Thou Shalt Be a Selfish Overlay Neighbor

Implications of Selfish Neighbor Selection on the Design and Performance of Overlay Networks

Azer Bestavros

Computer Science Department
Boston University

Joint work with

G. Smaragdakis (BU→Deutsche Telecom), N. Laoutaris (BU→Telefonica), J. Byers (BU), P. Michiardi (Eurecom), V. Lekakis (U. of Crete), and M. Roussopoulos (Harvard→FORTH)

http://www.cs.bu.edu/groups/wing

Eurecom, Sophia-Antipolis, France
November 12, 2008

On the importance of neighbors

- Neighbor selection is a key building block for many applications – e.g., selecting
  - inter ISP peering relationships as in BGP
  - intra ISP router topology
  - neighbors in proxy caching networks
  - neighbors in P2P applications as in Bittorrent
- Performance depends largely on the quality of one's neighborhood.

Love thy neighbor as thyself…

… unless you can afford to move!

- In cyberspace, changing one's neighborhood is cheap – just rewire!
- Especially true for overlay networks.
- Implications?

Overlay networks

- Communication between any two nodes is possible; any node is a candidate as a neighbor
- Cost of overlay link could be arbitrary; e.g., delay, physical hops, distance, …

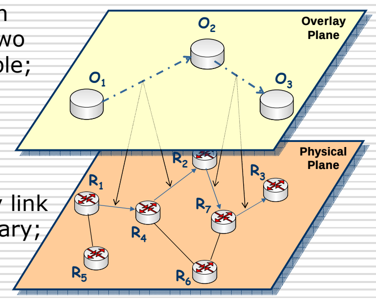

Example uses of overlays as …

- Routing Networks (e.g., Skype):
  - Send unicast traffic from one overlay node to another
  - Node's objective is to minimize its average (or maximum) routing cost to all destinations
- Broadcast Networks (e.g., MS update):
  - Send data from one node to all nodes in the overlay
  - Node's objective is to minimize its average (or maximum) broadcast cost to all destinations
- Query Networks (e.g., Gnutella):
  - Find content available in some (unknown) overlay node
  - Node's objective is to query the most number of overlay nodes using scoped flooding

Choosing thy neighborhood game

- Given an established overlay network
  - A node evaluates the advantage (if any) from picking a different set of neighbors
  - If rewiring is warranted, the node changes its (outbound) neighbors accordingly
  - This rewiring may trigger more rewiring by other nodes

and the "Selfish Neighbor Selection" (SNS) game goes on…

How we depart from prior work?

- Selfish routing†
  - Game input: Fixed network topology
  - Game outcome: Selfishly constructed source-based routes over the topology
- ➔ Our SNS work:
  - Game input: Shortest-path routing
  - Game outcome: Selfishly constructed network topology

† References: [Roughgarden & Tardos, JACM'02] [Qiu *et al*, Sigcomm'03]

SNS Game: Interesting questions

- What is the optimal strategy for playing the SNS game?
- How does it compare to empirical ones (e.g., random, nearest neighbor, …)?
- Under what conditions will neighborhoods stabilize (i.e., reach Nash-like equilibrium)?
- What do the resulting Nash-equilibrium overlay structures look like?
- What is the impact of partial/incomplete knowledge on optimal strategies?
- What is the price of anarchy?

SNS Game: Interesting questions

- What is the effect of node churn on stability and performance?
- What is the effect of changing costs due to changes in physical network?
- What if some (most) nodes are naïve? malicious? adversarial?
- How does this all scale with the size of the network?
- Could answers to the above questions inform systems/protocol design?
- …

Formulation of SNS for routing

- Notation:
  - S_{-i} is the residual wiring graph defined by the local wirings of all nodes except node v_i
  - S is the global wiring graph obtained by adding v_i's choice of neighbors s_i to S_{-i}

$$S = S_{-i} \cup \{s_i\}$$

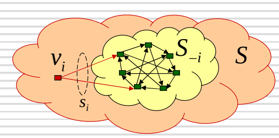

Defining an overlay neighborhood

- Assumptions by prior works
  - No cap on number of neighbors
    - Impractical – think about implications on scoped flooding in P2P, link state for routing, OS socket overheads, up-link bandwidth fragmentation, …
  - Neighbor relationships are symmetric
    - Presumptuous – communication is directed and costs are often asymmetric
- ➔ Our assumptions:
  - Nodes have a small bounded-degree $k << n$
  - Neighboring relationship is directed

Selfish Neighbor Selection (SNS)

- Players:
  The set of overlay nodes, $V=\{v_1,\ldots,v_n\}$
- Strategies:
  A strategy $s_i \in S_i$ for v_i amounts to selecting k_i outgoing overlay links; $|S_i|=(n\text{-}1 \text{ choose } k_i)$
- Outcome:
  $S=\{s_1,\ldots,s_n\}$ is the "global wiring" composed of all "individual wirings" s_i

Formulation of SNS for routing

- The objective of node v_i is to find the local wiring s_i that minimizes

$$C_i(S) = C_i(S_{-i} \cup \{s_i\}) = \sum_{v_j \in V_{-i}} p_{ij} \cdot d_S(v_i, v_j)$$

$$|s_i| \leq k$$

where

- p_{ij} is the preference of v_i for v_j as destination
- $d_S(v_i, v_j)$ is the cost of routing from v_i to v_j in S

How we depart from prior work?

- Prior work assume undirected links, unbounded degree, and uniform destination preferences
  - In [Fabrikant *et al*, PODC'03], a node may "buy" as many undirected links as it wants, each at cost α, so as to minimize the purchase + access cost

$$C_i(S) = \underbrace{\alpha \cdot |s_i|}_{\text{Purchase Cost}} + \underbrace{\sum_{v_j \in V_{-i}} d_S(v_i, v_j)}_{\text{Access Cost}}$$

  - In [Chun *et al*, Infocom'04], effect of non-uniform link costs α_{ij} is empirically evaluated.
- Appropriate for telecom networks, but not overlays; results in preferential attachment…

Neighbor selection strategies

- Best-Response (BR) is the optimal local neighbor selection strategy for node v_i:
  - BR leverages knowledge of topology and link costs of residual graph S_{-i} to minimize $C_i(S)$
- Empirical local strategies that do not use global information:
  - k-kandom does not use any link information
  - k-closest uses only local information

BR for SNS (for routing) is NP hard

- Theorem:
  Under uniform overlay link weights (e.g., hop-count), finding the BR to S_{-i} is equivalent to solving the asymmetric k-median on S_{-i} with reversed distance cost

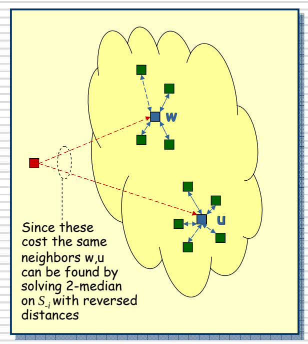

BR for SNS (for routing) is NP hard

- Theorem:
  Under uniform overlay link weights (e.g., hop-count), finding the BR to S_{-i} is equivalent to solving the asymmetric k-median on S_{-i} with reversed distance cost
- Corollaries:
  - BR is NP hard; constant approximation for *metric* k-median does not apply
  - $O(1)$-approximation with $O(\log n)$ blow-up in number of medians [Lin and Vitter, '92] is possible

Game theoretic results for SNS†

- Theorem:
  All games with uniform node preference, node degree, and link costs have pure Nash equilibria (stable graph).
  - In any such stable graphs, the cost of any node is at most $2 + k^{-1} + O(1)$ that of any other node.
  - The diameter of the stable graph for a uniform game is $O(\mathrm{sqrt}(n \log_k n))$.
- Theorem:
  There exist non-uniform games with no pure Nash equilibria.

† Proofs, constructions, and more results in *Laoutaris, Rajaraman, Sundaram, Teng, "A bounded-degree network formation game"*, from arXiv-CoRR cs.GT/0701071.

Empirical evaluation of SNS (routing)

- Obtain BR wiring for SNS game as follows

```
start with an arbitrary wiring ;
until wiring is stable or within threshold {
   for each node v_i {
      BR(v_i) ← heuristically† solve asymmetric k-median;
   }}
```

➔ Vary problem inputs/parameters and evaluate resulting wirings *w.r.t.* topological features, individual node cost, and overall social cost

† Two heuristic implementations:
- ILP using Simplex method (Cplex Tomlab toolbox)
- Local search (with r-link swap, $r = 1, 2, \ldots, k$; $O(n^r)$ *complexity*)

Results under complete uniformity

- Under unit link costs and uniform routing preference to all destinations, we know that a Nash-equilibrium exists.
- What are the characteristics of the resulting wiring graphs?
  - Are they random?
  - Do they exhibit a uniform in-degree distribution?

Results under complete uniformity

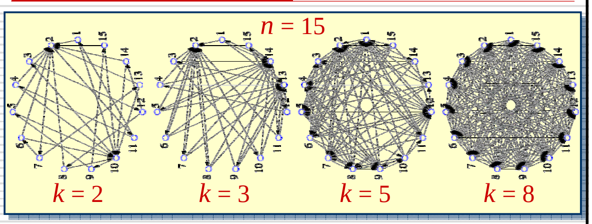

- Not uniform, but skewed in-degree distribution
- Selfishness yields preferential attachment to "accidentally" popular nodes
- Phenomenon more evident for small k/n – why?

Effect of skewed routing preference

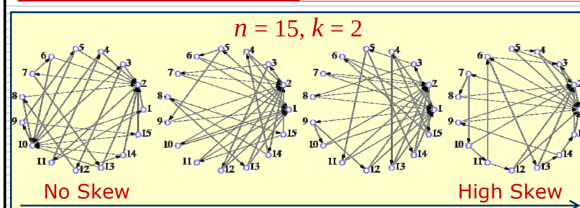

- Preferential attachment to "inherently" popular nodes satisfies selfishness' need for popular nodes for small k
- What happens with larger k ?

The two sources of in-degree skew

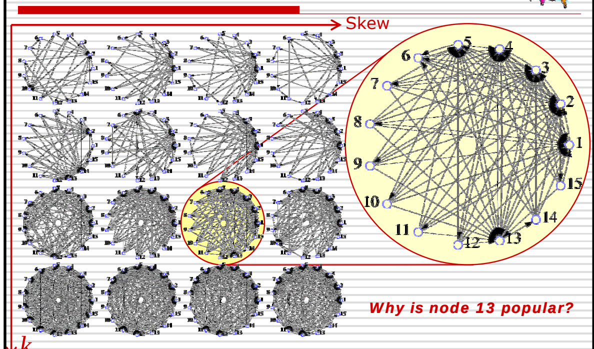

Effect of heterogeneous link costs

- Link cost generation
  1. Synthetically using BRITE:
     - Barabasi-Albert (BA) model with heavy-tailed 2D placement
     - Euclidean distance used to derive cost of overlay links
  2. Empirically from PlanetLab:
     - 300-node PlanetLab topology
     - All-pair ping traces used to derive cost of overlay links
  3. Empirically from AS-level maps
     - 12/2001 Rocket-Fuel data of the Internet topology
     - AS-level hop-count used to derive cost of overlay links
- Control parameter
  - Bound on out-degree (k) ≈ link density (β)

Experimental setting

- **Neighbor selection strategy**
  a. The k-random heuristic
  b. The k-closest heuristic, a.k.a. greedy
  c. SNS Best Response (BR) wiring using ILP
- **Experiments done in nine permutations**
  - Three strategies for a new comer, each assuming residual graph was wired using one of the three strategies
- **Performance metrics**
  - Individual Cost = Average cost for a newcomer
    - ➔ Cost ratio for strategy x = C(x)/C(BR)
  - Social Cost = Sum of cost for all nodes
    - ➔ Social Cost ratio for strategy x = SC(x)/SC(BR)

SNS over random residual networks

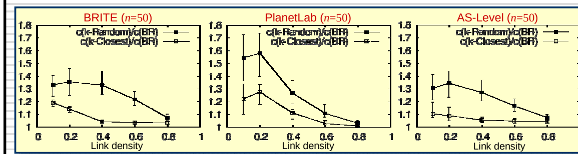

➔ BR is dominant, with k-closest decidedly better than k-random. BR's benefit pronounced for small k – why?

If your neighbors are naïve, it pays to be selfish

SNS over greedy residual networks

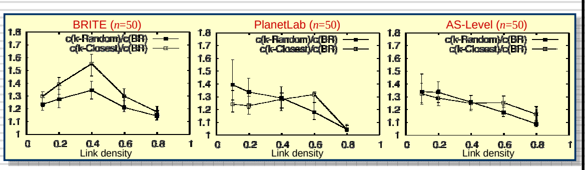

➔ BR is dominant, with k-random slightly better than k-closest – why?

If your neighbors are greedy, it pays to be selfish

SNS over selfish residual networks

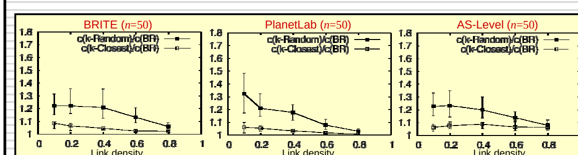

➔ BR is dominant, but not by a significant margin, with k-closest being quite competitive – why?

If your neighbors are selfish, it's OK to be naïve

Social cost benefit from SNS

	$\beta = 0.1$		$\beta = 0.2$	
	k-Random/BR	k-Closest/BR	k-Random/BR	k-Closest/BR
BRITE	1.44	1.53	1.52	1.84
PlanetLab	2.23	1.48	1.75	1.23
AS-level	2.04	1.90	1.83	1.61

➔ Adopting BR as a neighbor selection strategy results in a significant reduction in the social cost (by 30-60%) over naïve (random/greedy) approaches.

The network is better off with selfish nodes!

Almost Utopia!

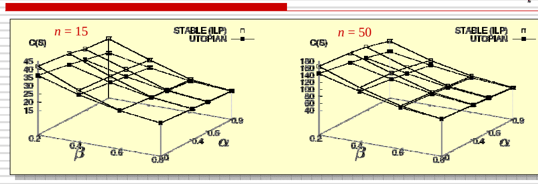

➔ Not much difference between social cost of SNS wiring and that of a Utopian wiring over wide ranges of preference skew and link density.

The network is almost a utopia with selfish nodes!

EGOIST: SNS prototype

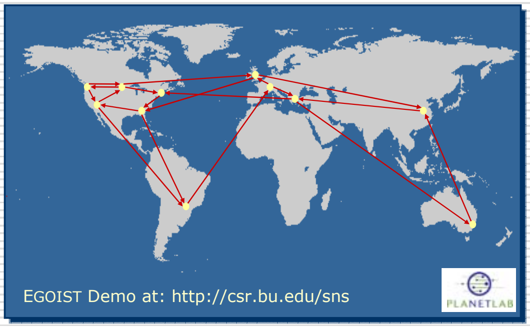

EGOIST: Implementation

Protocol for EGOIST overlay node v_i

1. Bootstraps by connecting to arbitrary neighbors
2. Joins link-state protocol to get residual graph
3. Measures cost to candidate neighbors
4. Wires according to chosen strategy (default: BR)
5. Monitors and announces overlay links

† We have also implemented a light-weight version of this protocol, in which steps 2, 4, and 5 are implemented on a central server.

EGOIST: Features

- **Supported metrics:**
  - Delay (actively/passively monitored with **ping**/**Pyxida**)
  - Available bandwidth (monitored with **pathChirp**)
  - Node load (monitored with **loadavg**)
- **Supported wiring strategies:**
  - *k*-random
  - *k*-closest
  - *k*-regular
  - Best-Response (Delay and AvailBw formulations)
  - Hybrid Best-Response (subset of links donated to the network)
- **BR Computation:**
  - By using the full residual graph
  - By sampling the residual graph

EGOIST: Baseline results (n=50)

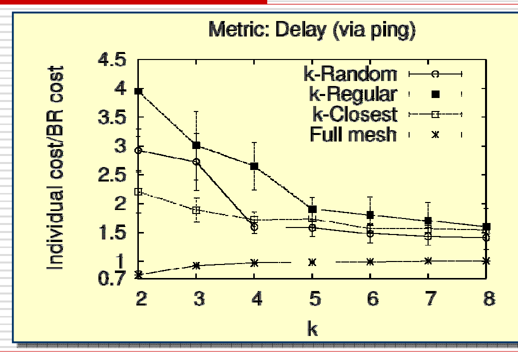

EGOIST: Active *vs.* passive

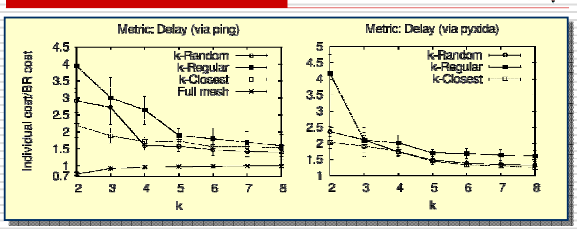

- Passive approaches deliver comparable results (across strategies) with much less overhead!
- Greedy indistinguishable from random; regular

EGOIST: Other metrics

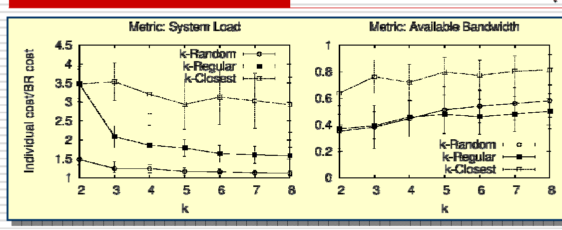

- Significant gains possible with BR
- Greedy's performance is lagging – why?

EGOIST: Re-wiring frequency

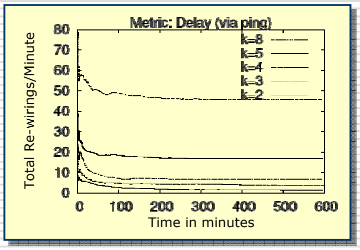

- Overlay fairly stable, especially for small k
- Re-wirings increase quite rapidly with k – *why?*

EGOIST: Marginal utility of re-wiring

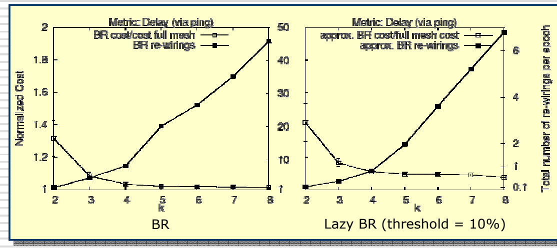

- Most of the benefit achieved with k ~ 3-4
- Re-wirings could be reduced using "lazy" BR

EGOIST: Effect of churn

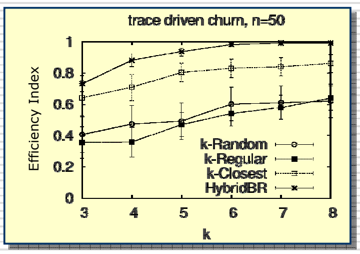

- HybridBR delivers much of the efficiency of BR
- Greedy strategy less susceptible to churn than random and regular strategies

EGOIST: Effect of churn

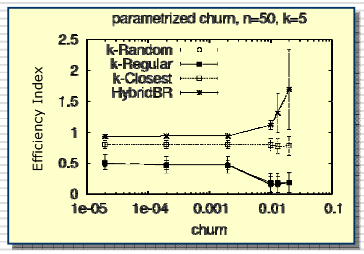

- BR dominates non-BR wirings strategies
- At very high churn, using HybridBR pays off

EGOIST: Vulnerability to abuse

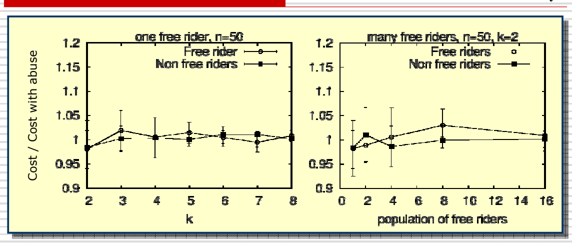

- Free riders avoid being chosen as neighbors by inflating cost of their outgoing links (*2 above)
- EGOIST is robust to abuse by free riders (not the case with greedy neighbor selection)

EGOIST: Effect of partial knowledge

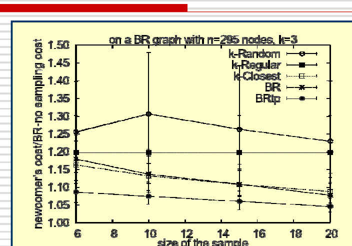

- Sampling rate affects BR and greedy strategies
- Topology-based biased random sampling significantly improves BR's performance

Other SNS objectives

- Routing Networks (e.g., Skype):
  - Send unicast traffic from one overlay node to another
  - Node's objective is to minimize its average (or maximum) routing cost to all destinations
- Broadcast Networks (e.g., MS updates):
  - Send data from one node to all nodes in the overlay
  - Node's objective is to minimize its average (or maximum) broadcast cost to all destinations
- Query Networks (e.g., Gnutella):
  - Find content available in some (unknown) overlay node
  - Node's objective is to query the most number of overlay nodes using scoped flooding

The *n*-way broadcast problem

- Each node needs to send a file to all others
  - Exchange of large scientific data-sets in grid computing
  - Distribution of traffic log files for network-wide IDS
  - Synchronization of distributed databases
  - Distributed backup
- Use swarming to reduce link stress
  - How do we create the underlying torrent topology?
  - Could SNS lead to better overlay on which to swarm?
  - What would constitute a selfish objective?
    - ➔ Maximize the average bandwidth over all nodes
    - ➔ Maximize the minimum bandwidth across all nodes

Swarming over SNS overlays

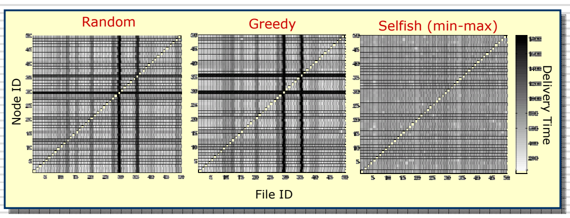

Thou shalt swarm over selfishly-constructed overlays!

Query routing over SNS overlays

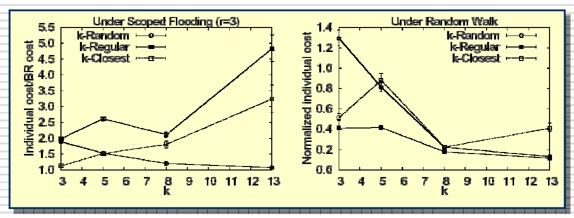

Thou shalt query over selfishly-constructed overlays!

Take home messages

- Performance of overlays depends highly on neighbor selection strategy
- Framing neighbor selection as a strategic game yields highly optimized overlays
- Implementing SNS is practical and yields overlays that are robust to churn/abuse

➔ Papers, demos, traces, and code available from **`http://csr.bu.edu/sns`**

Publications

"EGOIST: Overlay Routing using Selfish Neighbor Selection"
Georgios Smaragdakis, Vassilis Lekakis, Nikolaos Laoutaris, Azer Bestavros, John W. Byers and Mema Roussopoulos.
ACM CoNEXT 2008.

"Swarming on Optimized Graphs for n-way Broadcast"
Georgios Smaragdakis, Nikolaos Laoutaris, Pietro Michiardi, Azer Bestavros, John W. Byers and Mema Roussopoulos.
IEEE INFOCOM 2008.

"Implications of Selfish Neighbor Selection in Overlay Networks"
Nikolaos Laoutaris, Georgios Smaragdakis, Azer Bestavros and John W. Byers.
IEEE INFOCOM 2007.