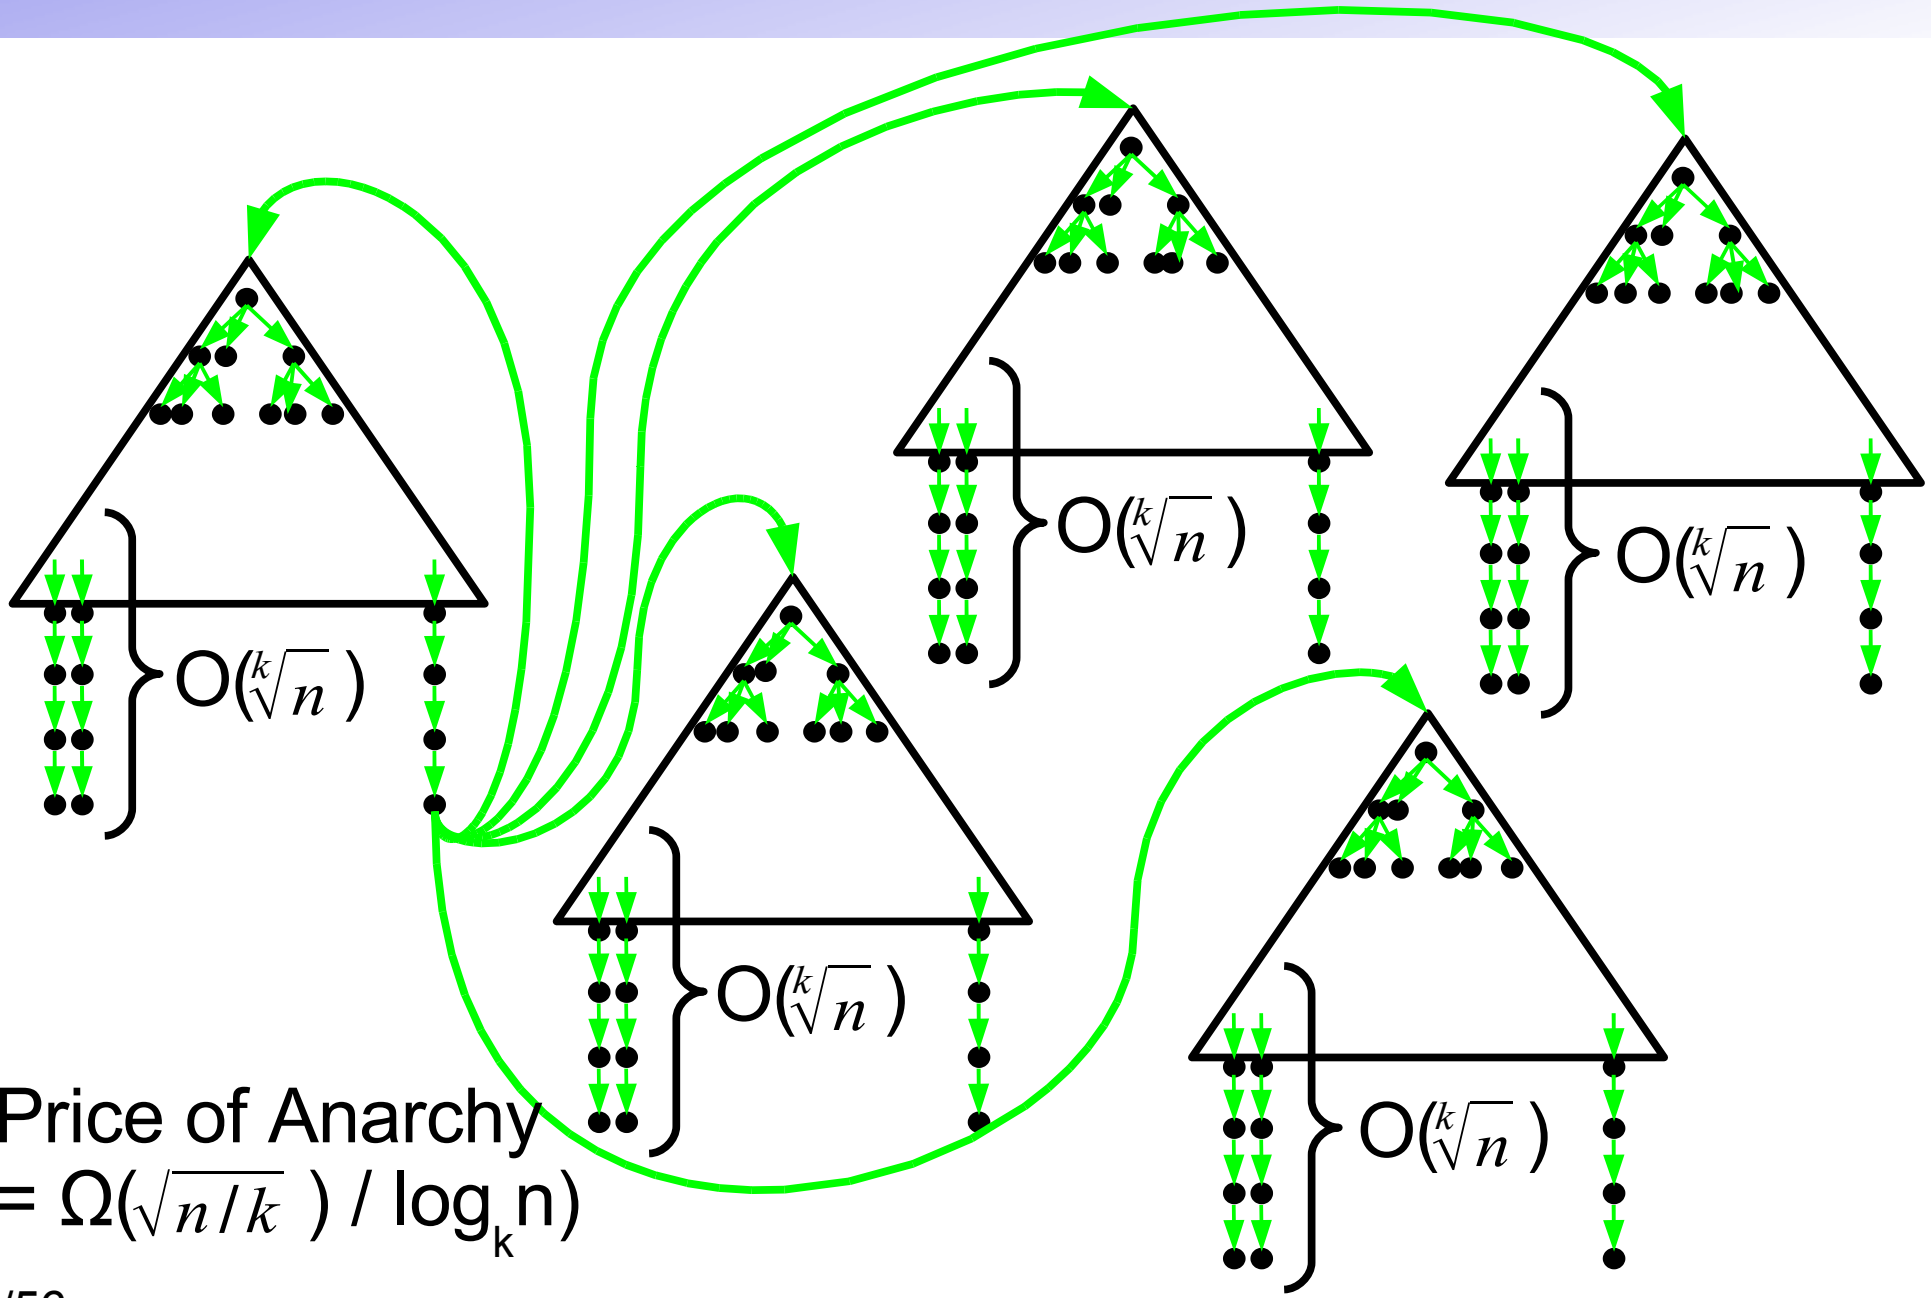

Price of Anarchy
= $\Omega(\sqrt{n/k}\ ) / \log_k n)$

k-Connection Game

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?

  - Stable or oscillating?

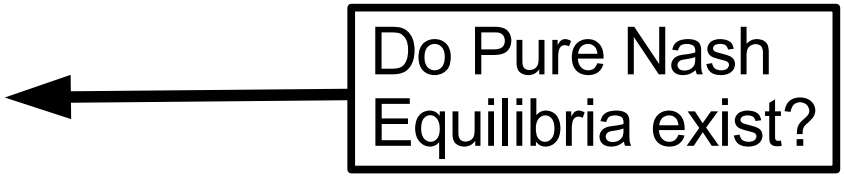

stable

  - Converging or cycling?

k-Connection Game

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?

Price of Anarchy = Worst Nash / Social Optimum

$\sqrt{n/k}$ / $\log_k$ n

Do Pure Nash Equilibria exist?

stable

Will best response trials converge to an equilibrium?

Experimental Results

- Best response walks
  - If all decisions are left to selfish nodes, will the system converge to one of these equilibria?
  - Which one?

Experimental Results

- Best response walks

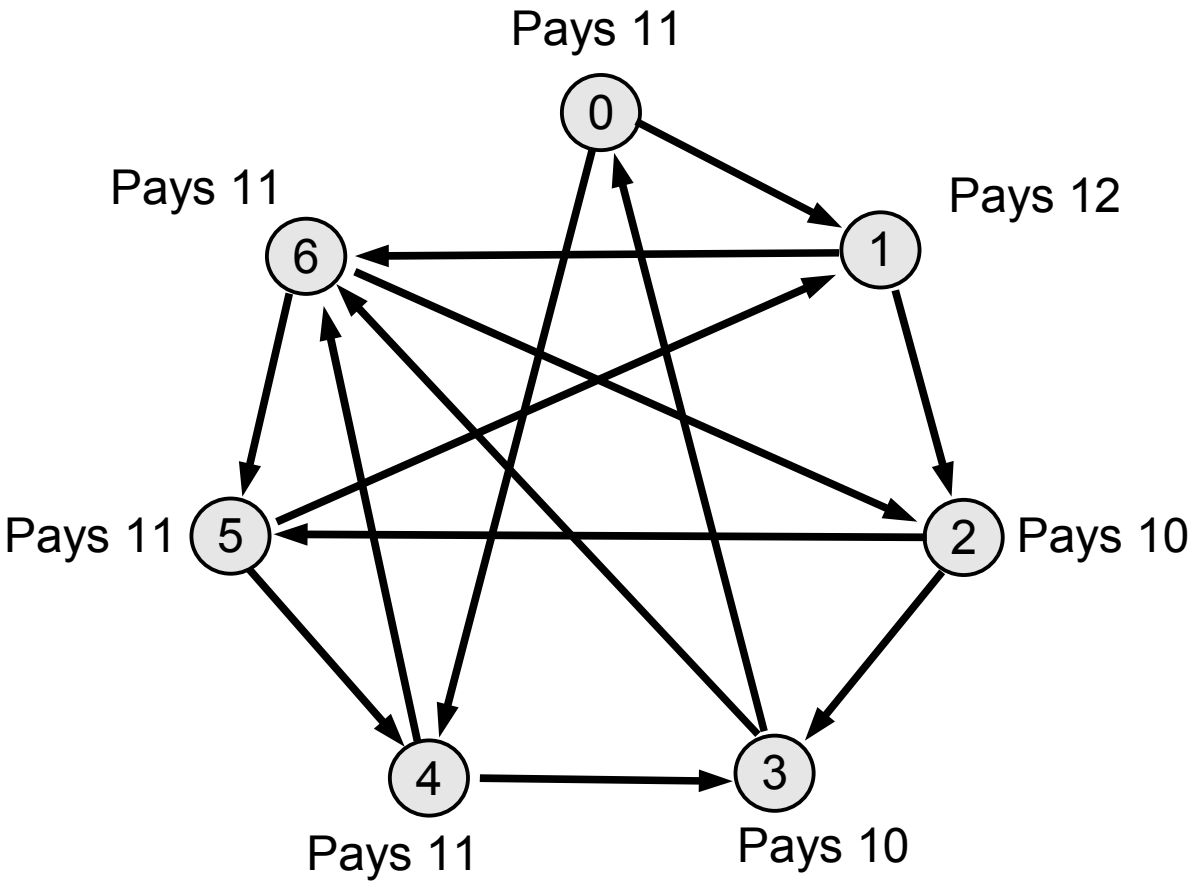

Experimental Results

- Best response walks

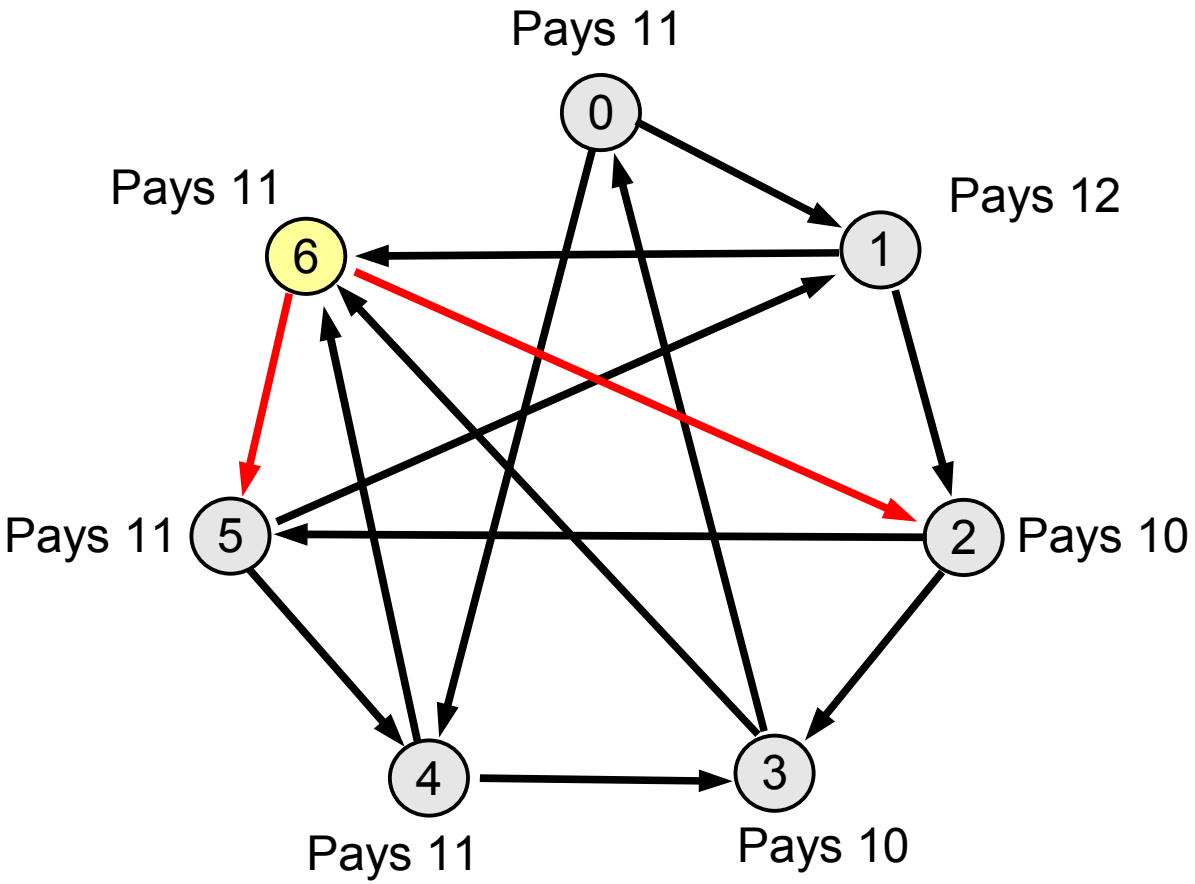

Experimental Results

- Best response walks

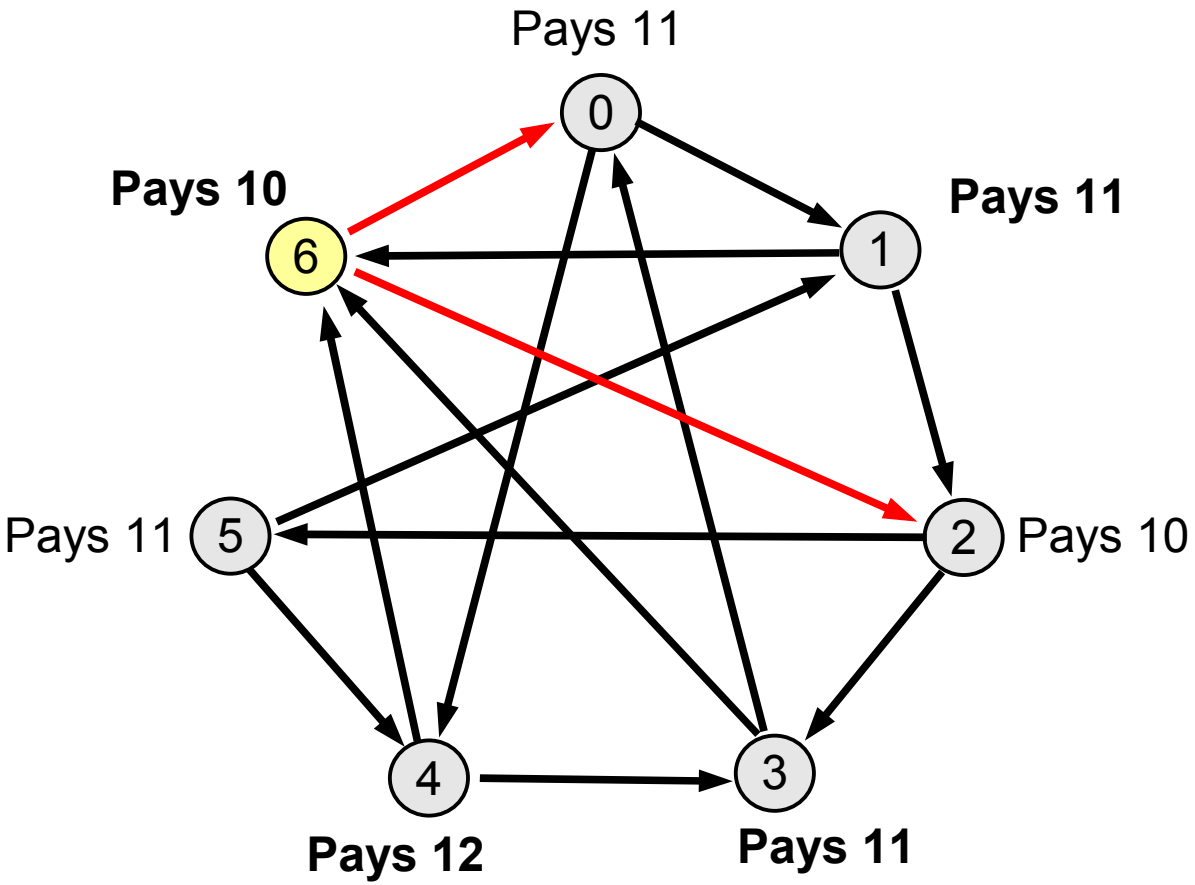

Experimental Results

- Best response walks

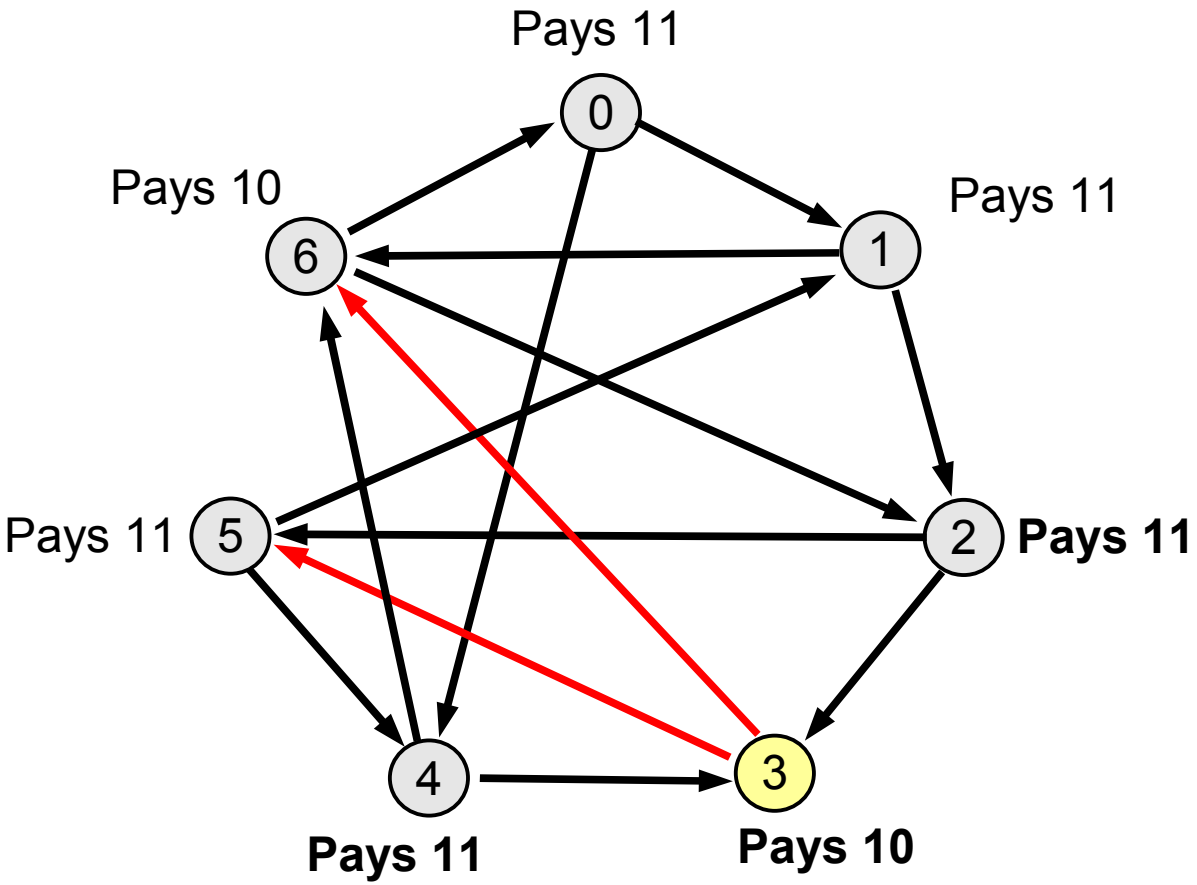

Experimental Results

- Best response walks

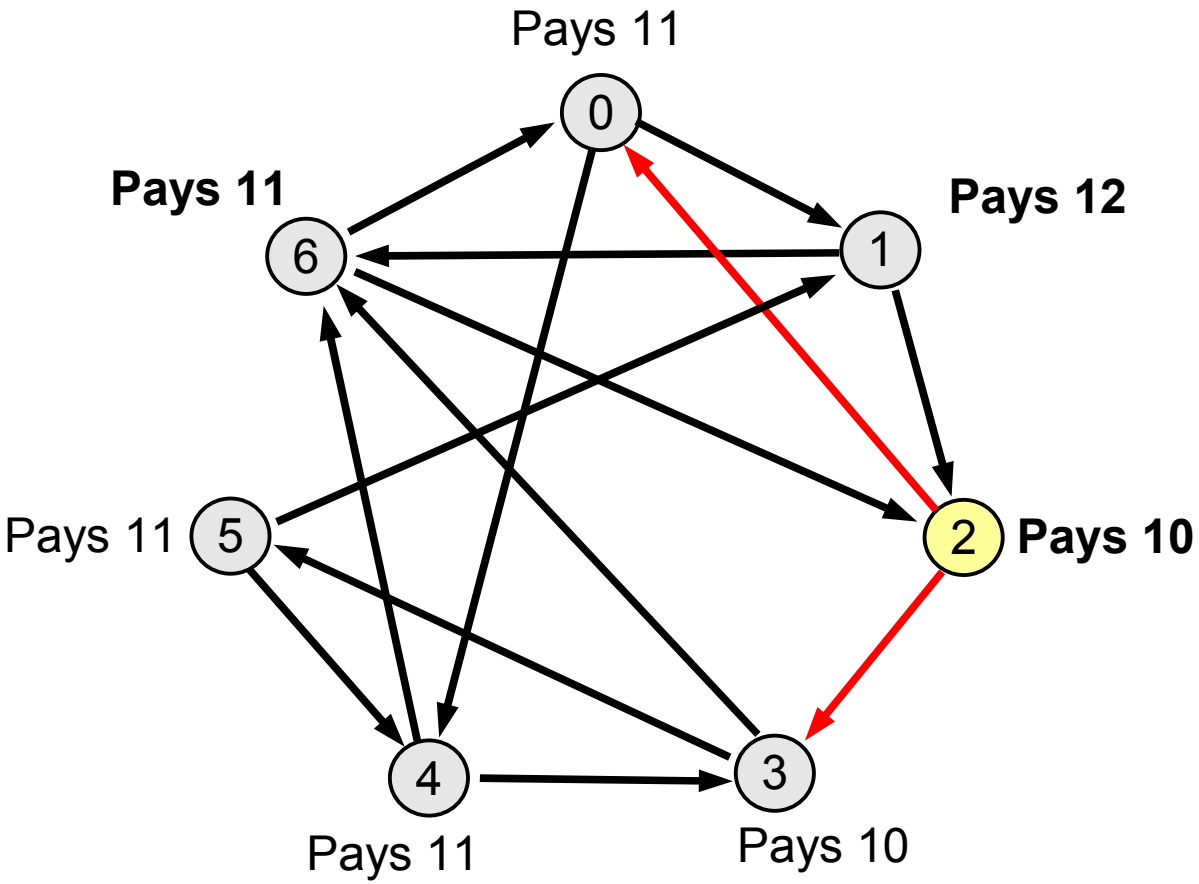

Experimental Results

- Best response walks

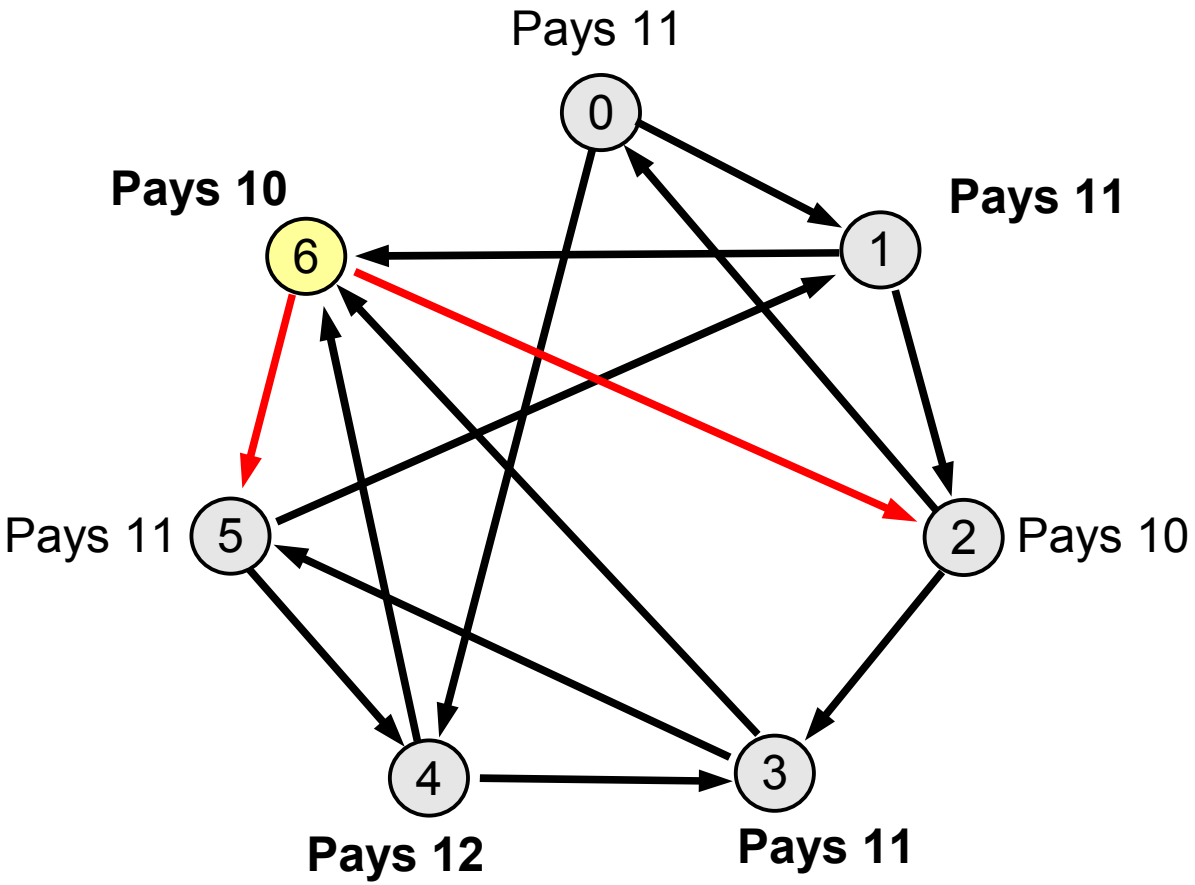

Experimental Results

- Best response walks

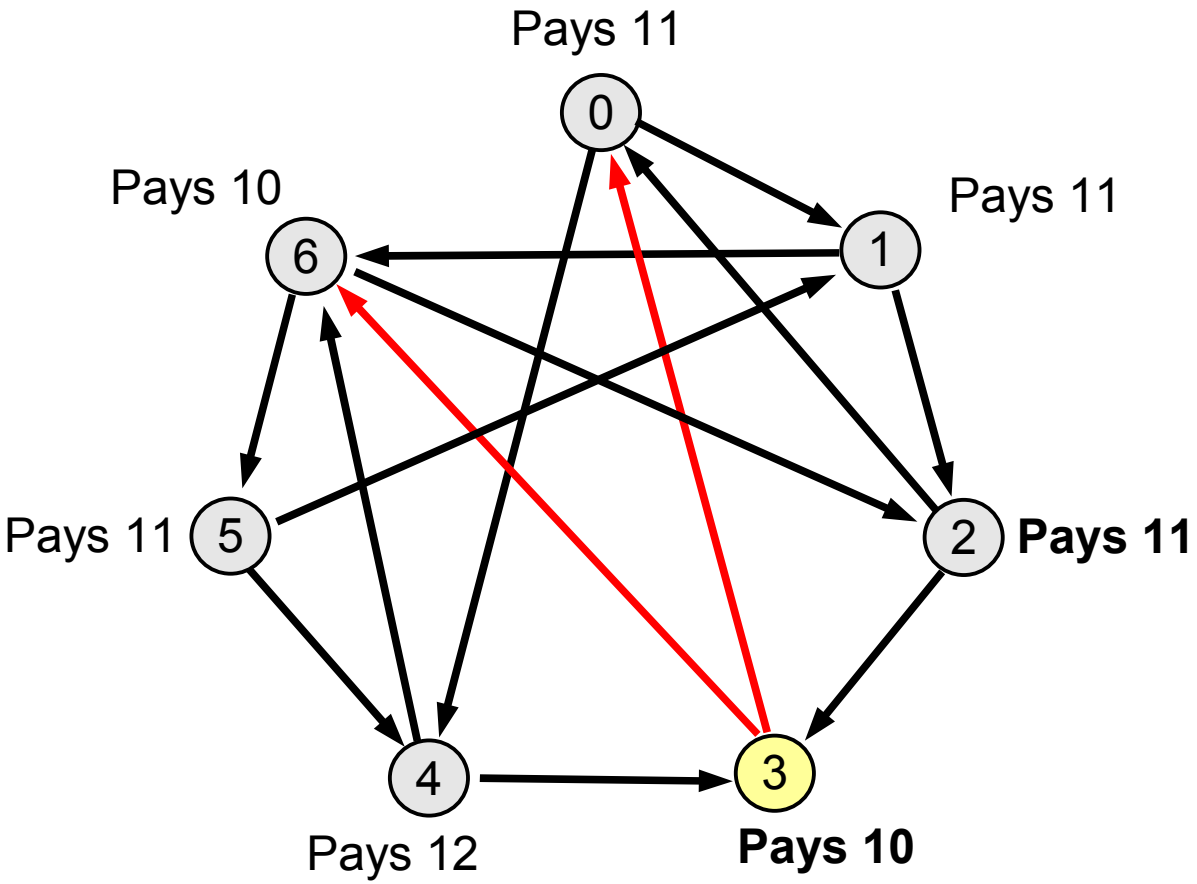

Experimental Results

- Best response walks

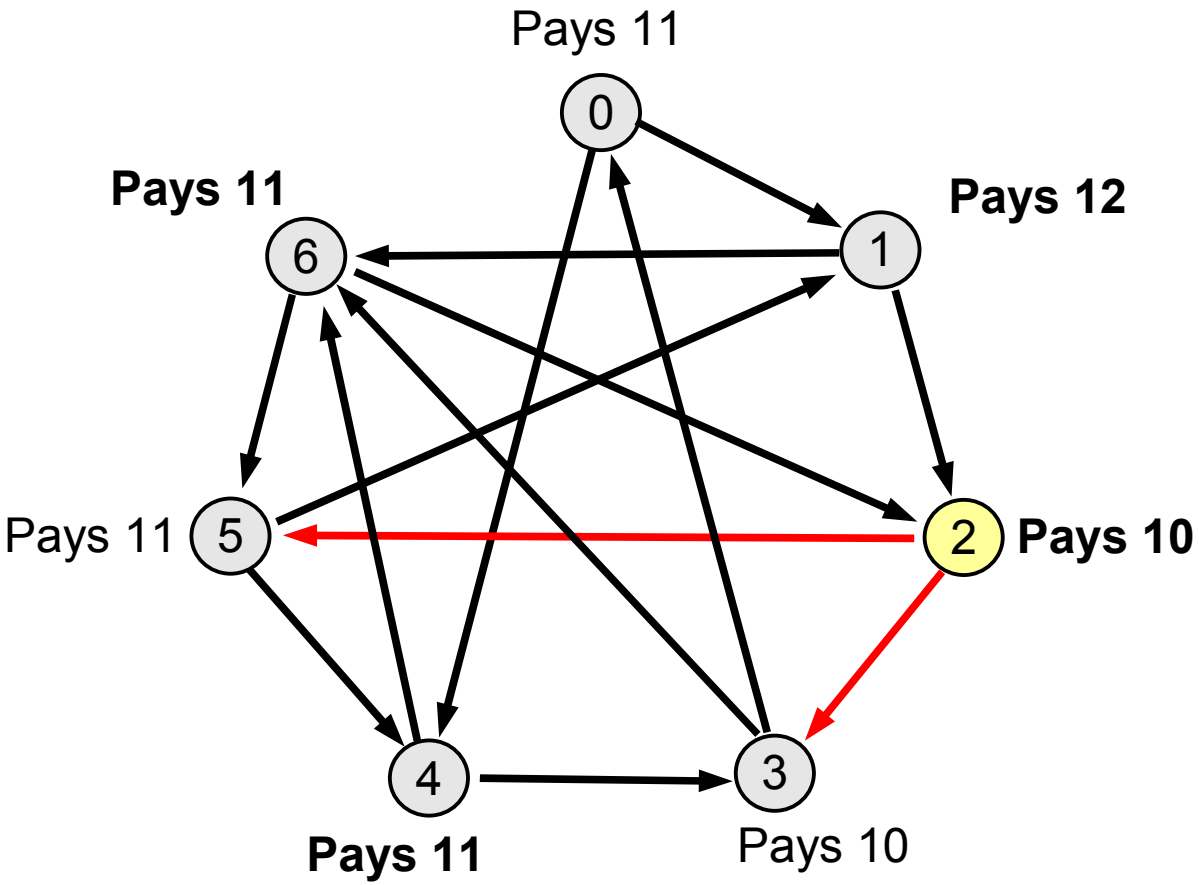

Experimental Results

- Best response walks

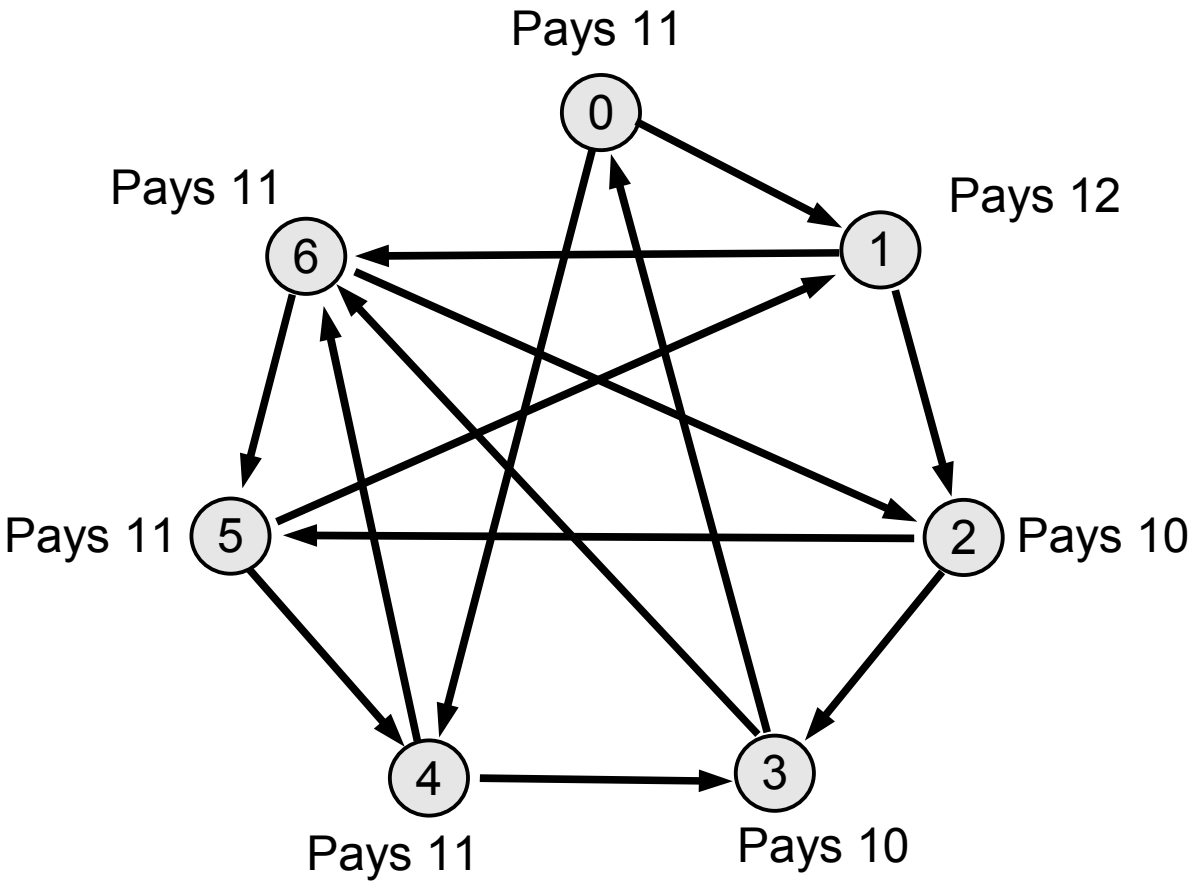

Experimental Results

- Best response walks
- More attempts at convergence

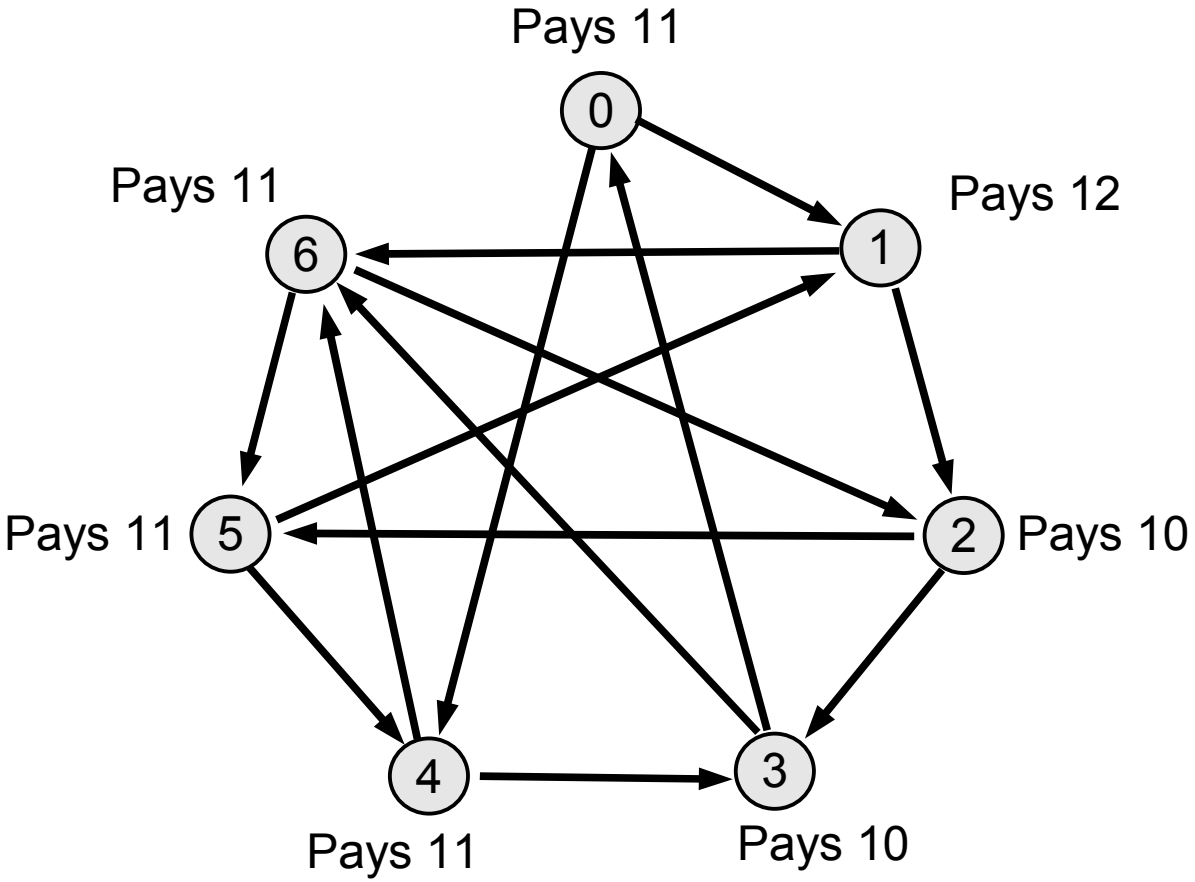

k-Connection Game

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?
  - Stable or oscillating?
  - Converging or cycling?

Price of Anarchy = Worst Nash / Social Optimum

$\sqrt{n/k}$ / $\log_k n$

Do Pure Nash Equilibria exist?

stable

Will best response trials converge to an equilibrium?

May cycle

k-Connection Games

- Number of nodes
- Budget of allowed links per node, *k*
- Want to minimize hop count distance.
- Each node picks ≤ *k* links to minimize

$$\sum_{\text{other nodes}} (\text{hop count distance to node})$$

Bounded Budget Connection (BBC) Games

- Number of nodes
- Link cost for each directed pair of nodes
- Budget of allowed link cost per node, $k(v)$
- Want to minimize hop count distance.
- Each node picks ≤ k links to minimize

$$\sum_{\text{other nodes}} (\text{hop count distance to node})$$

Bounded Budget Connection (BBC) Games

- Number of nodes
- Link cost for each directed pair of nodes
- Budget of allowed link cost per node, $k(v)$
- Length metric from the perspective of each node
- Each node picks ≤ k links to minimize

$$\sum_{\text{other nodes}} (\text{shortest path distance to node})$$

Bounded Budget Connection (BBC) Games

- Number of nodes
- Link cost for each directed pair of nodes
- Budget of allowed link cost per node, *k(v)*
- Length metric from the perspective of each node
- Affinity for each directed pair of nodes
- Each node picks ≤ *k* links to minimize

$$\sum_{\text{other nodes}} (\text{affinity} * \text{shortest path distance to node})$$

Example

Edge Cost = 1,
Budget = 2

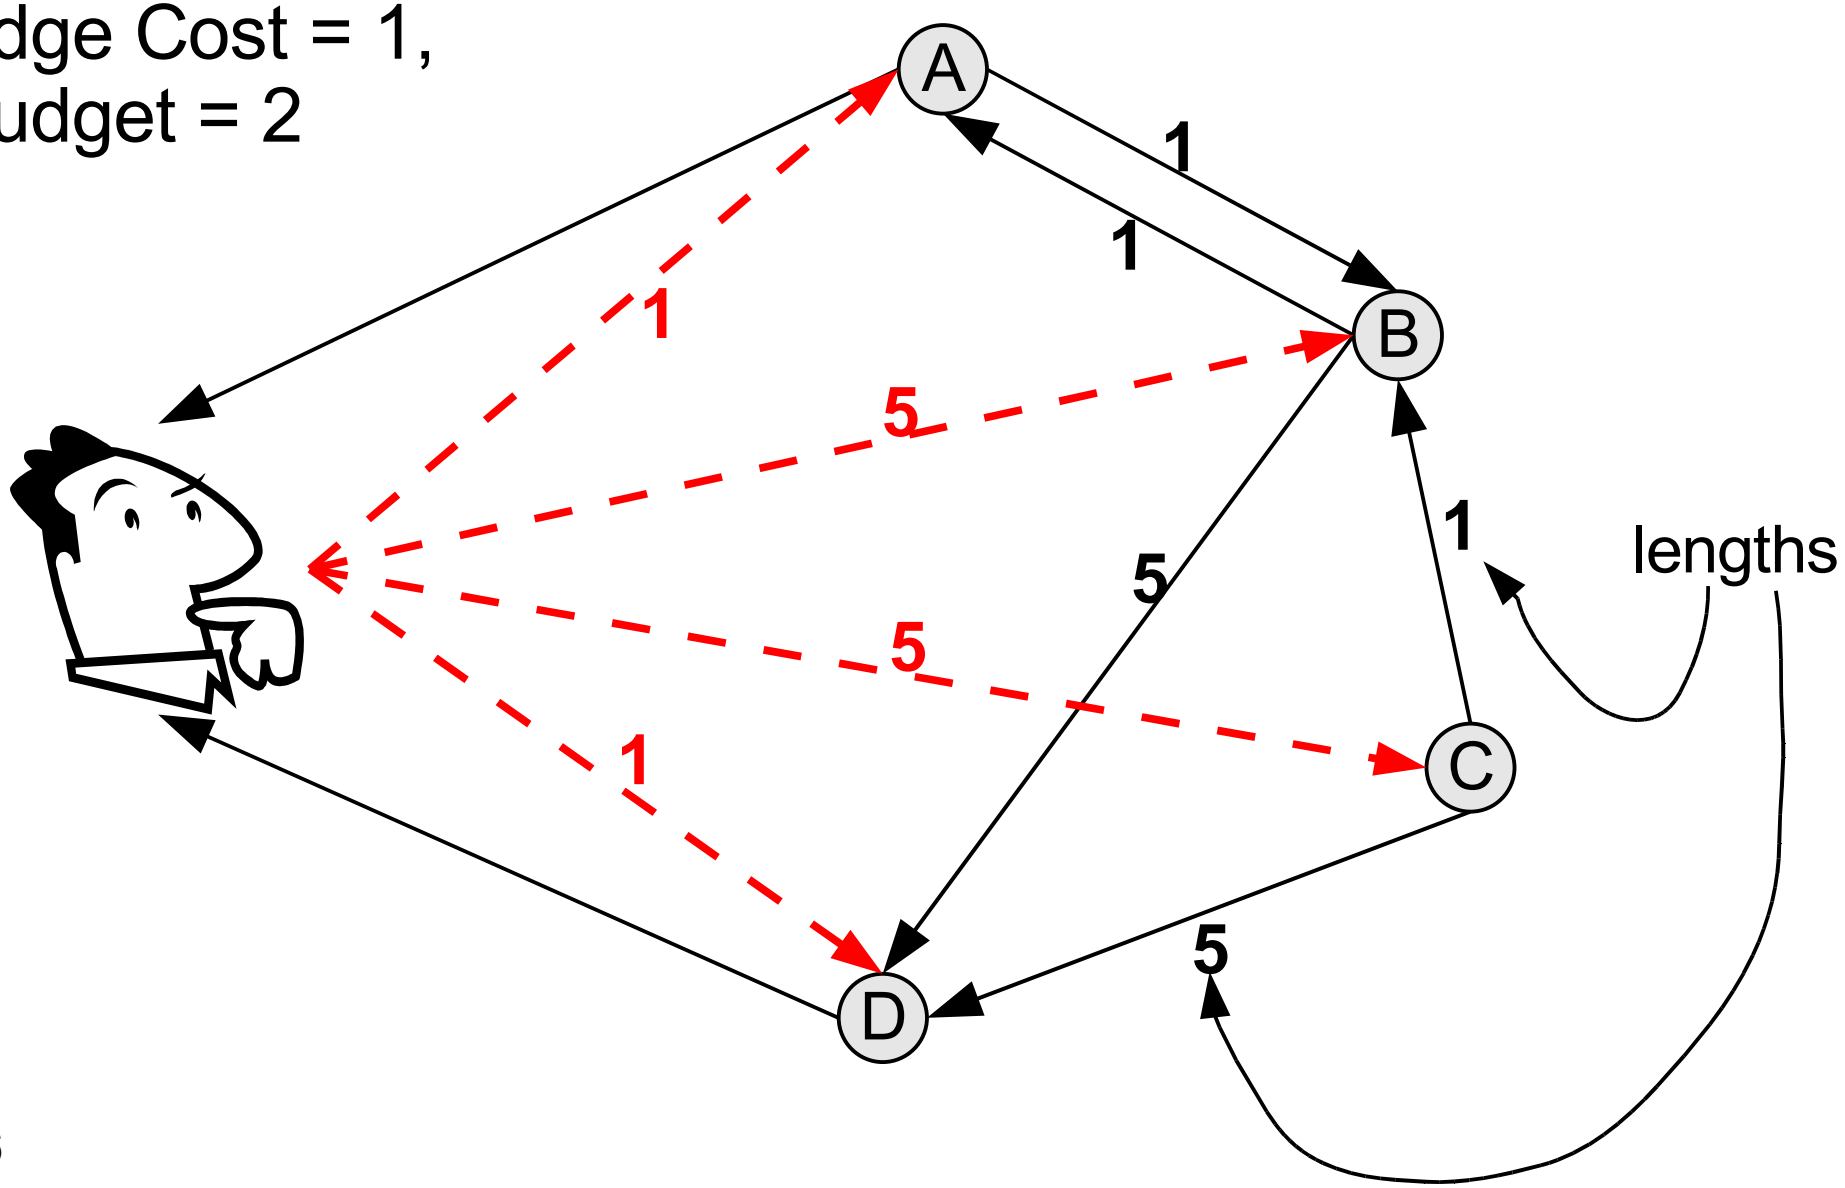

Example

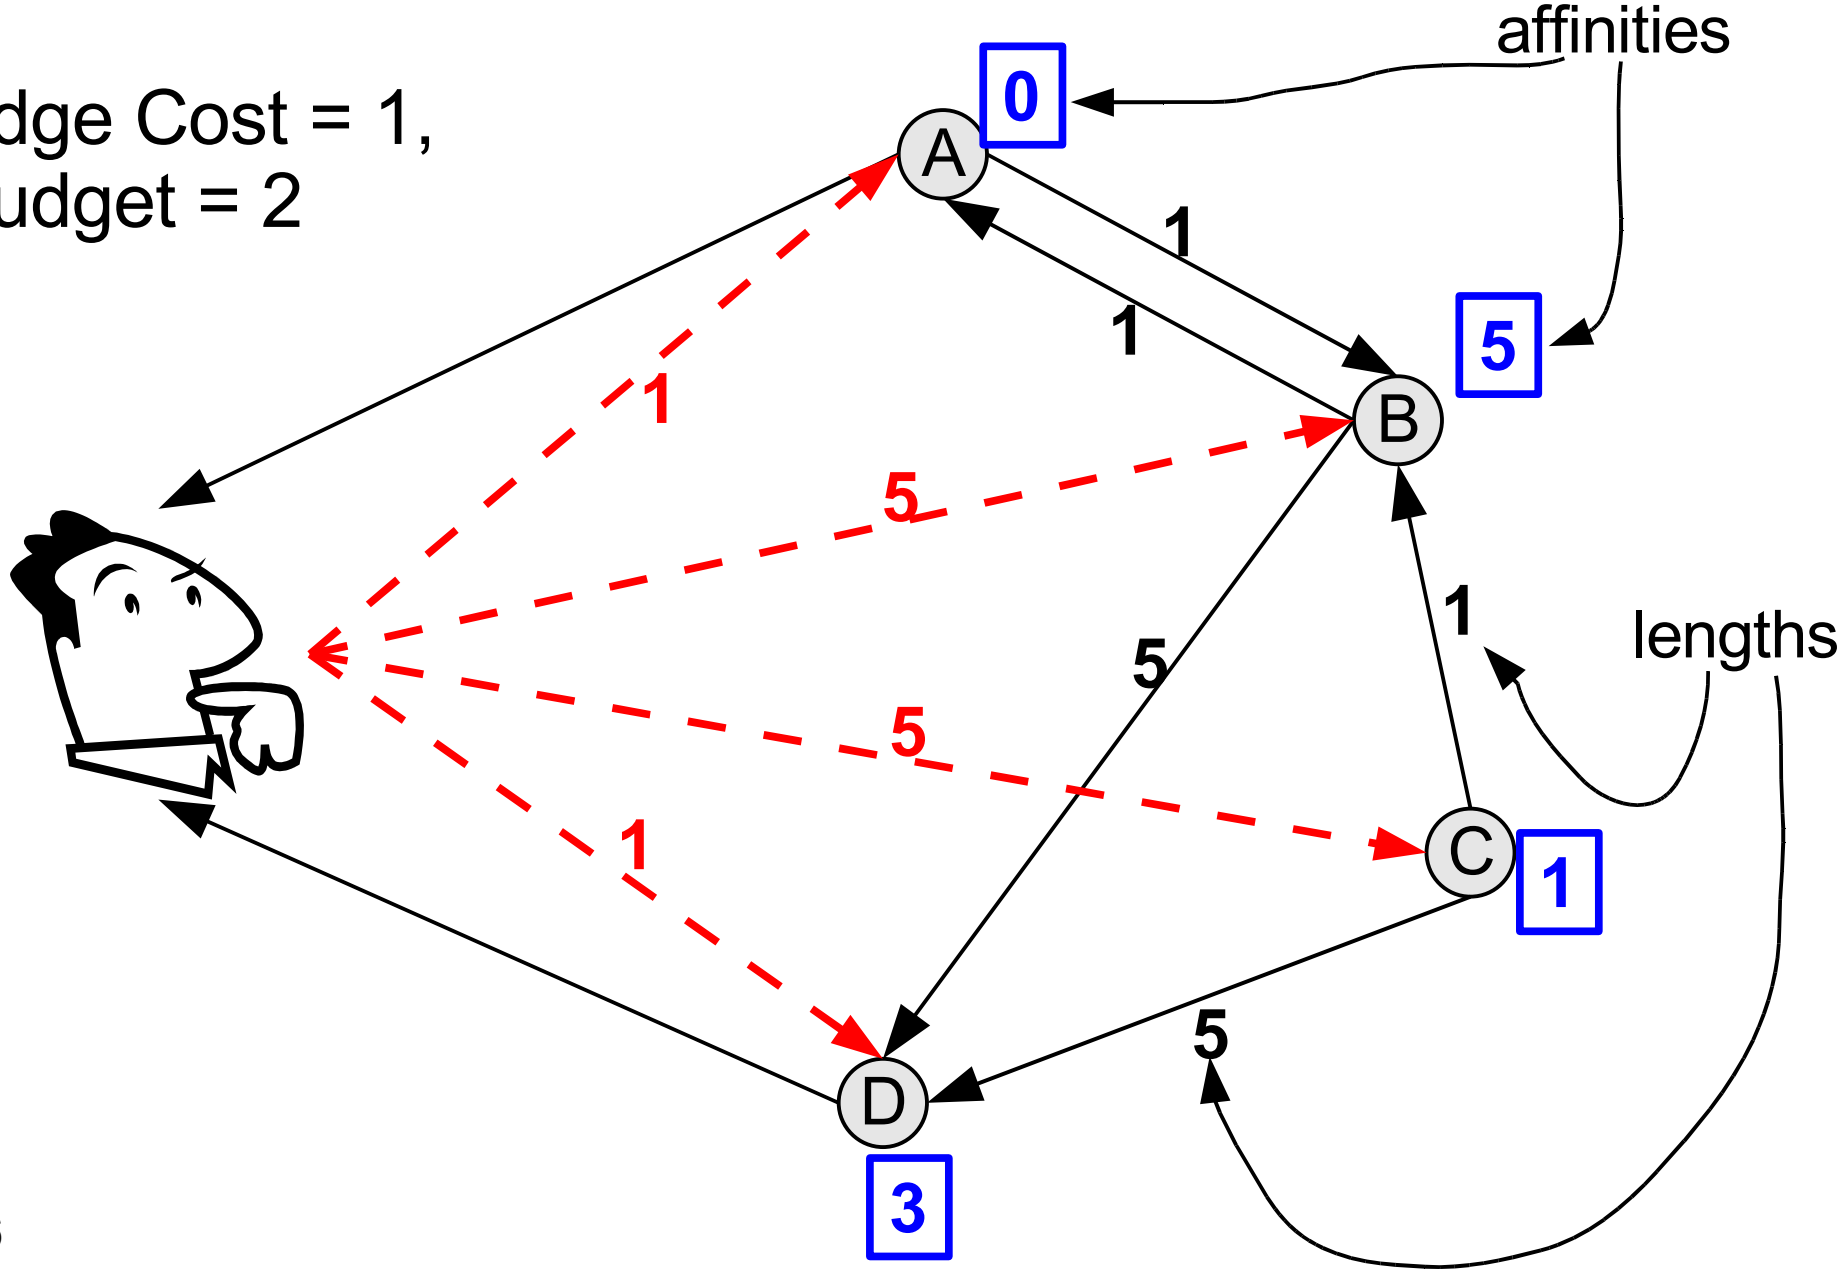
Edge Cost = 1,
Budget = 2
affinities
0
A
1
1
5
B
1
5
5
1
lengths
5
1
C
1
5
D
3

Example

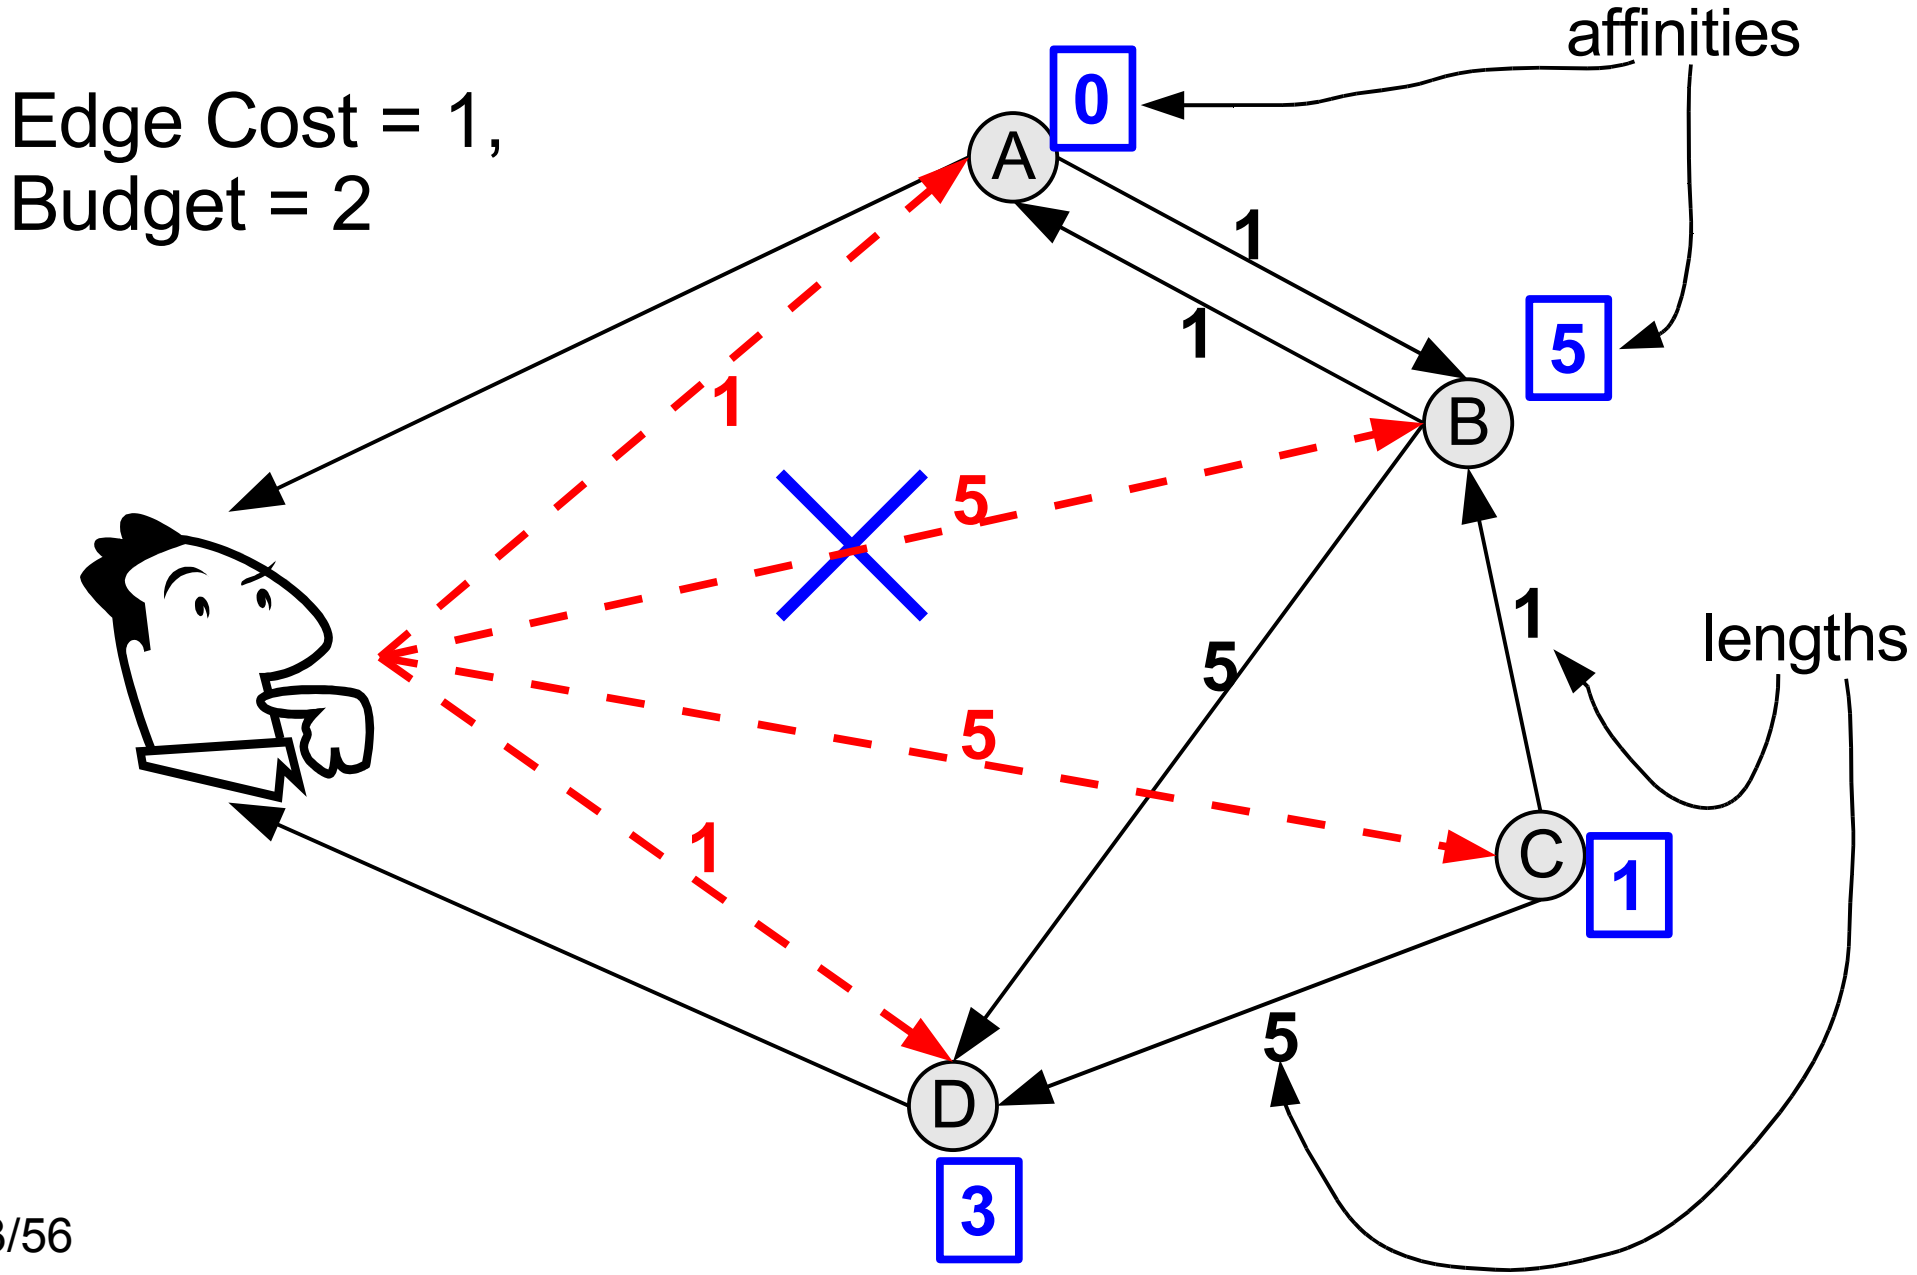
Edge Cost = 1,
Budget = 2
affinities
0
A
1
1
5
B
1
5
5
1
lengths
5
5
1
C
1
5
D
3

Example

Edge Cost = 1,
Budget = 2

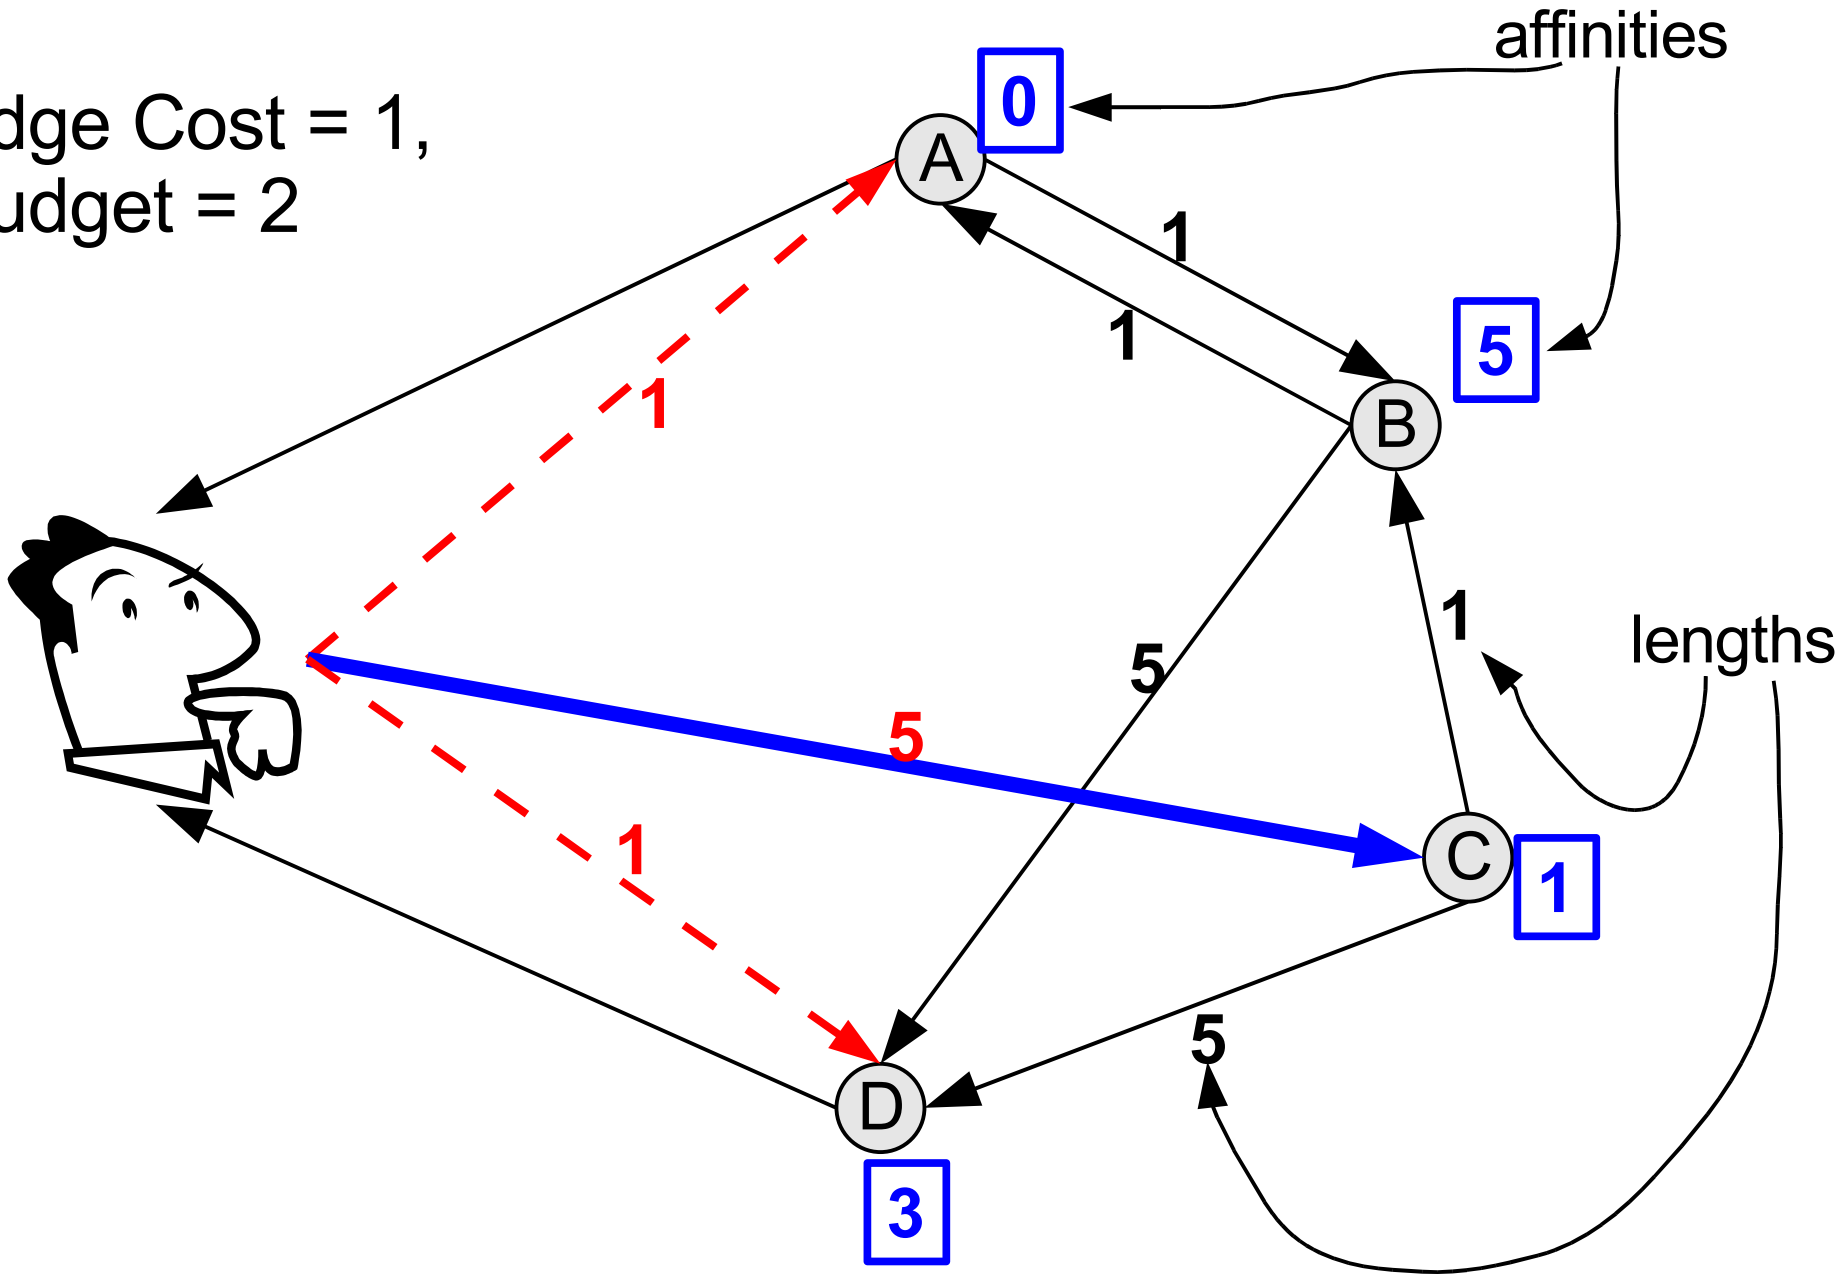

Example

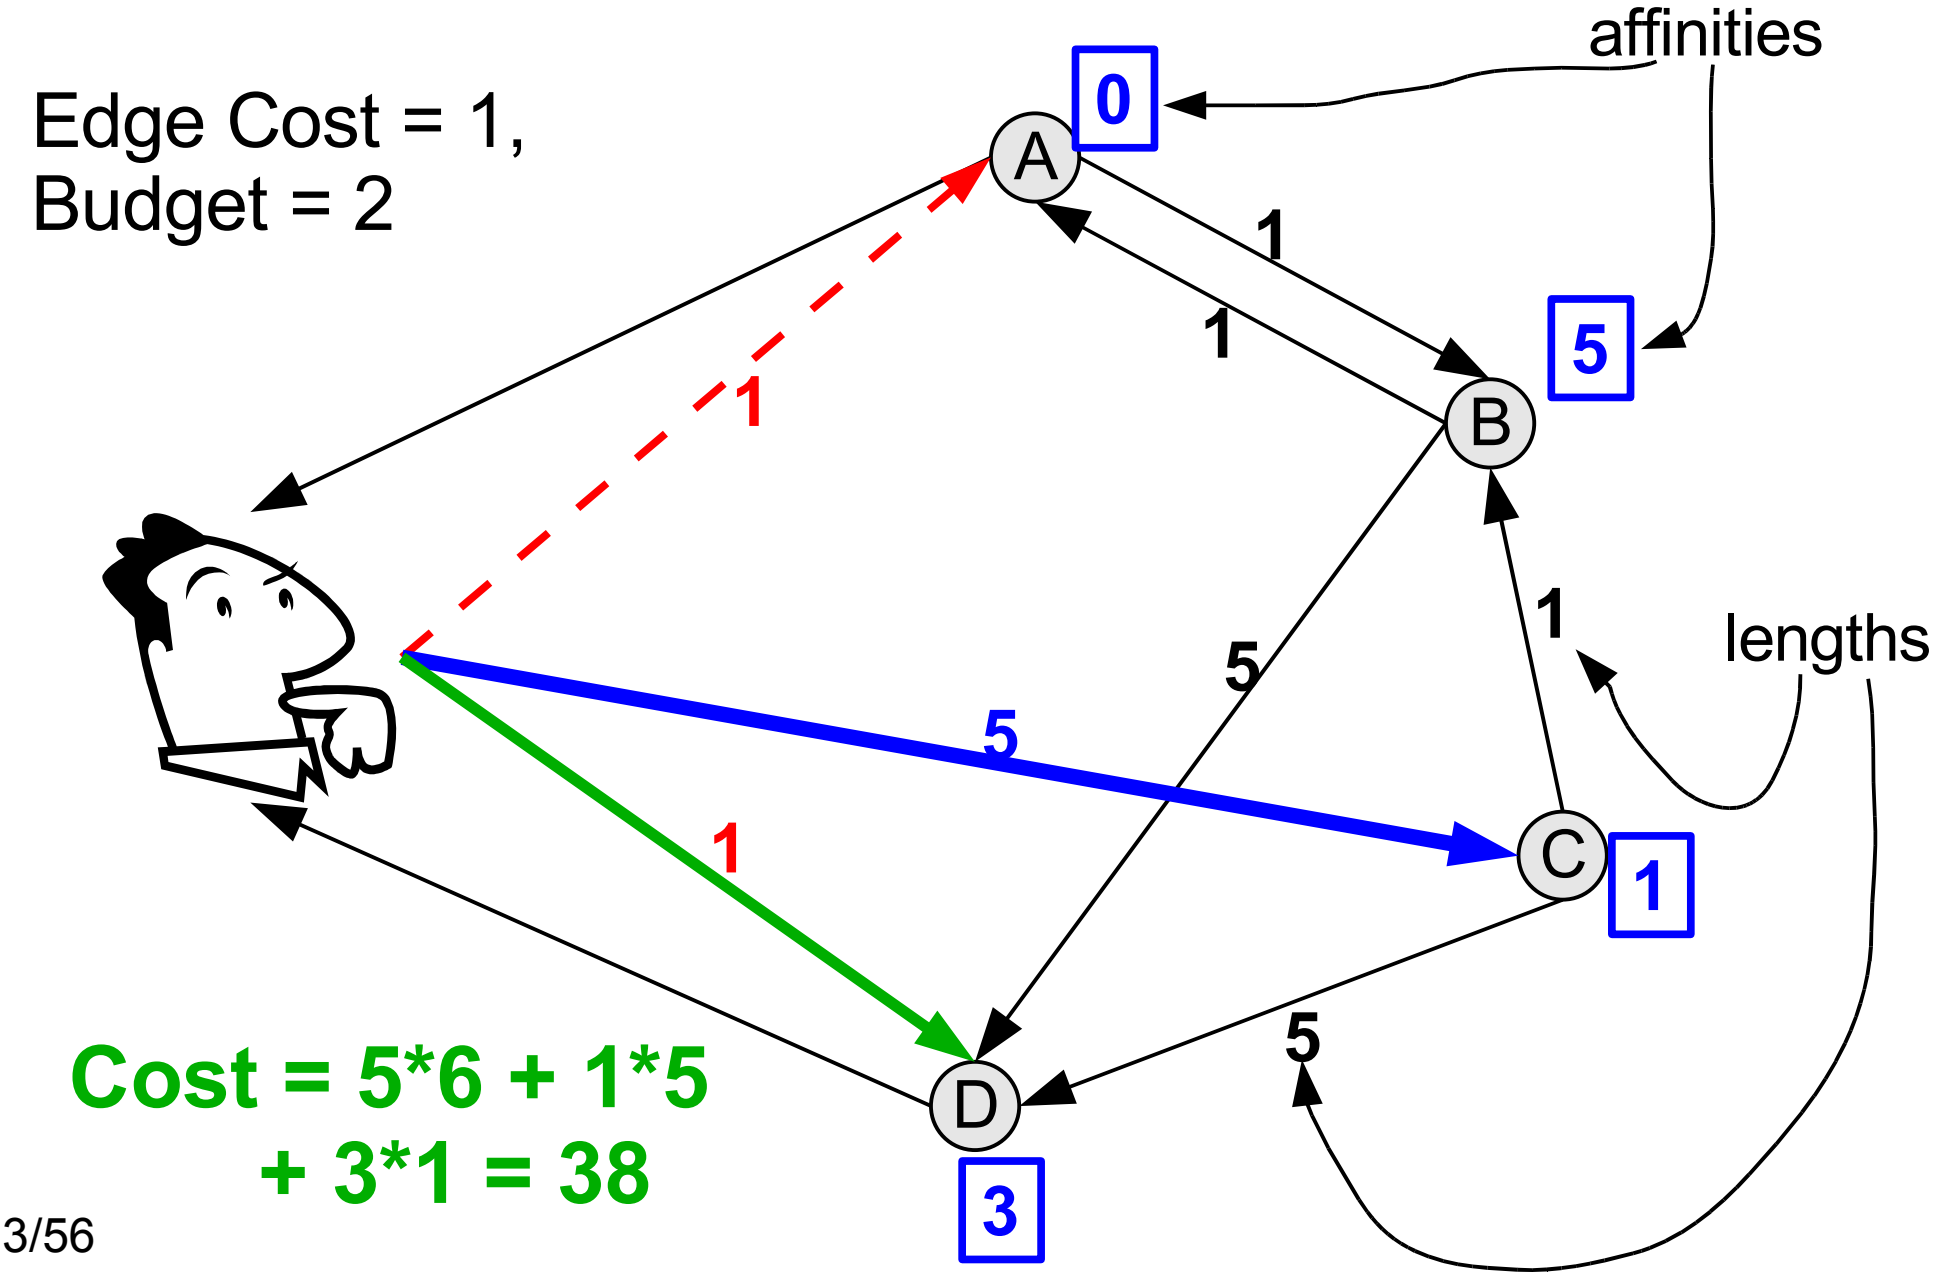
Edge Cost = 1,
Budget = 2
affinities
0
5
1
A
B
C
D
1
1
1
5
5
1
lengths
5
1
3
Cost = 5*6 + 1*5
+ 3*1 = 38

Example

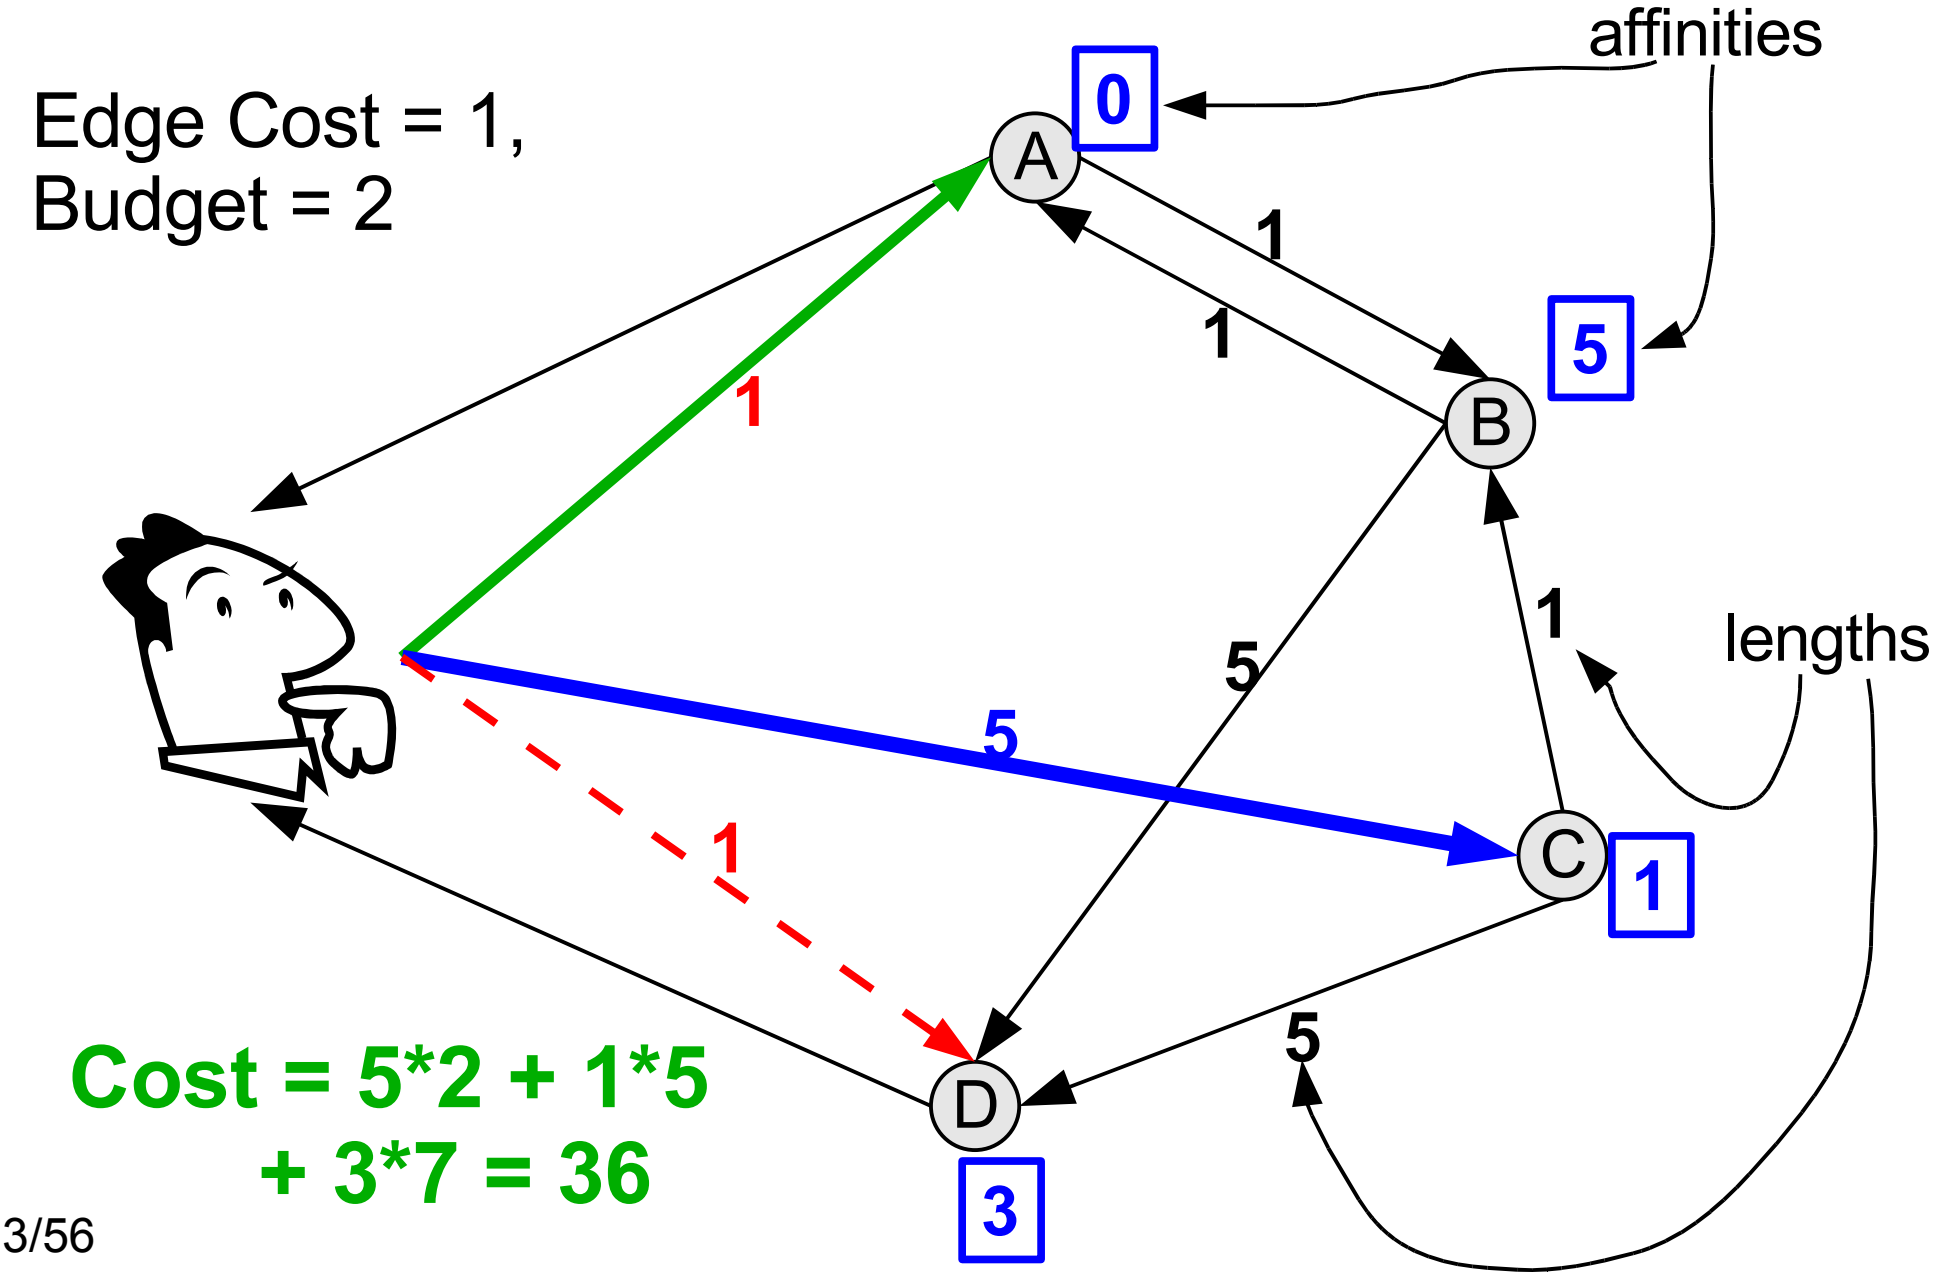
Edge Cost = 1,
Budget = 2
affinities
0
5
1
A
B
C
D
1
1
1
5
5
1
lengths
5
1
3
Cost = 5*2 + 1*5
+ 3*7 = 36

Other Network Connection Games

- Edge cost built into utility instead of a budget built into actions. Undirected edges.
  - Fabrikant, Luthra, Maneva, Papadimitriou, and Shenker. On a network creation game. PODC, 2003.
  - Albers, Eilts, Even-Dar, Mansour, and Roditty. On Nash equilibria for a network creation game. SODA, 2006.
  - Demaine, Hajiaghayi, and Mahini. The Price of Anarchy in Network Creation Games. PODC, 2007.
- Price of anarchy:
  - $\Theta(1)$ upper bound if edge cost is $O(n^{(1-\varepsilon)})$.
  - Otherwise, bounded by $2^{O(\sqrt{\lg n})}$.

Experimental Results - edge costs built into utility

- Chun, Fonseca, Stoica, and Kubiatowicz. Characterizing selfishly constructed overlay routing networks. INFOCOM, 2004.
  - Resulting graphs are usually star graphs.
  - Utility for nodes depends heavily on link costs.
  - Interesting results are when link costs depend on in-degree.

Experimental Results - bounded links budget

- Laoutaris, Smaragdakis, Bestavros, and Byers. Implications of selfish neighbor selection in overlay networks. INFOCOM, 2007.
  - Approximate best response found using asymmetric k-median approximation.
  - Uniform popularity: stable solutions close to optimum, resulting topology has hubs.
  - Non-uniform popularity: If degree bound is low, popular nodes are hubs. Otherwise, hubs are anywhere.

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

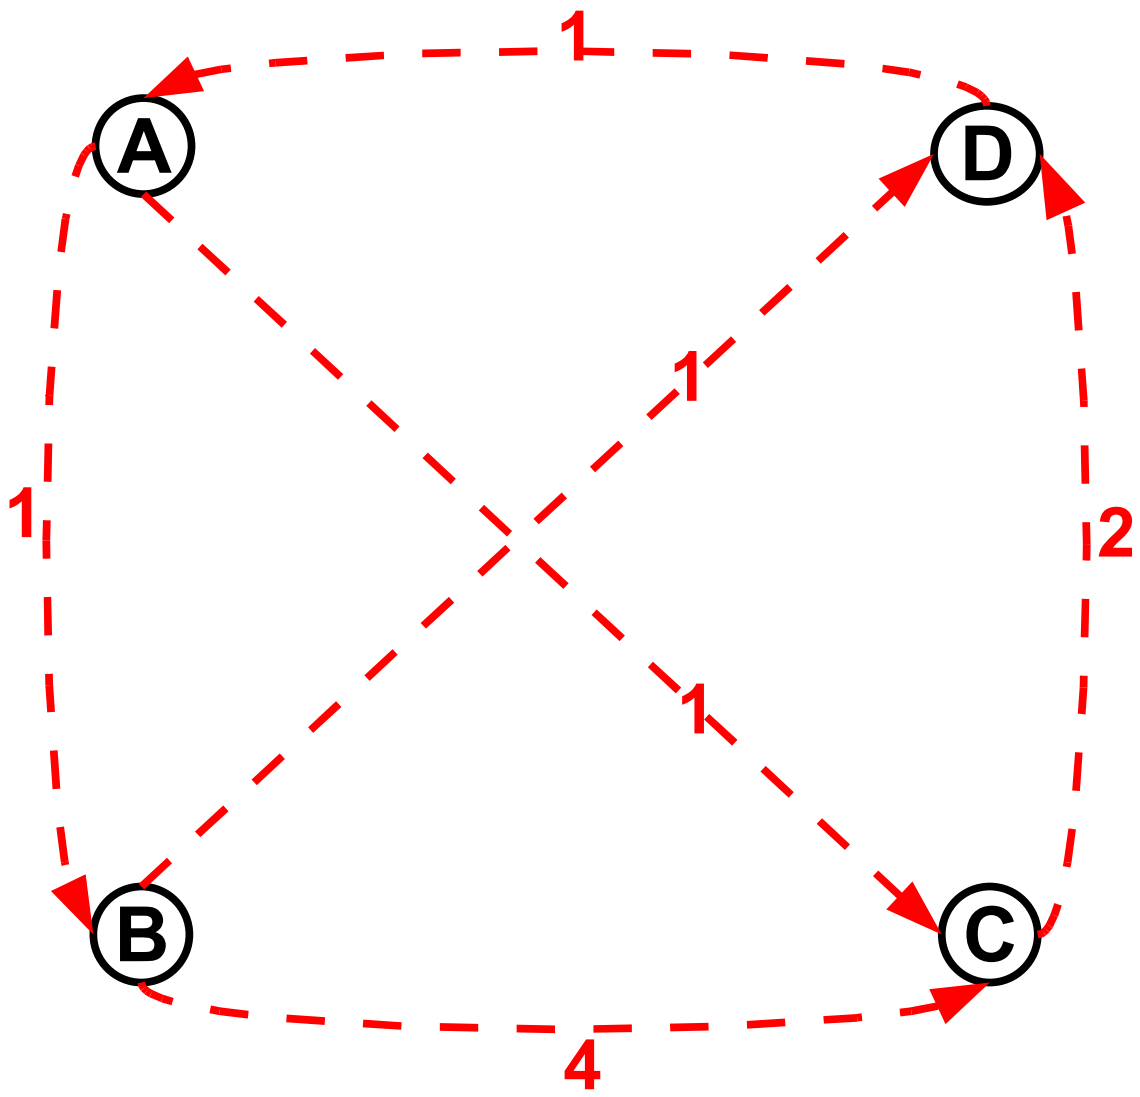

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

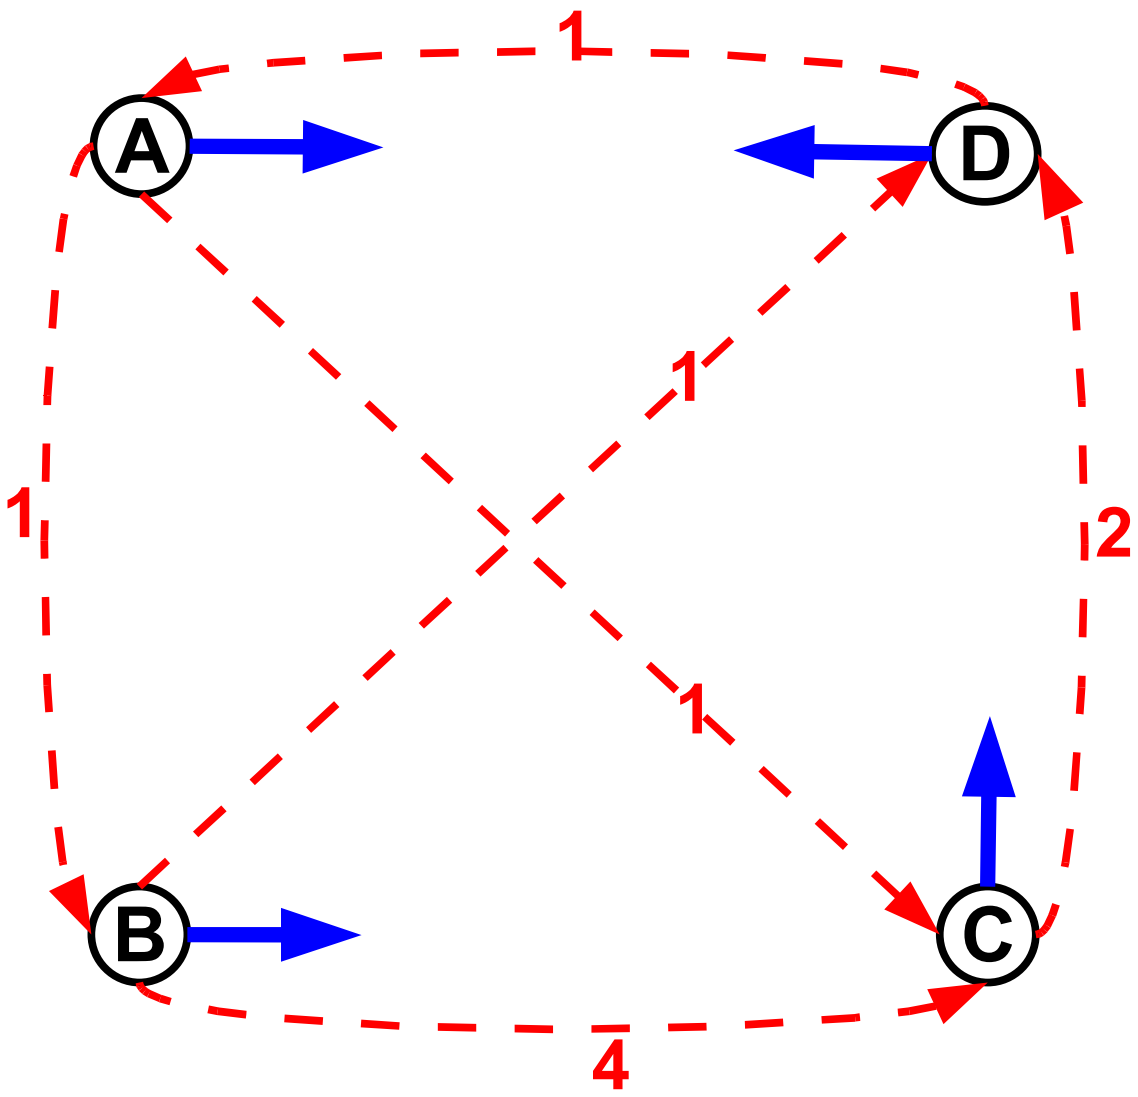

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

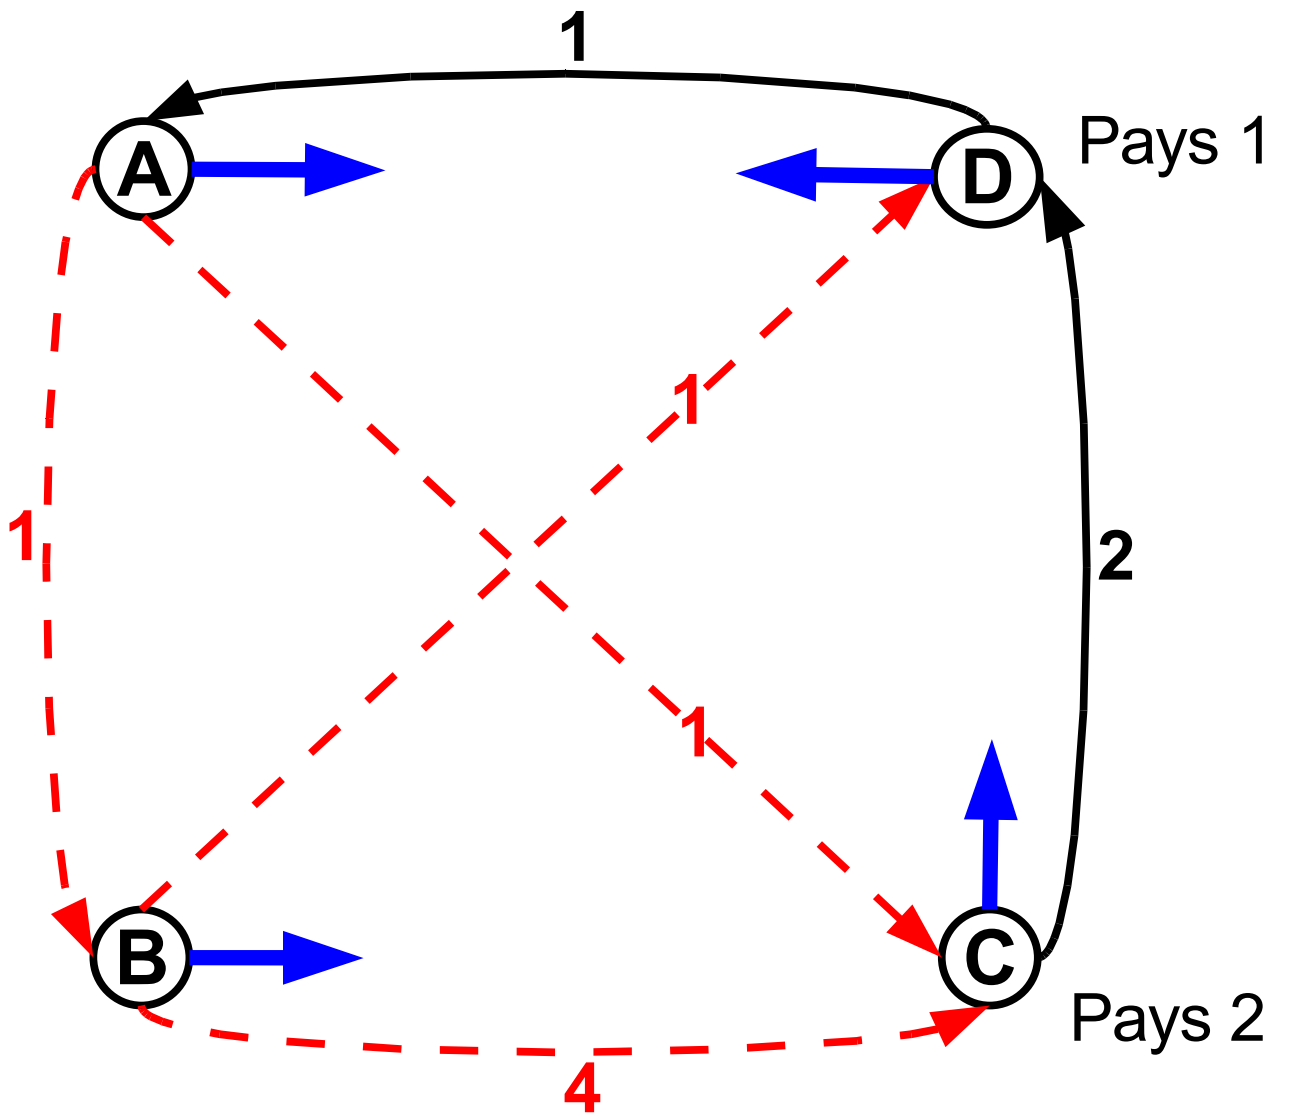

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

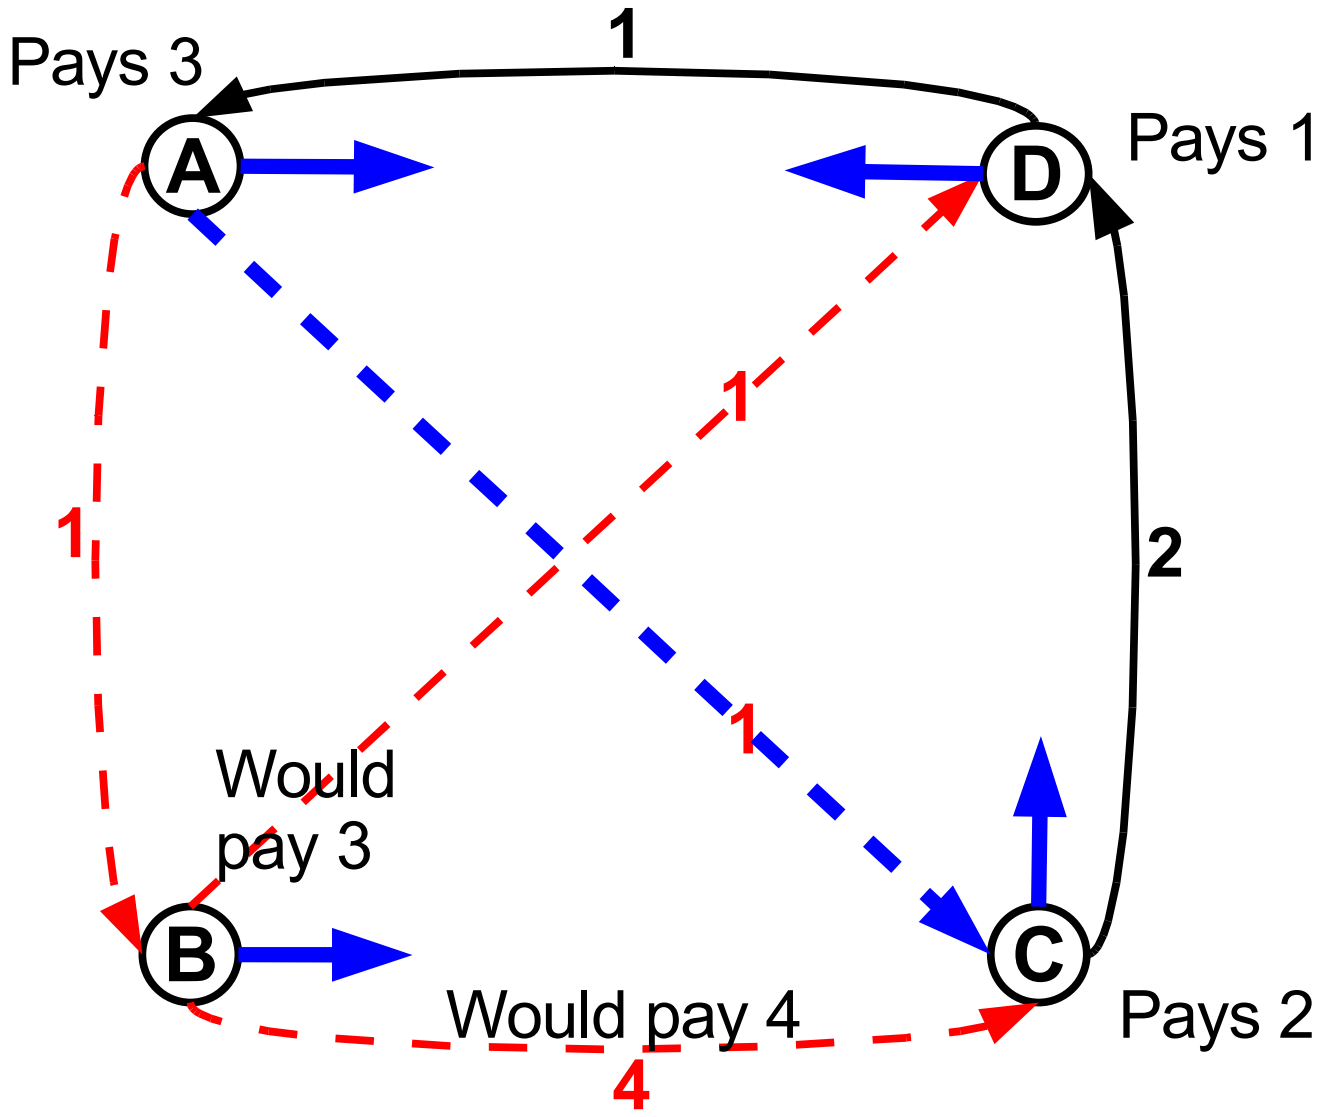

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

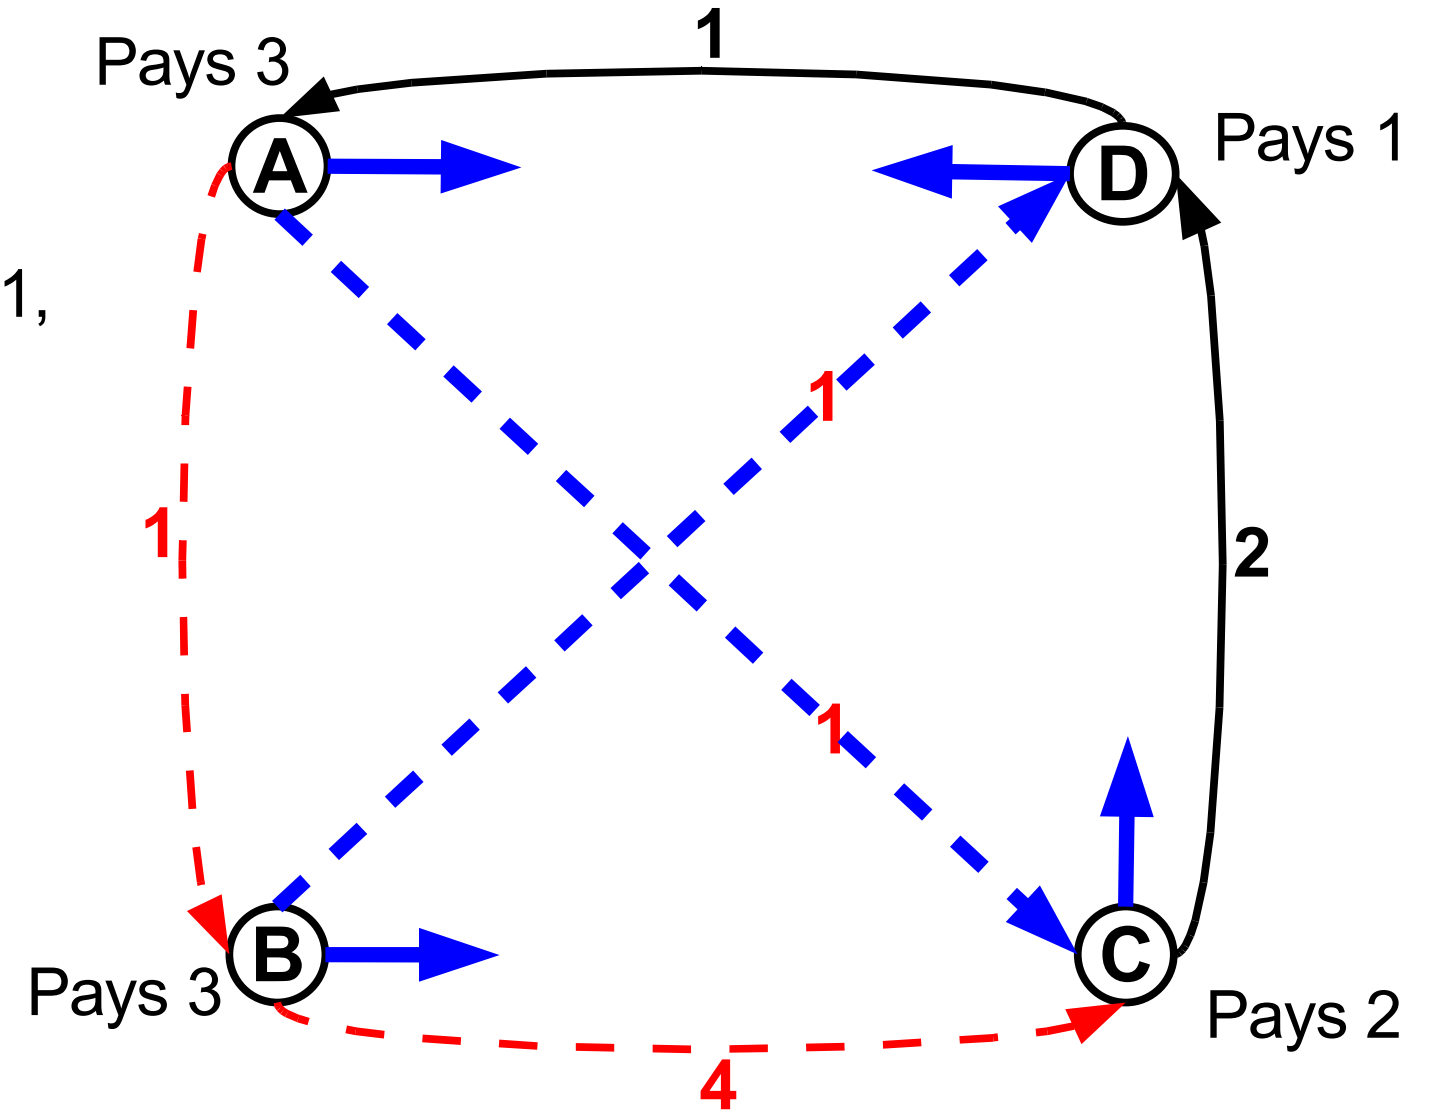

Sometimes no pure Nash equilibrium

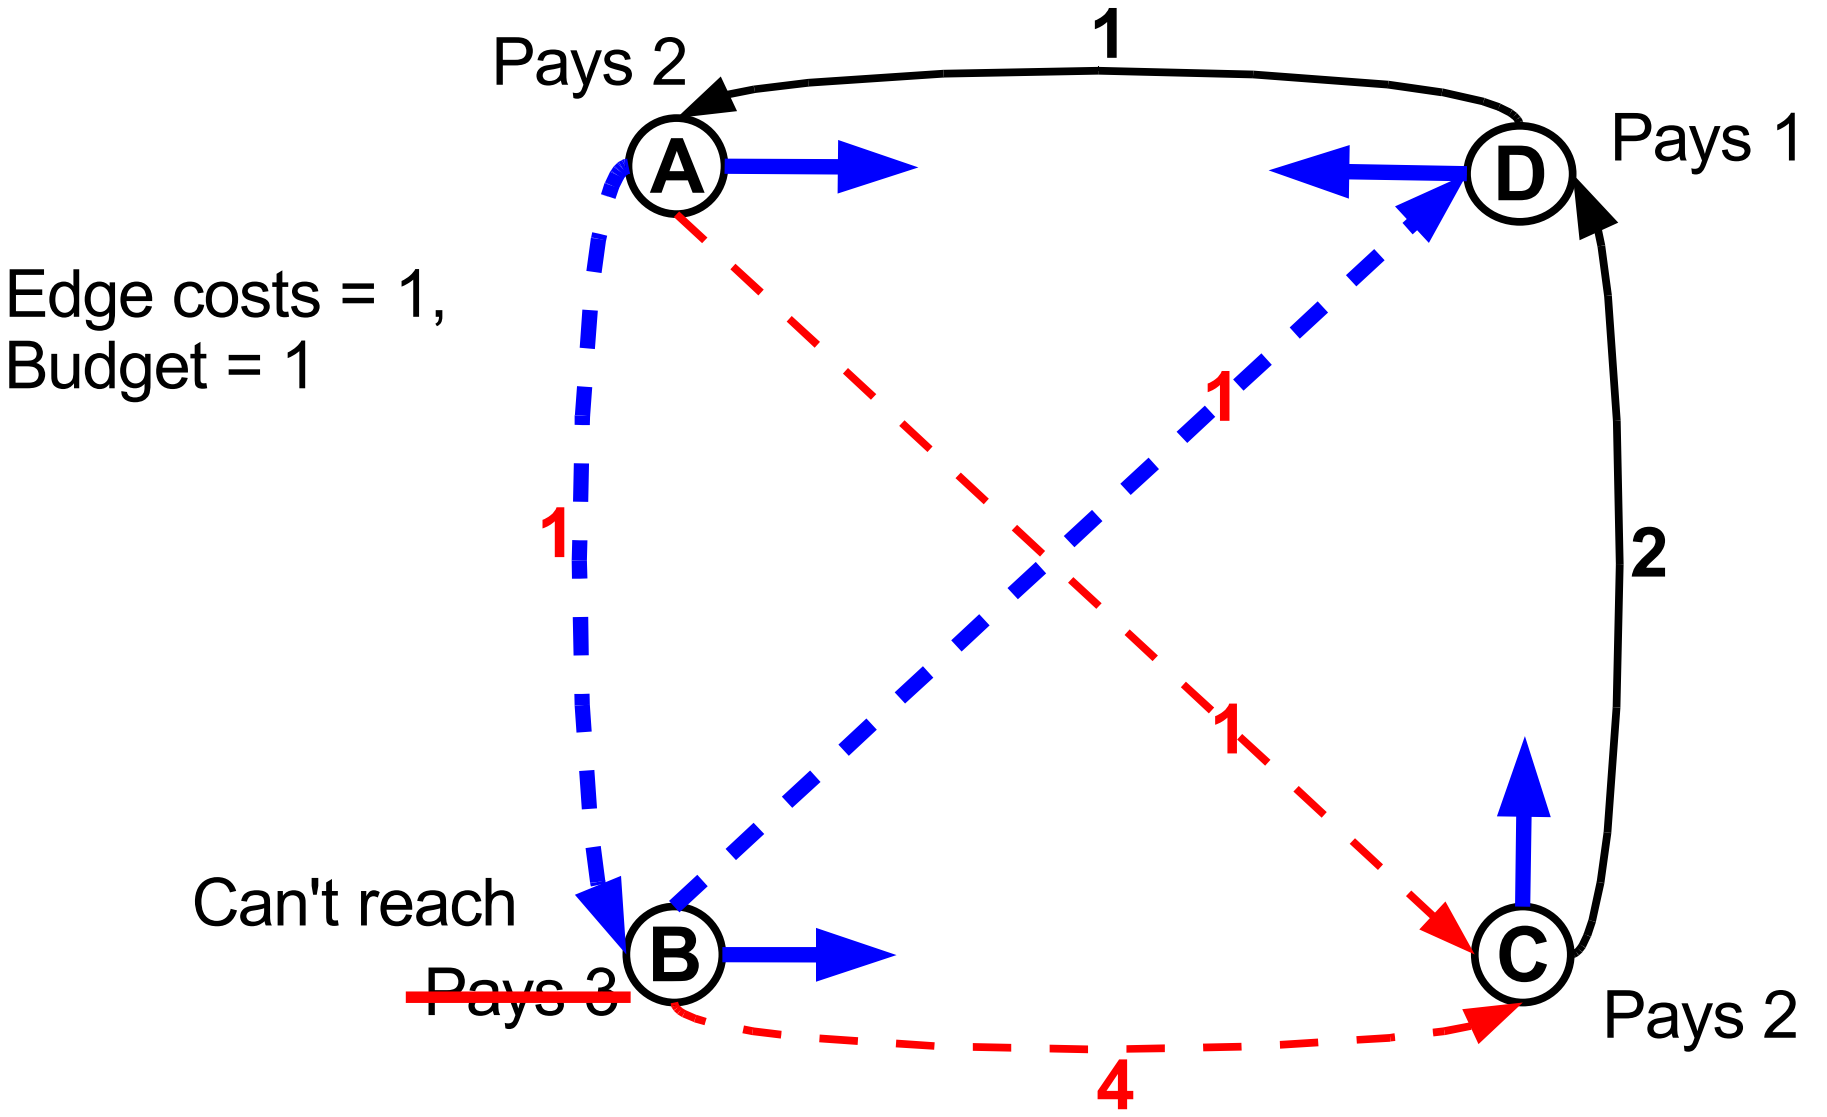

Sometimes no pure Nash equilibrium

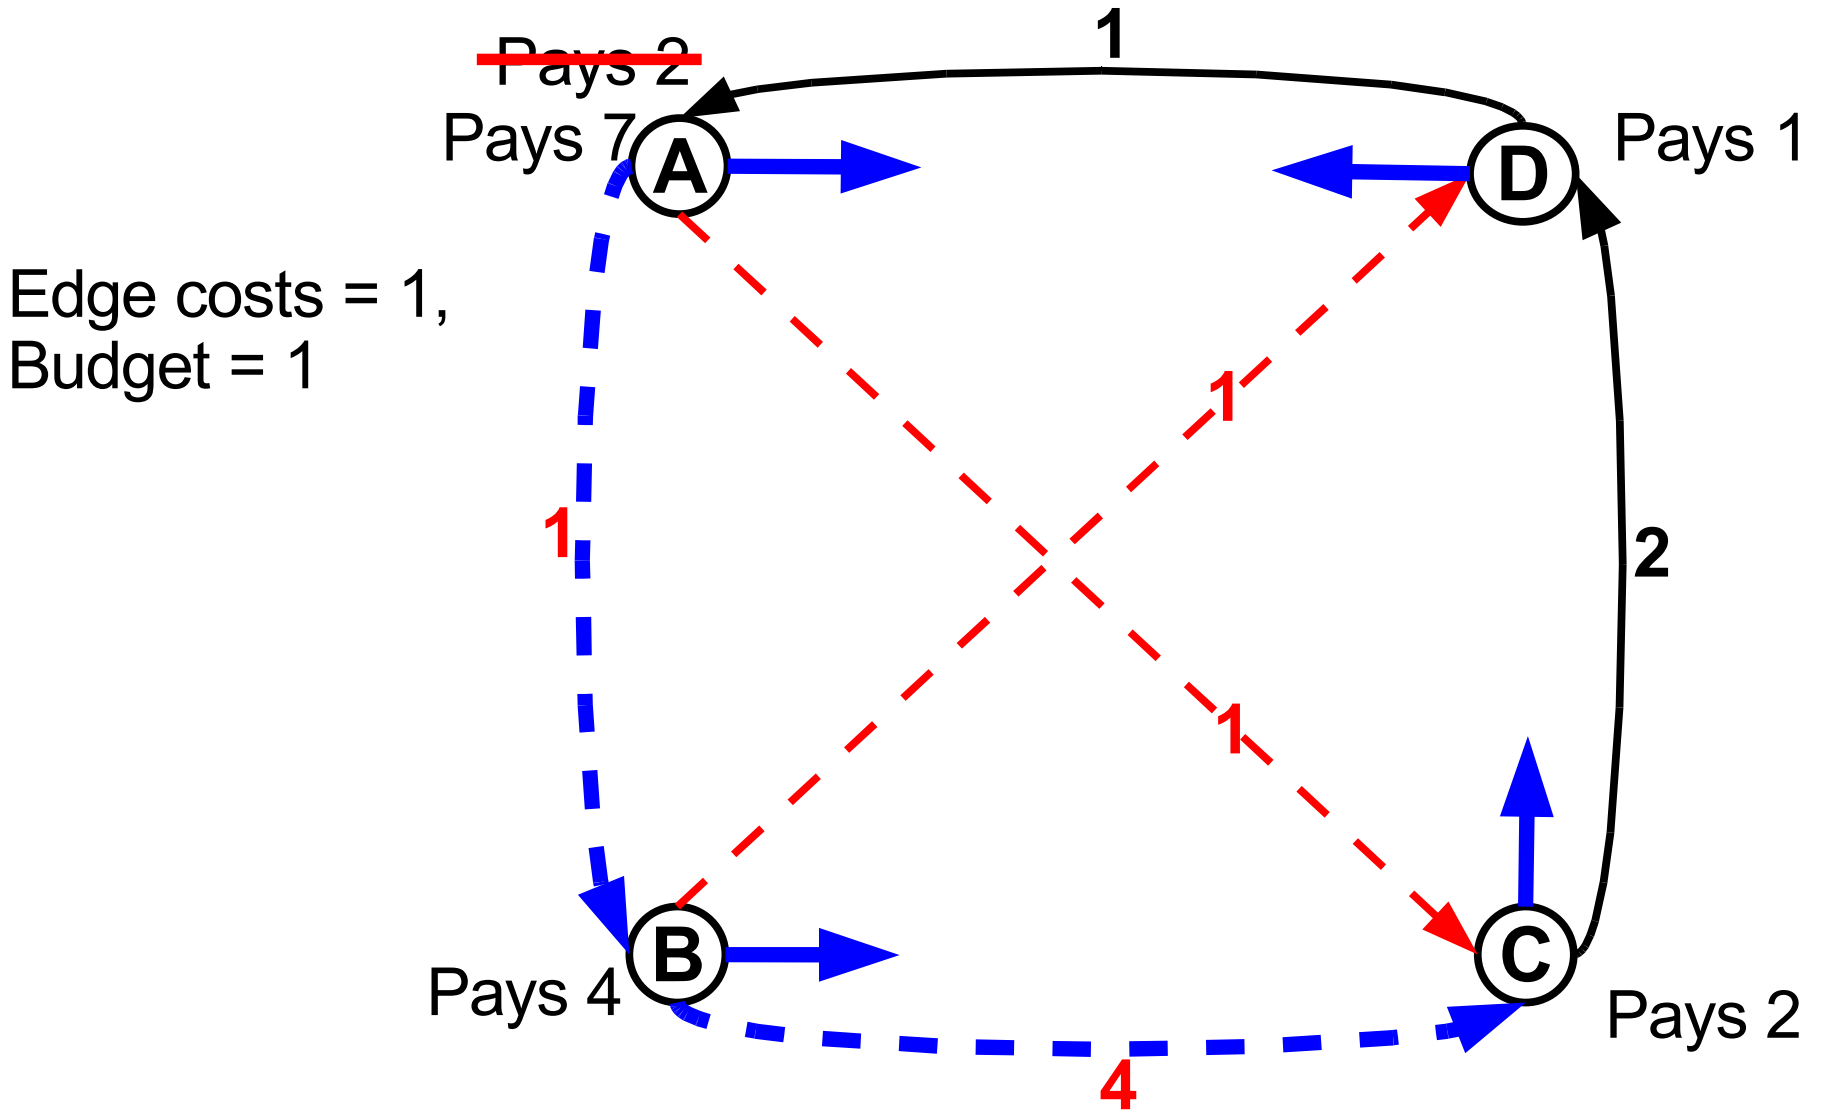

Sometimes no pure Nash equilibrium

Edge costs = 1,
Budget = 1

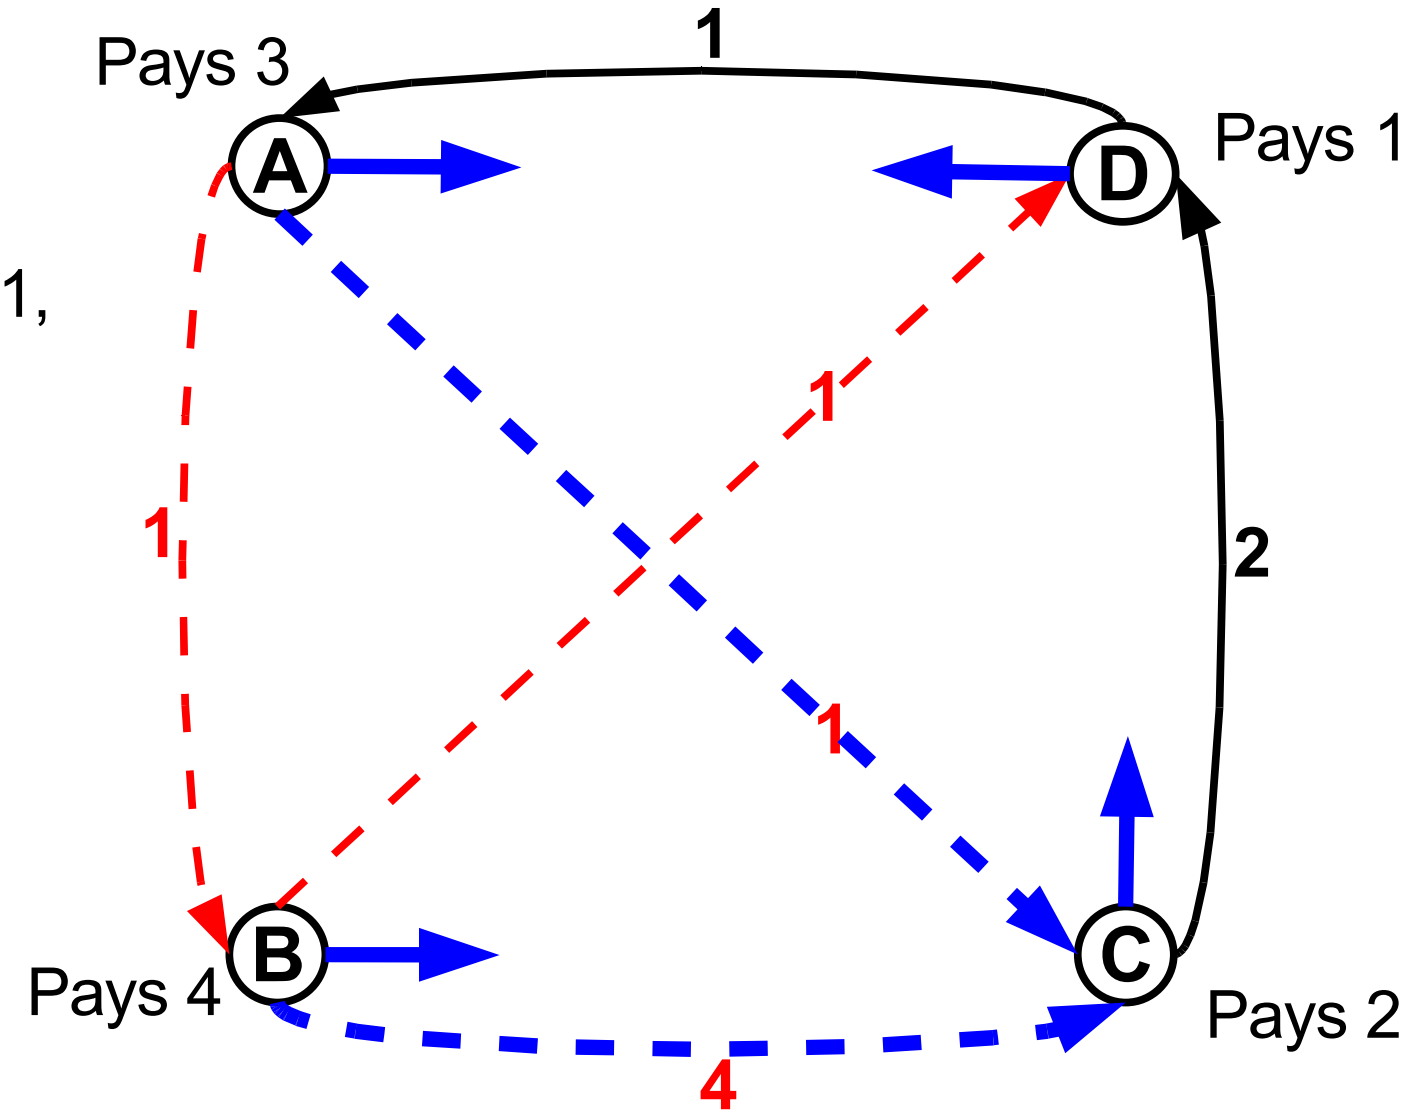

BBC Games

- NP hard to determine whether or not an equilibrium exists.
  - Even if only costs are non-uniform
  - Even if only affinities are non-uniform

BBC Games

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?

  - Stable or oscillating?

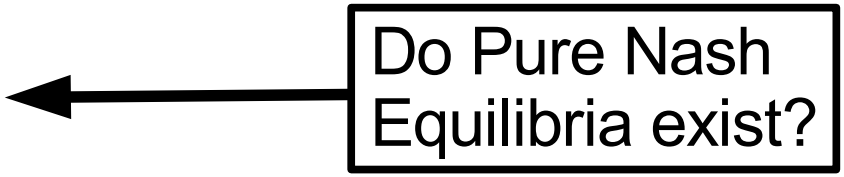

oscillating

  - Converging or cycling?

BBC Games

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?

  - Stable or oscillating?

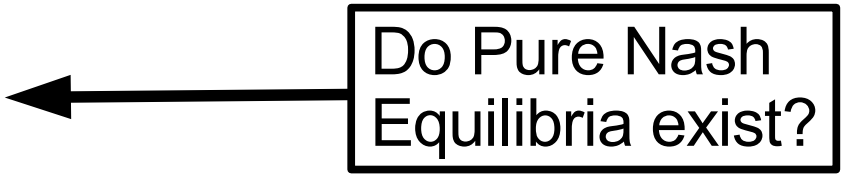

oscillating

  - Converging or cycling?

May cycle

Is there ever a stable solution?

- Including budget as part of the cost (make it depend heavily on degree)

Is there ever a stable solution?

- Including budget as part of the cost (make it depend heavily on degree)
- What if we allow fractional edges?
  - Stable solution always exists!

Fractional BBC Game

- Number of nodes
- Affinity for each directed pair of nodes
- Link cost for each directed pair of nodes
- Budget of allowed link cost per node, *k(v)*
- Length metric from the perspective of each node
- Each node *v* spends ≤ *k(v)* on fractions of links to minimize

$$\sum_{\text{other nodes}} (\text{affinity} * \text{cost of min cost 1 unit flow})$$

Fractional BBC Games (budget=1)

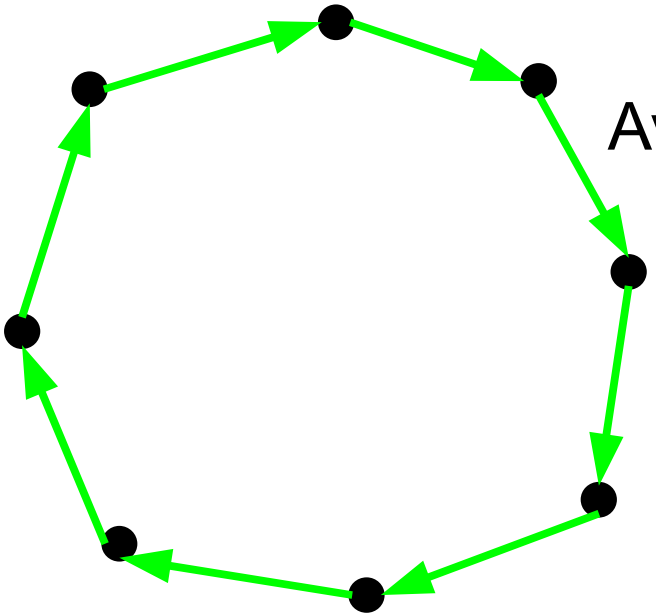

Each edge weight = 1
Average distance per node = (n-1) / 2

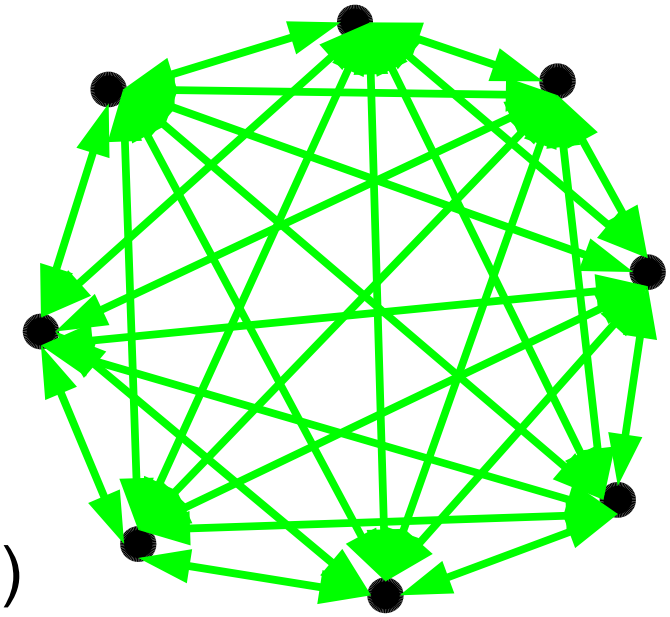

Each edge weight = 1/(n-1)
Average distance per node = 2 - 1/(n-1)

Fractional BBC Games

- Stable solution always exists
- Rational stable solution always exists

- More about this later....

Border Gateway Protocol

- Rehkter, Li. A Border Gateway Protocol (BGP version 4). RFC 1771, 1995.

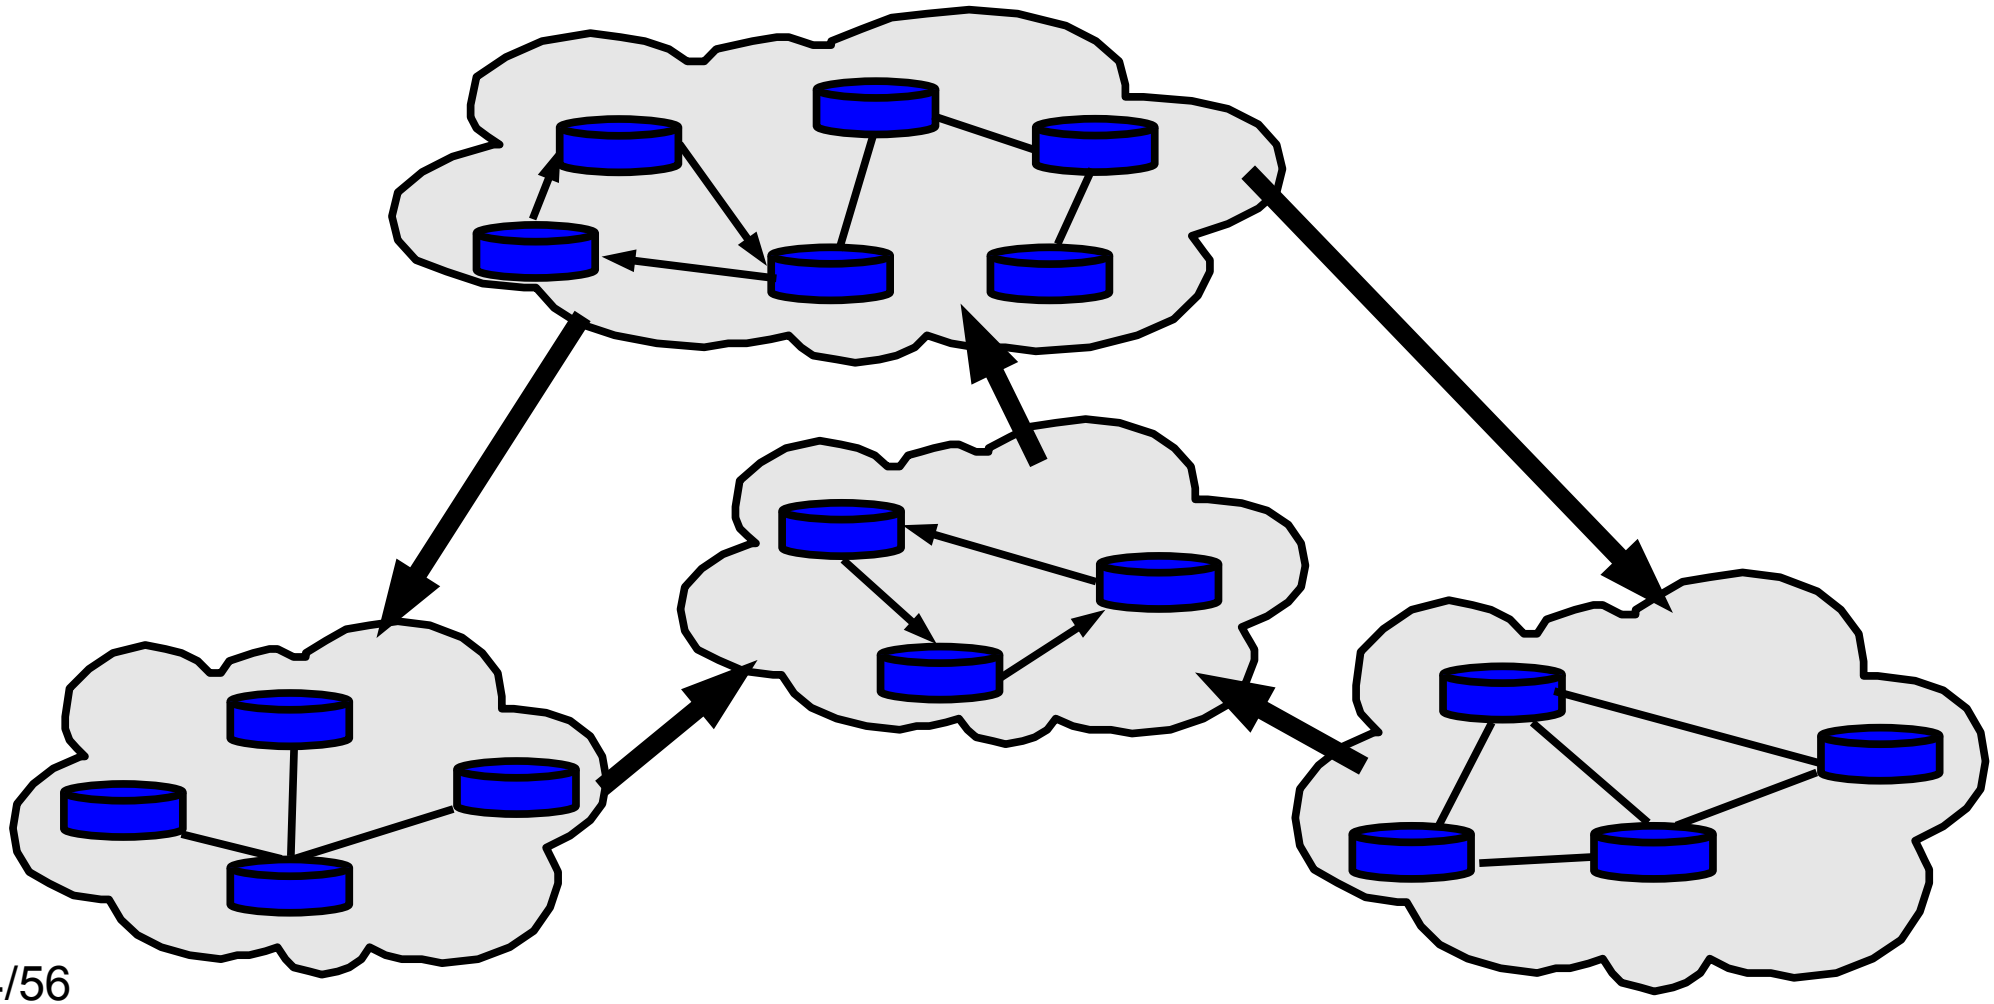

Border Gateway Protocol

- Varadhan, Govindan, and Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report USC CS TR 96-631, Dept of Computer Science, USC, 1996.

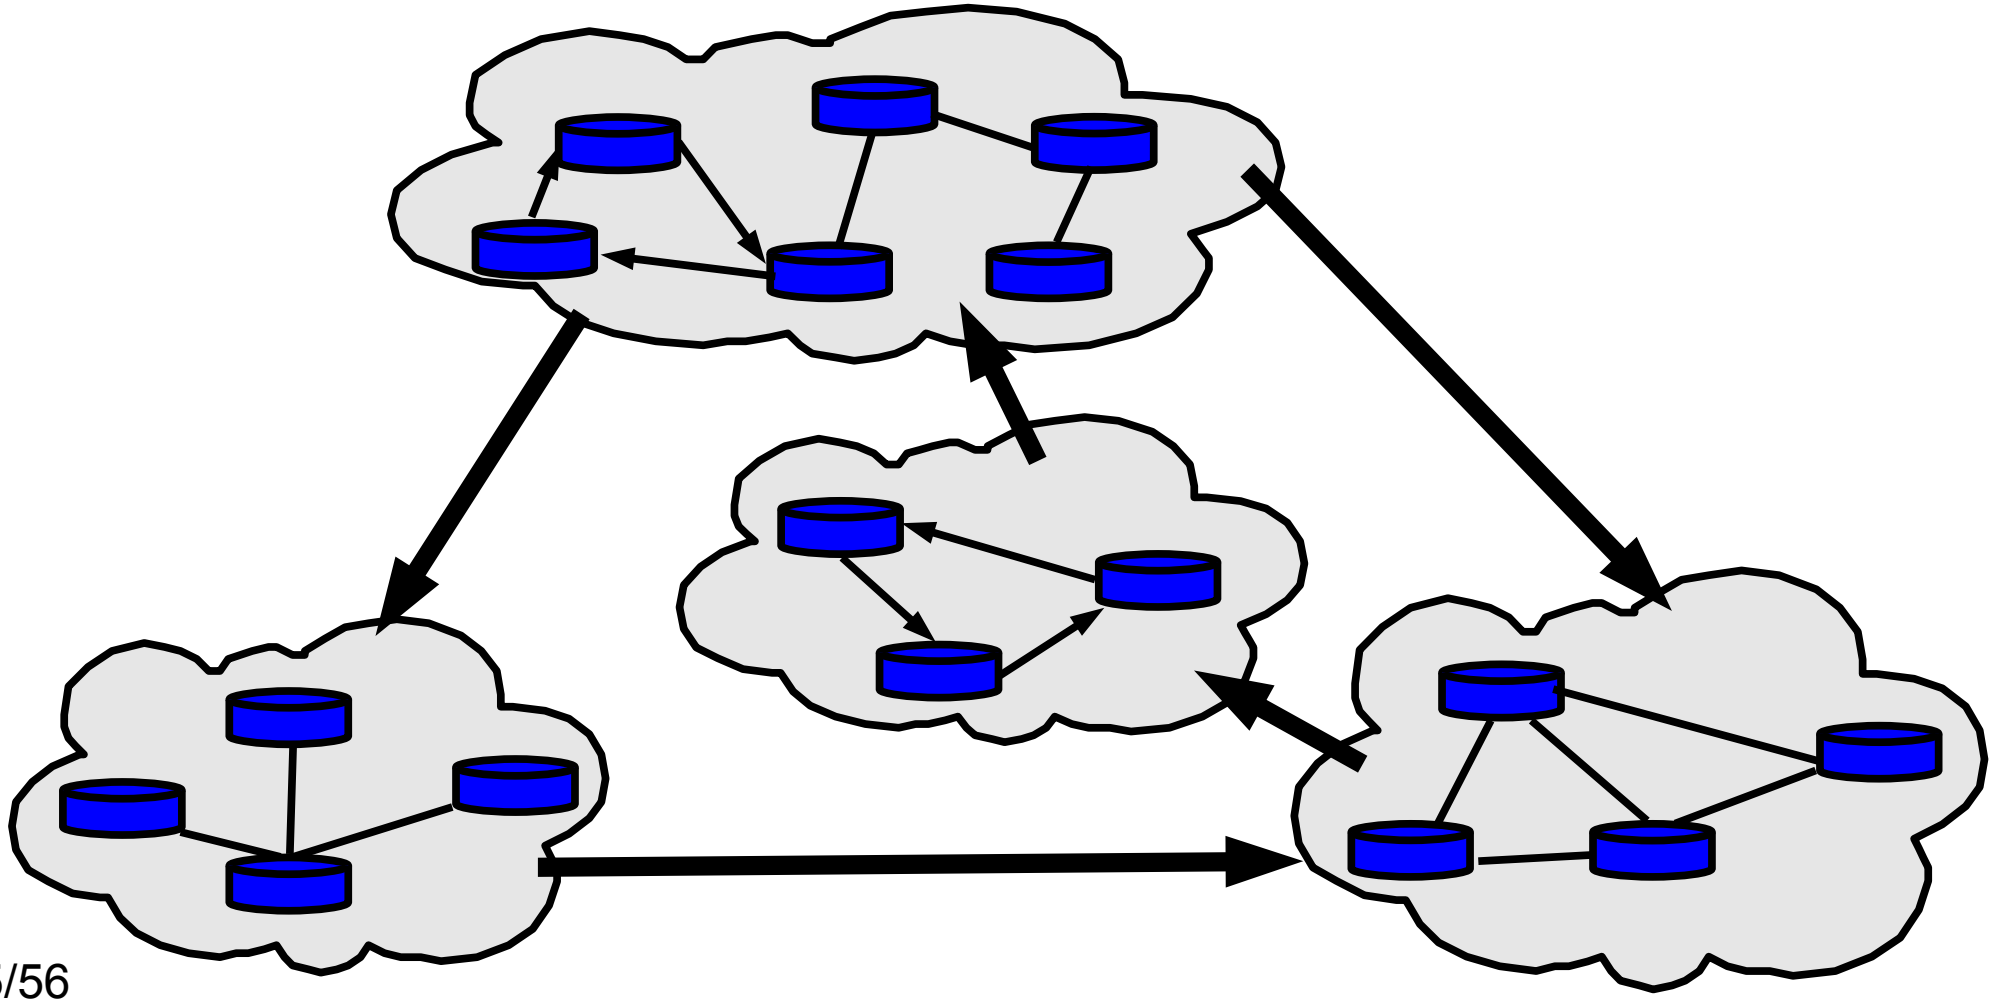

Border Gateway Protocol

- Varadhan, Govindan, and Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report USC CS TR 96-631, Dept of Computer Science, USC, 1996.

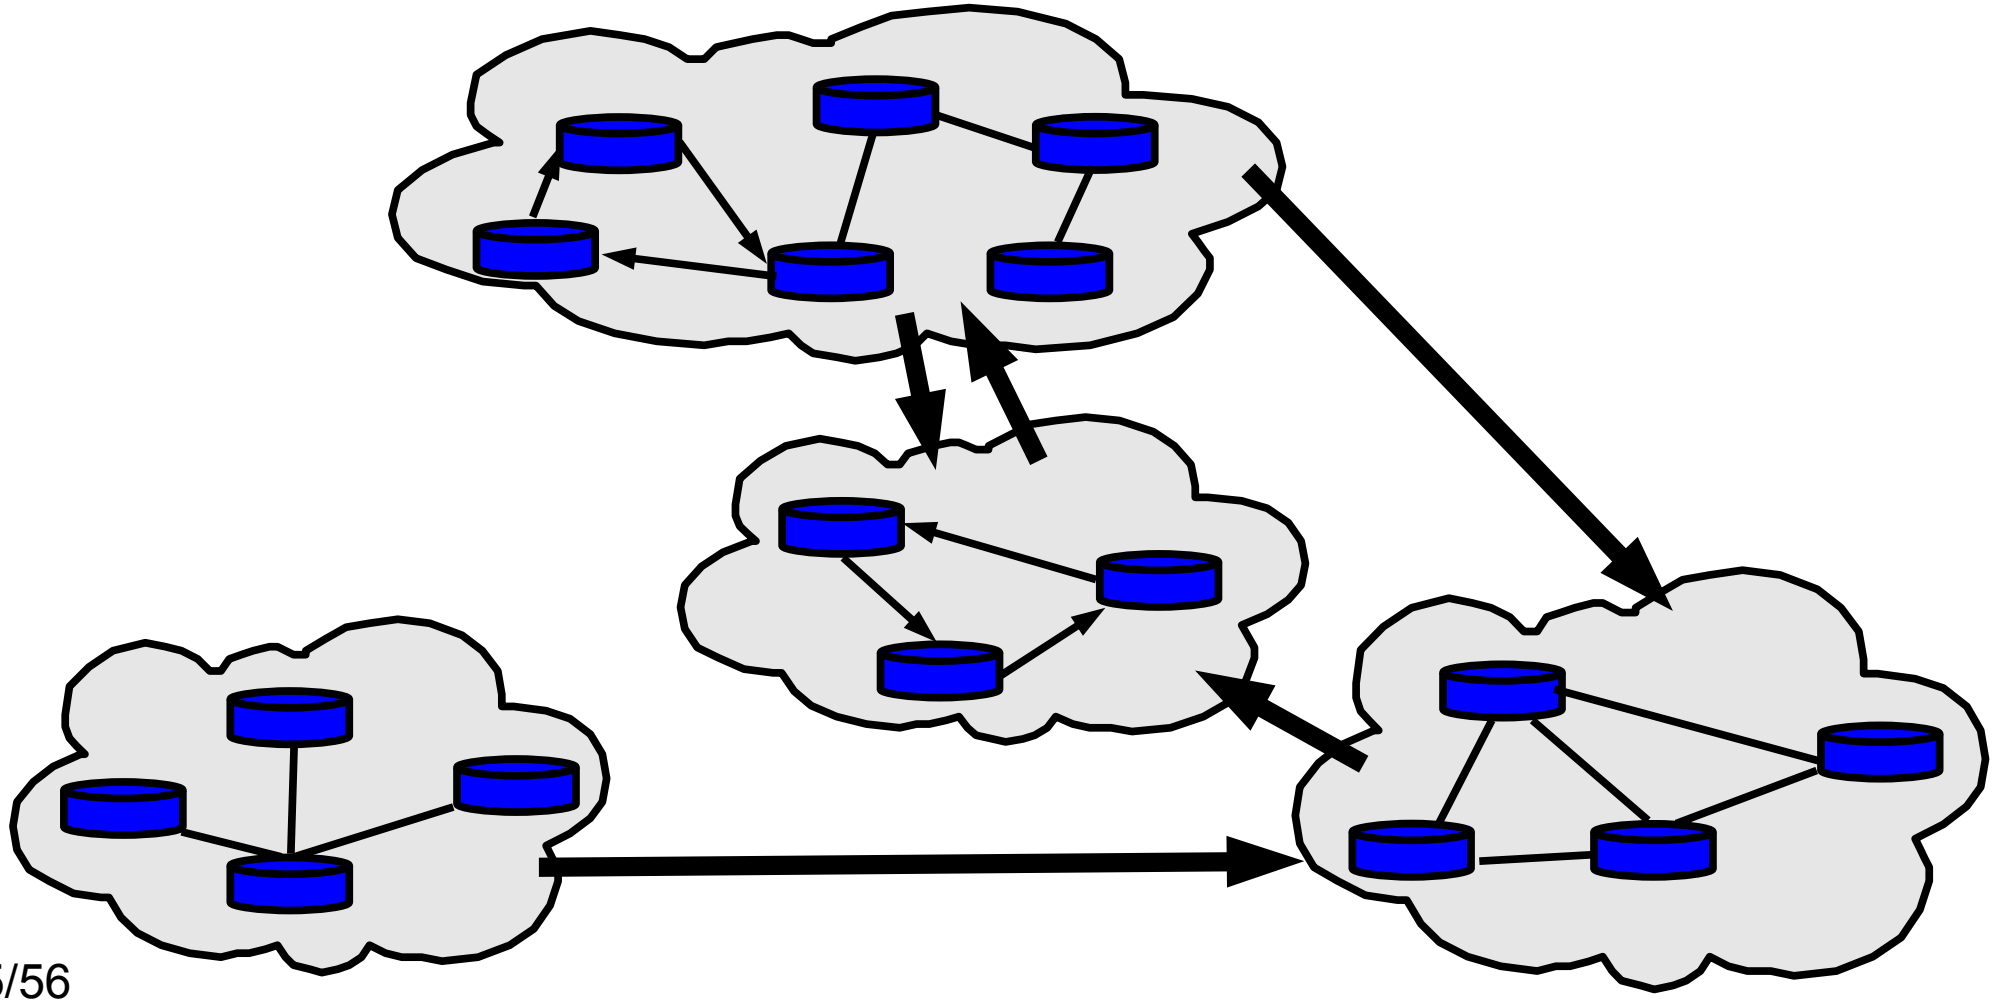

Border Gateway Protocol

- Varadhan, Govindan, and Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report USC CS TR 96-631, Dept of Computer Science, USC, 1996.

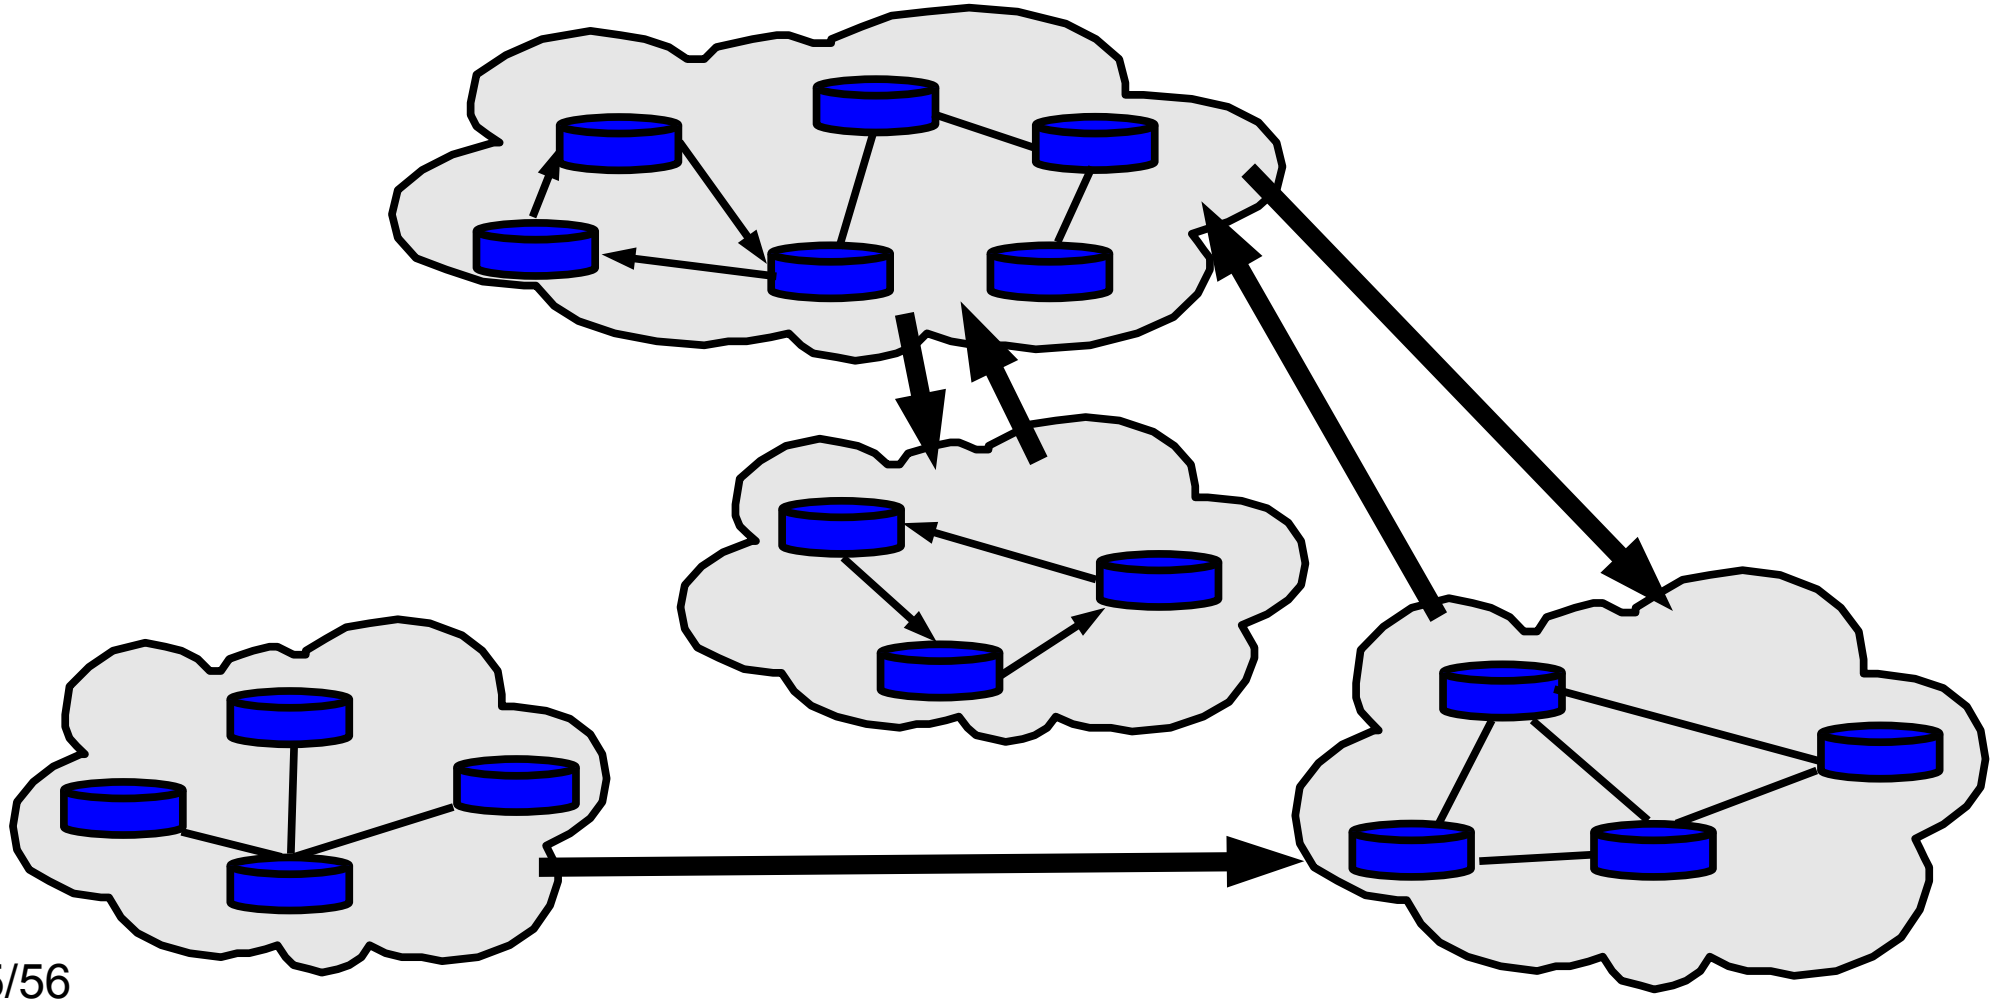

Border Gateway Protocol

- Varadhan, Govindan, and Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report USC CS TR 96-631, Dept of Computer Science, USC, 1996.

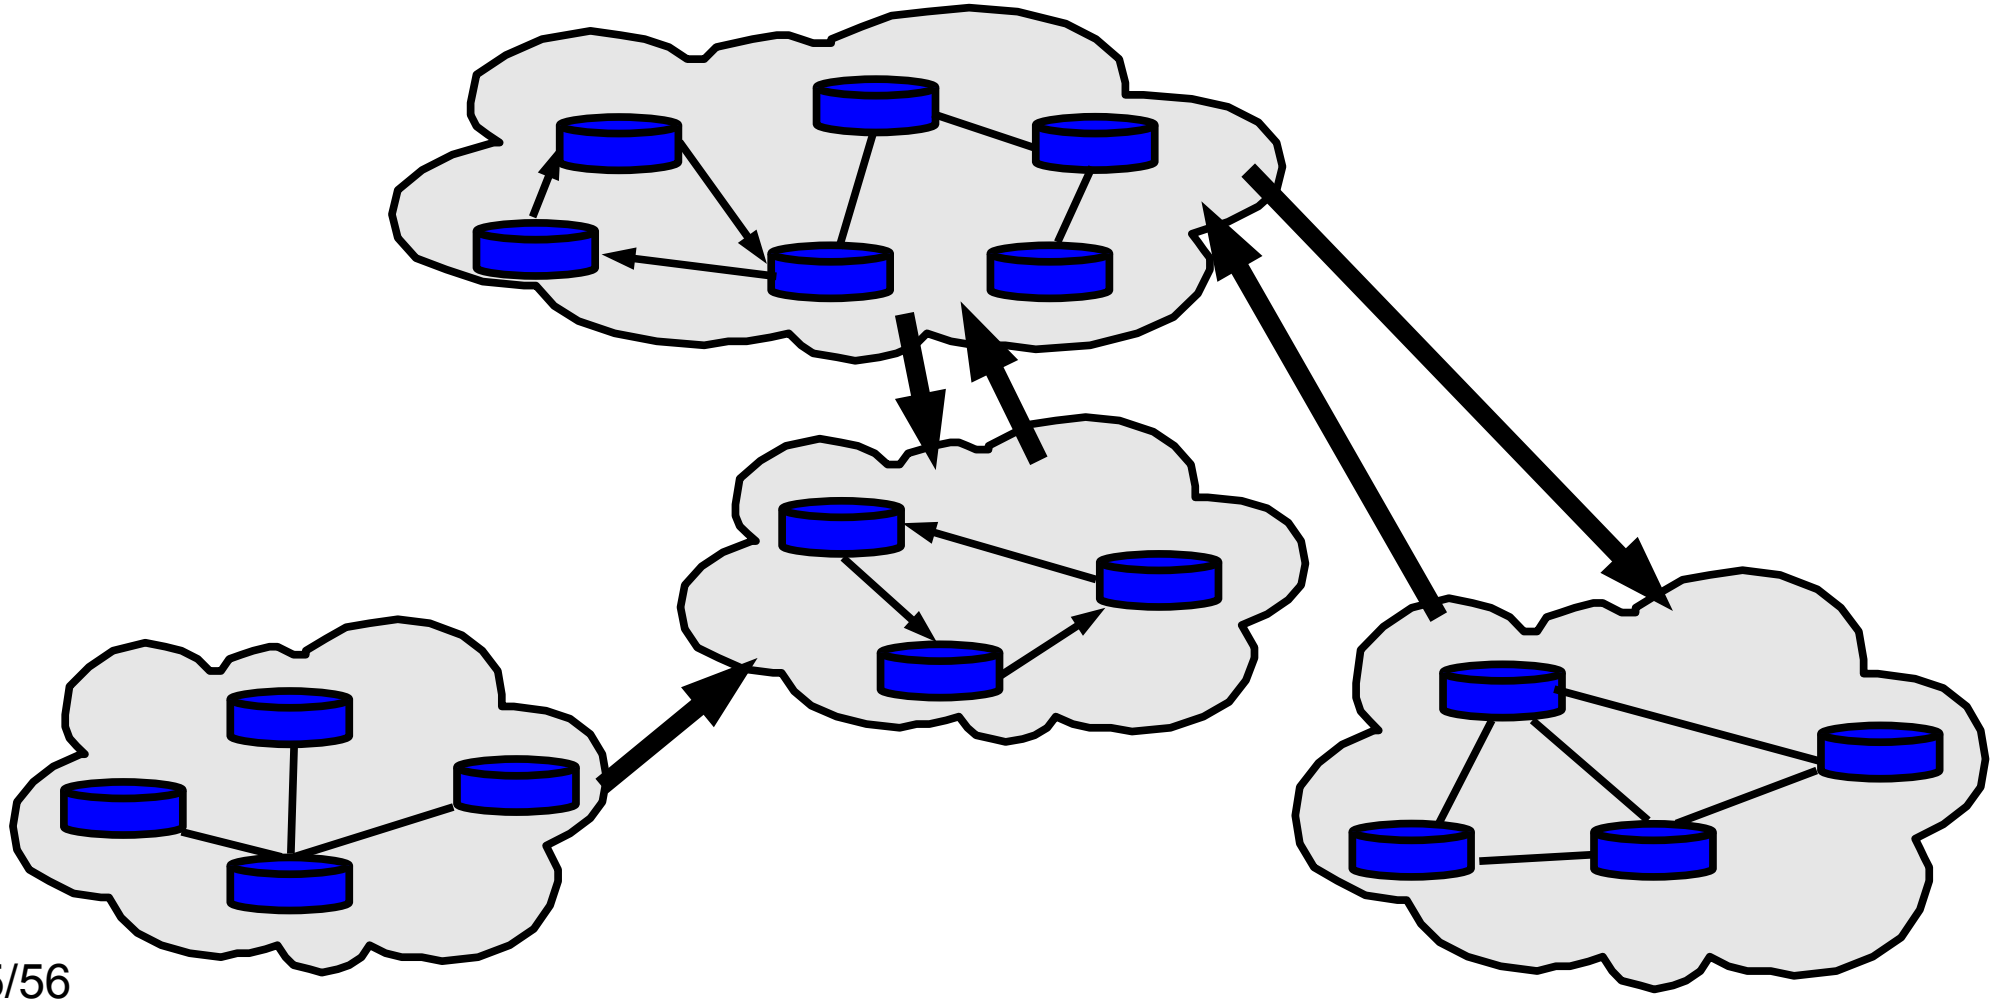

Border Gateway Protocol

- Varadhan, Govindan, and Estrin. Persistent Route Oscillations in Inter-Domain Routing. Technical Report USC CS TR 96-631, Dept of Computer Science, USC, 1996.

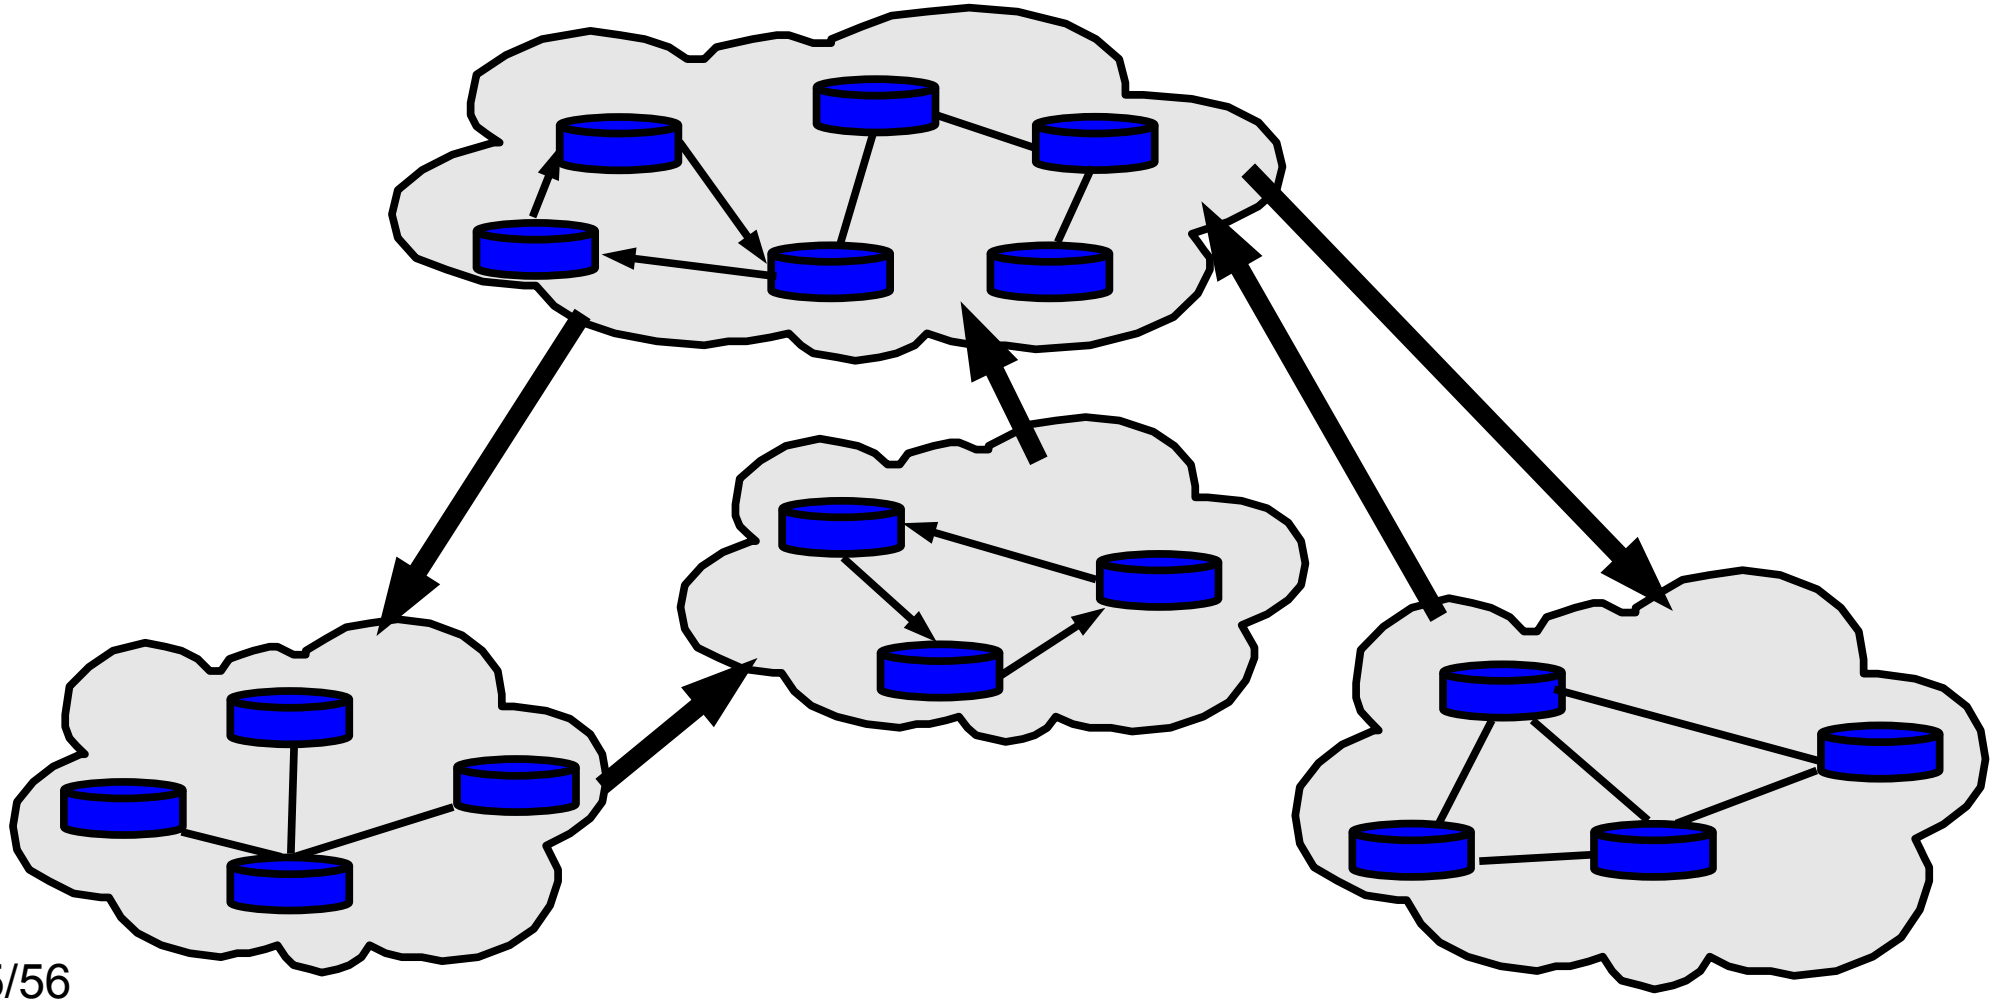

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

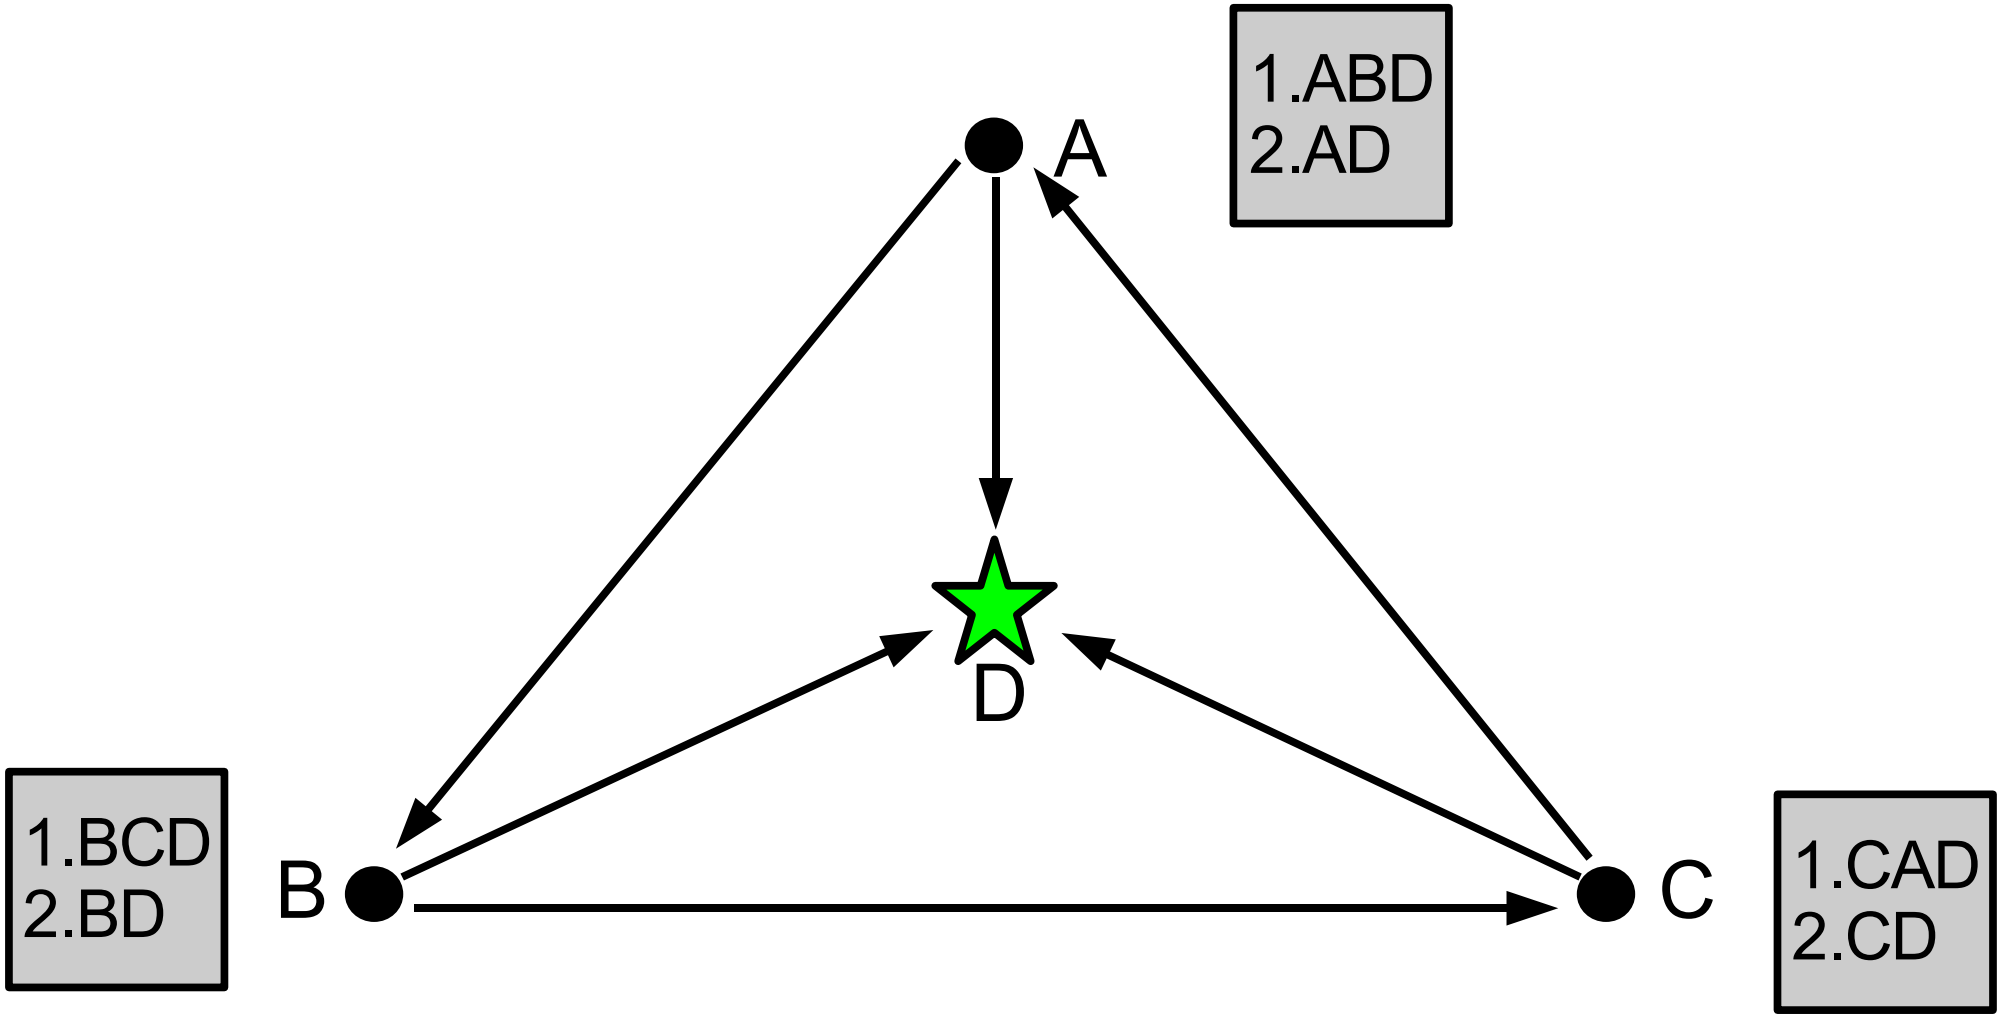

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

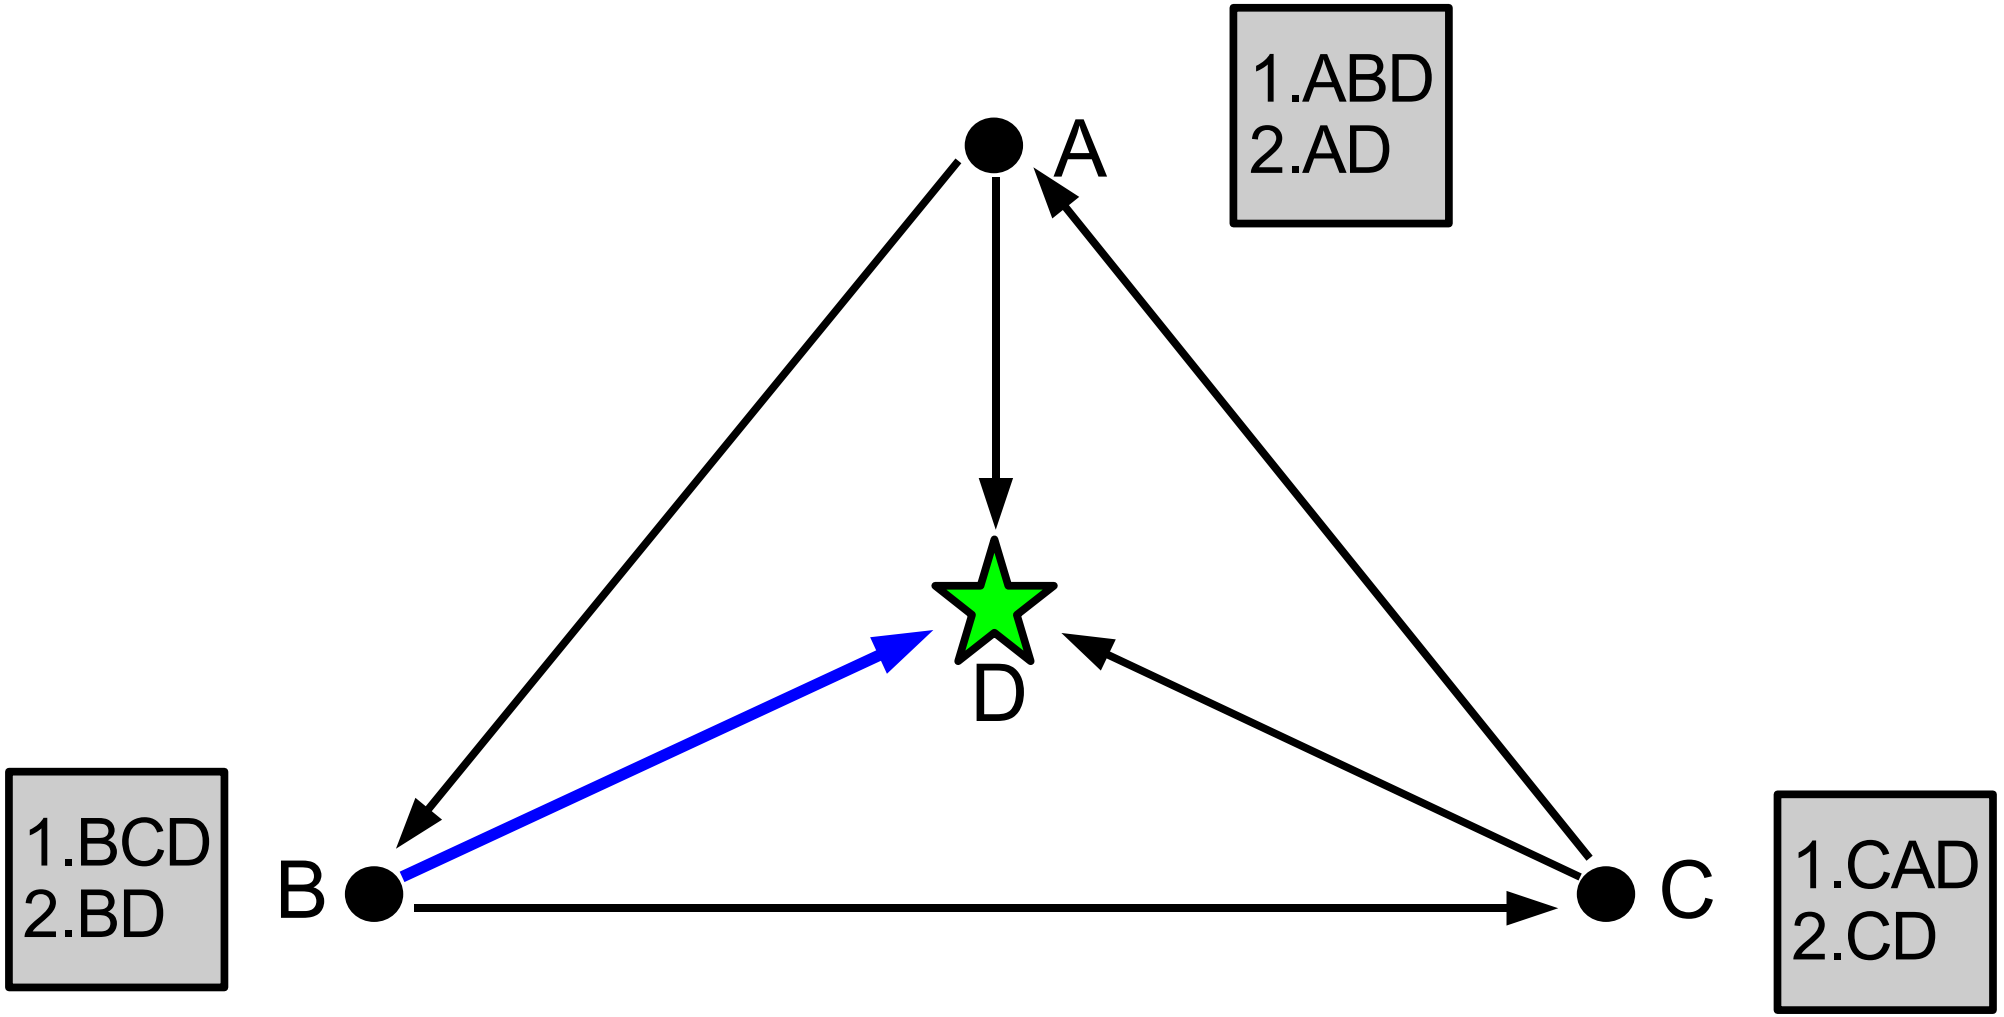

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

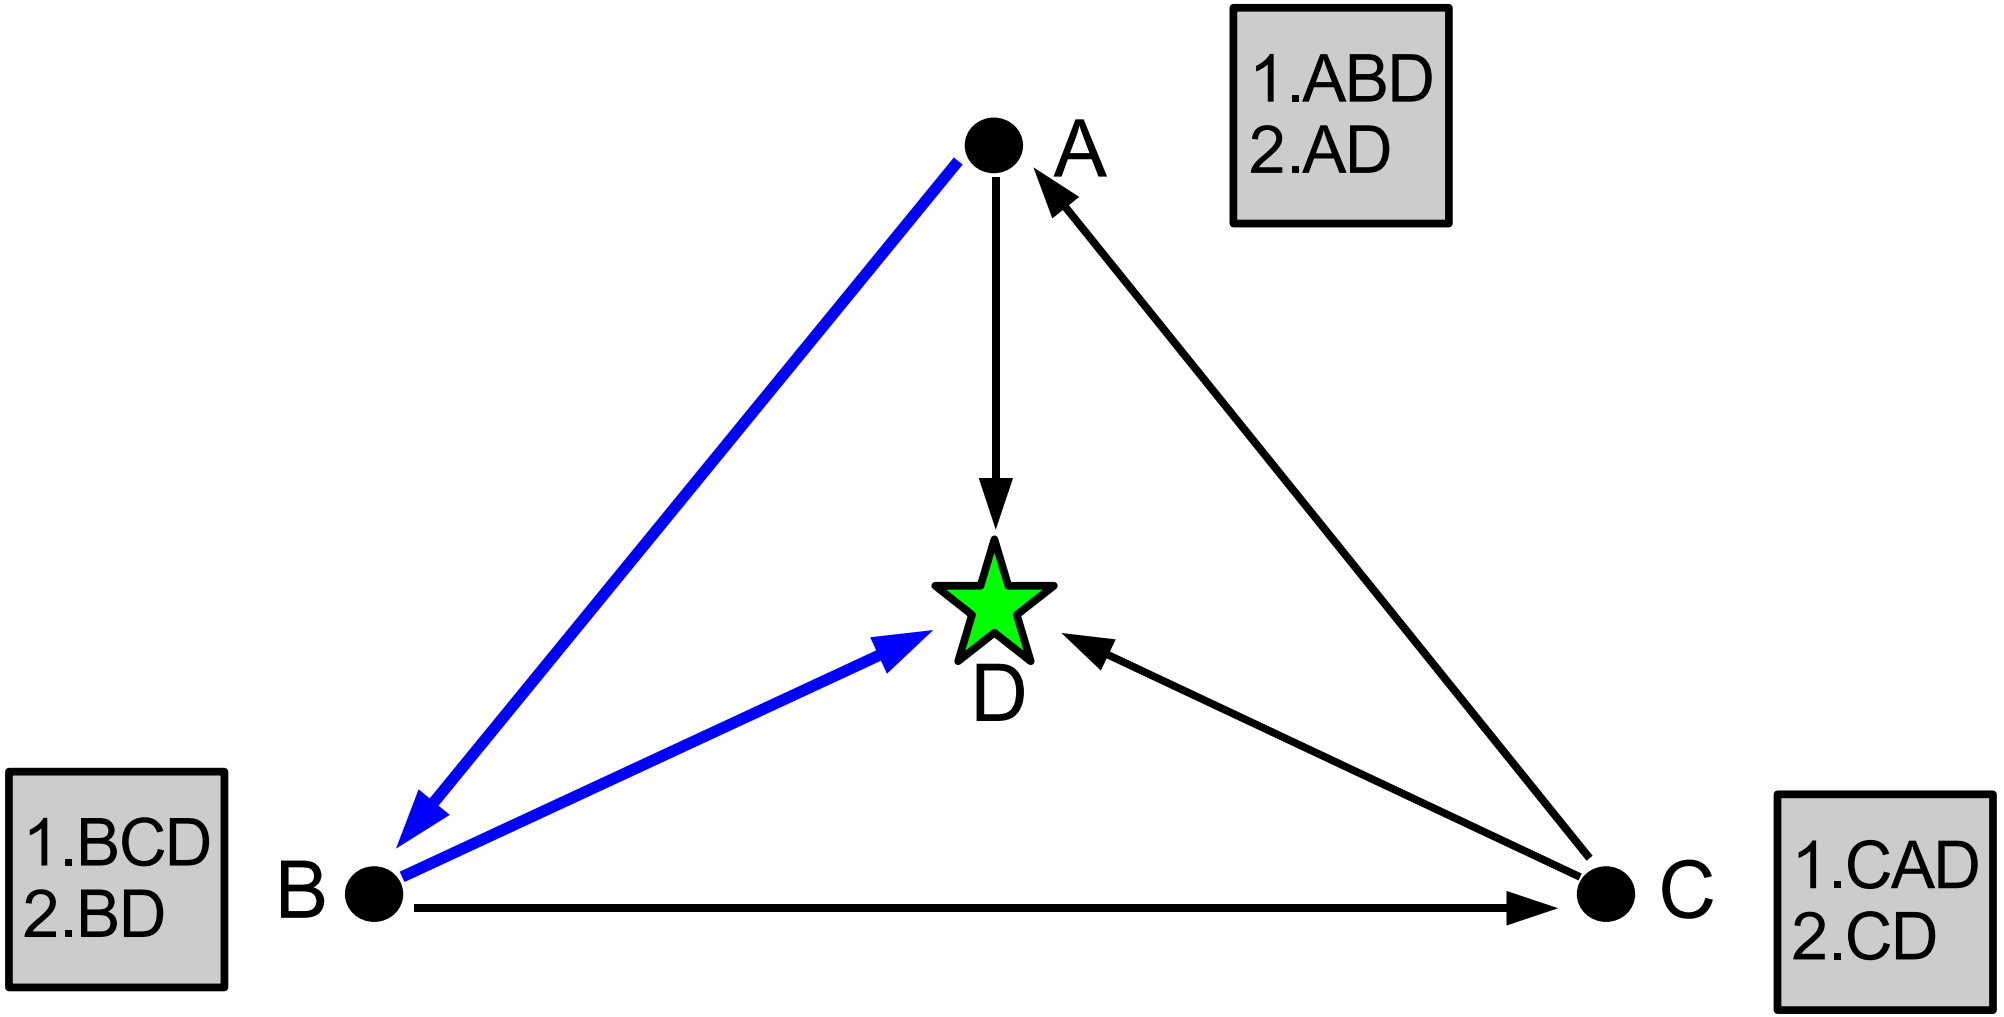

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

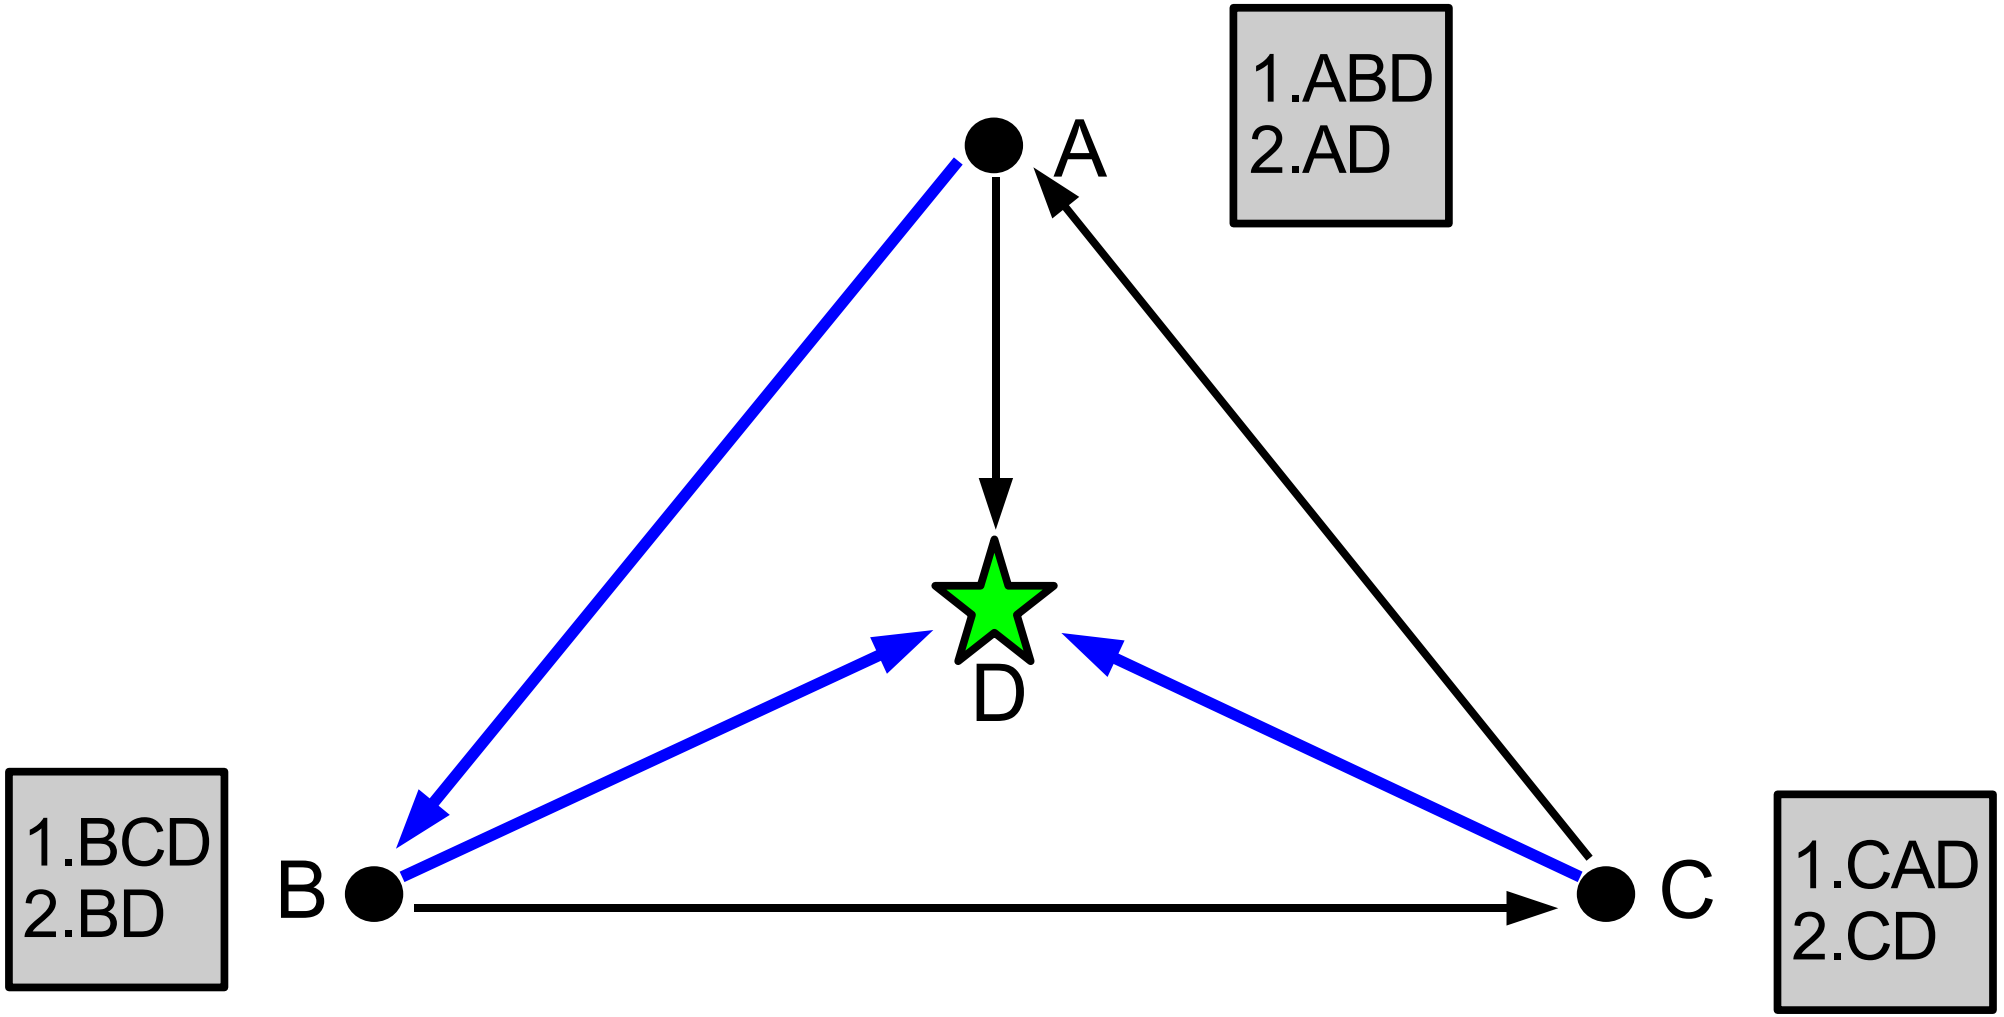

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

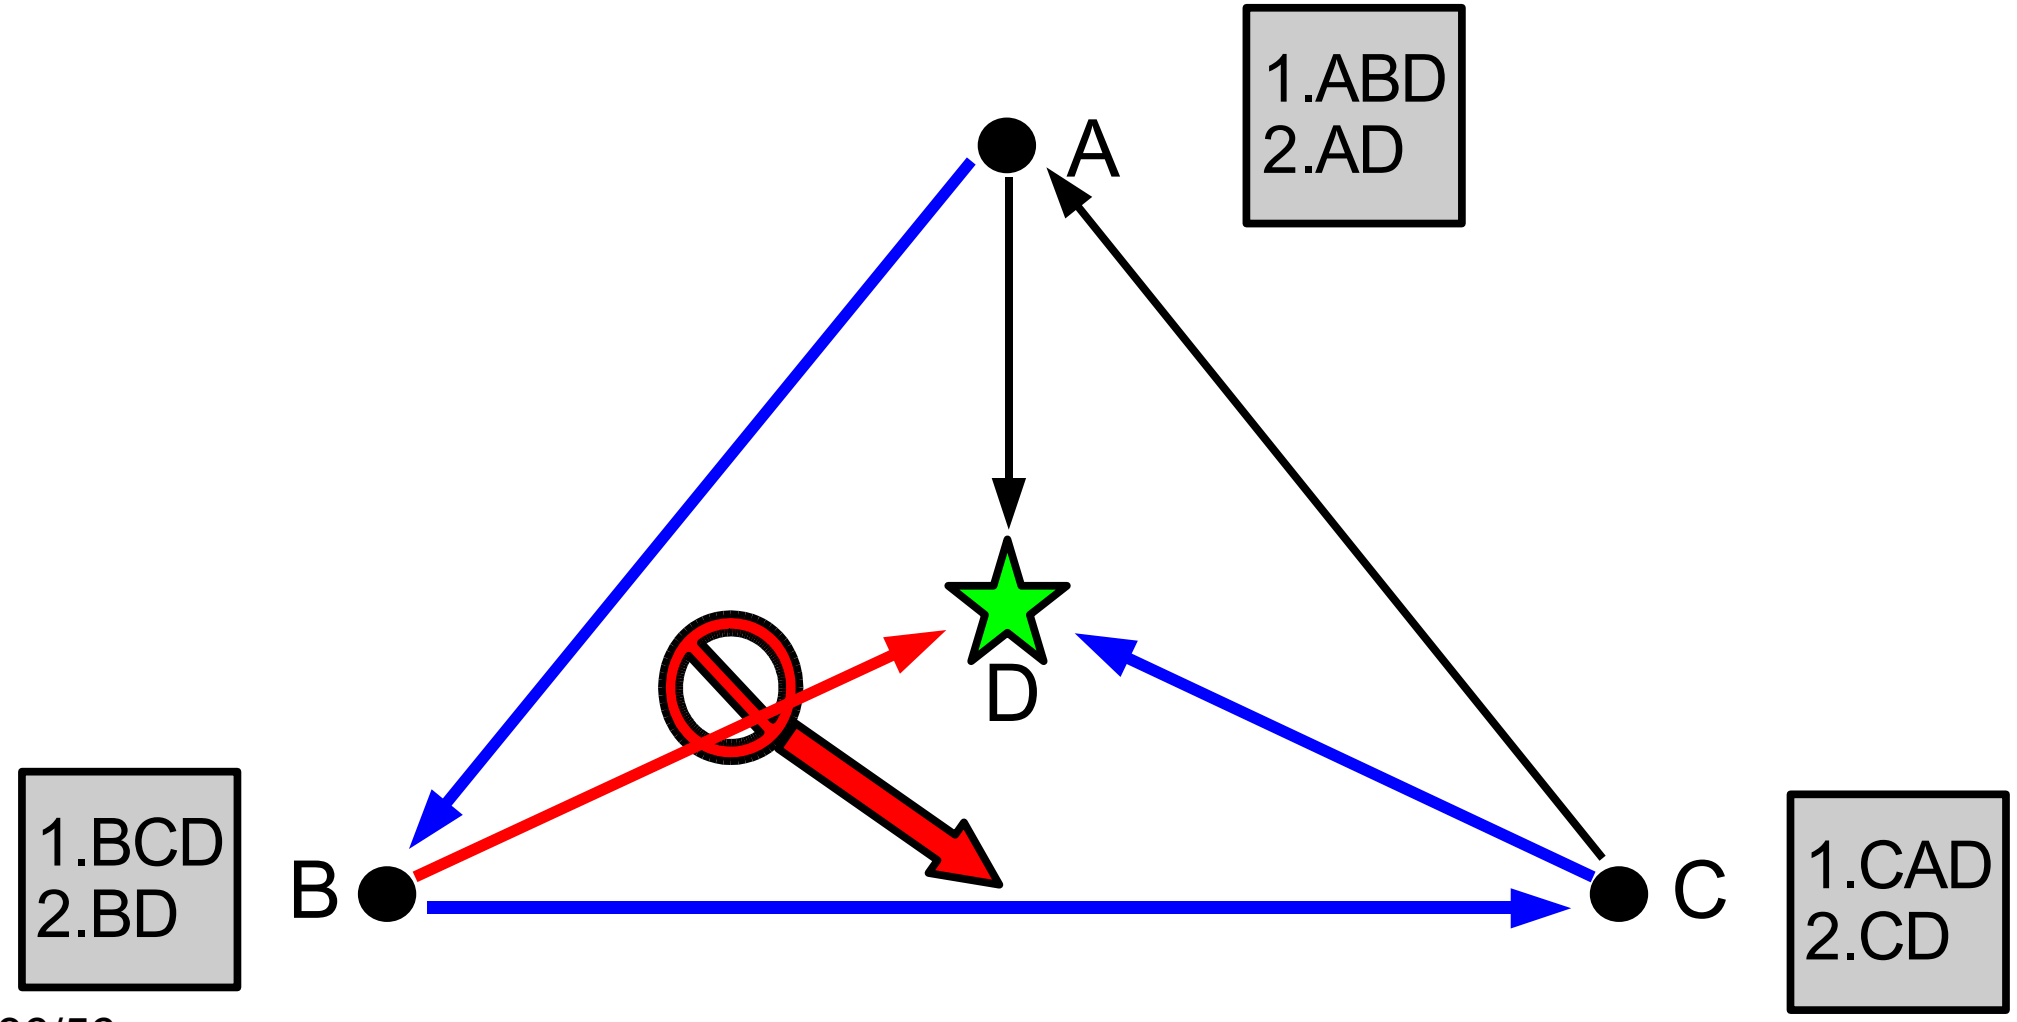

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

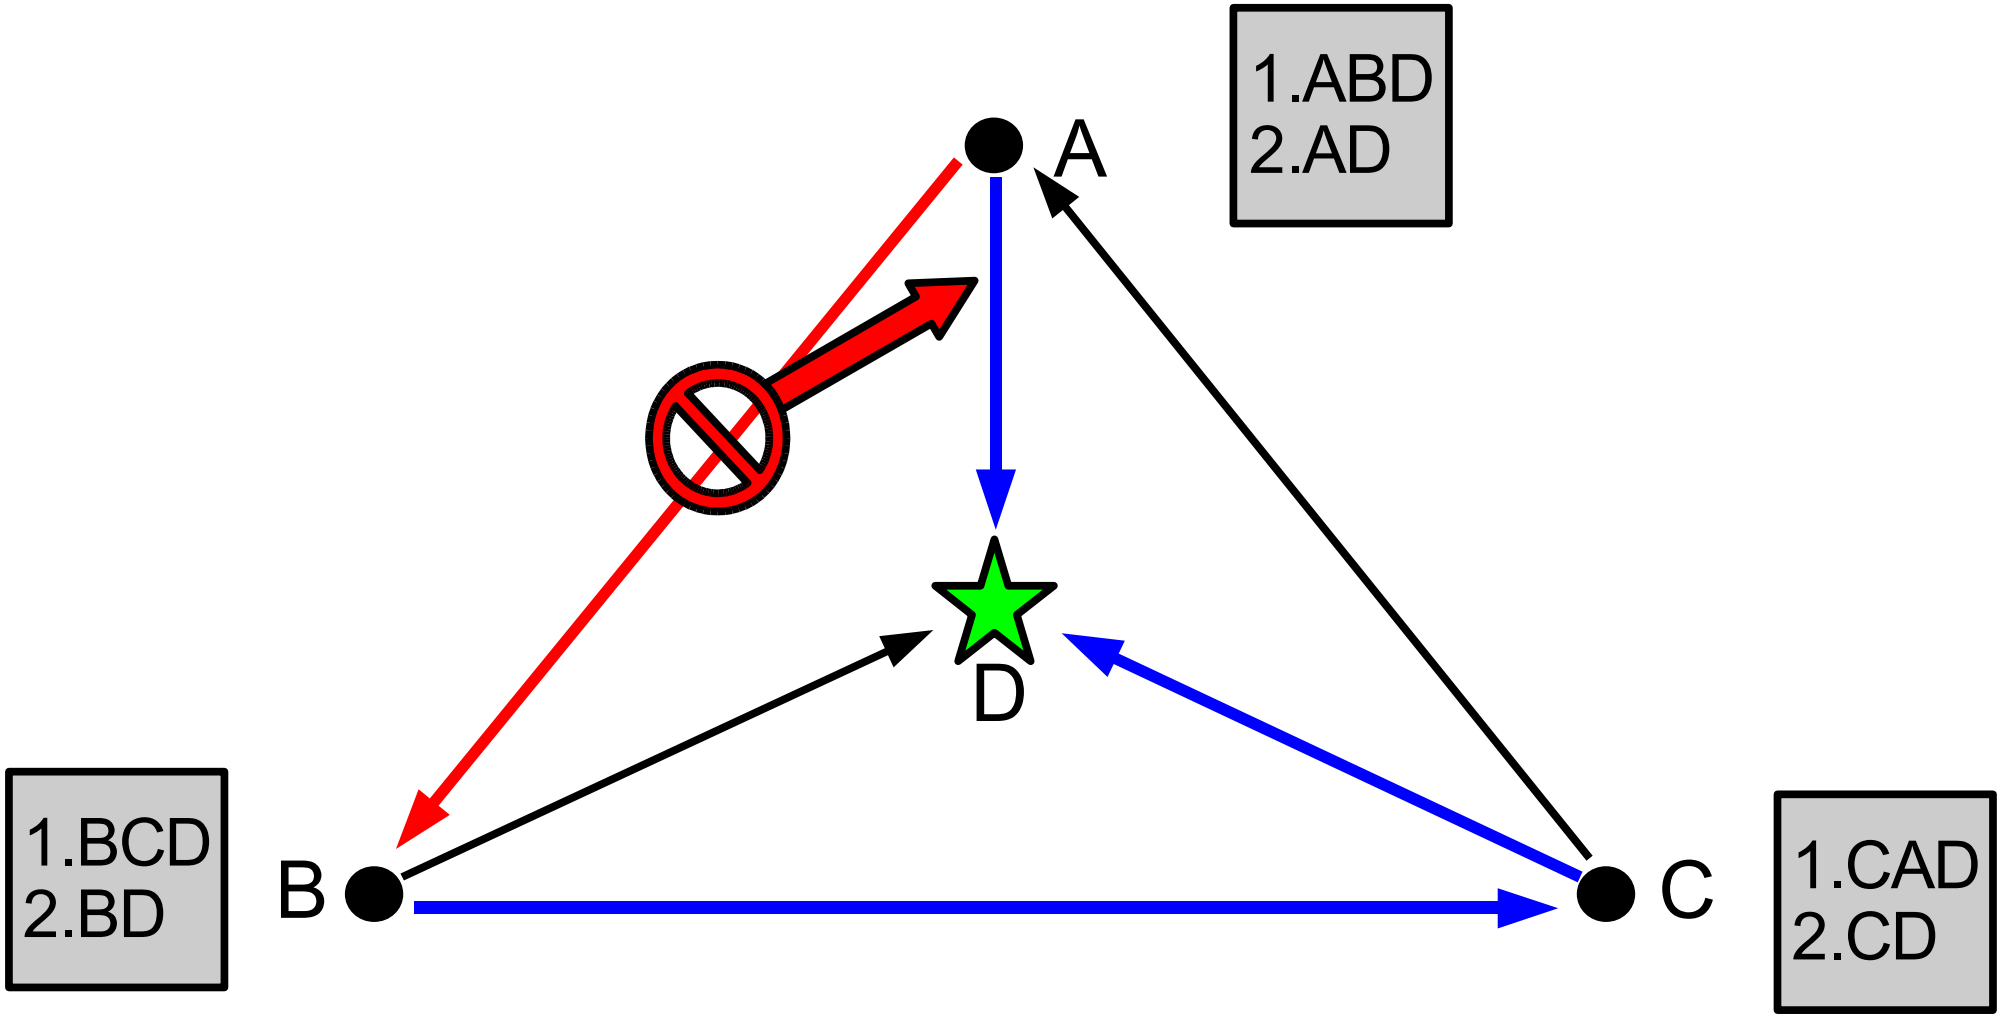

Stable Paths Problem

- Griffin, Shepherd, and Wilfong. The stable paths problem and interdomain routing. Transactions on Networking, 2002.

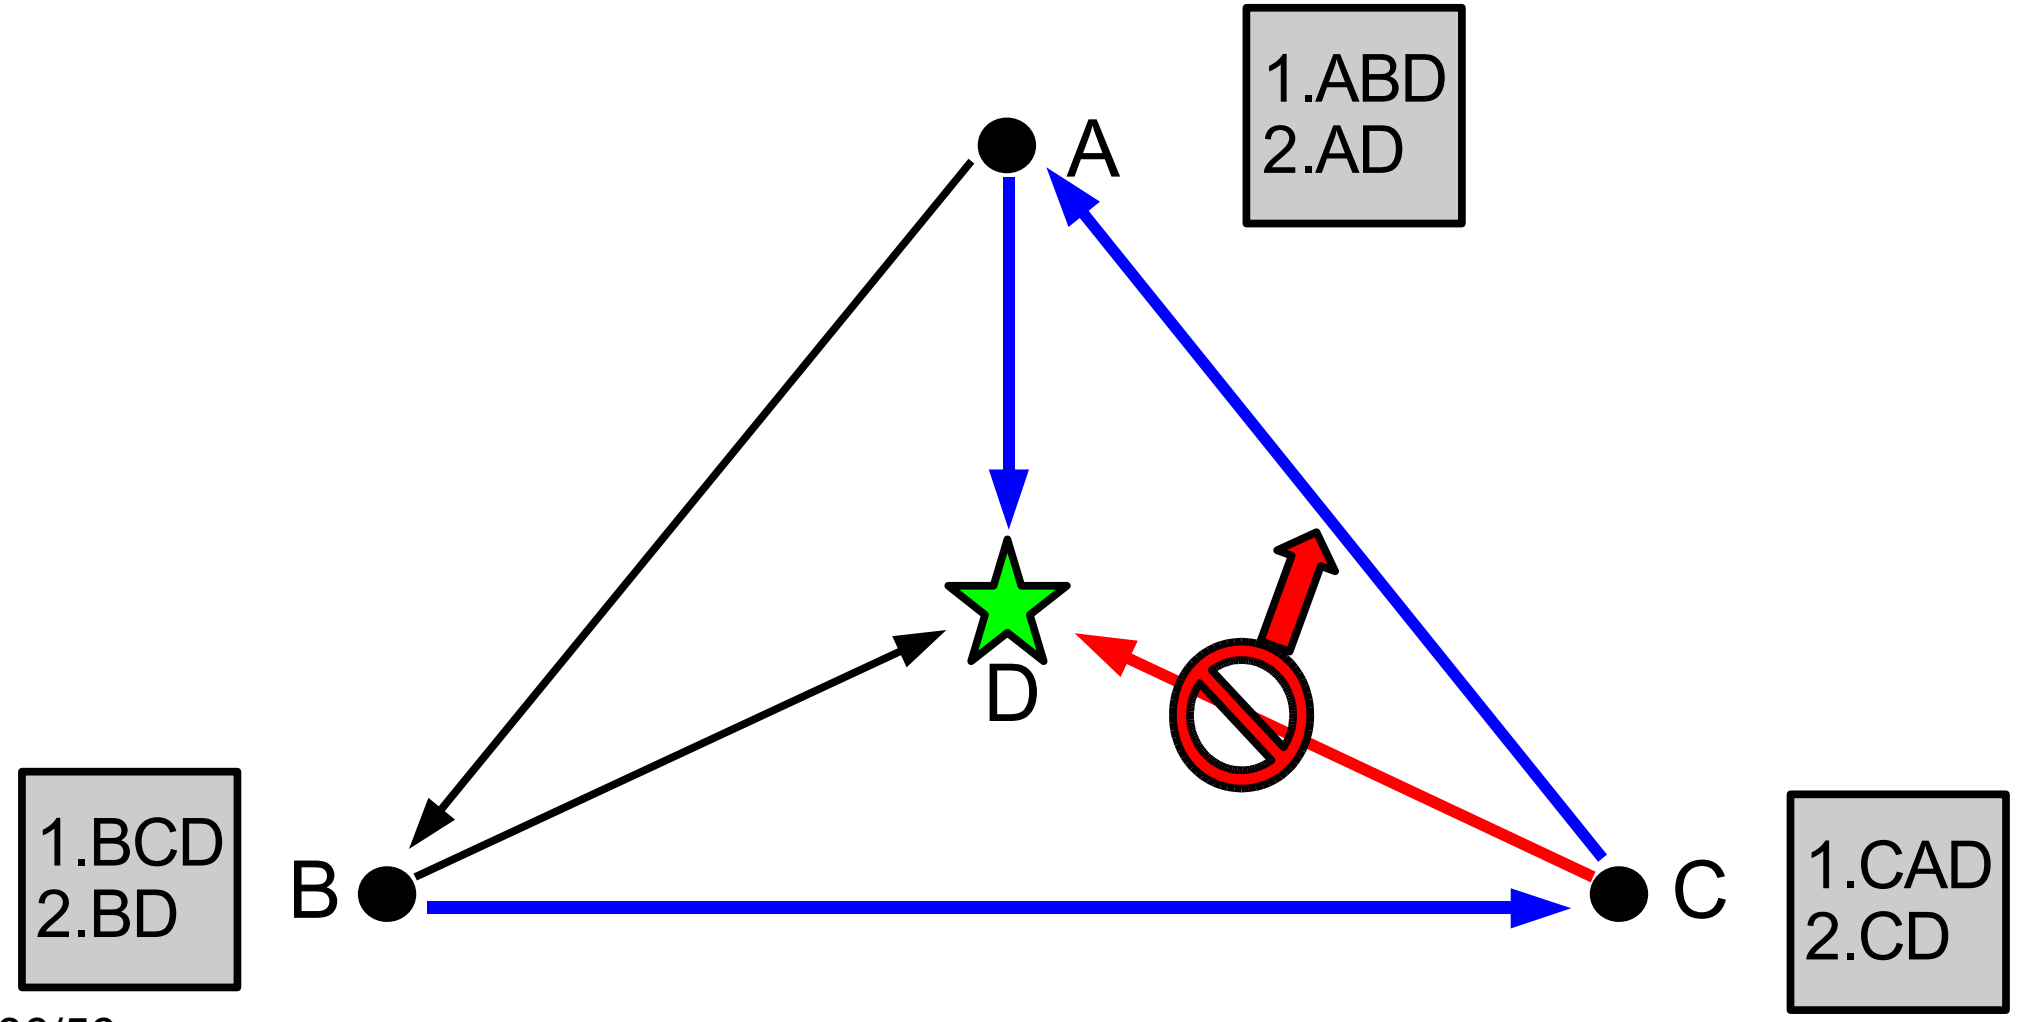

Fractional Stable Paths Problem

- Haxell and Wilfong. A fractional model of the border gateway protocol (BGP). SODA, 2008.

- Total paths must sum to 1
- Can't use a path more than a node ahead of you uses its portion.

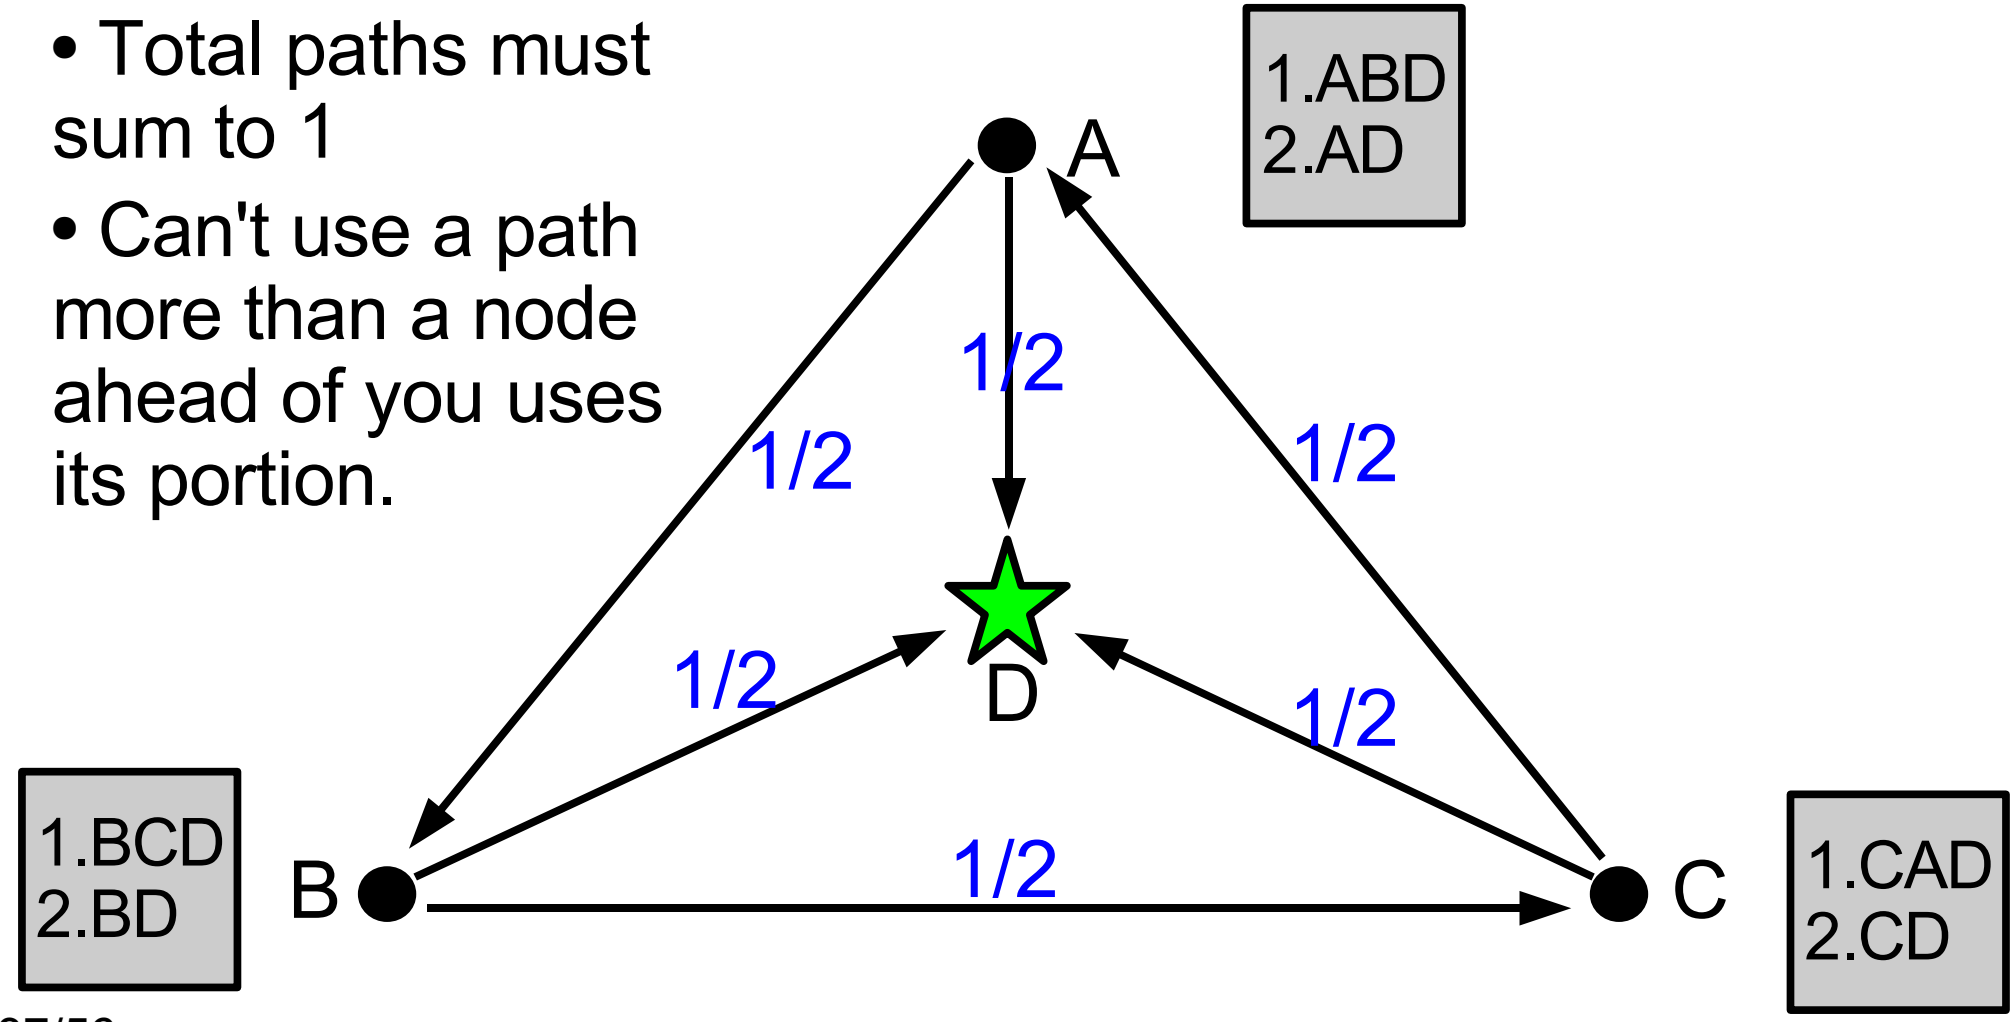

Fractional Stable Paths Problem

- Think of this as a possible “improved” BGP.
- Choosing fractional paths – fraction of bandwidth allowed.
- Can't use a path more than the fraction purchased.
- Peak rate: assume each node has a different peak time, each can use full fraction.

Fractional Networks

Fractional BBC:

- Multiple destinations
- Numerical utility (based on path length)

Fractional Stable Paths:

- Can reduce to single destination
- Ordinal utility (based on path preferences)

Fractional Networks

Fractional BBC:

- Choosing fractional edges

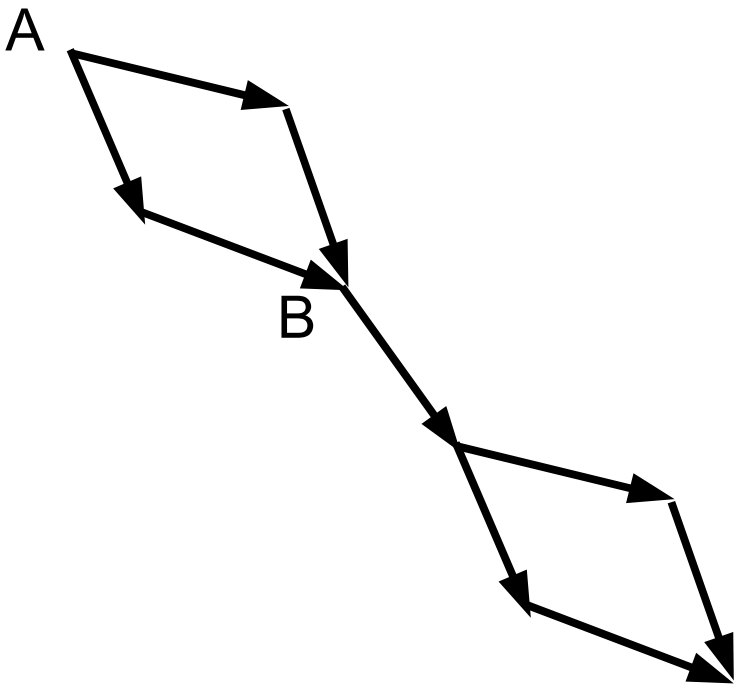

Fractional Stable Paths:

- Choosing fractional paths

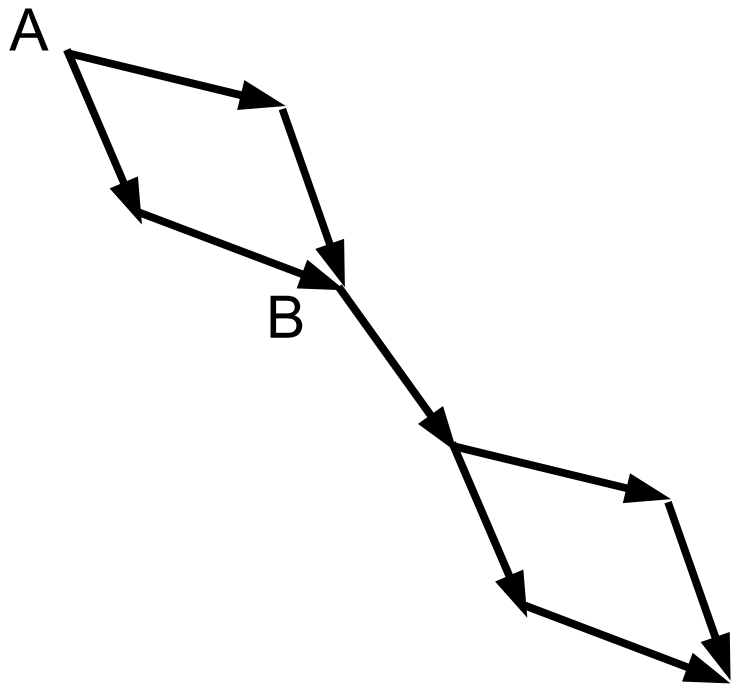

Fractional Networks

Fractional BBC:

- Choosing fractional edges

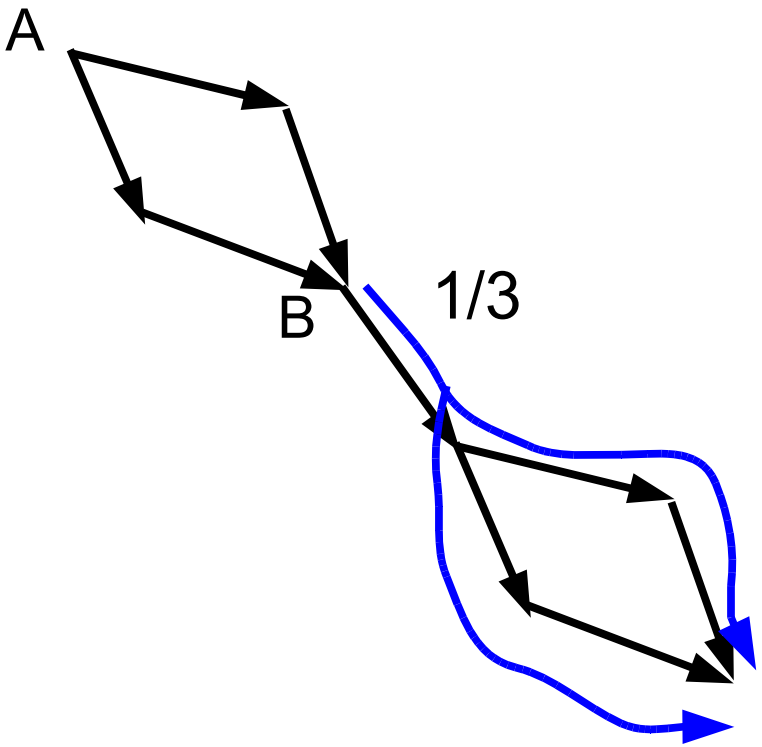

Fractional Stable Paths:

- Choosing fractional paths

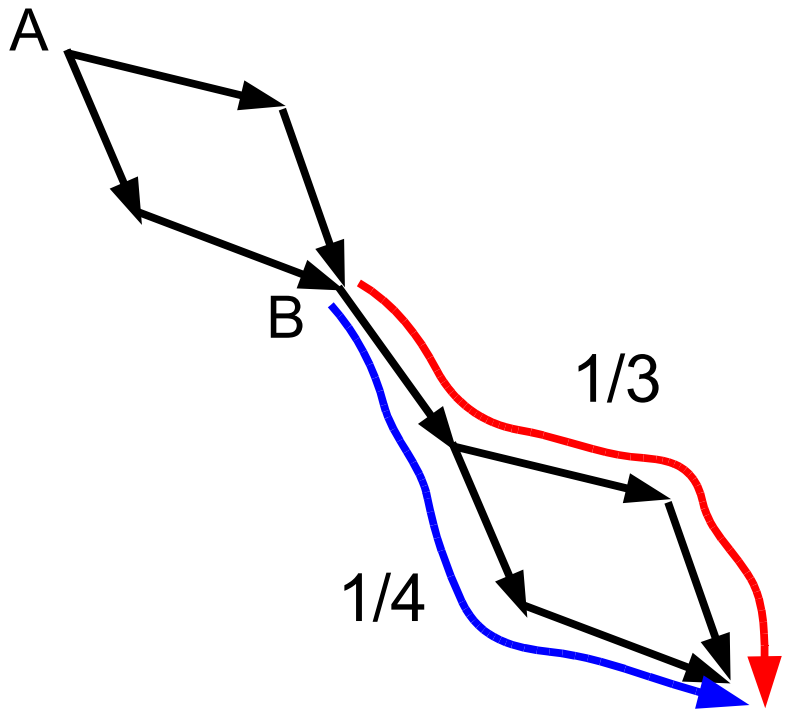

Fractional Networks

Fractional BBC:

- Choosing fractional edges

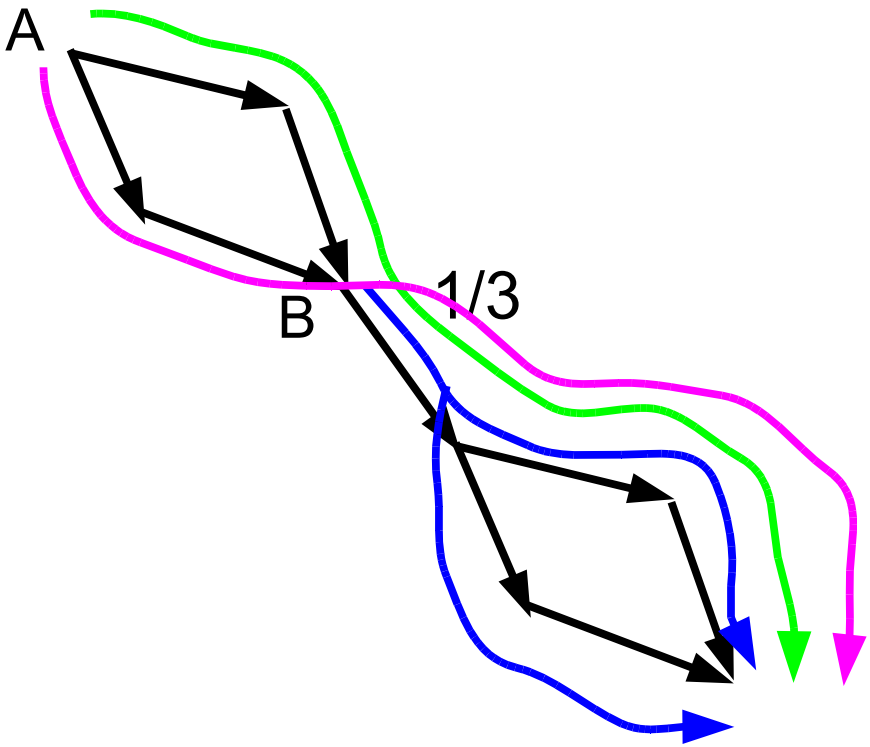

Fractional Stable Paths:

- Choosing fractional paths

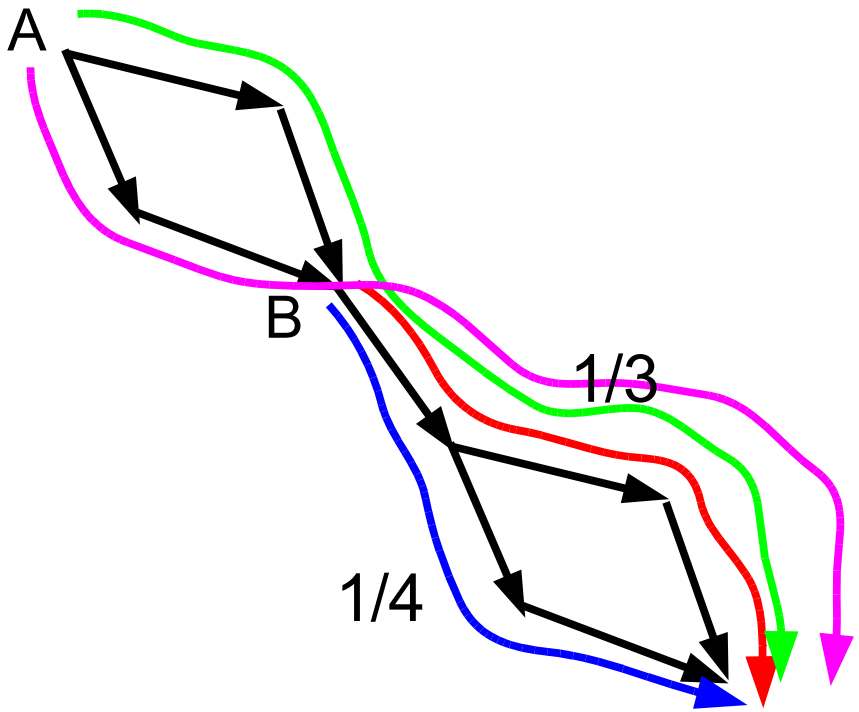

Fractional Networks

Fractional BBC:

- Choosing fractional edges

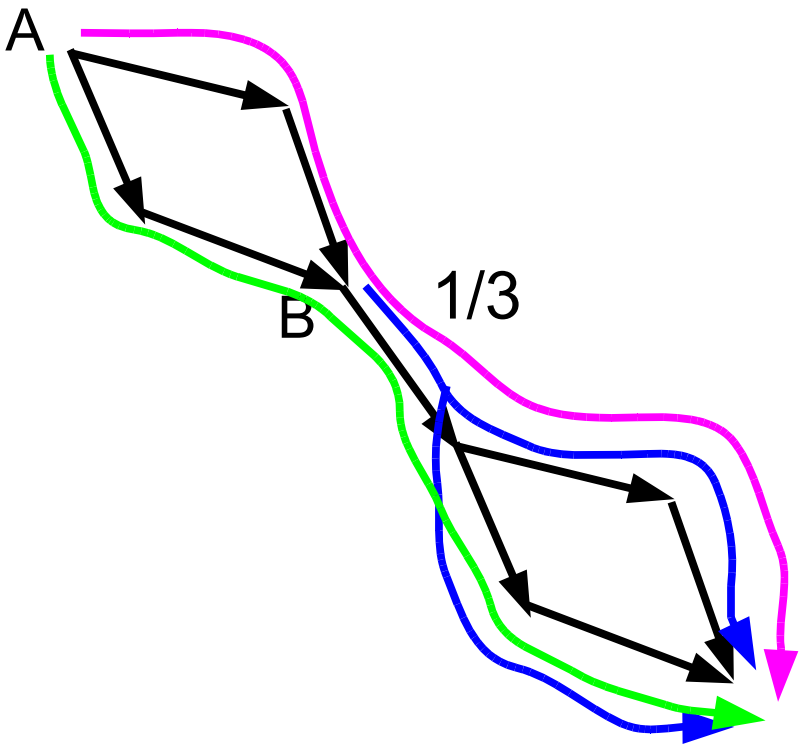

Fractional Stable Paths:

- Choosing fractional paths

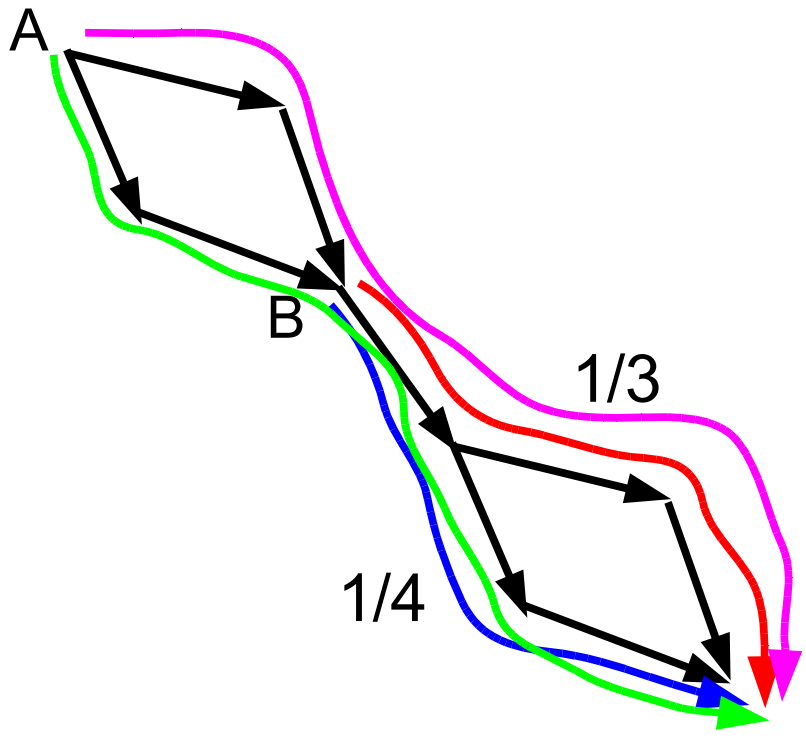

Fractional Networks

- Fractional BBC and Fractional Stable Paths are equivalent if:
  - BBC with single destination, uniform costs and lengths, budget 1
  - All paths are 1 or 2 hops

Fractional Networks

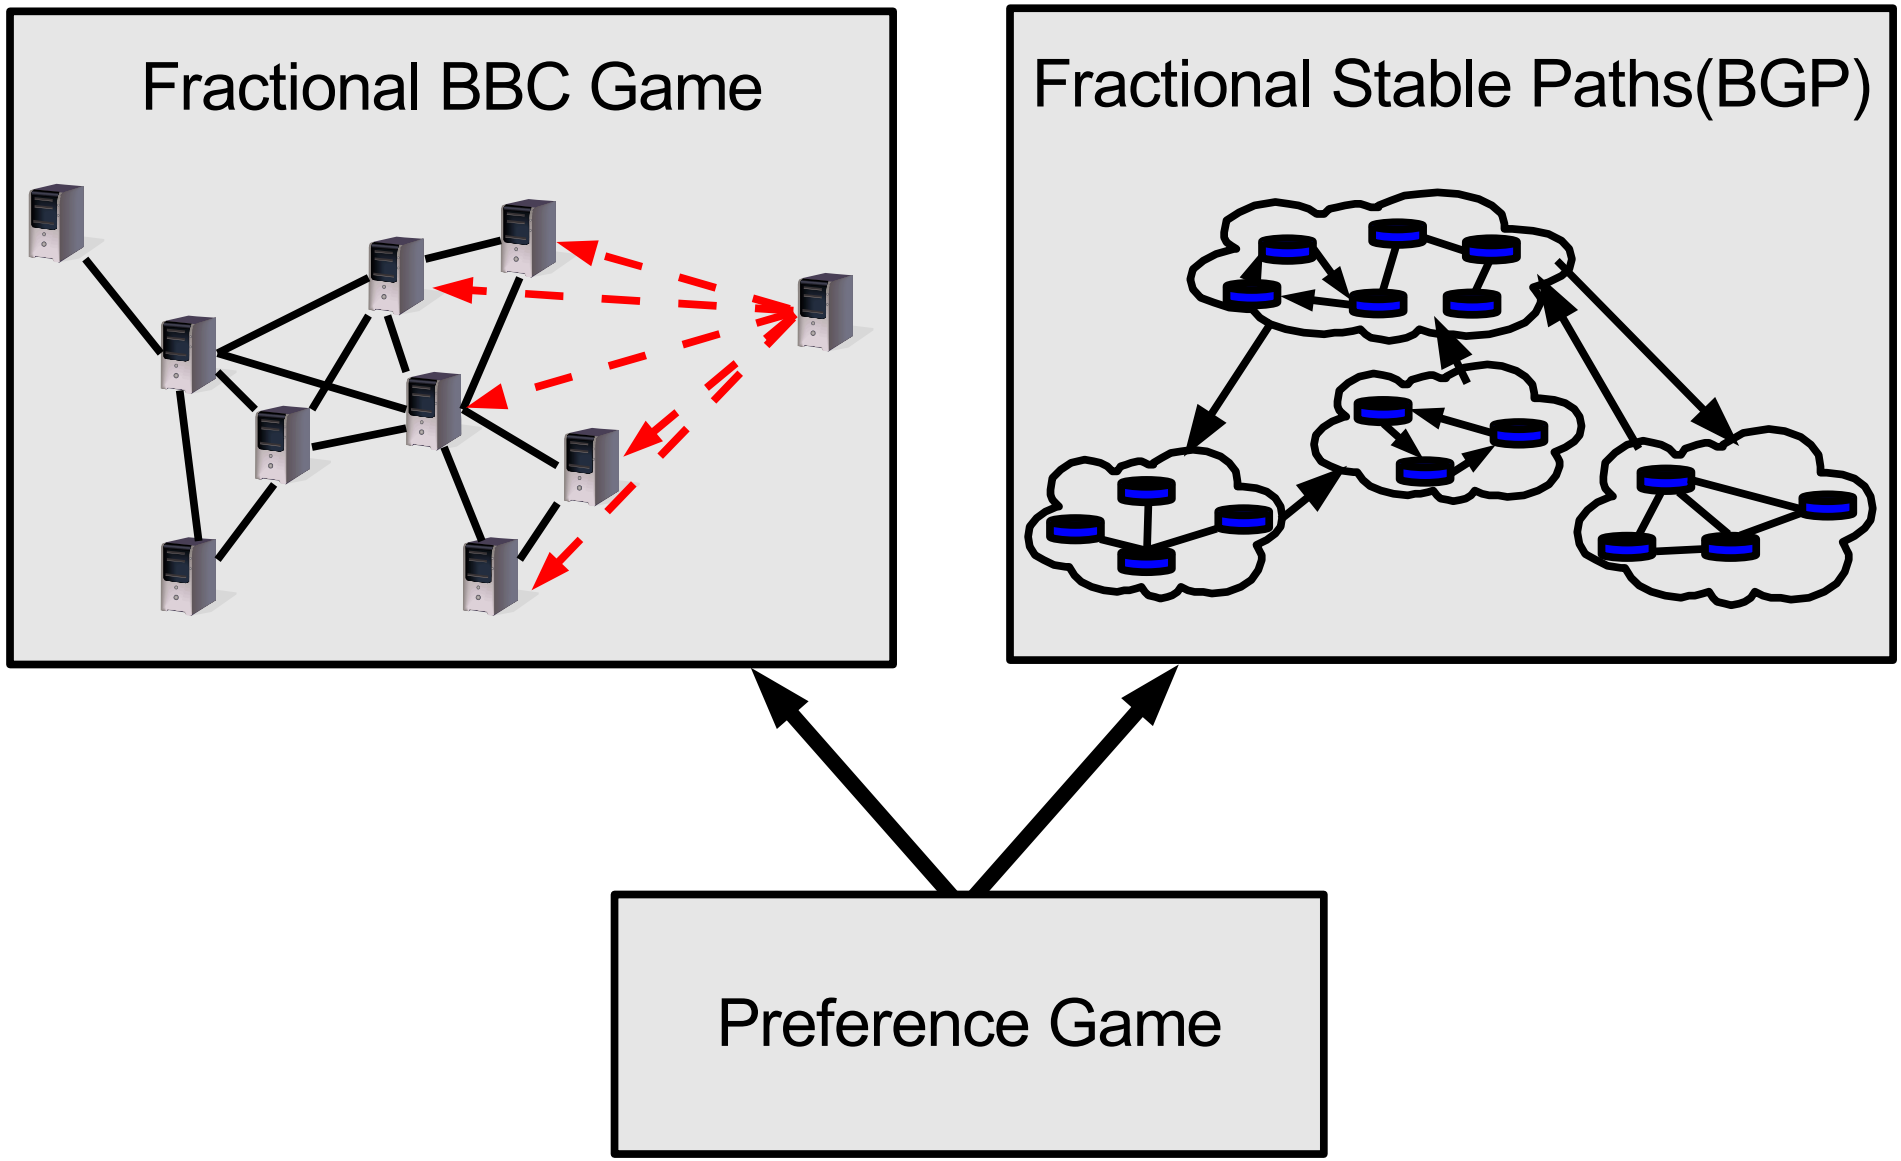

The Preference Game

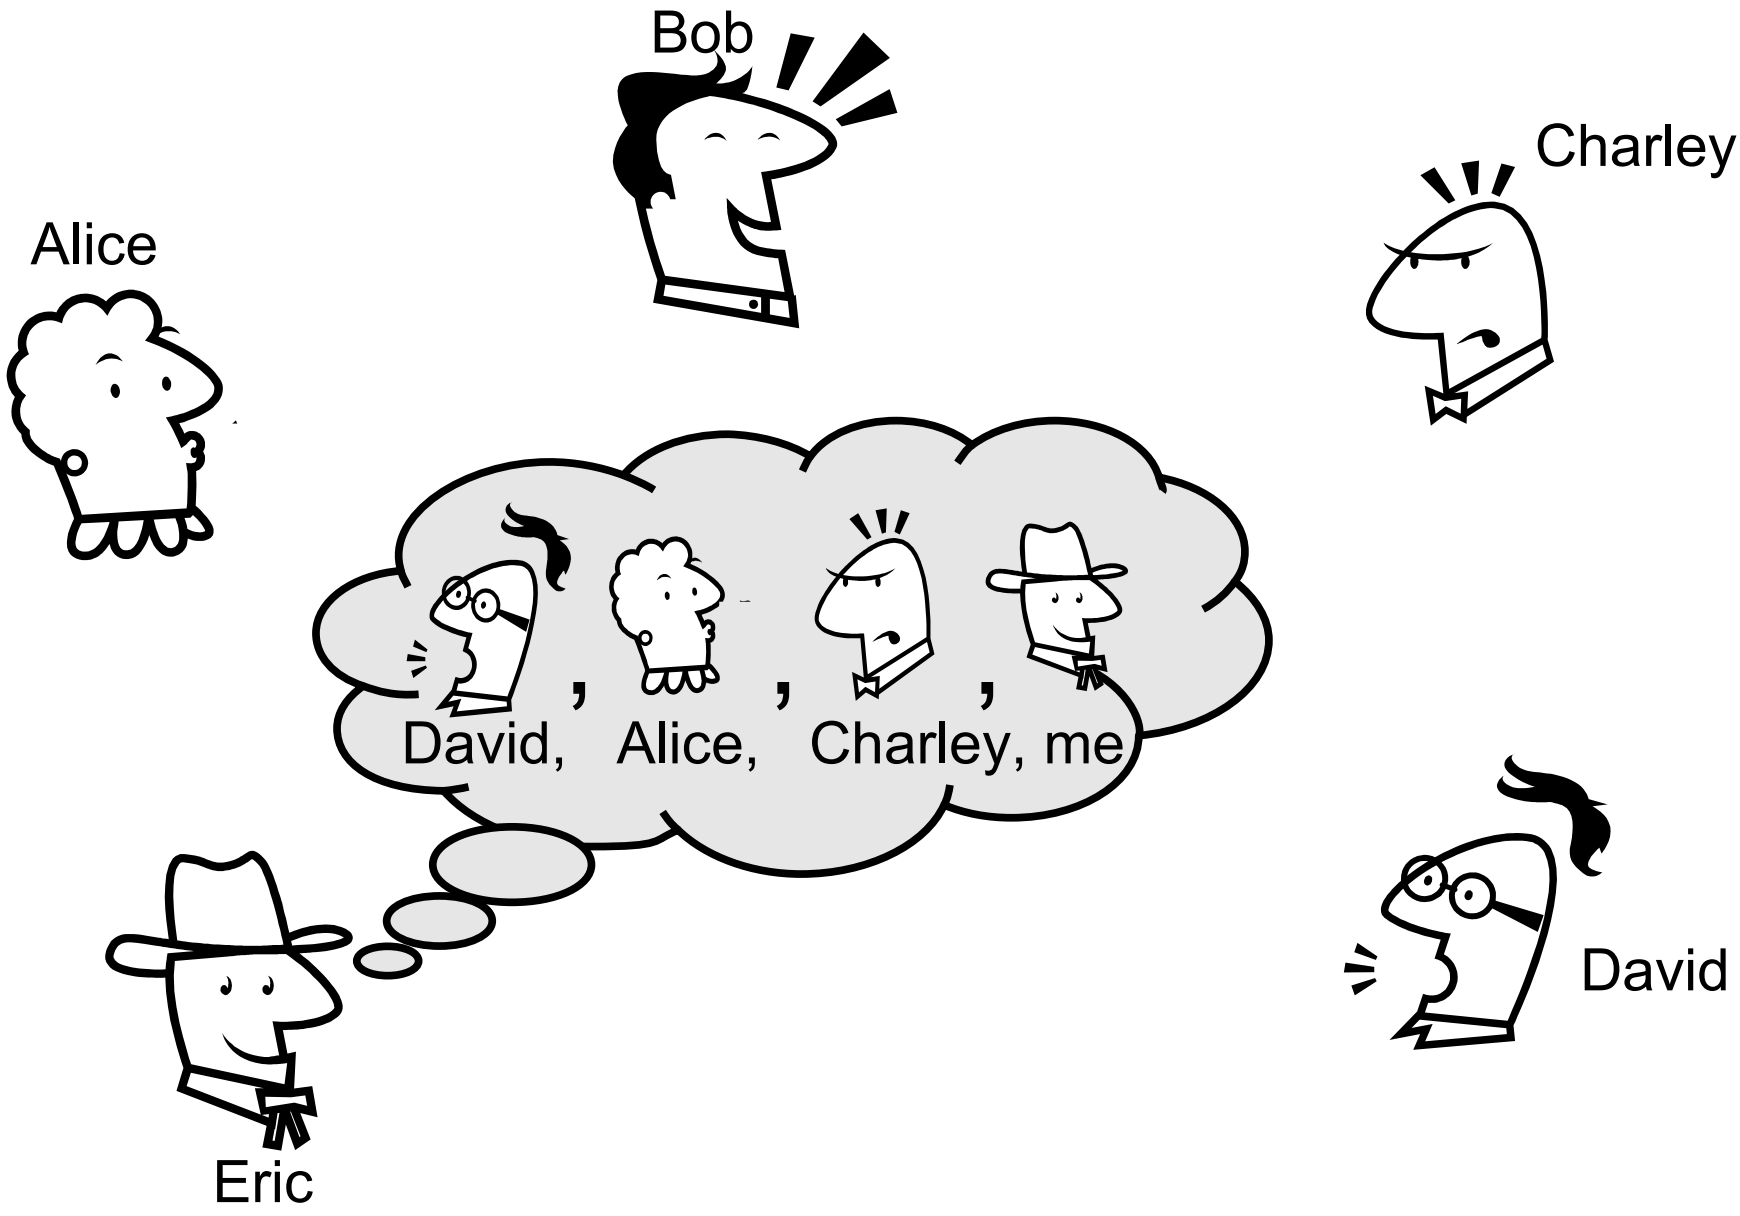

The Preference Game

- Assign weight to each player
- Assign total weight 1
- Assign weight at most the amount that player assigns to himself
- Maximize weight on highest preference players

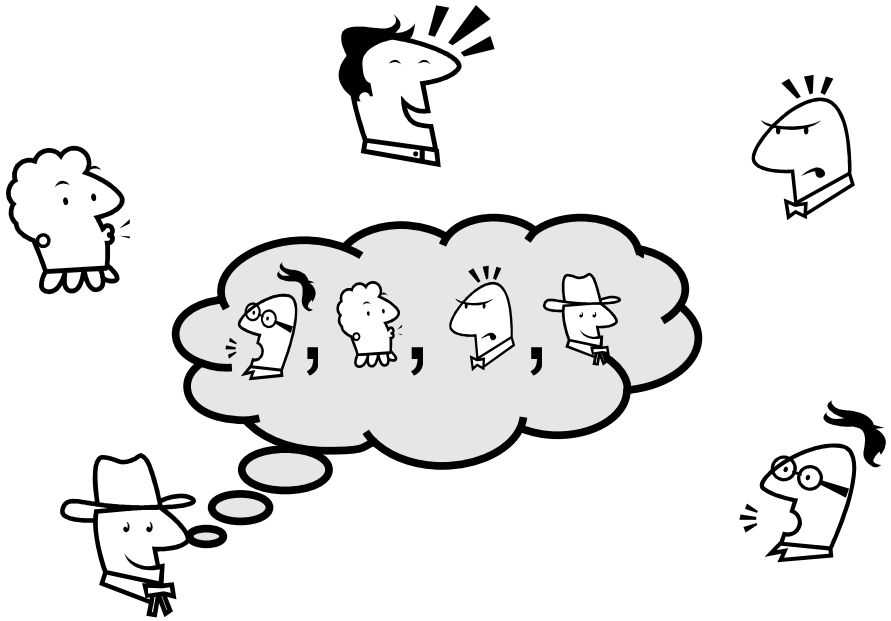

The Preference Game

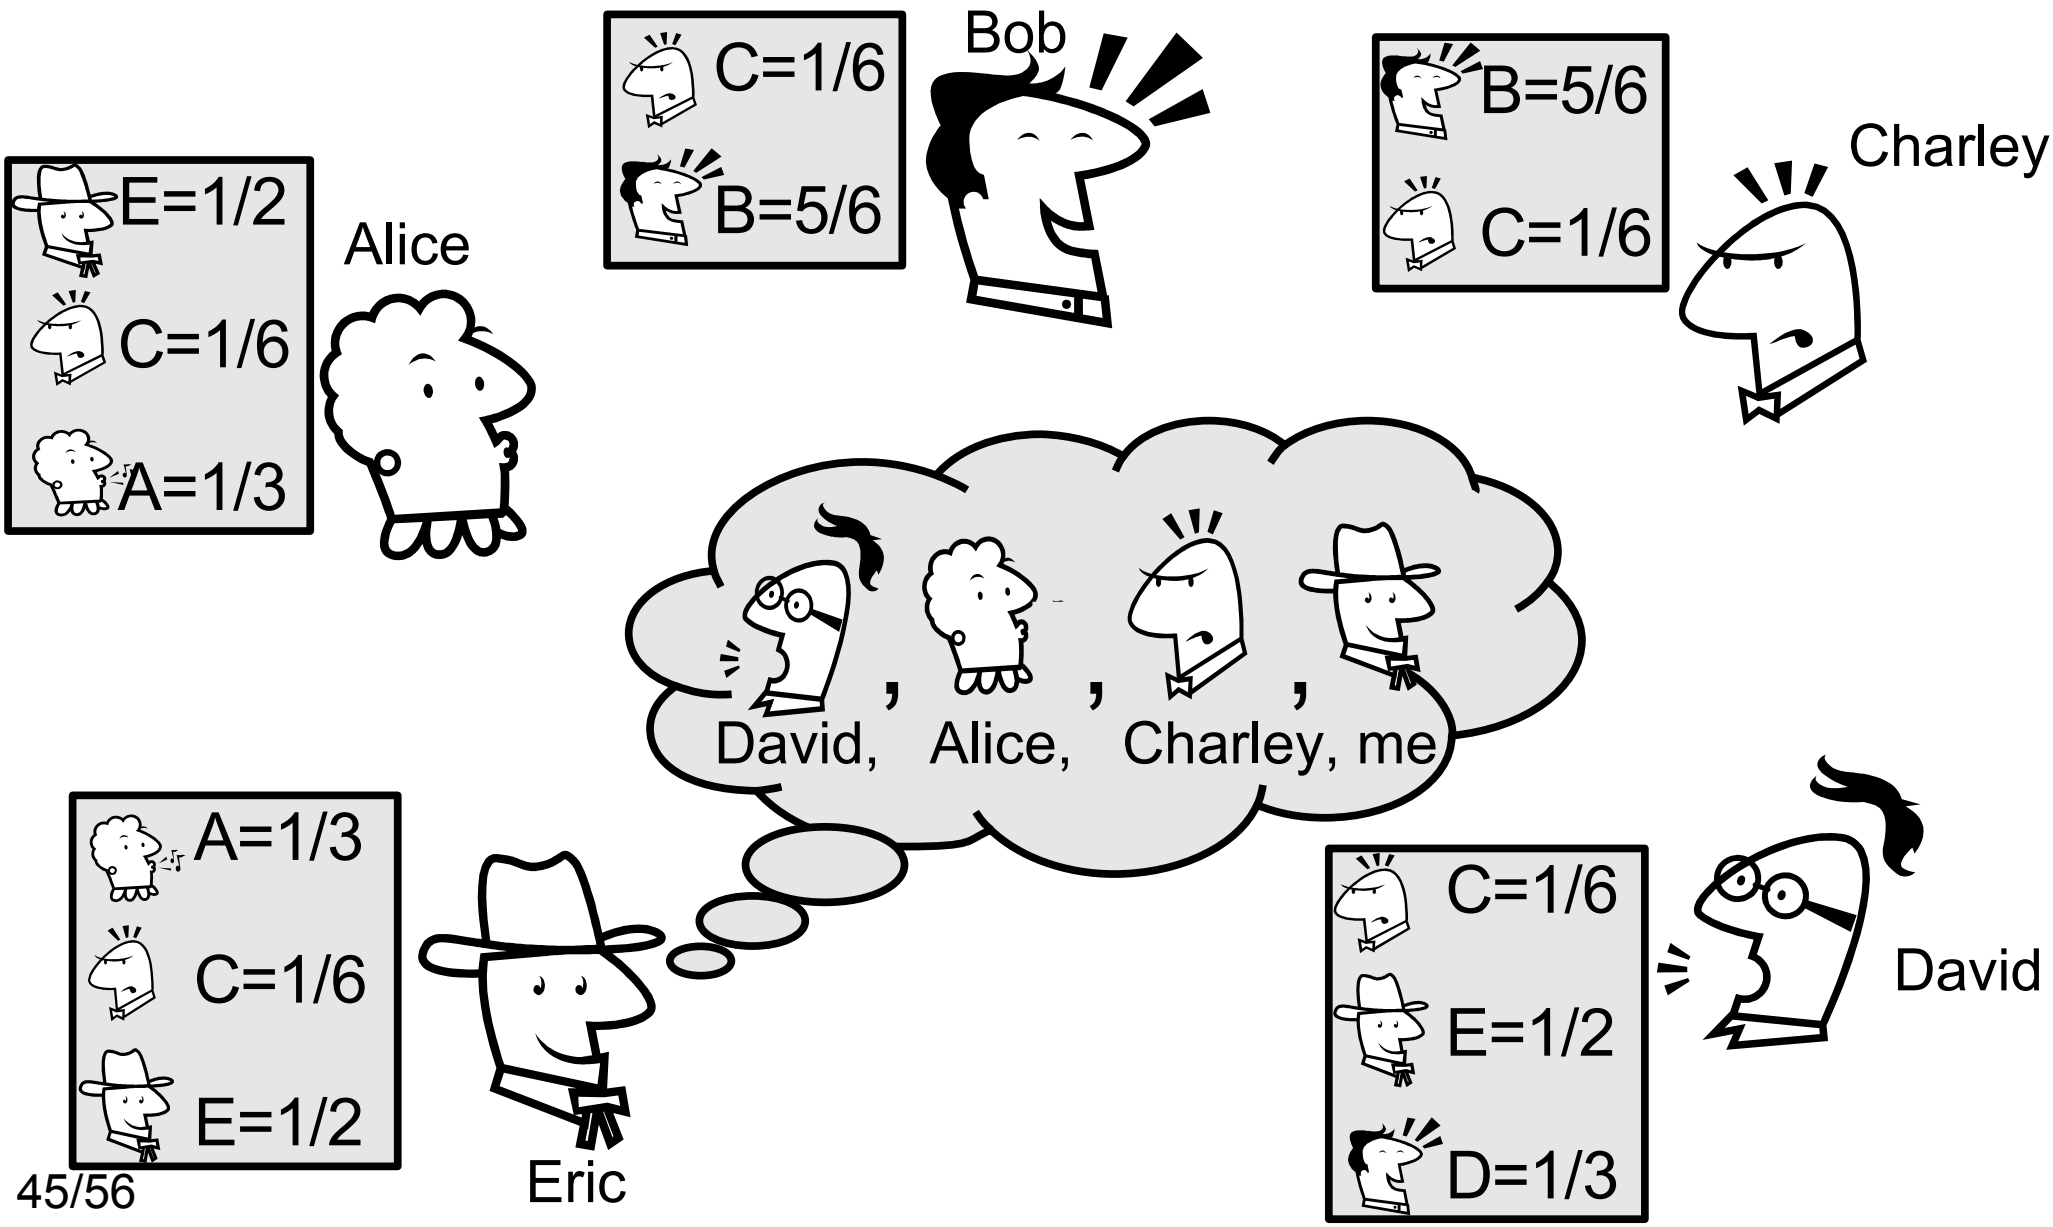

The Preference Game

- Reduces to both fractional BBC and fractional Stable Paths Problem (as well as many others).
- A pure Nash equilibrium always exists.
- In fact, a *rational* pure Nash equilibrium always exists.
- Seems like it should be easy to “solve”
- If preferences follow some rules, it is easy to solve.

The Preference Game

- In general: PPAD hard to find an equilibrium (even an approximate equilibrium)
  - PPAD = Same as “end of the line”
    - Papadimitriou. *On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence*. JCSS 1994.
  - As hard as finding mixed Nash in general games
    - Daskalakis, Goldberg, Papadimitriou. *The Complexity of Computing a Nash Equilibrium*. STOC, 2006.
    - Goldberg, Papadimitiou. *Reducibility Among Equilibrium Problems*. STOC, 2006.
    - Chen, Deng, and Teng. *Computing Nash Equilibria: Approximation and Smoothed Complexity*. FOCS, 2006.

The Preference Game

- In general: PPAD hard to find an equilibrium (even an approximate equilibrium)

- All algorithms unlikely to converge in polynomial time.
- Can't expect a real-world system to converge.

Fractional BBC, Fractional SPP

- Is any network resulting from decentralized design:
  - "Good" or "Bad"?

Price of Anarchy =
Worst Nash / Social Optimum

  - Stable or oscillating?

Do Pure Nash Equilibria exist?

stable exists

  - Converging or cycling?

Will best response trials converge to an equilibrium?

Fractional BBC, Fractional SPP

- Is any network resulting from decentralized design:
  - "Good" or "Bad"? ← Price of Anarchy = Worst Nash / Social Optimum
  - Stable or oscillating? ← Do Pure Nash Equilibria exist?

    stable exists
  - Converging or cycling? ← Will best response trials converge to an equilibrium?

    Likely to oscillate

Recap

- Even simple models may have “bad” equilibria

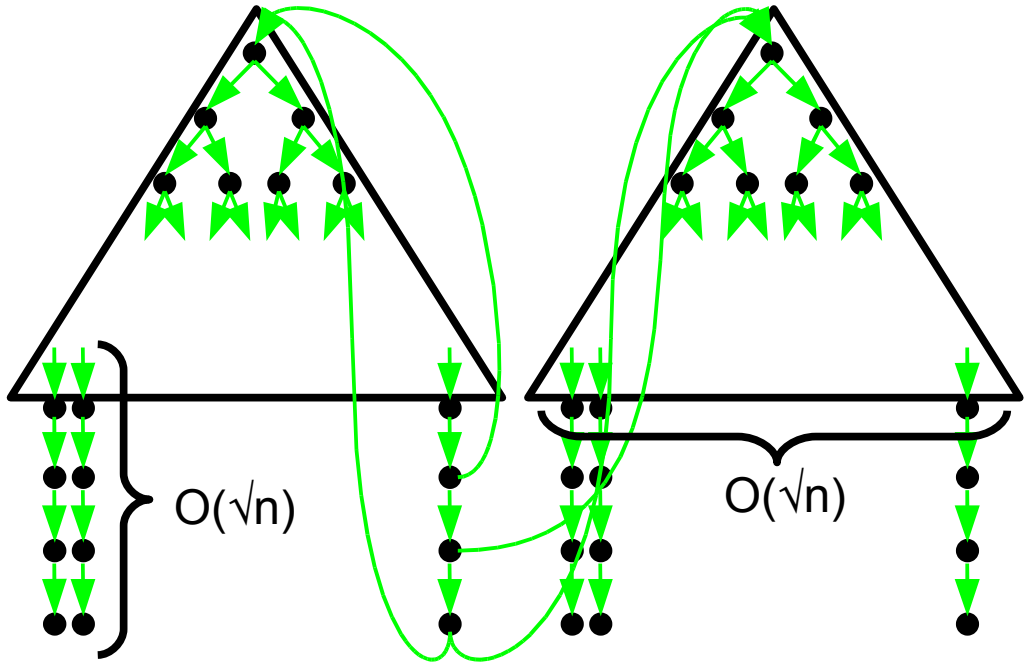

Recap

- Stable solutions don't imply convergence

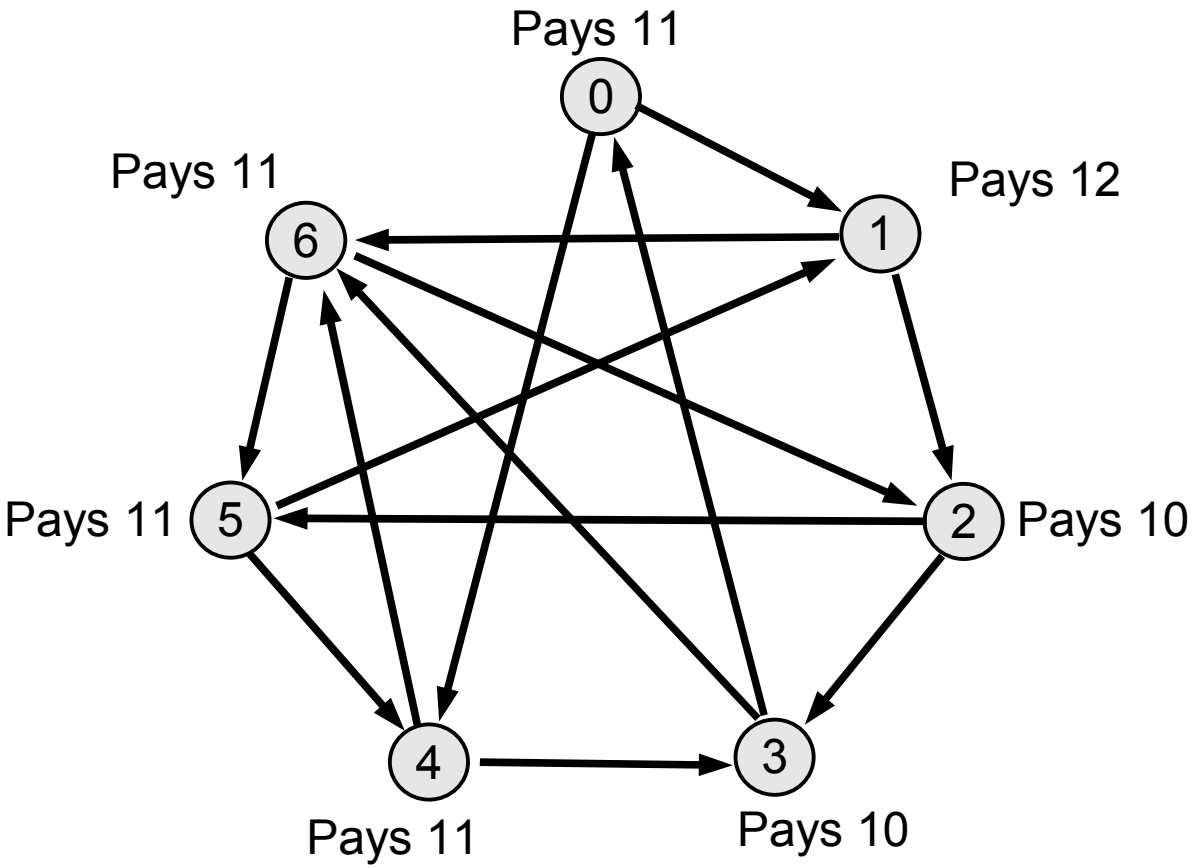

Recap

- More complicated (realistic) models may have no stable solution.

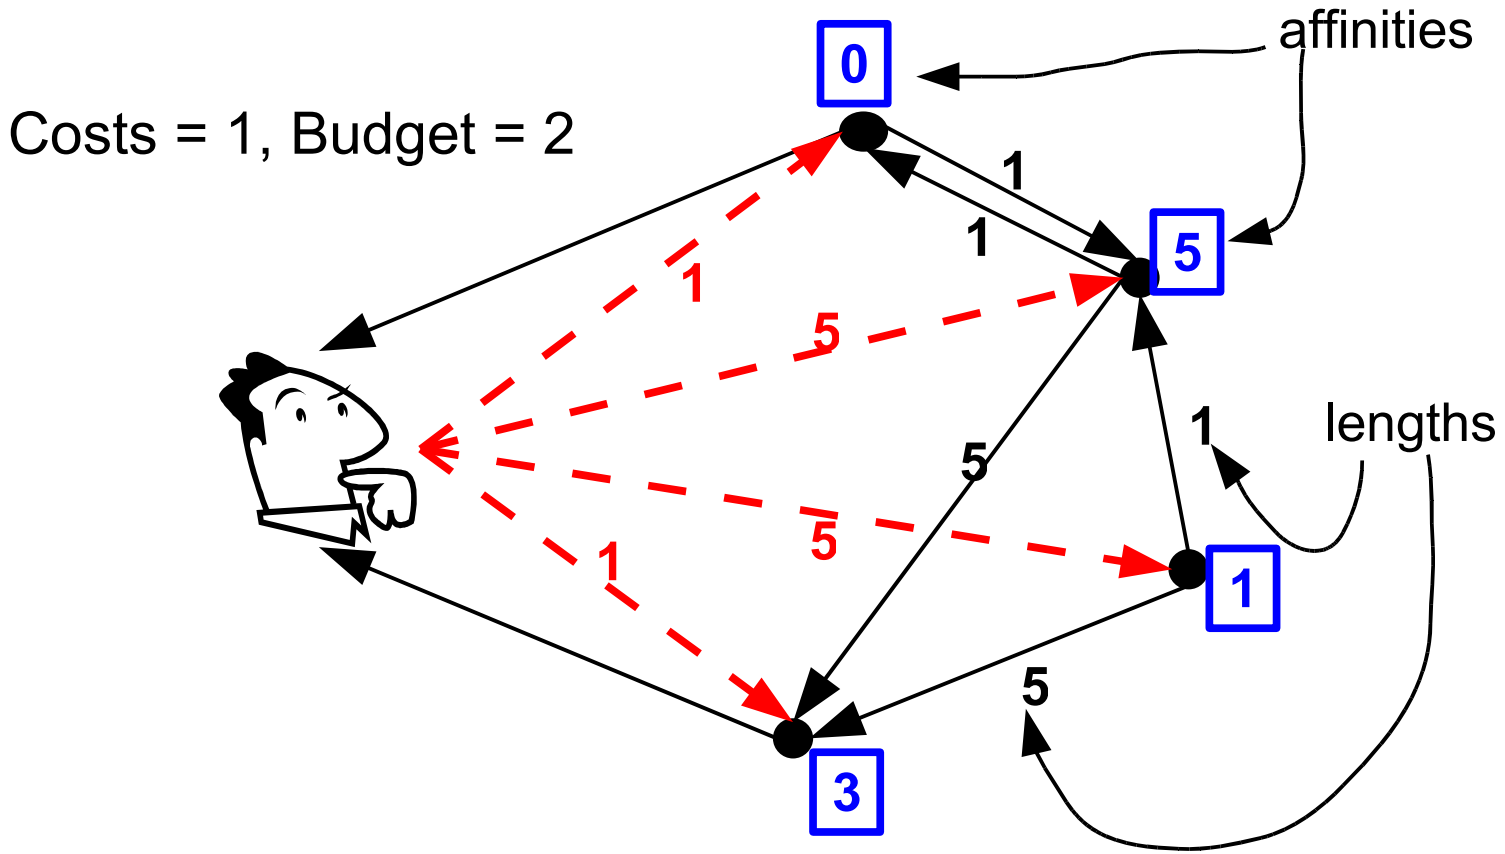

Recap

- Allowing fractional edges implies stable solutions, but may be hard to find.

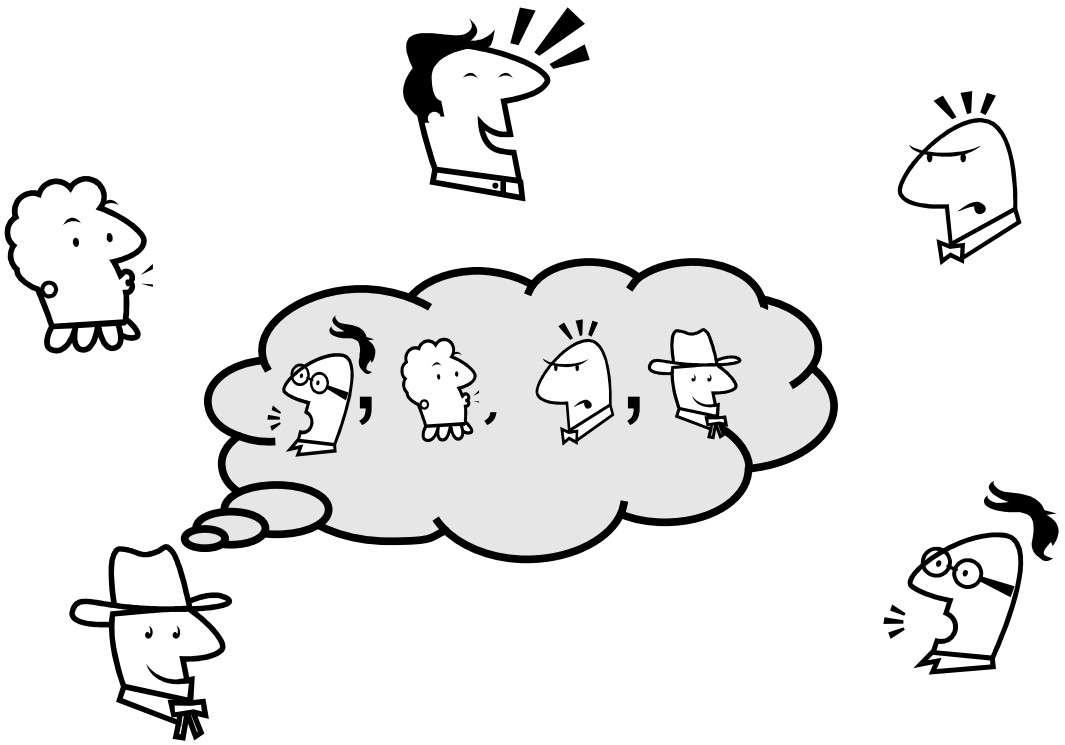

What does this mean?

- Stable solutions may not exist
- Stable solutions don't immediately imply convergence
- Stable solutions may be hard to find

so...

- If we allow nodes to choose their edges for an overlay network with bounded degree, the network is likely to oscillate.

What can we do? (open questions)

- Approximation – “close enough” to a best response

What can we do? (open questions)

- Approximation – “close enough” to a best response
- Include cost of changing edges into the utilities

What can we do? (open questions)

- Approximation – “close enough” to a best response
- Include cost of changing edges into the utilities
- Give the nodes a push towards a good solution

What can we do? (open questions)

- Approximation – “close enough” to a best response
- Include cost of changing edges into the utilities
- Give the nodes a push towards a good solution
- Repeated game payoffs
  - strategy is a series of moves
  - payoff based on an average, possibly weighted over time.

What can we do? (open questions)

- Approximation – “close enough” to a best response
- Include cost of changing edges into the utilities
- Give the nodes a push towards a good solution
- Repeated game payoffs
  - strategy is a series of moves
  - payoff based on an average, possibly weighted over time.
- Hierarchical instead of fully distributed

Results mentioned are from:

- N. Laoutaris, L. Poplawski, R. Rajaraman, R. Sundaram, S.-H. Teng. *Bounded Budget Connection (BBC) Games or How to Make Friends and Influence People, on a Budget*. PODC '08
- S. Kintali, L. Poplawski, R. Rajaraman, R. Sundaram, S.-H. Teng. *Reducibility Among Fractional Stability Problems.* arXiv:0904.1435v1 [cs.CC]

More about me:

- laura.poplawski@gmail.com
- http://www.ccs.neu.edu/home/ljp