On Characterizing Affinity and Its Impact On Network Performance

Gabriel Lucas
Abhishek Ghose
John Chuang

University of California, Berkeley

School of Information Management and Systems
Department of Electrical Engineering and Computer Science

Refer to our paper for more complete details:

Gabriel Lucas, Abhishek Ghose and John Chuang, On Characterizing Affinity and Its Impact on Network Performance. Proceedings of ACM SIGCOMM Workshop on Models, Methods and Tools for Reproducible Network Research, Karlsruhe Germany, August 25 2003.

Overview of presentation

What is affinity?

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

Major findings

Research questions we addressed

Going forward

What is affinity?

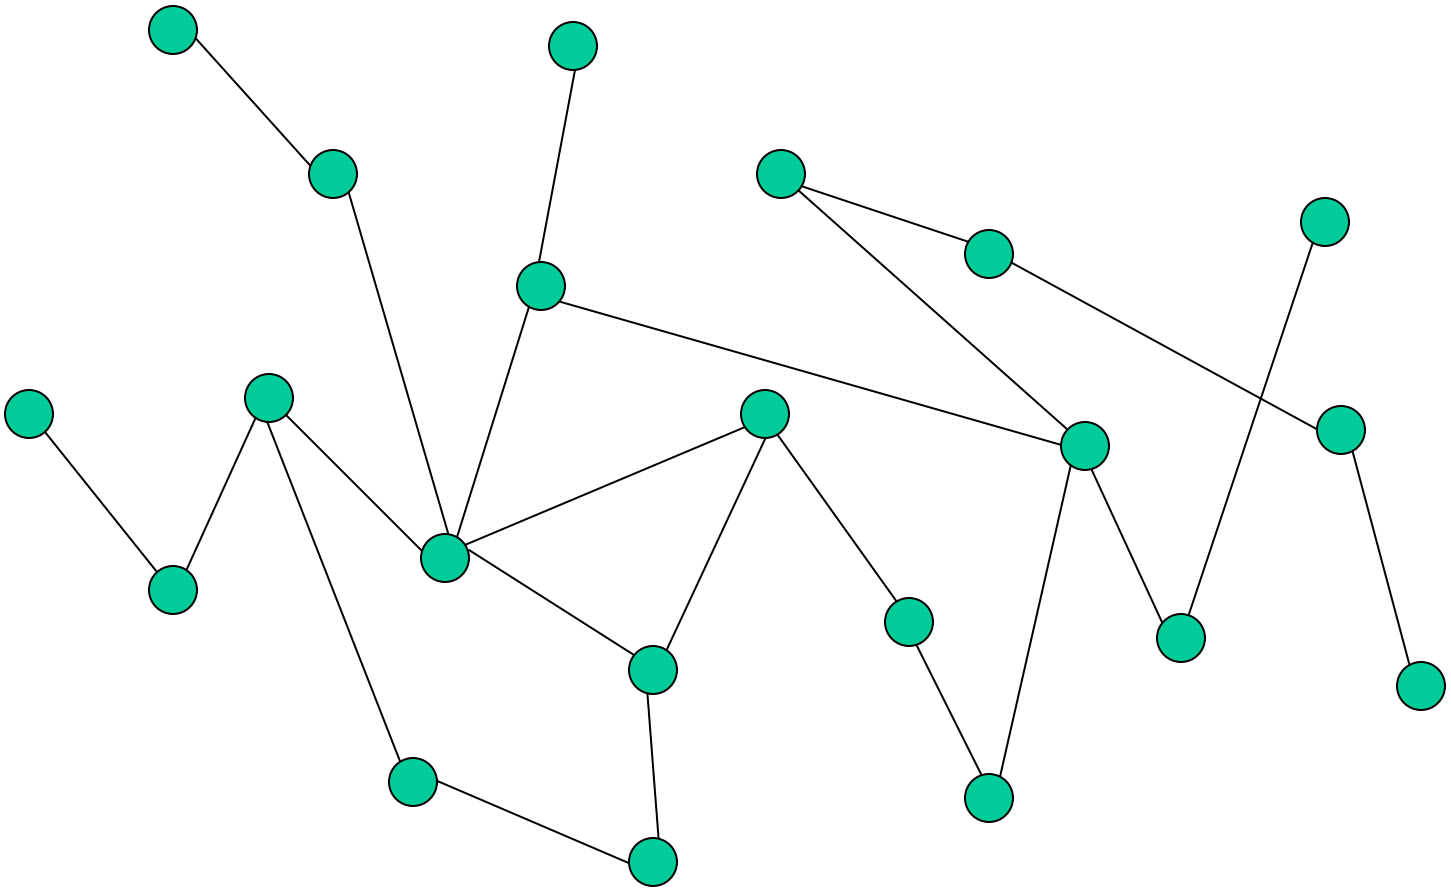

What is affinity?

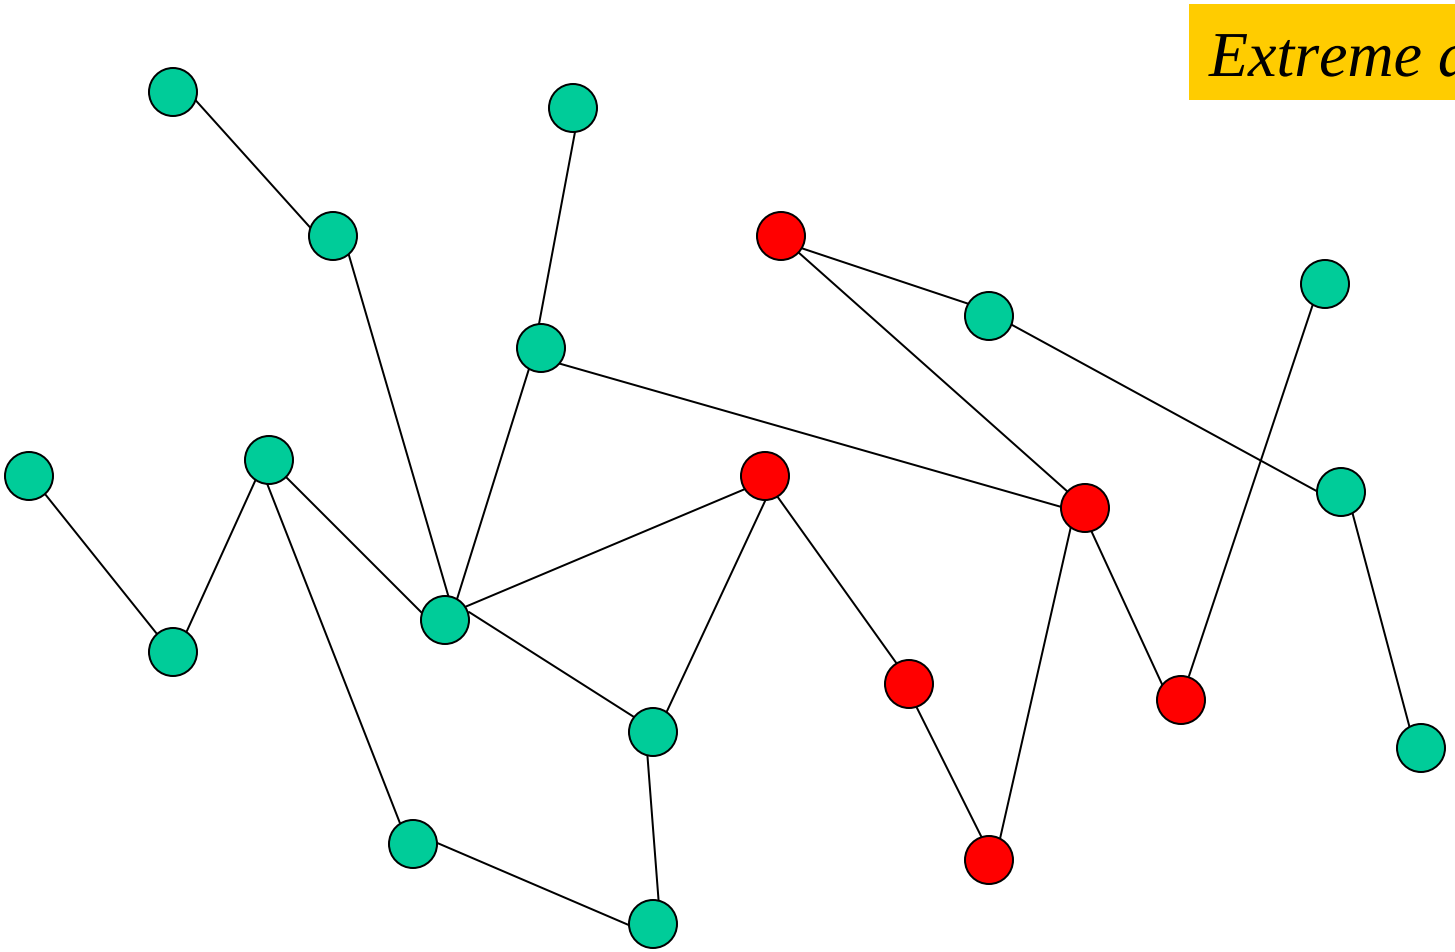

What is affinity?

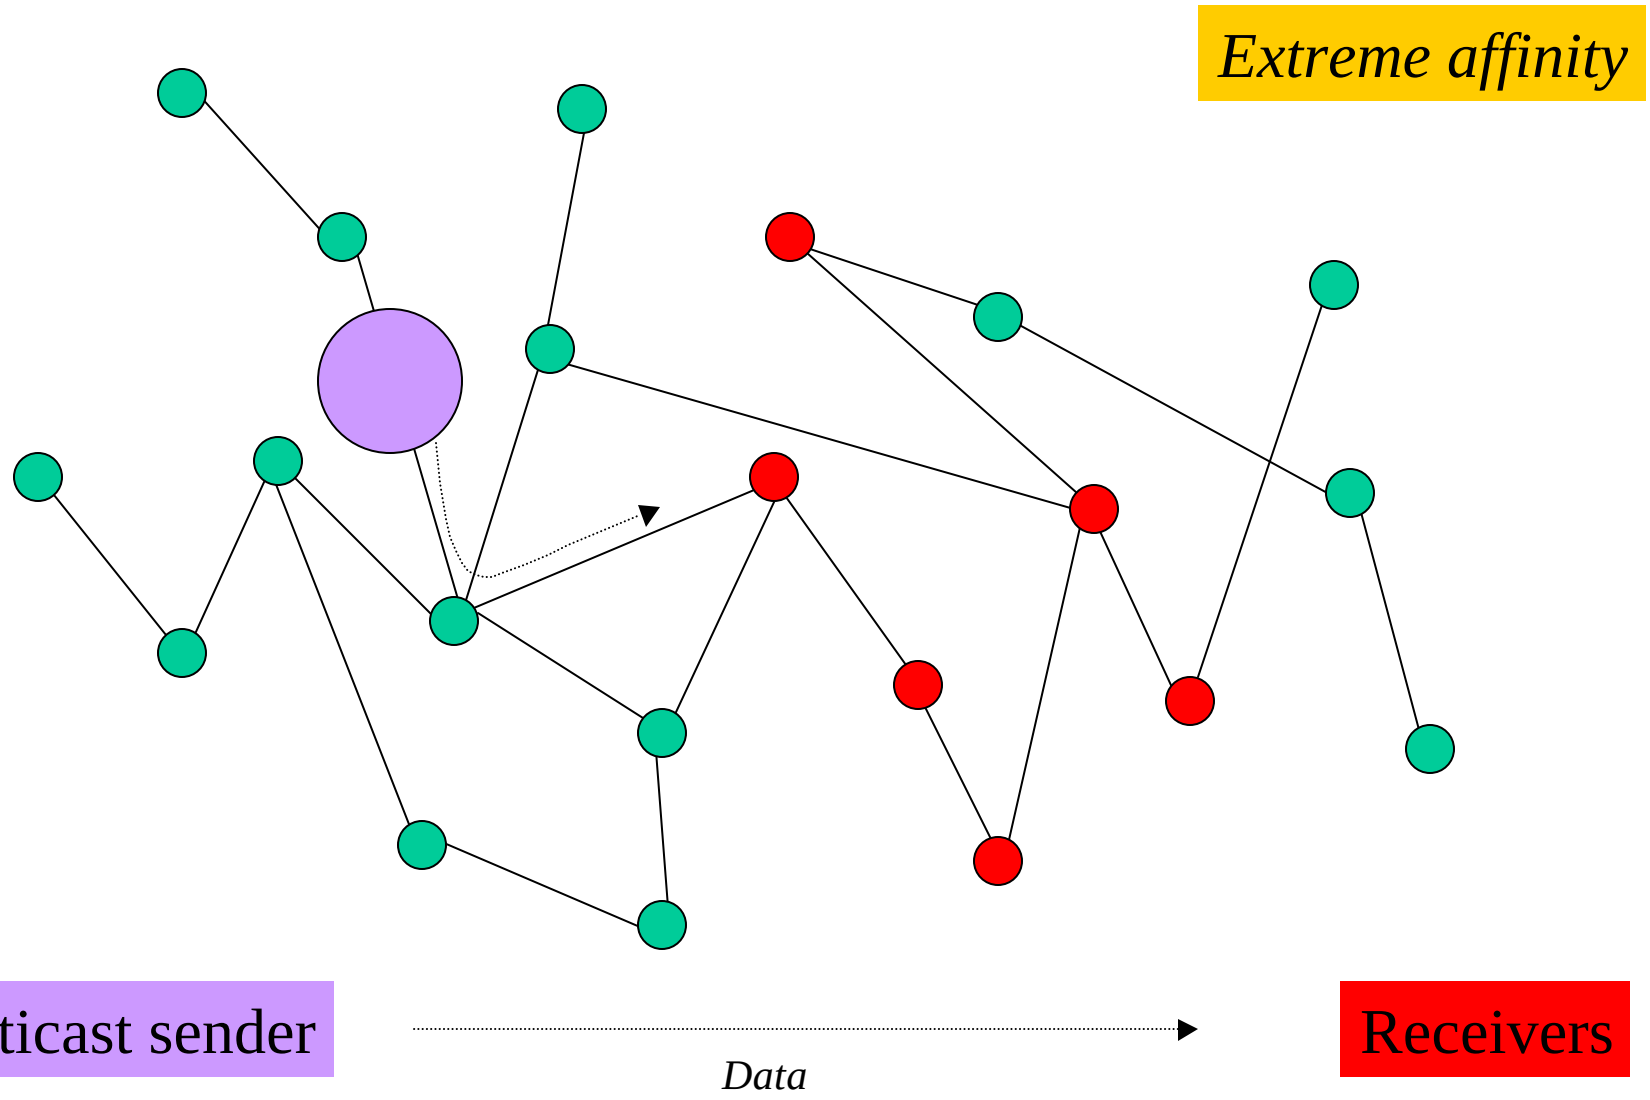

What is affinity?

Extreme disaffinity

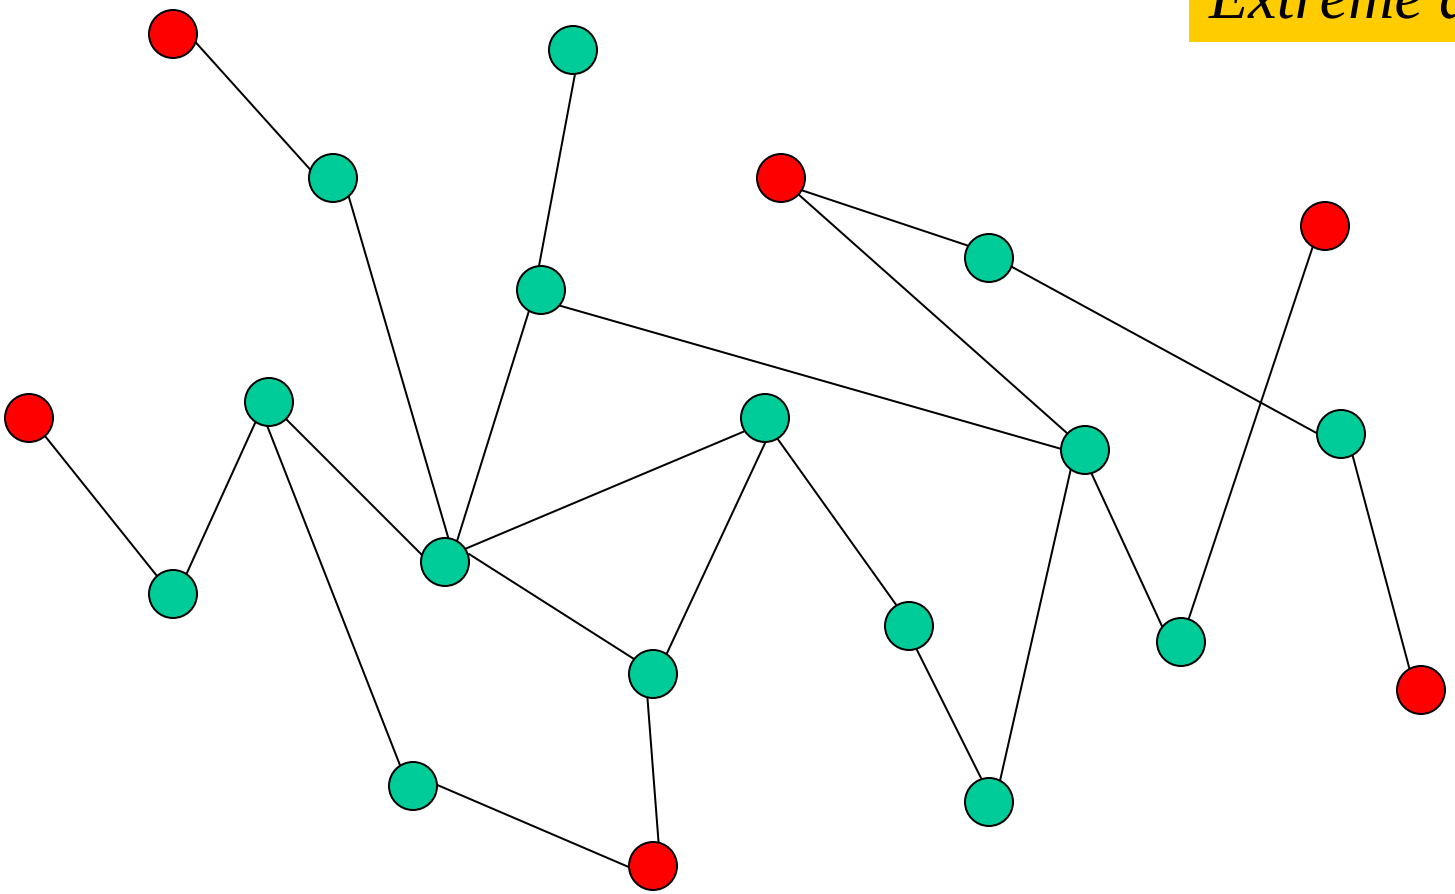

What is affinity?

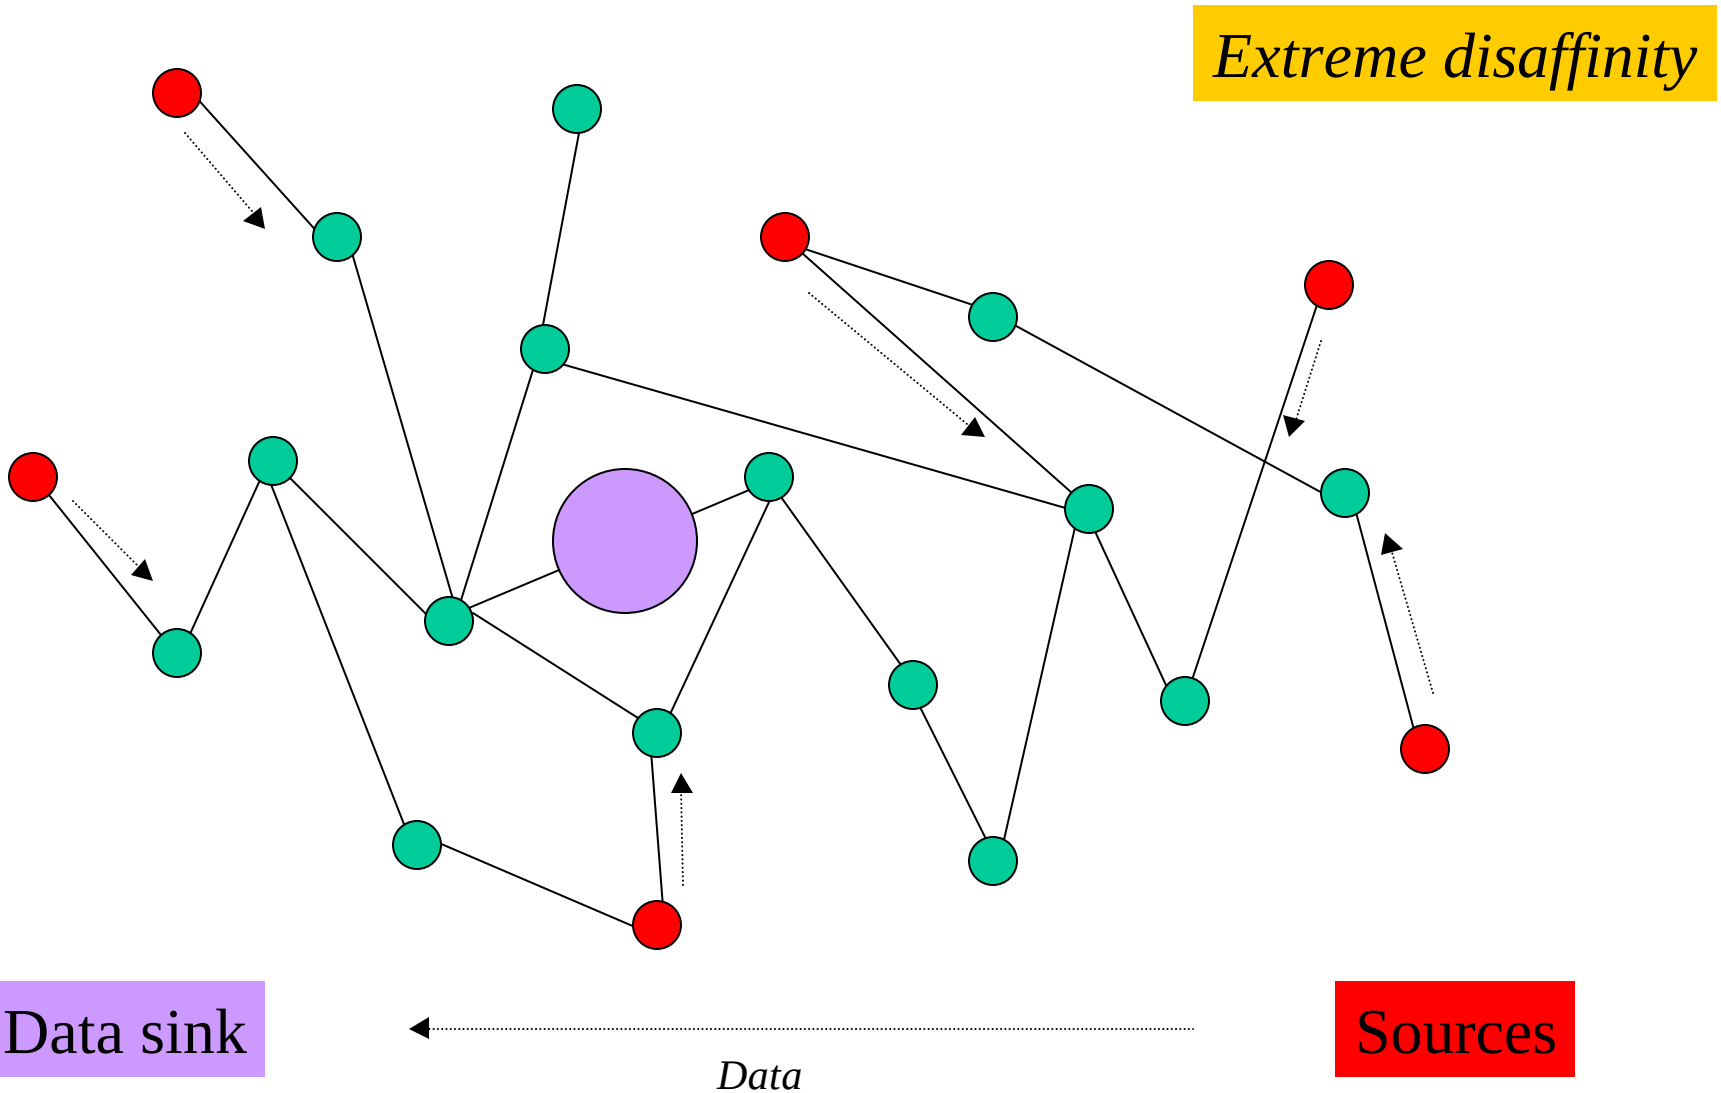

What is affinity?

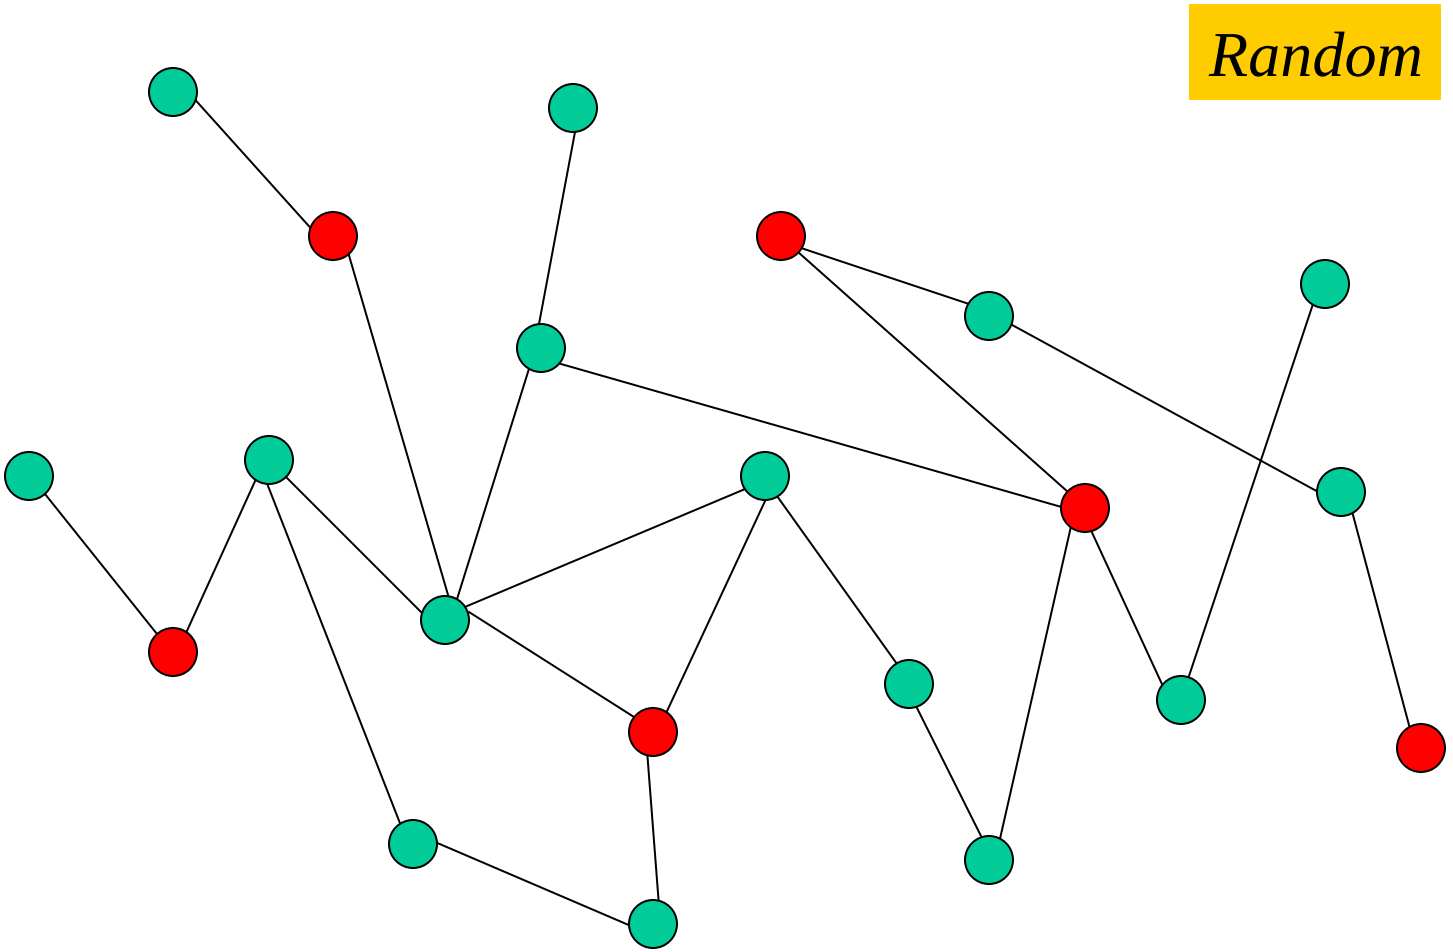

What is affinity?

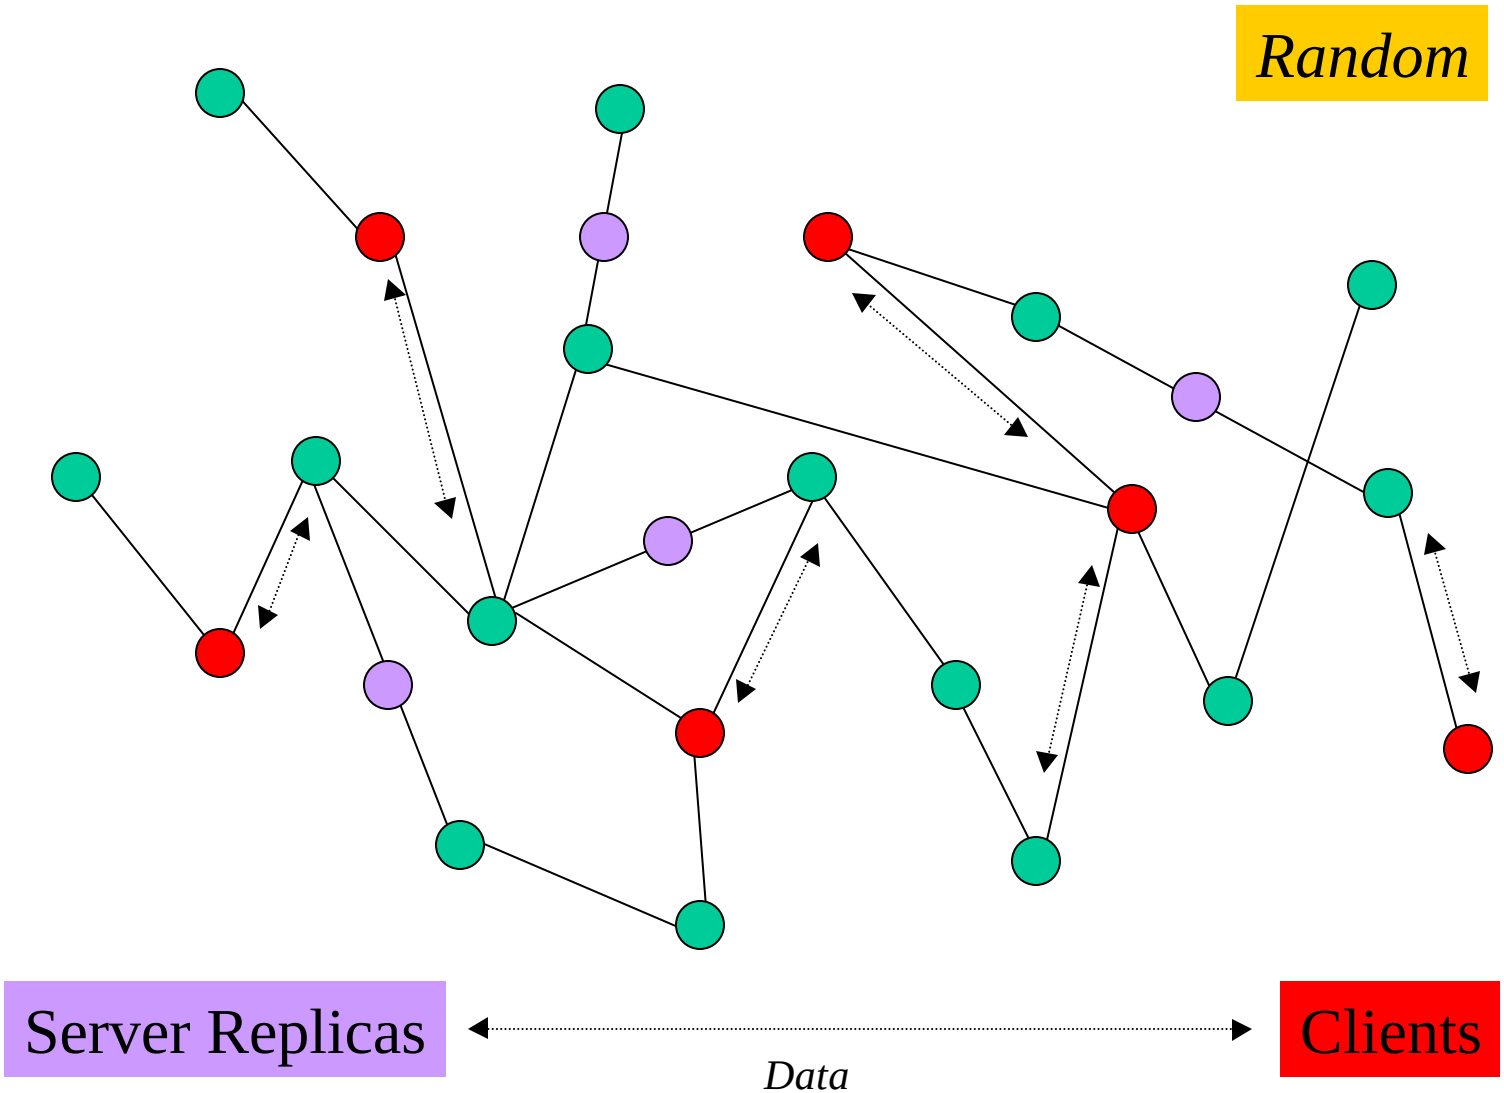

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

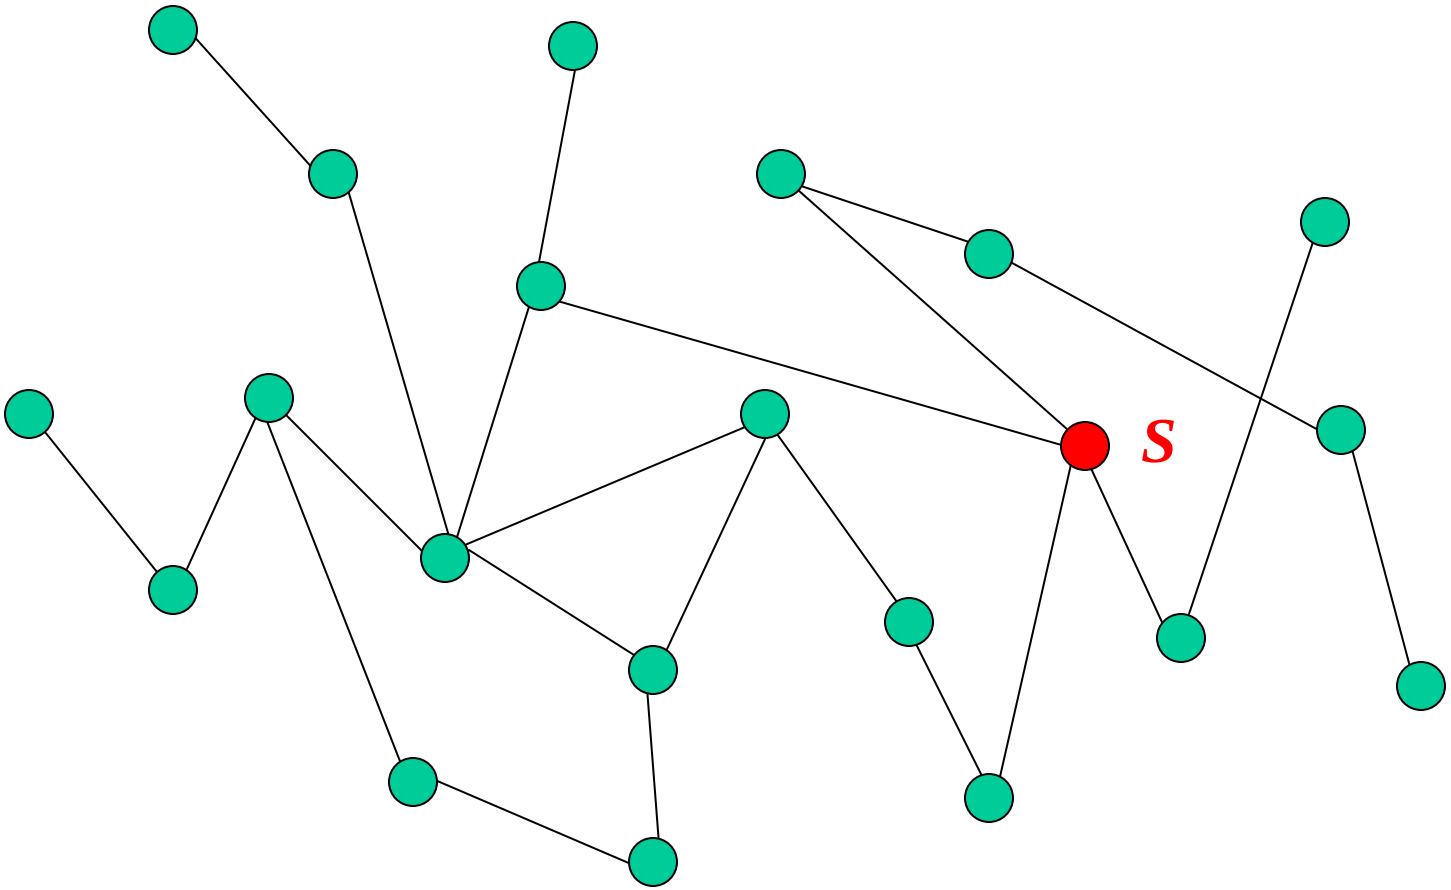

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

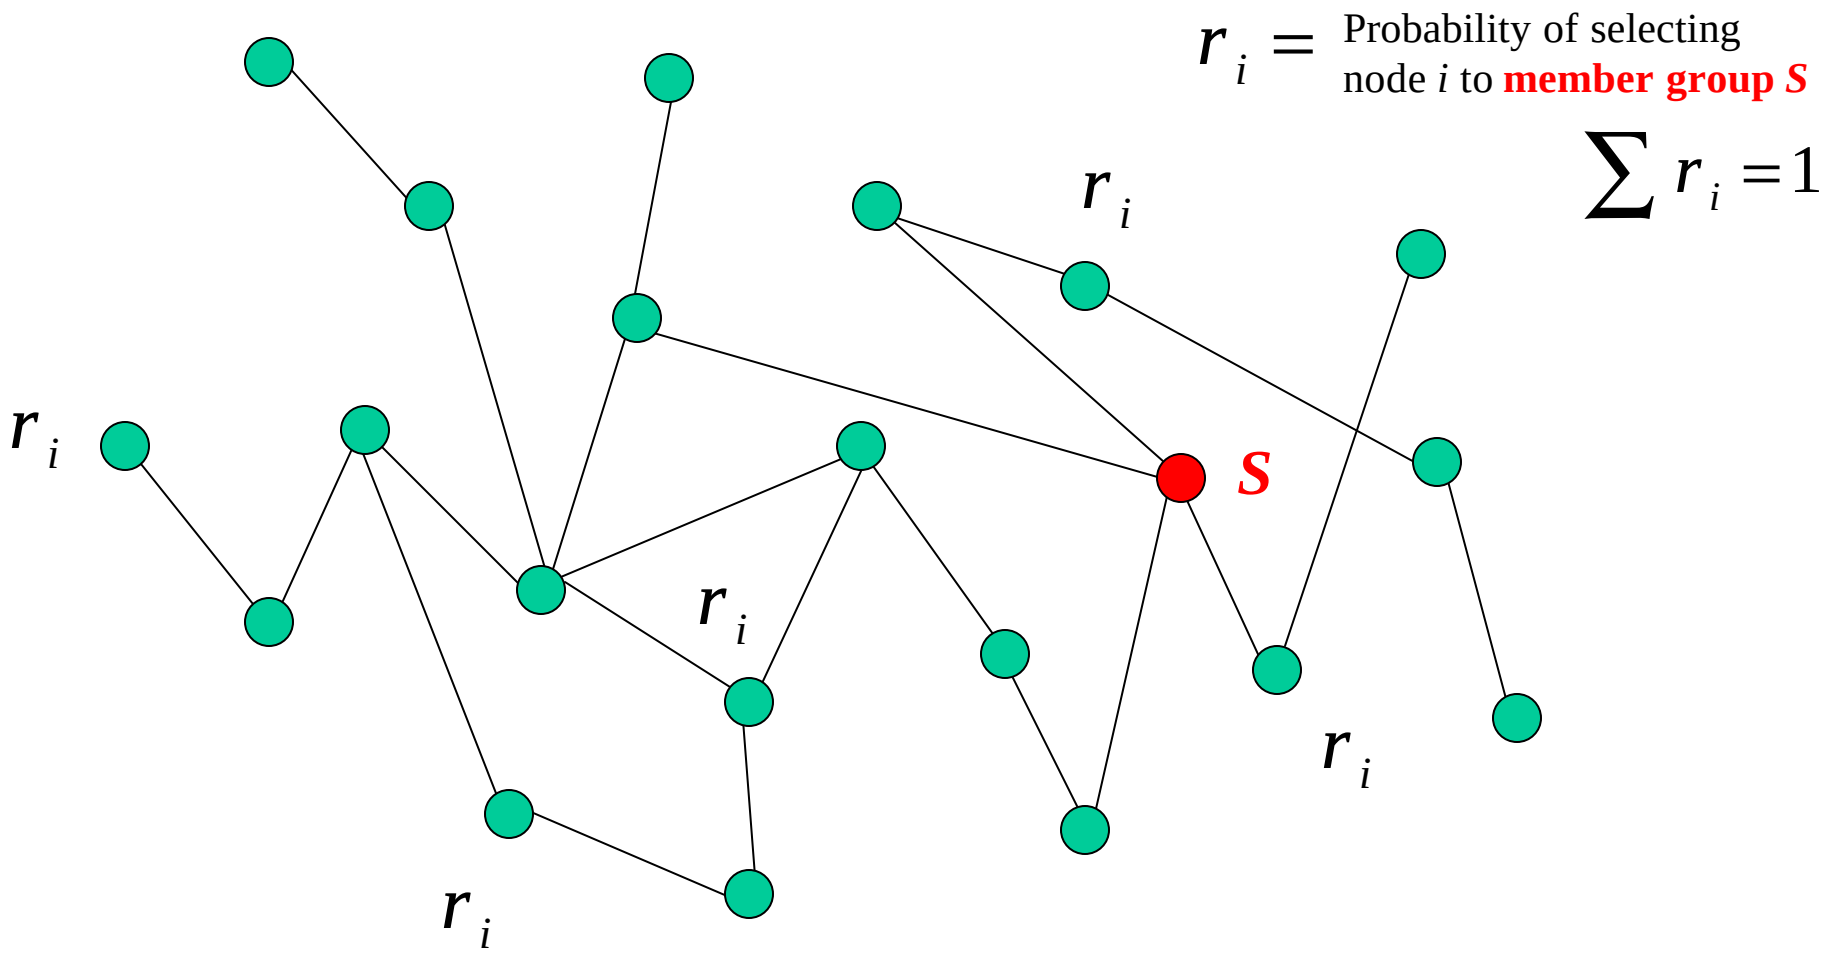

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

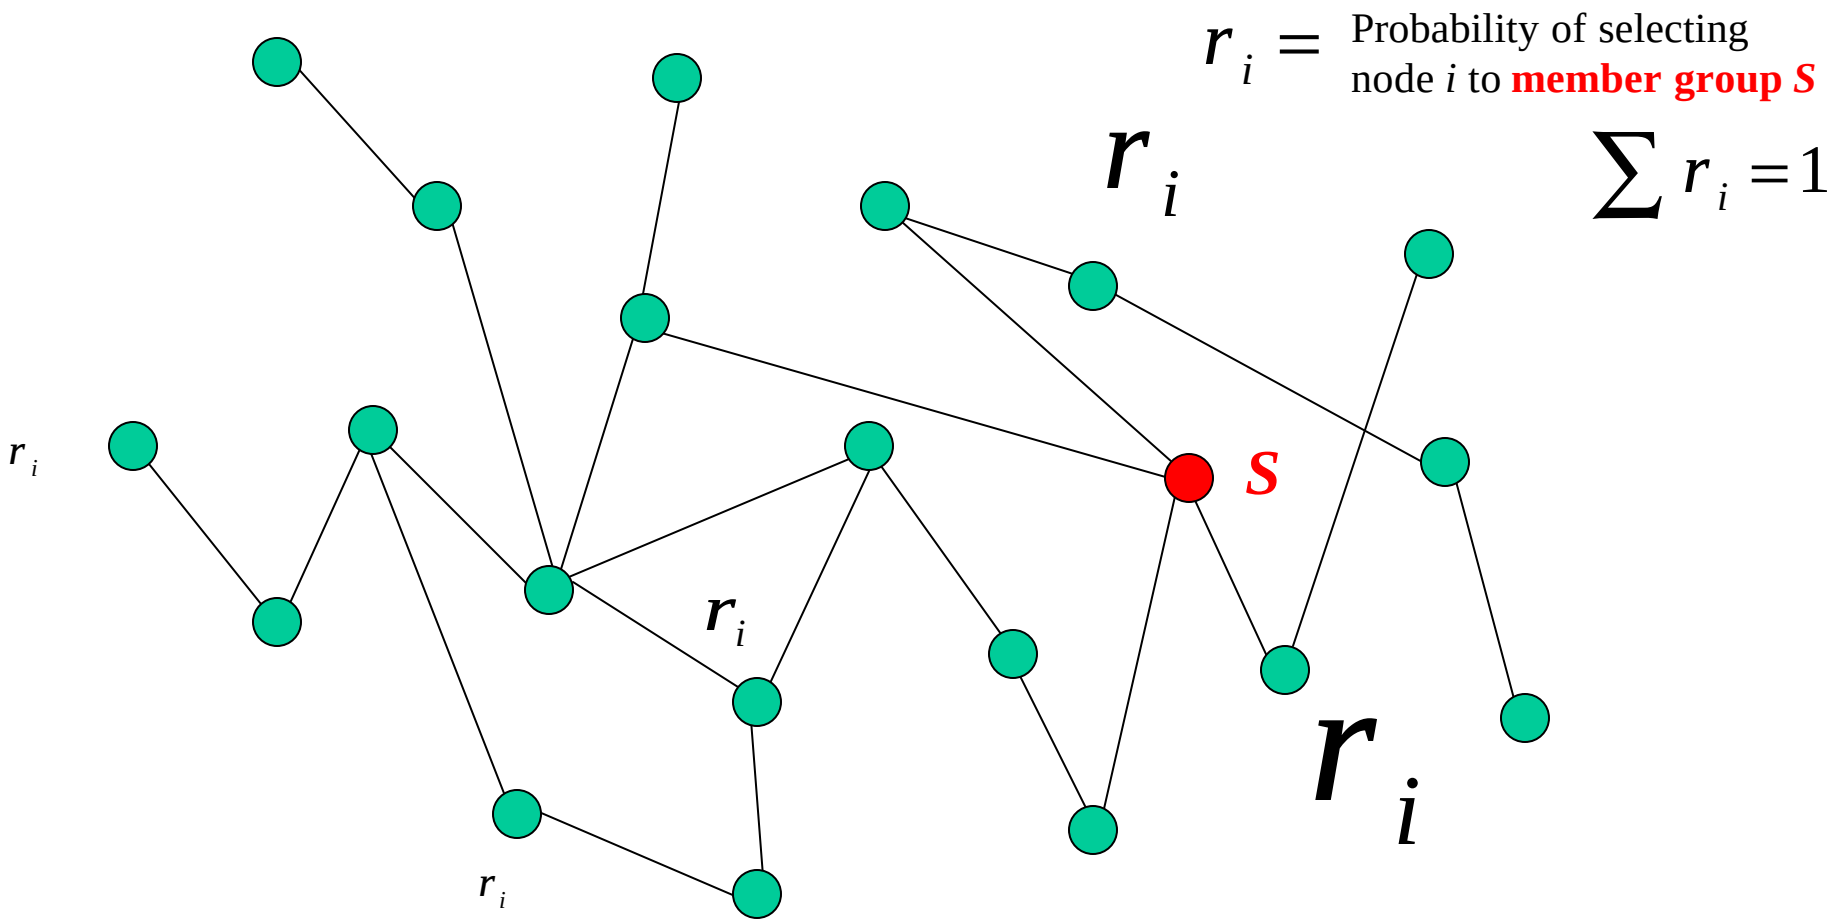

Extreme affinity

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

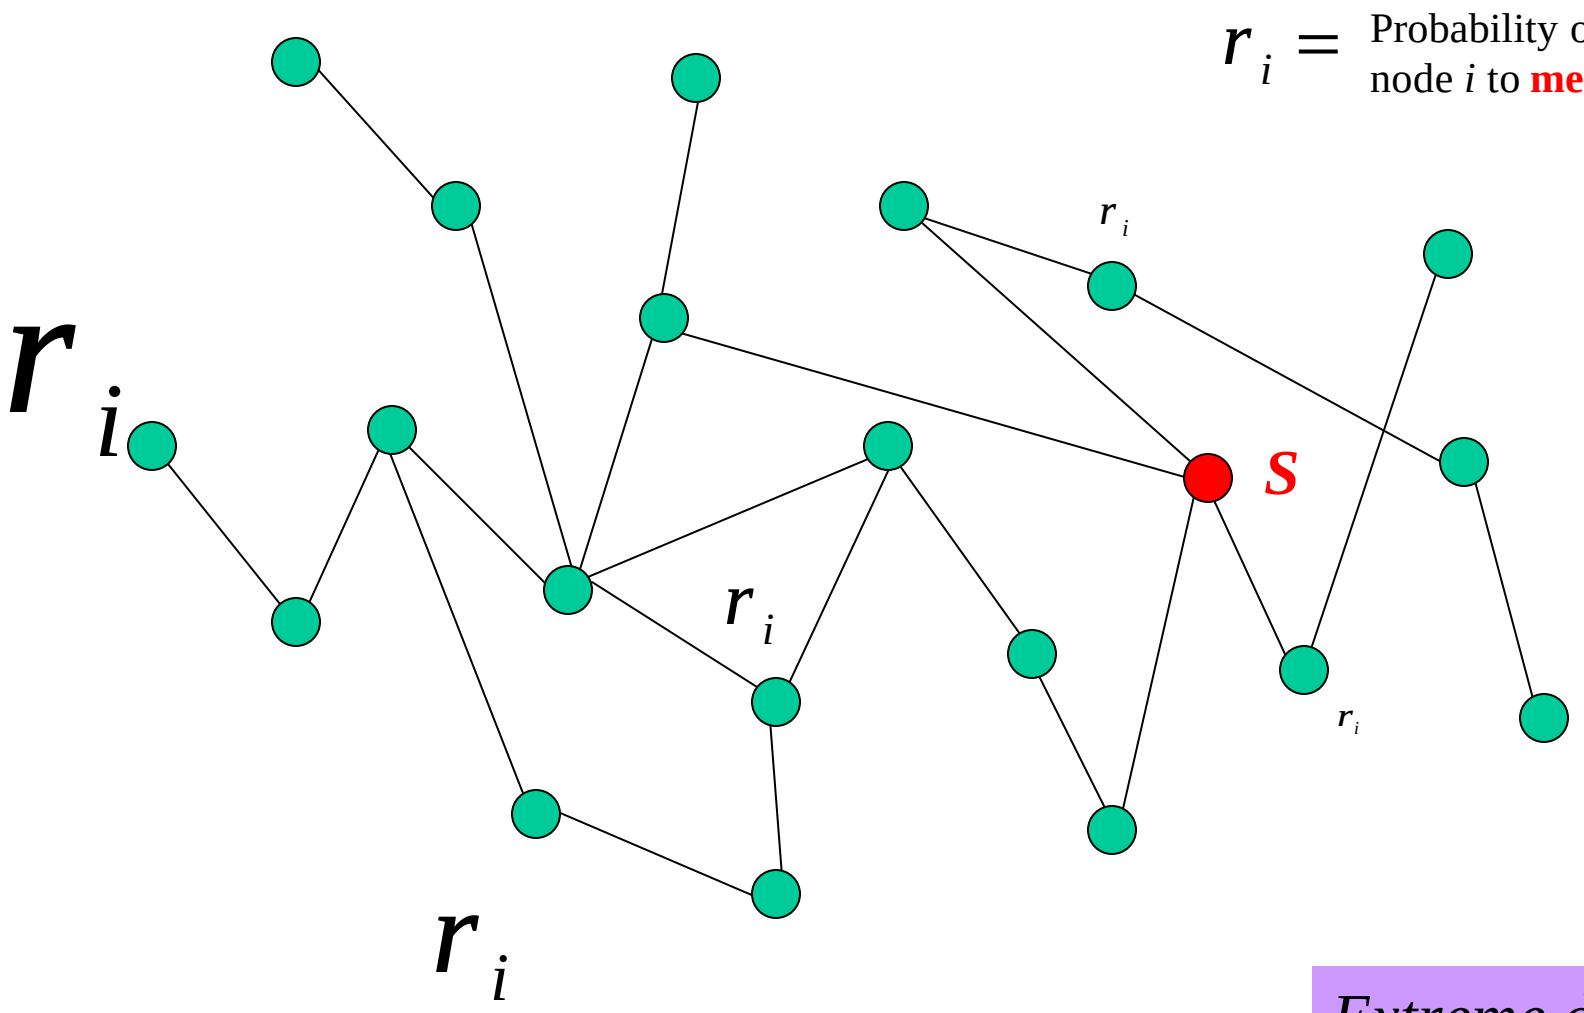

Extreme dissafinity

Affinity selection

Phillips, Shenker and Tangmunarunkit (2000)

Constant such that $\sum r_i = 1$

$$r_i = \frac{a}{(d_{iS})^b}$$

Affinity selection input parameter

Minimum distance from node$_i$ to ***S***

$b > 0 \rightarrow$ Affinity

$b < 0 \rightarrow$ Disaffinity

Major findings

Changes in affinity produce network variations

Prior work

Wong, *et al*. (2000): Multicast
Radoslavov, *et al*. (2001): Replica placement
He and Papadopolous (2000): Routing services

Characterizing affinity is possible

Affinity effects and limits are network specific

Our work

Affinity analysis improves
network performance predictions

Open questions

1. How should we analyze a member group to determine its *affinity level*?

   Can we do better than classification by b ?

2a. How does network topology affect affinity selection?

2b. Are extreme affinity and disaffinity network-specific boundaries?

3. Can affinity analysis refine network performance predictions?

Goal: Describe affinity level with a single number.

1. Development of affinity metrics

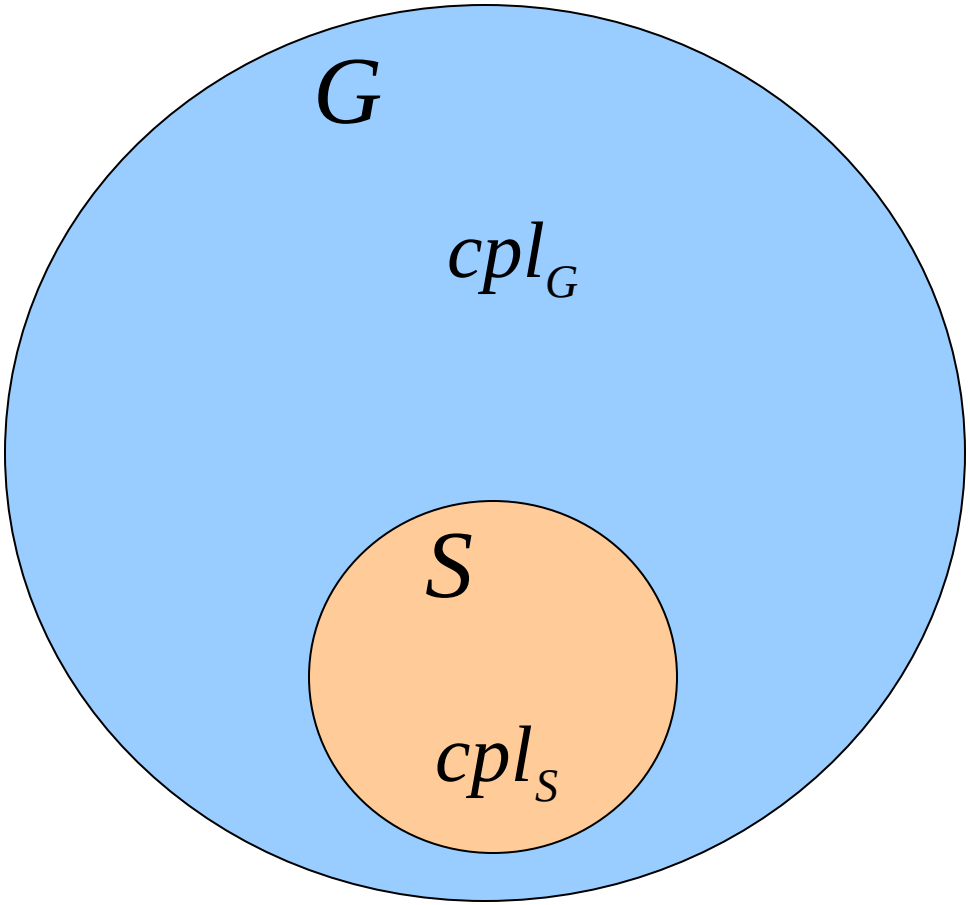

$$l = \frac{cpl_S}{cpl_G}$$

cpl = Characteristic path length = Average path length

1. Development of affinity metrics

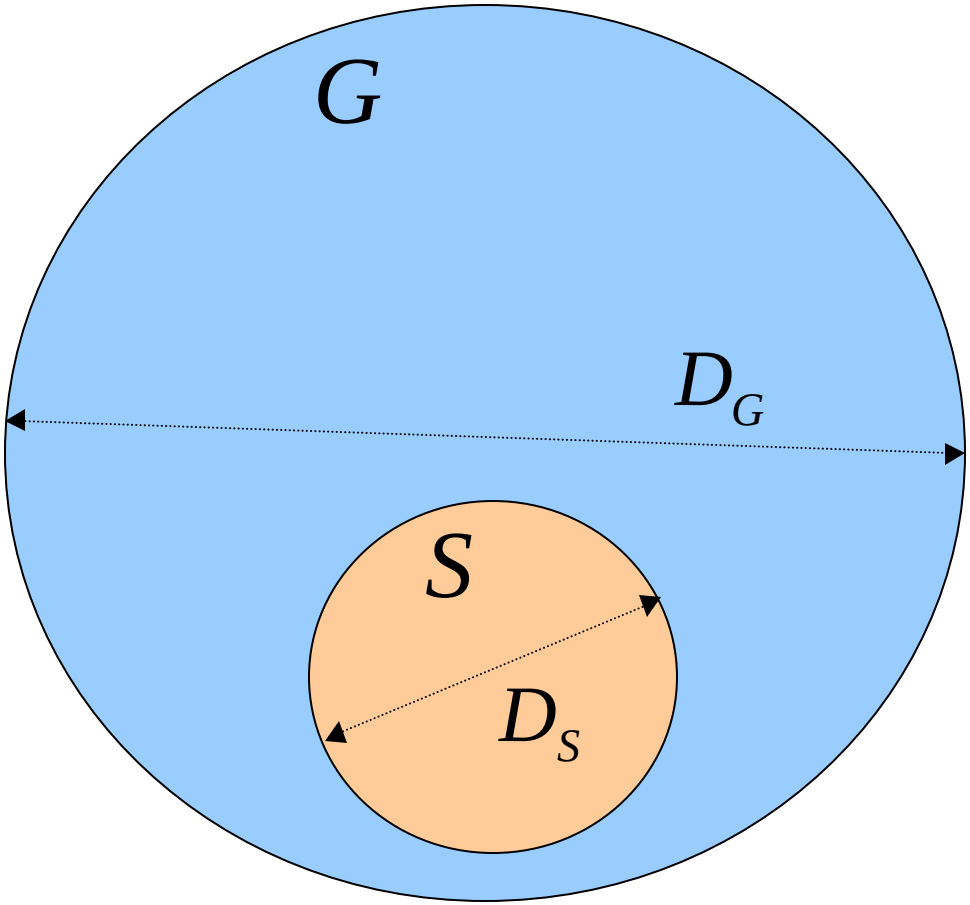

$$l = \frac{cpl_S}{cpl_G}$$

$$y = \frac{D_G}{D_S}$$

...and others

cpl = Characteristic path length = Average path length

1. Development of affinity metrics

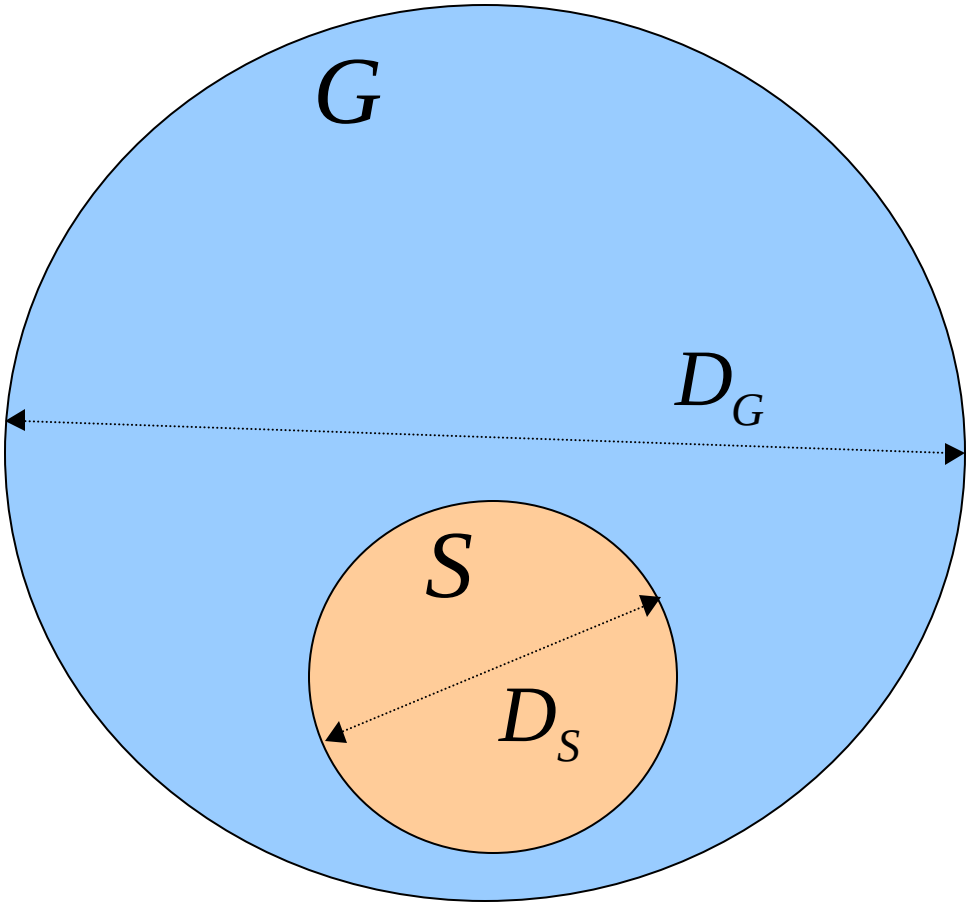

$$l = \frac{cpl_S}{cpl_G}$$

$$y = \frac{D_G}{D_S}$$

...and others

cpl = Characteristic path length = Average path length

2. Network impact on affinity

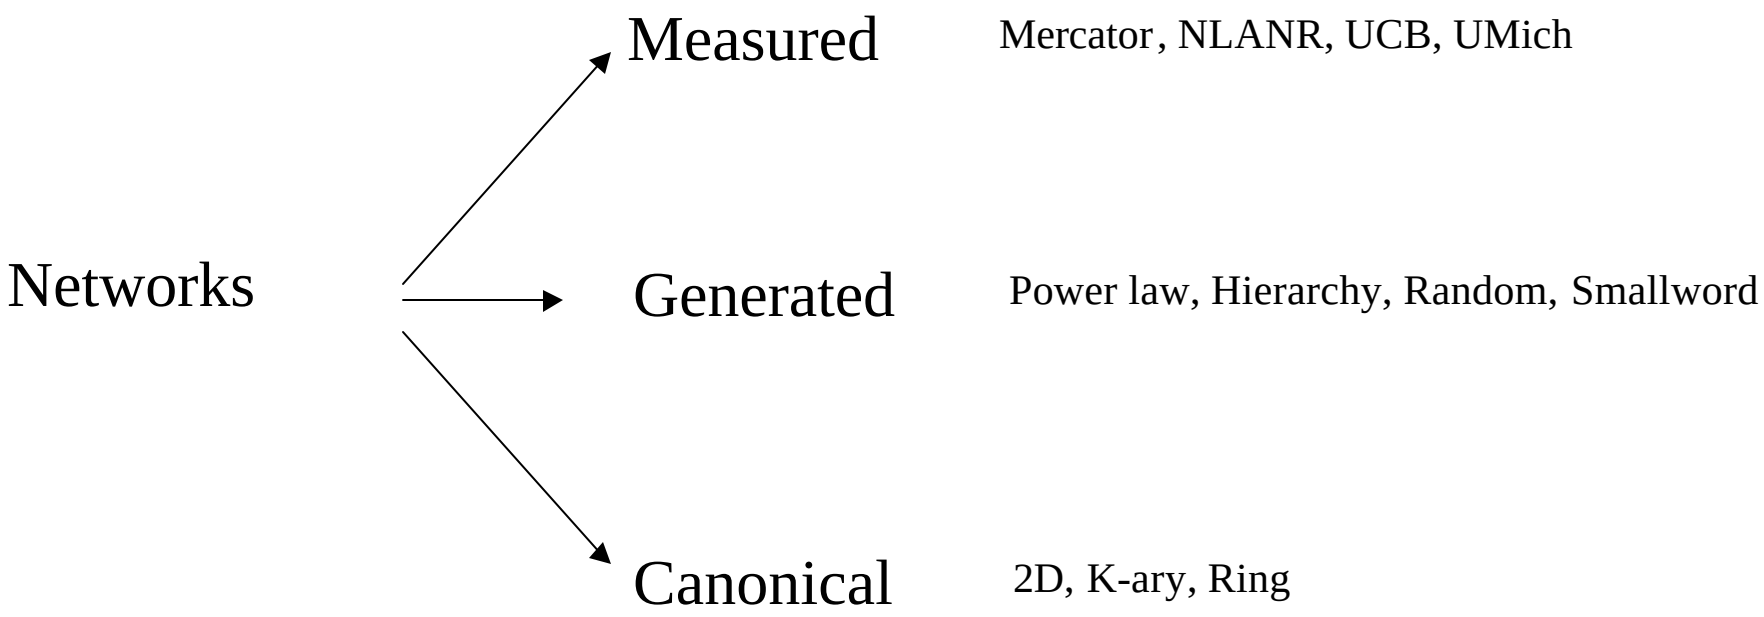

Between 5,000 and 30,000 nodes

Average degree between 2.00 and 6.00

Member groups sized between 0.1% and 10% of network size

Affinity levels between $b = 15$ and $b = -15$

2. Network impact on affinity

Member group size = 0.1% of network

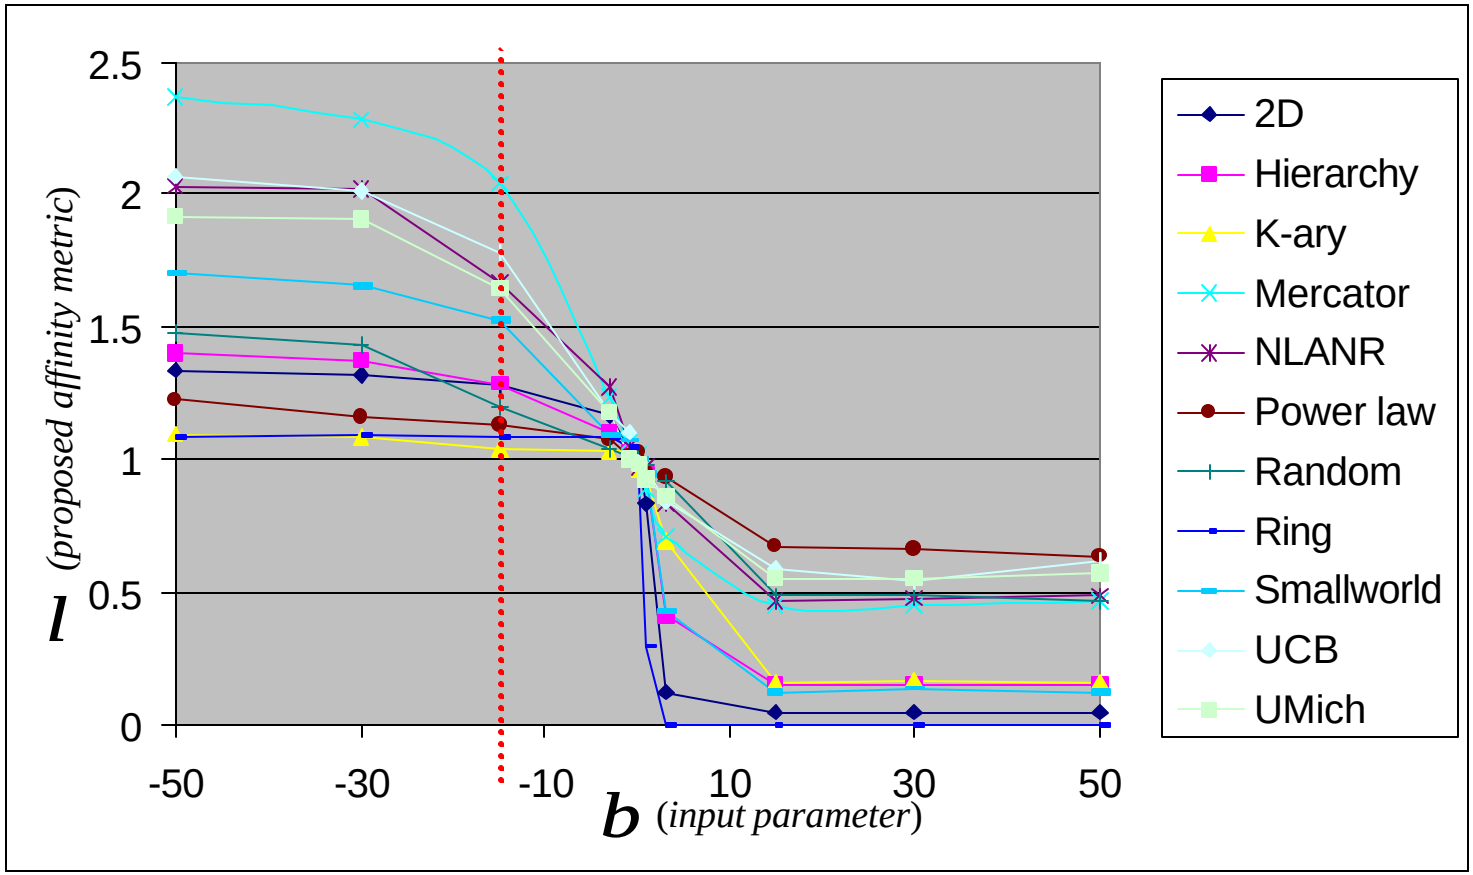

Affinity varies widely for a given value of b

2. Network impact on affinity

Member group size = 0.1% of network

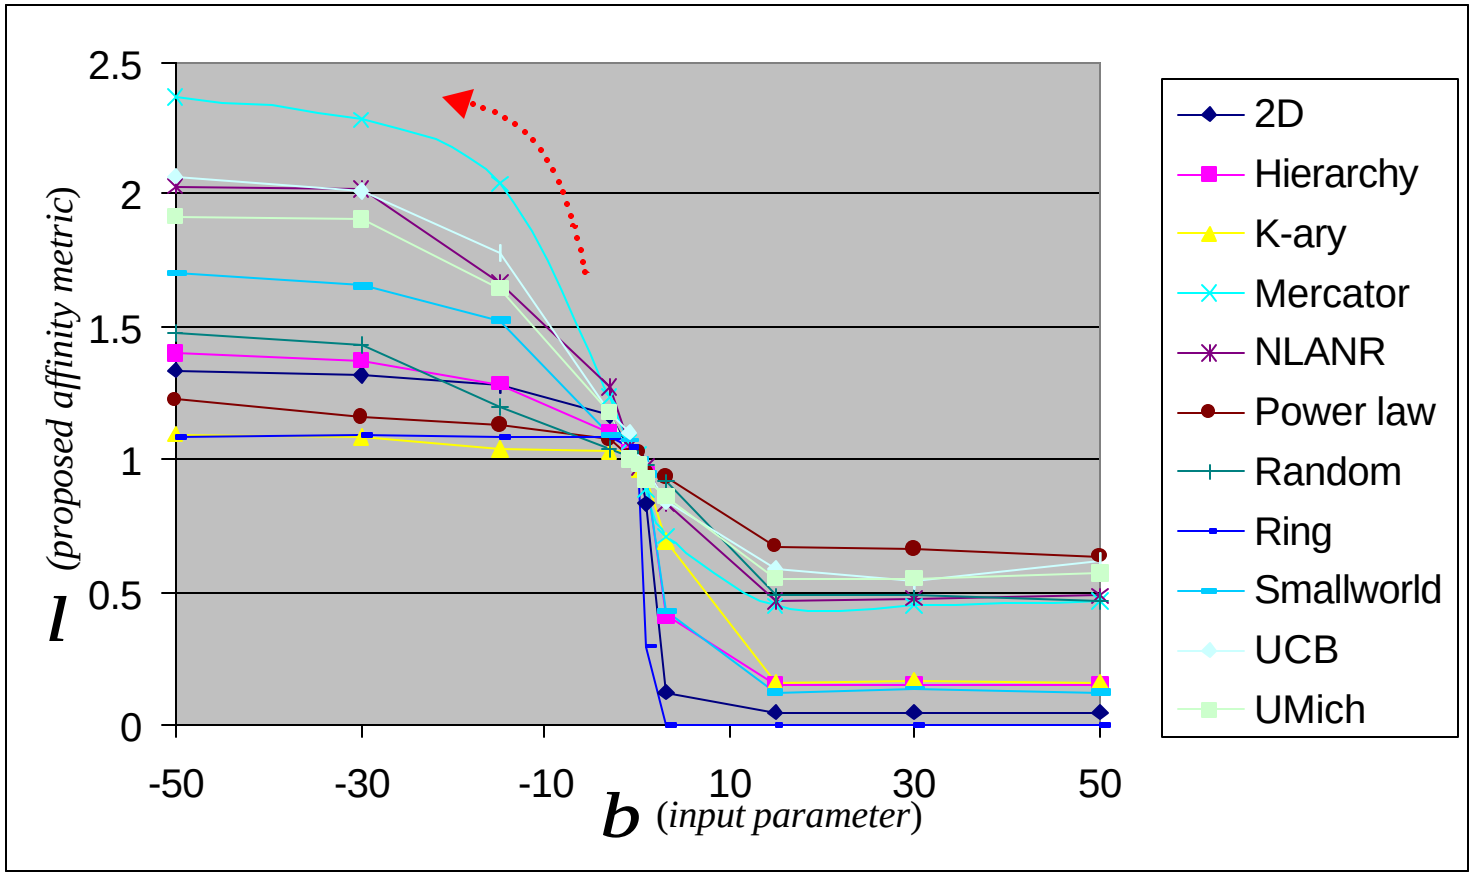

Affinity varies widely for a given value of b

$b = -15$ is not always the right lower bound

2. Network impact on affinity

Member group size = 0.1% of network

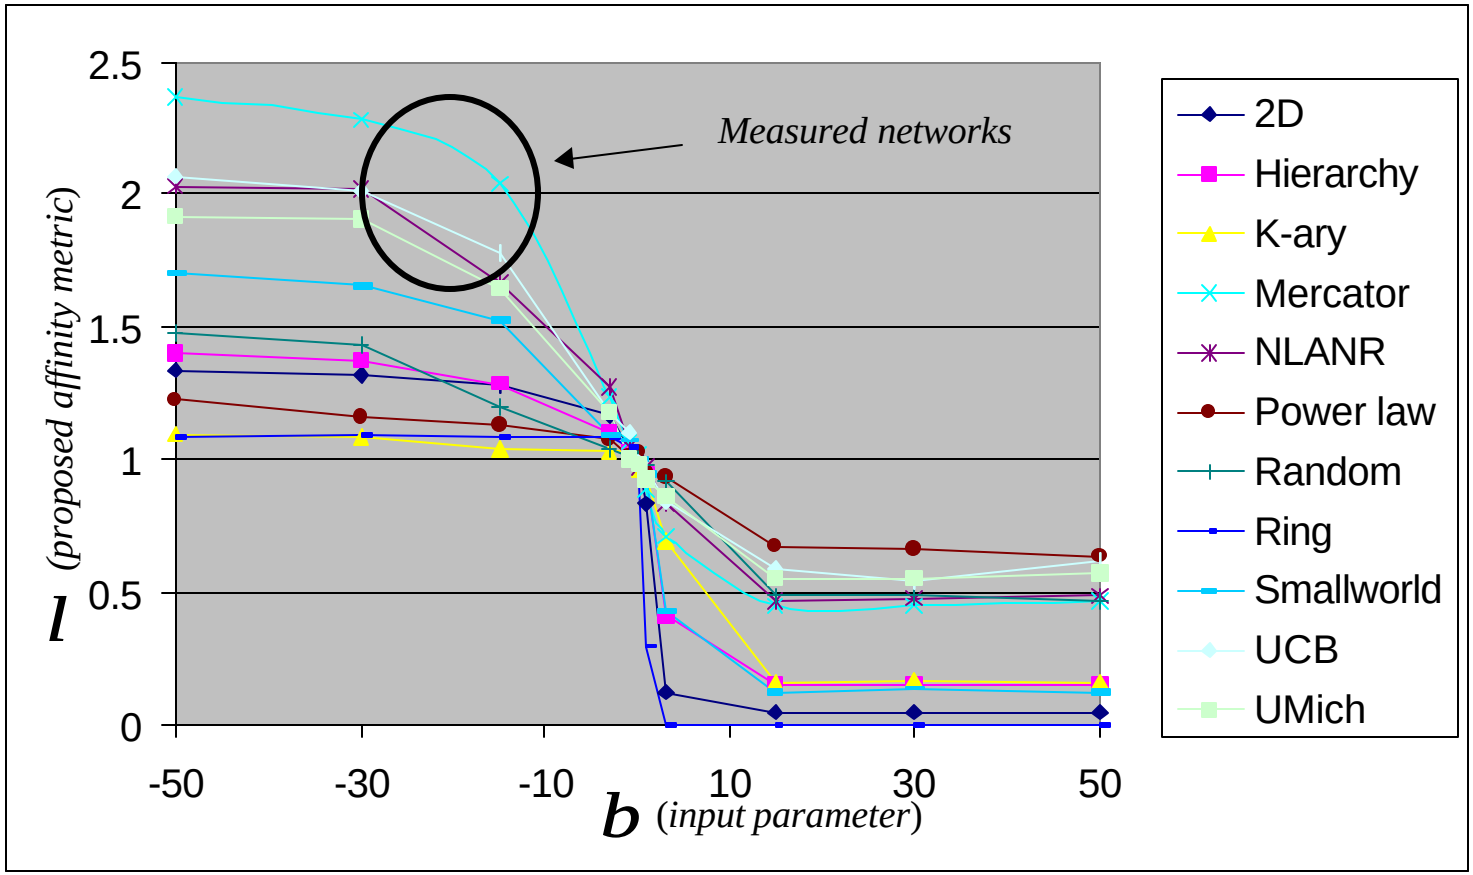

Affinity varies widely for a given value of b

$b = -15$ is not always the right lower bound

3. Case study: Multicast efficiency

Chalmers and
Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

L_m → (# hops in shortest path tree)

L_u → ($\sum$ unicast paths)

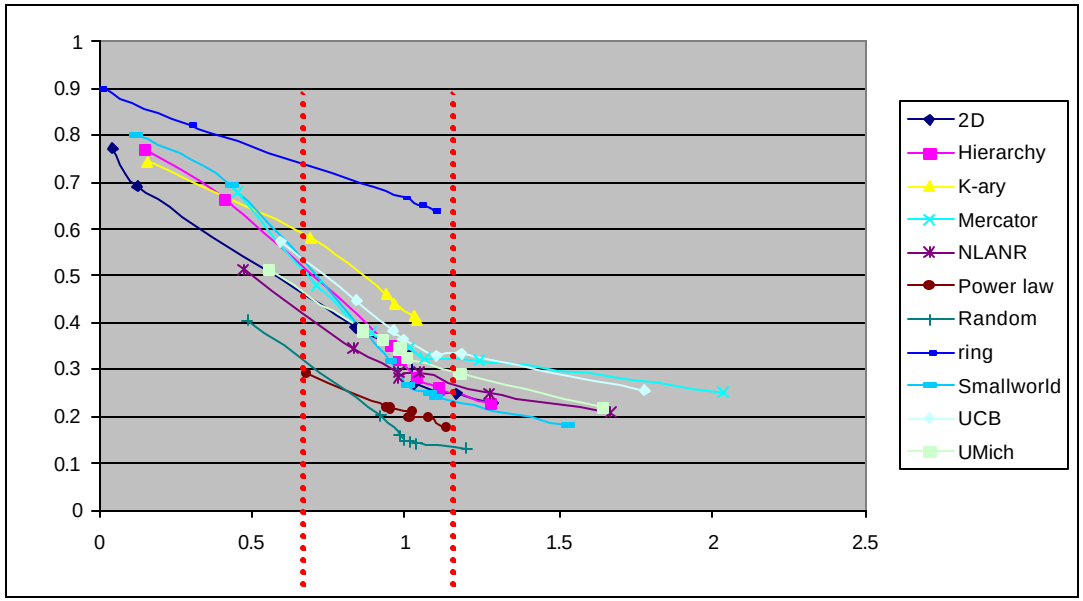

d (*multicast efficiency*)

l (*proposed affinity metric*)

Chalmers and Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

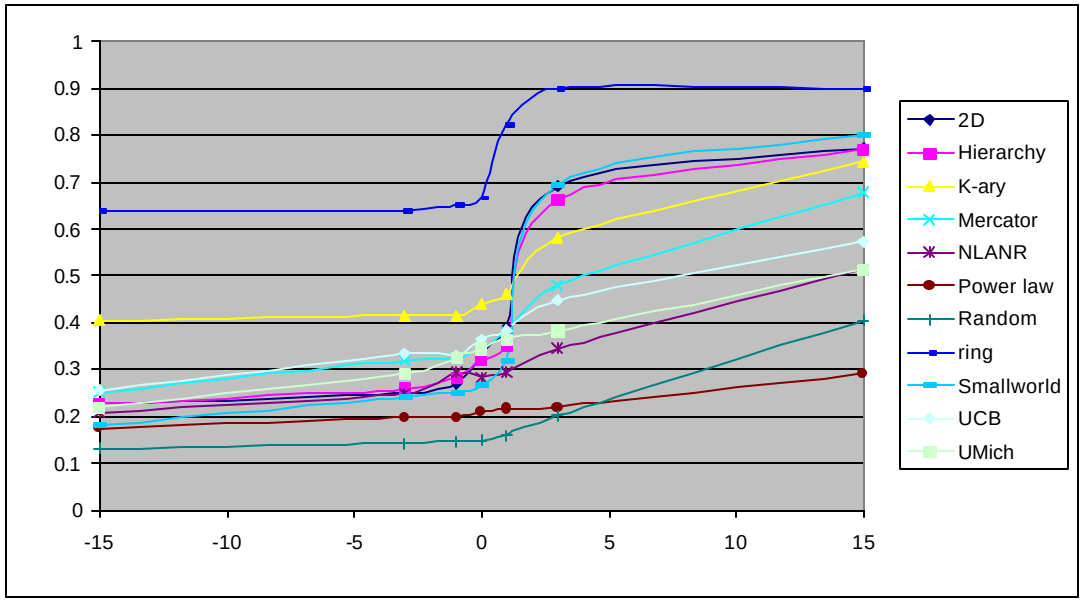

d (*multicast efficiency*)

b (*input parameter*)

Observations:

Range of l varies

magnitude

value

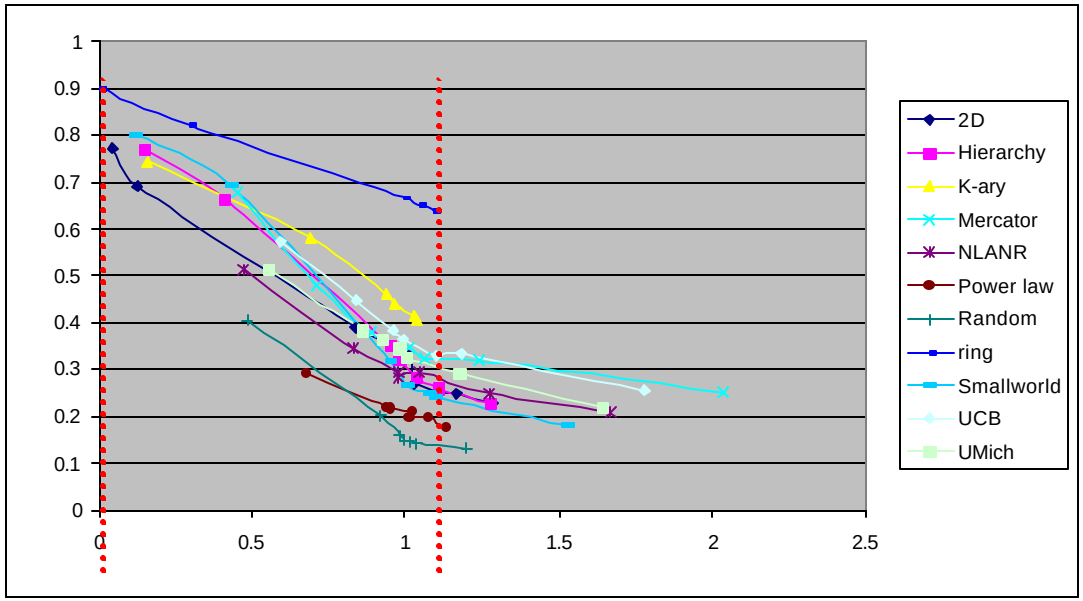

d (*multicast efficiency*)

l (*proposed affinity metric*)

Chalmers and Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

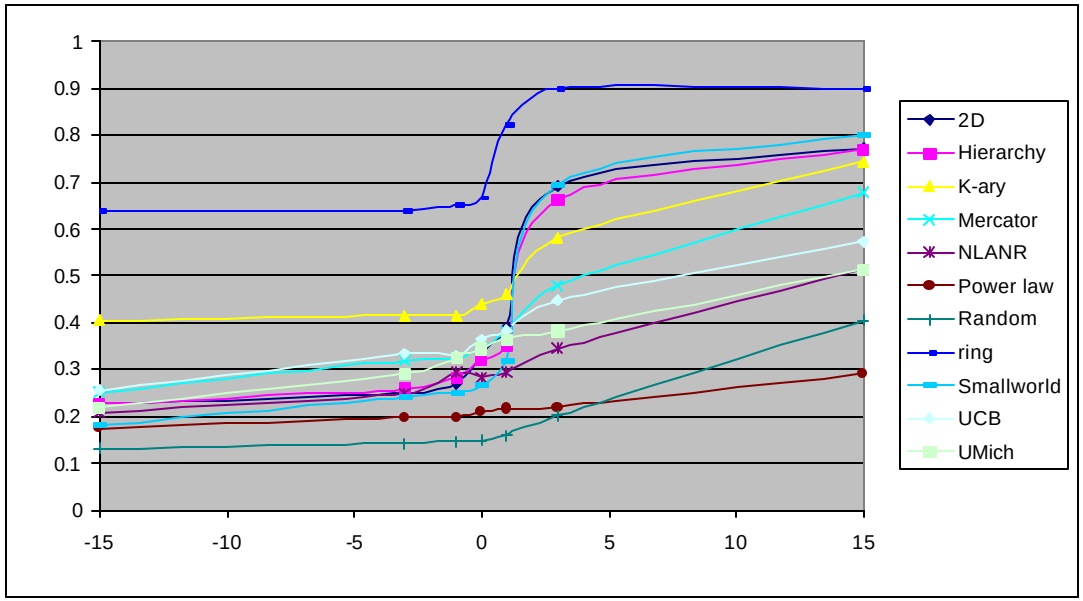

d (*multicast efficiency*)

b (*input parameter*)

Observations:

Range of l varies

magnitude

value

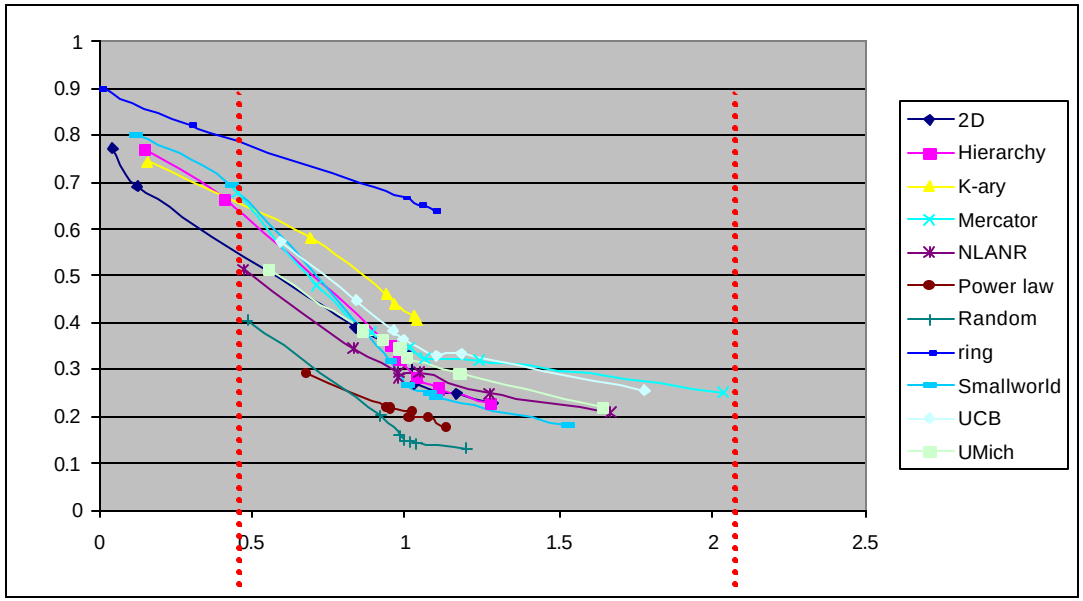

d (*multicast efficiency*)

l (*proposed affinity metric*)

Chalmers and Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

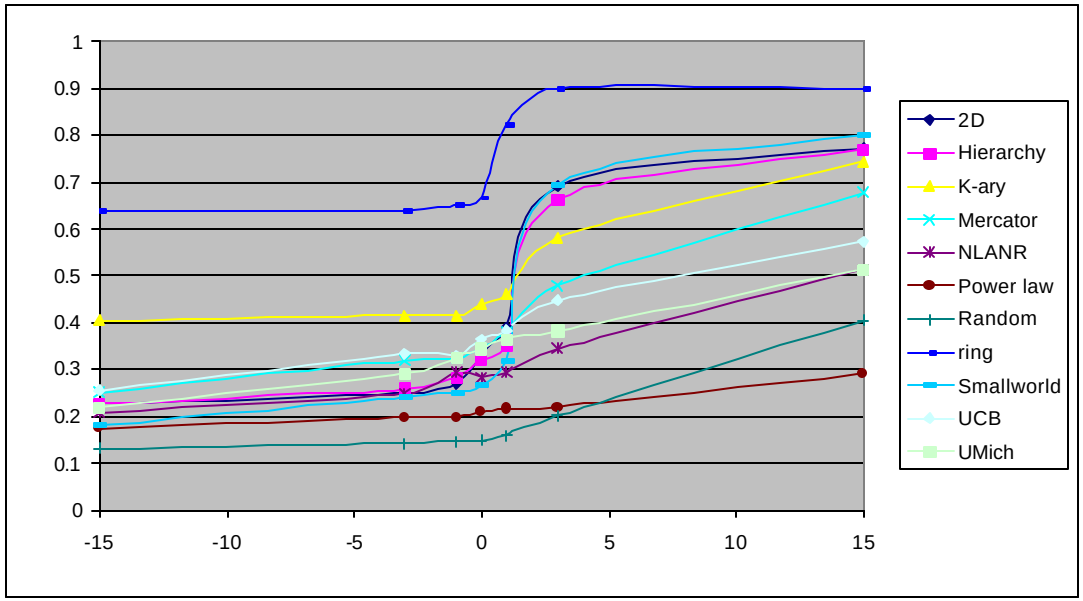

d (*multicast efficiency*)

b (*input parameter*)

Observations:

Range of l varies

magnitude

value

d (*multicast efficiency*)

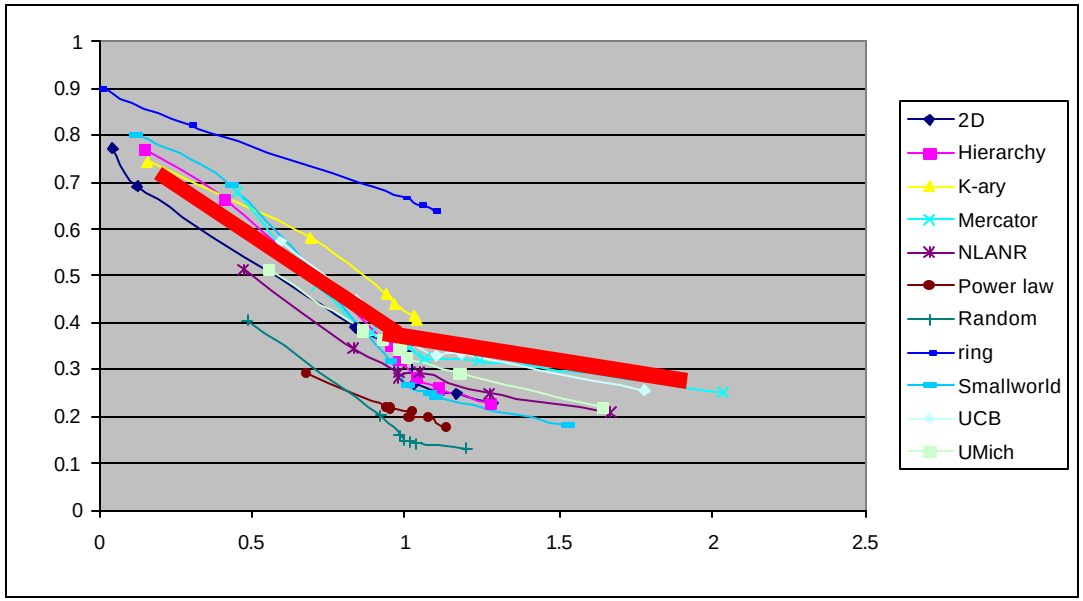

l (*proposed affinity metric*)

Chalmers and Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

Observations:

Range of l varies

magnitude

value

Knee at $l = 1$

Better linear fit

d (*multicast efficiency*)

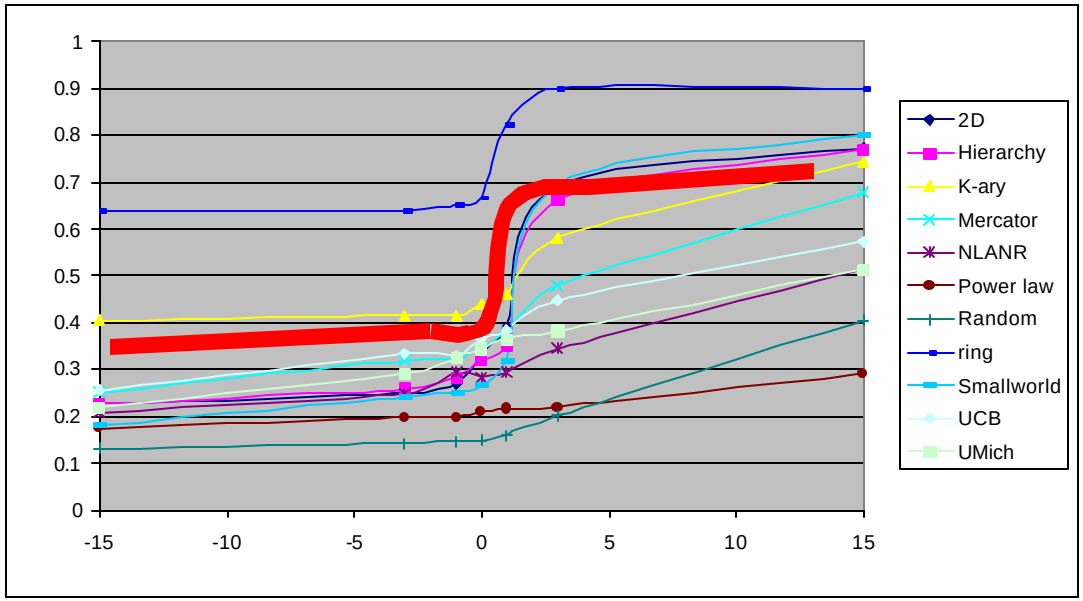

b (*input parameter*)

Chalmers and Almeroth (2001)

$$d = 1 - \frac{L_m}{L_u}$$

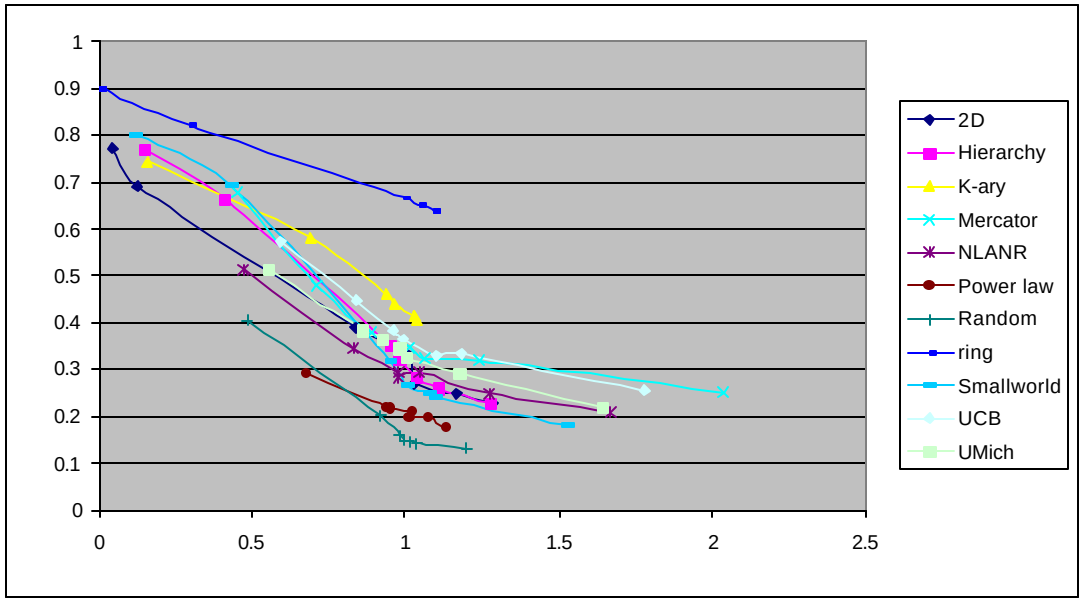

Regression:

$$d = a\,l + b$$

$r^2 = 0.5618$

$$d = a_i\,l + b_i$$

$r^2 = 0.9208$

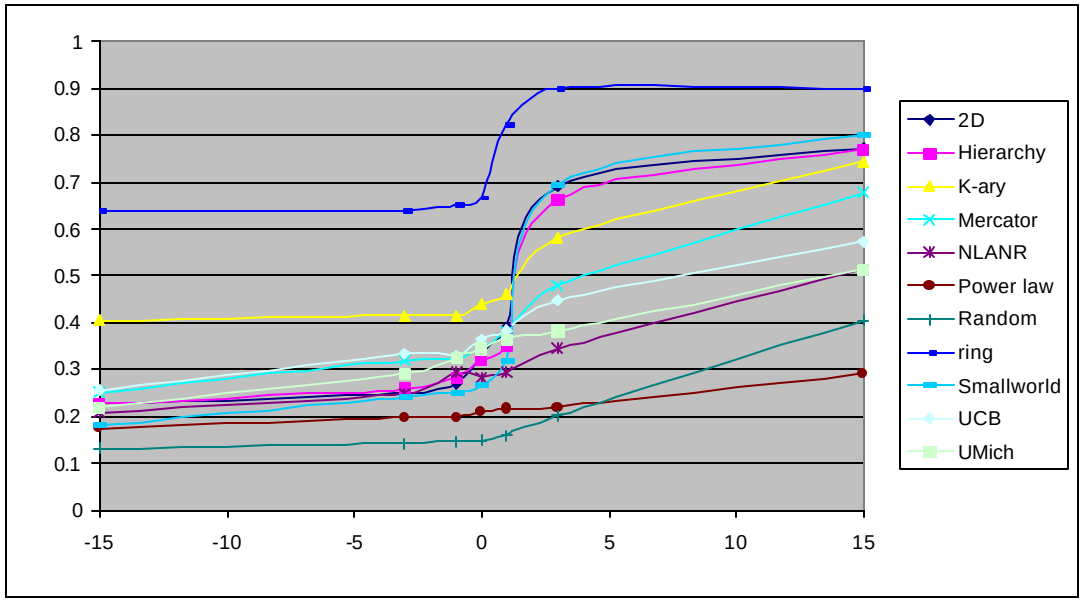

$$d = a\,b + b$$

$r^2 = 0.2375$

$$d = a_i\,b + b_i$$

$r^2 = 0.8178$

Going forward

Include affinity selection as part of simulations

- Variations in network performance
- Represents realistic scenarios

Be mindful of network topology

- Constraints on affinity selection

Examine actual affinity of subgroups

- l
- Refine and develop other metrics

Develop an affinity utility

- Analogous to topology generators
- Aid with selection and analysis