

# **Game-Theoretic Aspects of Diffuse Computing: Results and Future Directions**

**Joan Feigenbaum**

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

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

# Two Views of Multi-agent Systems

CS

Focus is on  
Computational &  
Communication  
Efficiency

Agents are  
Obedient,  
Faulty, or  
Adversarial

ECON

Focus is on  
Incentives

Agents are  
Strategic


# Internet Computation

- Both **incentives** and **computational and communication efficiency** matter.
  - “Ownership, operation, and use by numerous independent self-interested parties give the Internet the characteristics of an **economy** as well as those of a **computer**.”
- ⇒ Study of “Computational Game Theory” or “Incentive-Compatible Algorithms”

# Outline

- Multicast Cost Sharing
- Interdomain Routing
- Forwarding Packets in Wireless, *ad hoc* Networks
- Anonymity Systems
- Secret Sharing with Utilities
- Complexity of Information Markets

# Multicast Cost Sharing Mechanism-Design Problem



Users' types

Link costs

Receiver Set

Which users receive the multicast?

Cost Shares

How much does each receiver pay?

# Results for “Obedient” Network and Strategic Users

- **Feigenbaum-Papadimitriou-Shenker [STOC '00]**  
Welfare-maximizing mechanism has **good network complexity**.
- **Feigenbaum-Krishnamurthy-Sami-Shenker [FST-TCS '02]**  
Budget-balanced mechanisms have **bad network complexity**.
- **Feigenbaum-Krishnamurthy-Sami-Shenker [FST-TCS '02]**  
**Strategyproof** mechanism cannot be both approximately budget-balanced and approximately welfare-maximizing.

# Obedient Network, continued

- **Archer-Feigenbaum-Krishnamurthy-Sami-Shenker**  
Group-strategyproofness, “approximate” budget balance and approximately minimum worst-case welfare loss with better network complexity.  
(Not constant-factor approximation.)

## Open questions include:

- Constant-factor approximation with same network complexity as in AFKSS?
- “Approximate group-strategyproofness” and good network complexity?

# Results on “Strategic Network”

## Mitchell-Teague:

- Nodes collect signed messages so they can “prove” they paid correctly.
  - Sign data in FPS protocol.
  - $O(1)$  additional signed messages per link
- Content provider randomly “audits.”
- Cheaters are fined.
- Nodes that report cheaters are rewarded.

# Strategic Network, continued

## Theorem (MT02):


- (1) Cheater's expected welfare is negative.
- (2) If node  $N$ 's ancestors and children are welfare-maximizing, then  $N$ 's *only* welfare-maximizing strategy is to follow the protocol.

## Open questions include:

- Distributed auditing?
- Achieve same goal without PKI?

Note: Related work by Fiat *et al.* on digital-goods auctions.

# Lowest-Cost Routing Mechanism-design Problem



Agents: Transit ASs

Inputs: Transit costs

Outputs: Routes, Payments

# Problem Statement

Agents' types: Per-packet costs  $\{c_k\}$

(Unknown) global parameter: Traffic matrix  $[T_{ij}]$

Outputs:  $\{route(i, j)\}$

Payments:  $\{p^k\}$

Objectives:

- Lowest-cost paths (LCPs)
- Strategyproofness
- “BGP-based” distributed algorithm

# BGP-based Computational Model (1)


- Follow abstract BGP model of Griffin and Wilfong:  
Network is a graph with nodes corresponding to ASes and bidirectional links; intradomain-routing issues are ignored.
- Each AS has a routing table with LCPs to all other nodes:

Dest.	LCP				LCP cost
AS1	AS3	AS5	AS1		3
AS2	AS7	AS2			2

Entire paths are stored, not just next hop.

# BGP-based Computational Model (2)

- An AS “advertises” its routes to its neighbors in the AS graph, whenever its routing table changes.
- The computation of a single node is an infinite sequence of stages:



- Complexity measures:
  - Number of stages required for convergence
  - Total communication

# Results on Interdomain Routing

## Feigenbaum-Papadimitriou-Sami-Shenker [PODC '02]

- Unique **strategyproof** mechanism that satisfies one natural assumption.
- BGP-based **algorithm** that computes paths and prices with small performance penalty, *assuming nodes obey the protocol*.

## Mitchell-Sami-Talwar-Teague

- Use **signatures** to force nodes to obey the protocol.

**Open question:** Achieve same goal without PKI?

# Ongoing Work on Policy-Routing MD

## [F-Griffin-S-S]

- **Per-packet**  $c_k$  is an unrealistic cost model.
- AS route preferences are influenced by reliability, customer-provider relationships, peering agreements, *etc.*

### General Policy Routing:

- For all  $i, j$ , AS  $i$  assigns a **value**  $v^i(P_{ij})$  to each potential **route**  $P_{ij}$ .
- Mechanism Design Goals:
  - Maximize  $V = \sum_{i,j} v^i(P_{ij})$ .
  - For each destination  $j$ ,  $\{P_{ij}\}$  forms a tree.
  - **Strategyproofness**, **good network complexity**

# General Policy Routing is Hard

NP-hard even to approximate  $V$  closely

Approximation-preserving reduction from  
Maximum Independent Set

Possible approaches:

Restricted class of networks

Restricted class of valuation functions  $v^i()$



# Current Status

- (Easy to derive) **strategyproof mechanism**.
- Known MDST **distributed algorithms**.
- Conjecture: No “**BGP-based**” protocol.
- Interpretation for Mechanism Design: Can have BGP-like routing policies or a BGP-like protocol but not both.

Ongoing work: Alternative formulations of **incentive-compatible** interdomain routing.


# Ad-Hoc and/or Mobile Networks

- Nodes make same incentive-sensitive decisions as in traditional networks, *e.g.*:
  - Should I transit traffic?
  - Should I obey the protocol?
  - Should I connect to the network?
- These decisions are made more often and under faster-changing conditions than they are in traditional networks.
- Resources (*e.g.*, bandwidth and power) are scarcer than in traditional networks. Hence:
  - Global optimization is more important.
  - Selfish behavior by individual nodes is potentially more rewarding.

# Zhong-Chen-Yang [Infocom '03]


## The “Sprite” Payment System for Wireless *ad-hoc* Routing

- Nodes get **receipts** as evidence of forwarding.
- Whenever indoors: Use Internet to submit receipts and obtain **credits** (virtual currency).
- System deducts credits from source node and adds credits to forwarding nodes.
- Cryptographically based payment system ensures no credit without forwarding and no cheating (even with collusion) to reduce payment.

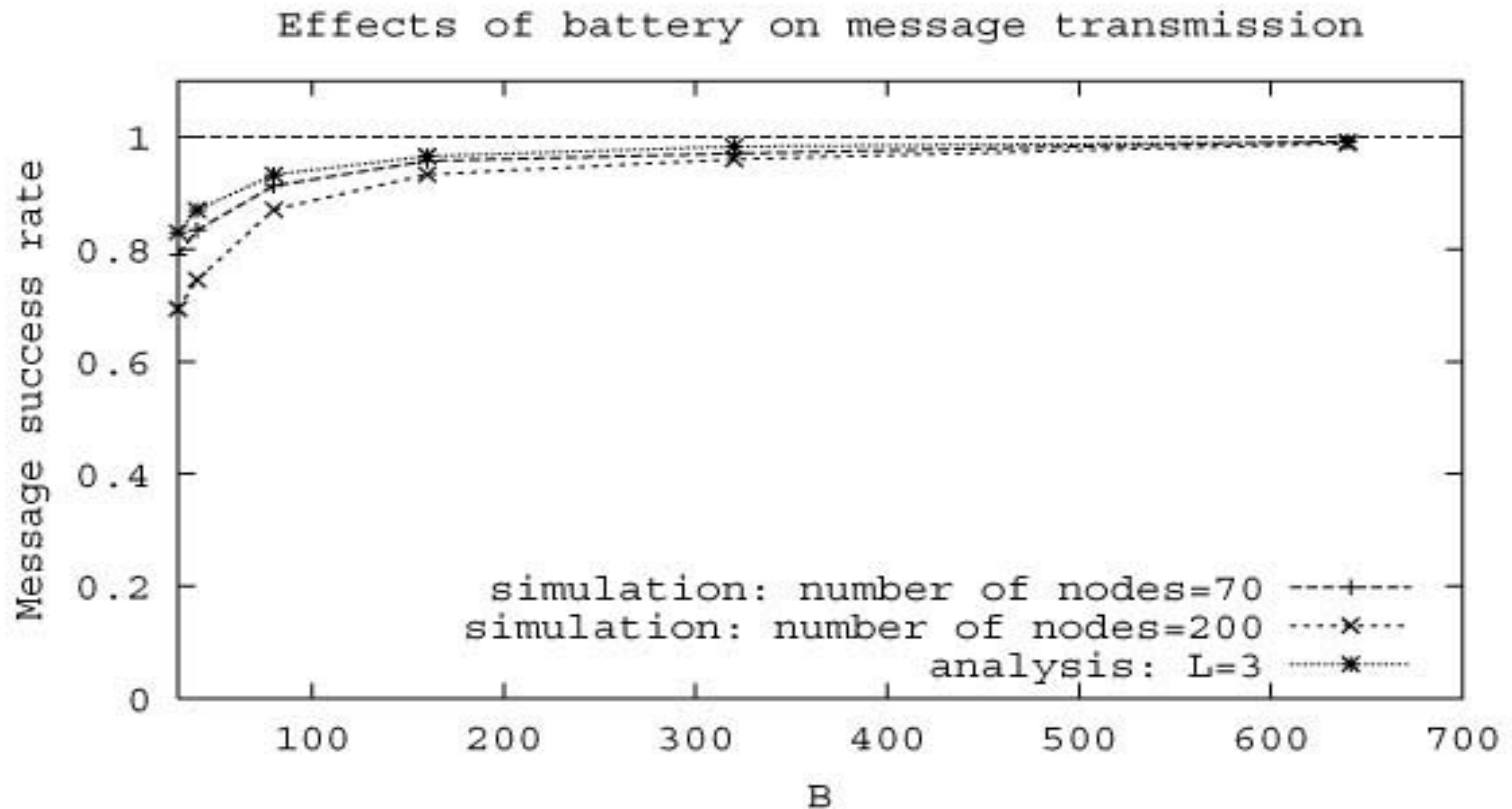


---

Wide-area wireless network



# Performance: OK unless battery is low



# Anonymity Systems: Basic Goals

Build a network

- where every message has a high probability of correct delivery
- where every message has a very low probability of anonymity compromise
- without heavyweight protocols (ZK proofs, *etc.*)
- without special trusted parties
- with minimal assumptions about the honesty and competence of intended participants

# Dingledine-Syverson [FC '02]

Two kinds of agents running network nodes:

- honest: wants reputation for running a reliable node.
- dishonest: wants to break anonymity and reliability of network.


## D-S Contributions:

- Designed reputation system to track reliability of networks made of these types of agents.
- “Creeping Death”:  
A modest-size coalition of dishonest agents can get themselves placed anywhere in the reputation spectrum.
- Modified design to be resilient against creeping-death attack:  
Choose all paths (or “cascades”) from a broader range of reputations.

# Acquisti-Dingledine-Syverson [FC '03]

- In anonymity systems, there can be a level of **free-riding** that **optimizes everybody's utility**.
- In anonymity systems, **price discrimination** doesn't work as well as it does in other settings, because **available anonymity is automatically reduced** when the market segment shrinks.
- **Current state of the economics of anonymity:**
  - No good trust metric currently well deployed.
  - No clear way yet to model dishonest nodes (who reveal what they know) as opposed to lazy nodes (who don't provide service).
  - **Adoption disincentives: Highly sensitive agents won't adopt first (because they need lots of users); low-sensitivity agents won't adopt without other incentives.**
  - Key open problem: exit node liability.

# Secure, Multiparty Function Evaluation



$$O = O(t^1, \dots, t^n)$$

- Each  $i$  learns  $O$ .
- No  $i$  can learn anything about  $t^j$  (except what he can infer from  $t^i$  and  $O$ ).

# Extensive SMFE Theory Developed by Cryptographic Researchers

- Agents are either “good” or “bad.”
  - “Good” is what’s called “obedient” in DAMD.
  - “Bad” could mean, e.g.,
    - Honest but curious
    - Byzantine adversary
- Typical Results
  - If at most  $r < n/2$  agents are honest but curious, every function has an  $r$ -private protocol.
  - If at most  $r < n/3$  agents are byzantine, every function has an  $r$ -resilient protocol.

([BGW '88] uses threshold- $r$  secret sharing and error-correcting codes.)

# Ongoing Work by Halpern and Teague

- Don't assume that a fixed fraction of agents follow the protocol.
- Instead, add utilities to the picture. **An agent will deviate from the protocol if and only if it leads to an increase in utility.**
- Theorem: If there is a known upper bound on the running time of the SMFE protocol, **the only rational thing for an agent to do at every step of the protocol is to “defect,” i.e., not to participate at all.**

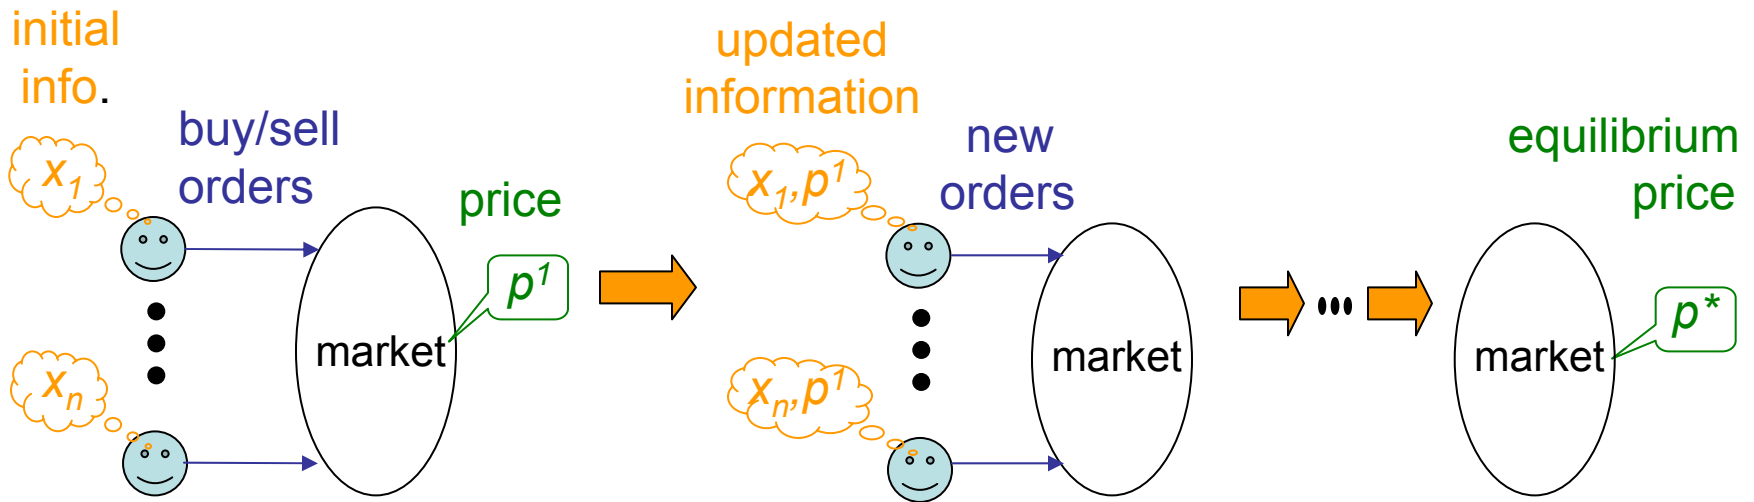
# Halpern and Teague (2)

- Shoham and Tennenholtz have shown that, even with a trusted party, **only certain functions can be computed “non-cooperatively.”**
- Theorem: Anything that can be **computed “non-cooperatively” with a trusted party** can be **computed non-cooperatively by a multiparty protocol**. (The running time of this protocol is a random variable.)

# Feigenbaum-Fortnow-Pennock-Sami [EC '03]: Computation in a Distributed Information Market

- Evidence from research markets, sports-betting markets, securities markets, *etc.* suggests that markets are very good at **aggregating information**.
- Markets are sometimes designed **explicitly for information aggregation**, e.g., the Hollywood Stock Exchange.
- Computational capacity of information markets:
  - **Which aggregation functions can be computed?**
  - How many securities must be traded?
  - **How fast does the market price converge?**

# Market as a Computational Device



Is equilibrium price  $p^*$  = desired aggregate  $f(x_1, x_2, \dots, x_n)$  ?

- Study **Boolean functions**, each  $x_i$  is a single bit.
- Trade in a single security  $F$  with payoff contingent on  $f$ .
- Use simplified *multiperiod Shapley-Shubik* market model.
- Assume that traders are common-prior Bayesians and that they bid **truthfully (not strategically)**.

# [FFPS '03] Results

## Computable Functions:

- If  $f$  can be expressed as a *weighted threshold function*, then, for any prior distribution, the **price of  $F$**  converges to the true value of  $f(x_1, x_2, \dots, x_n)$ .
- If  $f$  cannot be expressed as a weighted threshold function, there are priors for which **price of  $F$**  does not converge to  $f(\cdot)$ .

## Convergence Time:

- The market converges to equilibrium in **at most  $n$  rounds**.
- There are weighted-threshold functions and priors such that the market takes  **$n/2$  rounds** to reach equilibrium in the worst case.

# Future Directions on Information Markets

- Realistic data models
  - Real numbers
  - Noisy data
- Realistic economic models
  - Strategic traders
  - No common priors
- Other computational aspects
  - Complexity of *traders'* computations
  - Decentralized market
  - Connection to machine-learning theory?

# High-level Question about Incentives in Diffuse Computing

- Distinguish problems that are *hard to compute on the Internet* from those that are *easy*.
- Hardness derives from *interplay* of networked computation and incentive compatibility (e.g., budget-balanced, group-strategyproof multicast cost sharing).
- Develop general complexity theory
  - Efficient algorithms for interdomain routing, P2P file sharing, distributed resource allocation, *etc.*
  - Hardness results may lead to “secure algorithmic mechanisms.”