

# Incentive Compatibility in Diffuse Computing: Results and Future Directions

Supported by the DoD URI  
Program under ONR  
grant N00014-01-1-0795

Speaker:

Joan Feigenbaum

<http://www.cs.yale.edu/homes/jf>



# Why Incentive Compatibility Matters in Diffuse Computing

- Shift in focus from platform to network
- Previously “independent” actors are now part of a “continuously adapting” computational ecosystem.
- *Strategic* choices are important for adaptation or even survival in this ecosystem.

[Gebrowski & Garstka '98] [SPYCE proposal '00]

# Sample SPYCE Accomplishments on Incentives

- Multicast cost sharing [Yale Berkeley Stanford]
  - Welfare maximization is easy.
  - Budget balancing is hard.
- Economics of Anonymity Systems [NRL *et al.*]
  - Free riding can be beneficial.
  - Price discrimination doesn't work well.
- Incentivizing cooperation in Ad Hoc Networks [Yale]
  - Cryptography prevents cheating in *Sprite* system.
  - Performance can suffer if batteries are low.
- Computational capacity of information markets [Yale and NEC]
  - Weighted threshold functions of  $n$  variables are computable in  $n$  rounds.
  - No other functions are computable in this market model.

# New SPYCE Results on Incentives

- Spectrum sharing
- Interdomain routing
- Exchange-based mechanisms
- Paradoxical value of privacy

# Spectrum Sharing

In 802.11 wireless network:

- Each service provider (SP) owns a set of access points (APs).
- Each AP is assigned a channel (frequency).
- Limited number of channels available
- Neighboring APs must use different channels to avoid interference.

How should channels be assigned and by whom (e.g., by SPs or by the FCC)?

# Spectrum Sharing as a Game

[Haldorson, Halpern, Li, & Mirrokni; PODC '04]

SPs take turns assigning channels to APs.

Then they can bargain and trade channels.

Question: How far is a Nash Equilibrium (NE) of the game from an optimal assignment (OPT) of channels?

$$\text{PRICE OF ANARCHY} \triangleq \frac{\# \text{ users w/access in OPT}}{\# \text{ users w/access in NE}}$$

# Results Derived from Graph Coloring


Theorem 1 [HHLM]: If users are uniformly distributed and there are  $k$  channels, then

$$5 \leq \text{PoA} \leq 5 + \max\left(0, 1 - \frac{5}{k}\right)$$

Theorem 2 [HHLM]: If 2-buyer-1-seller bargains are allowed, then

$$3 \leq \text{PoA} \leq 3 + \max\left(0, 1 - \frac{3}{k}\right)$$

# Interdomain Routing



Agents: Transit Autonomous Systems (ASes)

Inputs: Routing Costs or Preferences

Outputs: Routes, Payments

# Previous SPYCE Results on IDR

- **Lowest-cost routing** [FPSS; PODC '02]  
Strategyproof, BGP-compatible mechanism
- **Policy routing** [FSS; PODC '04]
  - **Arbitrary preferences**  
Strategyproof mechanism, but NP-hard even to approximate
  - **Next-hop preferences**  
Strategyproof mechanism, polynomial-time computable, but not BGP-compatible

# New Formulation: Subjective-Cost Policy Routing

[Feigenbaum, Karger, Mirrokni, & Sami; Sept. 2004]

- Each AS  $i$  assigns a cost  $c_i(k)$  to every other AS  $k$ . "Subjective cost" of path  $P_{ij}$  is

$$c_i(P_{ij}) \triangleq \sum_{k \in P_{ij}} c_i(k)$$

- No req. that  $c_i(\cdot)$  and  $c_j(\cdot)$  be consistent
- Mech.-design goal: Minimize  $\sum_i c_i(P_{ij})$ .
- Generalization of many natural routing policies, e.g., lowest cost and "forbidden set"

# Results and Interpretation [FKMS]

- Unfortunately, almost all variants of **subjective-cost** routing that we studied, including **forbidden-set** policies, are NP-hard, even to approximate.
- Partial positive result if  $c_i(\cdot)$  has the form  $c_i(P) \triangleq \lambda_i \ell_1(P) + (1 - \lambda_i) \ell_2(P)$ , *i.e.*, a convex combination of two objective measures (think **price** and **latency**)
- Conjecture: Good distributed algorithms depend on significant **agreement** among ASes about which routes are good.

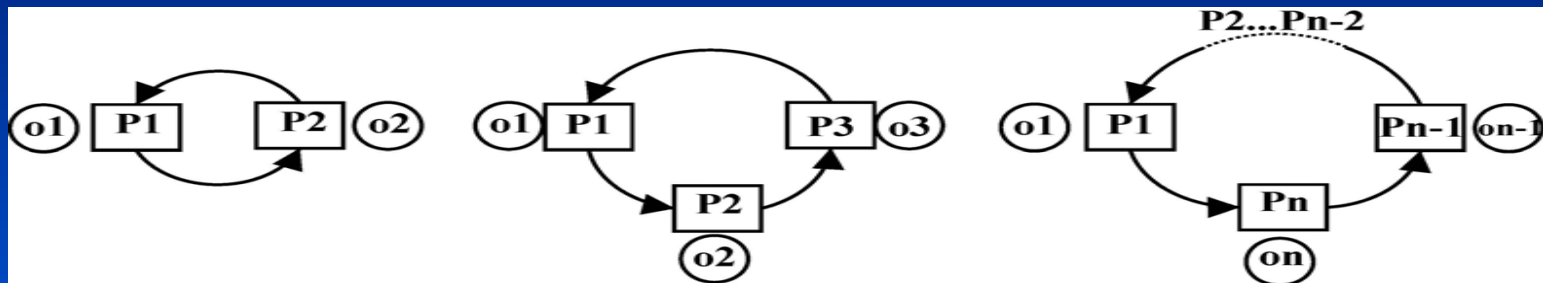
# Exchange-based Mechanisms

[Anagnostakis & Greenwald; ICDCS '04]

- Goal: Incentivizing agents in diffuse systems
  - Extremely limited "trusted third party" capabilities
  - For example, no micropayment services (yet?)
- Hypothesis: Sometimes, exchange (aka "barter") can support robust and efficient incentive mechanisms.
  - Evidence: Popularity of BitTorrent

# First Experiment

Cheat-proof N-way exchange rings in P2P file sharing



- Simulation and real-world measurement show that almost all file transfers can be served in rings of size 5.
- Efficient, simple, and robust
- Reveal flaws in BitTorrent and other fielded systems.

Next step: N-way bandwidth barter, focusing on wireless and neighborhood-area networks

# The Paradoxical Value of Privacy

[Syverson, EIS '03]

- People *say* that they value their privacy.
- People *act* in ways that compromise their privacy.
- Standard explanations  
[Varian, Odlyzko, Acquisti, *etc.*]
  - Principal-agent problems
  - Technology trends (*e.g.*, plummeting storage costs)
  - Immediate, tangible benefits vs. delayed, intangible benefits
- Killer app. for privacy reform: ID theft
- How *should* things work?

# Financial Scenario

- You say you're Bob, and I give you credit.
- You default.
- I say Bob defaulted.
  - Giving you credit was my choice, not Bob's.
  - If you're Bob, you should incur liability.
  - If not, I should incur liability. I'm the one who screwed up, damaging Bob's reputation.
  - Advantage: I can make the choice to authenticate, limit your credit, or assume greater risk.
  - Insurance industry will do the rest?

# Criminal Scenario

- You are arrested for a crime and claim to be Bob.
- I (*e.g.*, local court) assign arrest to Bob.
- You abscond.
- I say that Bob absconded.
  - *My report* of arrest and absconding should be subject to penalties of wrongful arrest.
  - So should any agency propagating my report.

# Conclusions

- ID theft is primarily an authentication problem and a reputation problem. It is *not* primarily a confidentiality problem.
- Current focus is on protecting the confidentiality of SSNs and other PII.
- Current costs imposed by the faulty infrastructure are borne by the victim.
- Benefits and costs of authentication and of reputation reporting should be borne by those who run the authentication protocols and report reputations.