Economics of Peer-to-Peer Systems

John Chuang

School of Information Management and Systems
University of California at Berkeley
chuang@sims.berkeley.edu
http://p2pecon.berkeley.edu/

Academia Sinica 2004 Summer Institute on P2P Computing
August 3 2004

Collaborators

- Nicolas Christin
- Yang-hua Chu (CMU)
- Michal Feldman
- Jens Grossklags
- Ahsan Habib
- Kevin Lai (HP)
- Christos Papadimitriou
- Ion Stoica
- Hui Zhang (CMU)

Economics of P2P?

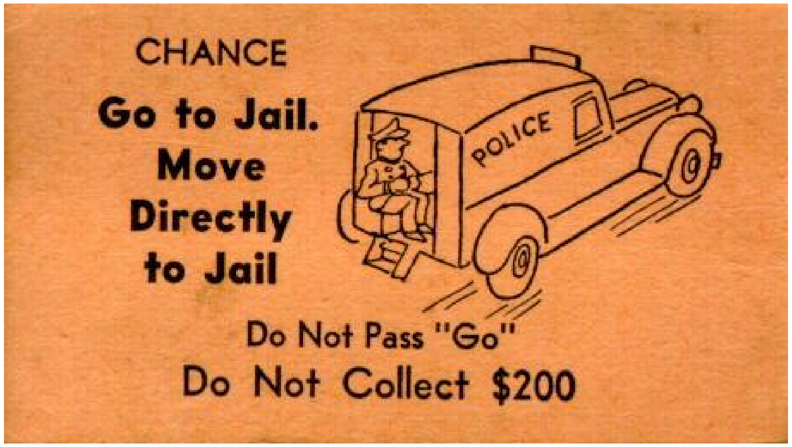

- This talk is <u>NOT</u> about the economic impact or legitimacy of P2P file sharing
- See:
  - Oberholzer & Strumpf, P2P's Impact on Recorded Music Sales.
  - Gopal, Bhattacharjee, Lertwachara, Marsden, Impact of Online P2P Sharing Networks on the Life Cycle of Albums on the Billboard Chart.

Economics of P2P

- This talk is about economics-informed design of P2P systems
  - Understanding system characteristics
    - Quantifying disincentives
    - Free-riding: individual rationality vs. collective welfare
    - Whitewashing: cheap pseudonyms
    - Information asymmetries: hidden info, hidden action
  - Designing incentive mechanisms
    - Tokens, reputation, taxation, contracts, etc.

Outline

- P2P system characteristics
  - Disincentives in sharing → free-riding
- Incentive mechanisms
  - Tokens, reputation, taxation, contracts, ...
  - Challenges: whitewashing, collusion, etc.
- Case study:
  - On-demand P2P streaming
  - Live event P2P streaming
- Information Asymmetry
  - Hidden action in multi-hop routing
- Conclusions

Diversity of P2P Systems

- Distributed storage, search, and retrieval
  - File-sharing: Napster, gnutella, kaZaA, Overnet, bitTorrent, ...
  - Anonymity/Persistence: Eternity, Freehaven, FreeNet, Publius, ...
  - DHTs: Chord, CAN, Pastry, Tapestry, OpenHash, ...
- Distributed computation
  - Globus (grid), Entropia, SETI@Home, etc.
- Communications
  - Connectivity: mobile wireless ad-hoc networks, “rooftop” networks
  - Redundancy: resilient overlay networks
  - Anonymity: onion-routing, MIX-net, Crowds
  - Distributed multimedia: skype (VoIP), ESM/Narada, Splitstream (live streaming), PROMISE (on-demand streaming)
- More at: http://www.openp2p.com/pub/q/p2p_category

P2P System Characteristics

- What do P2P systems have in common?
  - No infrastructure or service provider: rely on contributions by individual peers
  - Hidden action: difficult to monitor or enforce cooperation
  - Ad-hoc communities: highly dynamic memberships; interactions with strangers

Free-riding

- Fundamental tension between *individual rationality* and *collective welfare*
  - System utility derived solely from peer contributions
  - Contributions not costless → disincentives to share
- Rational peers choose to free-ride, i.e., consume but not contribute
- Free-riding prevalent in file-sharing networks [Adar00; Sariou02]
  - 66% of gnutella peers share no files
  - 10% of peers share 87% of files
  - 20% of peers share 98% of files
- [Adar00]: “Tragedy of digital commons”?

Questions

- What are the costs of participating in a P2P network? How significant are the disincentives for sharing (potential legal liability notwithstanding)?
- What are the effects of free-riding on P2P system performance? Are P2P systems doomed to failure due to non-cooperation?
- How do we design incentive mechanisms to encourage cooperation in P2P systems?

Disincentive for Sharing

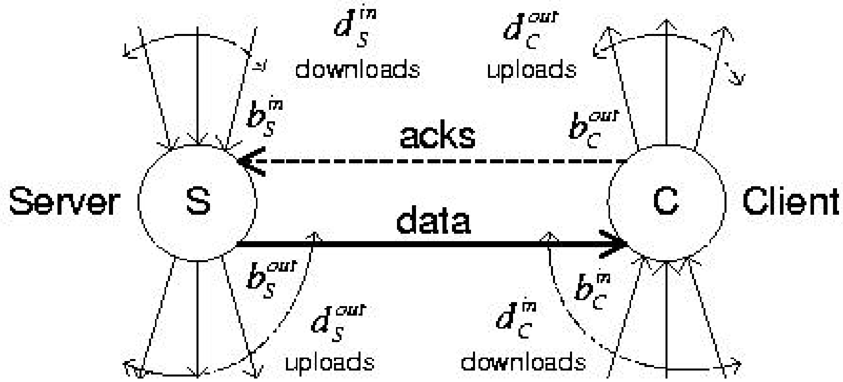

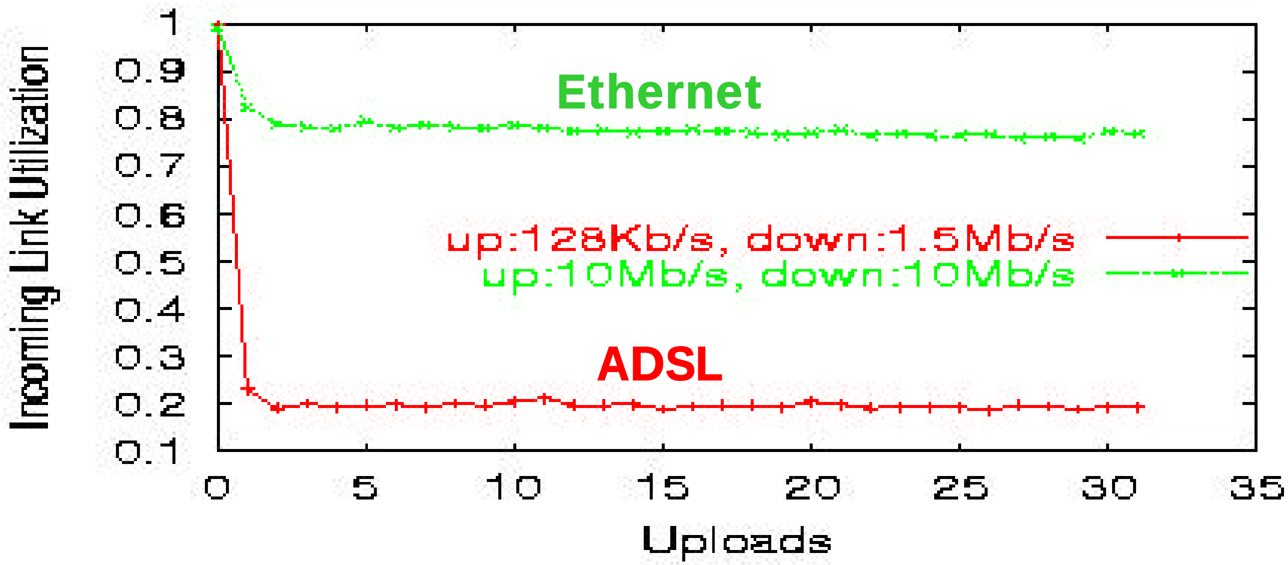

- Case 1: P2P file-sharing [Feldman03]
  - Incoming link utilization degrades by 20-80% when simultaneously uploading (ns-2 simulation)
  - Contention between TCP data and ACK

Disincentive for Sharing

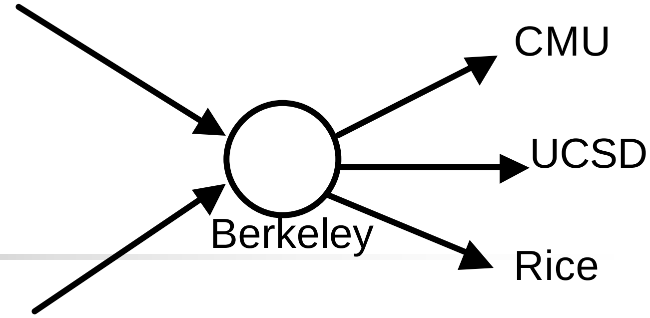

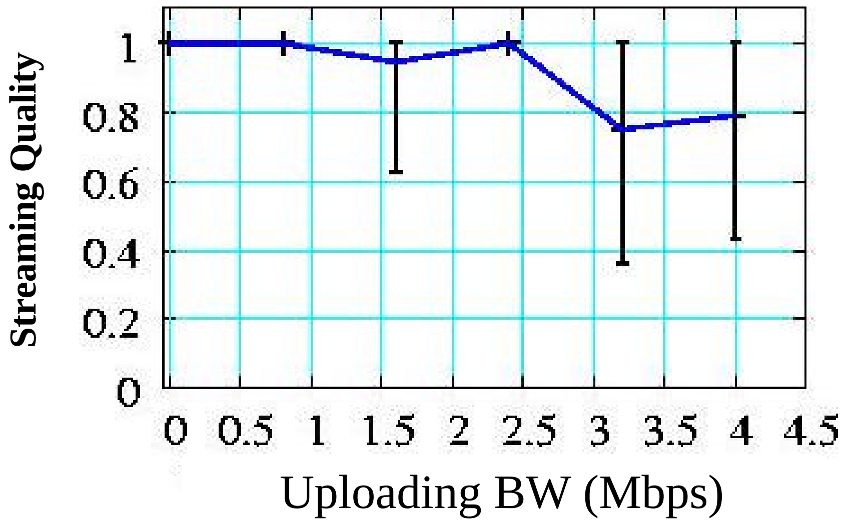

- Case 2: P2P media streaming [Habib04]
  - Streaming quality becomes highly variable as uploading bandwidth increases
    (planetlab experiment using PROMISE prototype)

General Cost Model [Christin04]

- A given node *u* requests an item, serves a request, or route requests between other nodes:
  - Latency cost (benefit) $$L_u = \sum_{v \in V} \sum_{k \in K_v} l_{u,k} t_{u,v} \Pr[Y = k]$$
  - Service cost $$S_u = \sum_{k \in K_u} s_{u,k} \Pr[Y = k]$$
  - Routing cost $$R_u = \sum_{v \in V} \sum_{w \in V} \sum_{k \in K_w} r_{u,k} \Pr[X = v] \Pr[Y = k] \chi_{v,w}(u)$$
  - Topology maintenance cost $$M_u = m_u \deg(u)$$

Participation Cost

- Cost can be highly variable, dependent on many factors, e.g., item popularity, network topology, routing algorithm, even node ID!

- Example: routing cost for various DHT overlay topologies [Christin04]

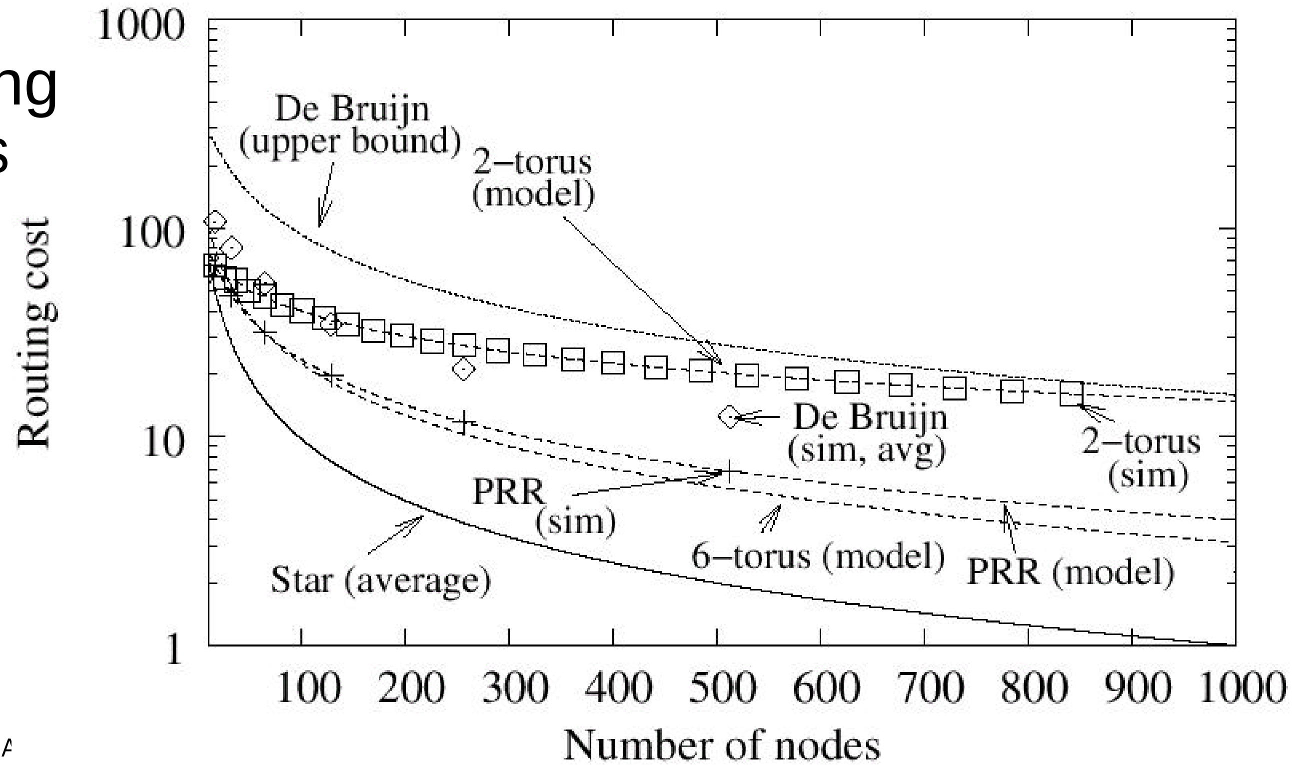

What can we do?

- Rely on altruism
  - No intervention necessary if *societal generosity* sufficiently high [Feldman04b]
  - Warm-glow theory: altruistic action may be part of rational behavior [Andreoni90]
- Enforcement
  - Obedient vs. malicious peers
  - Often circumvented by determined hackers
- Incentives
  - Rational users respond to reward and/or punishment
  - Security requirements still remain

Outline

- P2P system characteristics
  - Disincentives in sharing → free-riding
- Incentive mechanisms
  - Tokens, reputation, taxation, contracts, ...
  - Challenges: whitewashing, collusion, etc.
- Case study:
  - On-demand P2P streaming
  - Live event P2P streaming
- Information Asymmetry
  - Hidden action in multi-hop routing
- Conclusions

Incentive Mechanisms

- Tokens/currency
  - Appropriate for trading of multiple resource types
  - Examples: Mojonation [Wilcox-O'Hearn02], KARMA [Vishnumurthy03], tycoon [Lai04], ...
- Barter/taxation
  - Sometimes called “tit-fot-tat” or “bit-for-bit”
  - Appropriate for single commodity type
  - Examples: Bittorrent [Cohen03], ESM [Chu04]
- Reciprocity
  - Direct reciprocity (repetition)
  - Indirect reciprocity (reputation)

Direct Reciprocity

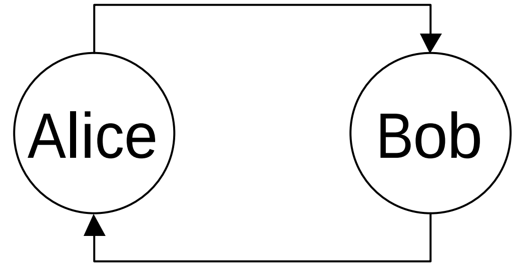

- Repetition encourages cooperation
  - e.g., Prisoners' Dilemma game:
    - one-shot game: mutual defection is dominant strategy
    - infinitely repeated game: mutual cooperation is dominant
- Simple tit-for-tat (TFT) strategy works very well in iterated prisoners' dilemma (IPD) tournaments [Axelrod84]
- Clustering (e.g., clubs [Asvanund03]) and server selection (e.g., CoopNet [Padmanabhan02]) may facilitate direct reciprocity

Direct Reciprocity

- But direct reciprocity can be difficult to achieve in P2P networks
  - Large populations and dynamic memberships → few repeat transactions
  - Asymmetries in interests
  - Asymmetries in capabilities

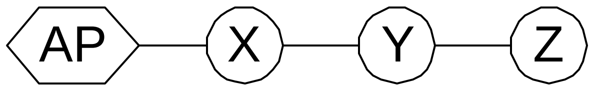

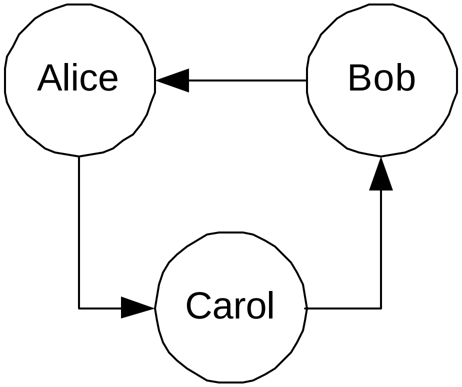

Indirect Reciprocity

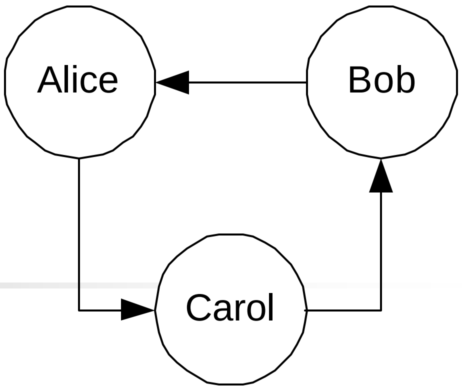

- Peers earn <u>reputation</u> via cooperation
- Reputable peers receive preferential treatment
- Implementation overhead for maintaining reputation information
- Various proposals
  - Image scoring [Nowak98],
    Free Haven [Dingledine90],
    Eigentrust [Kamvar03],
    Differentiated admission [Kung03],
    CONFIDANT [Buchegger02],
    ...

Tradeoffs and Challenges

- Design space for reciprocity-based schemes
  - Direct vs. indirect reciprocity?
    - Private vs. shared history
    - Server selection
    - Shared history: collusion resistance
  - Dealing with invisible defections
  - Dealing with strangers and whitewashers
  - Dealing with traitors
- Simulation-based study of robust incentive techniques in [Feldman04a]

Private History

- Corresponds to direct reciprocity
- Advantages
  - Implementation is simple and decentralized
  - Immune to collusion
- Disadvantages
  - Requires repeat transactions
    - e.g., low rate of turnover, small populations
  - Deals poorly with asymmetry of interest

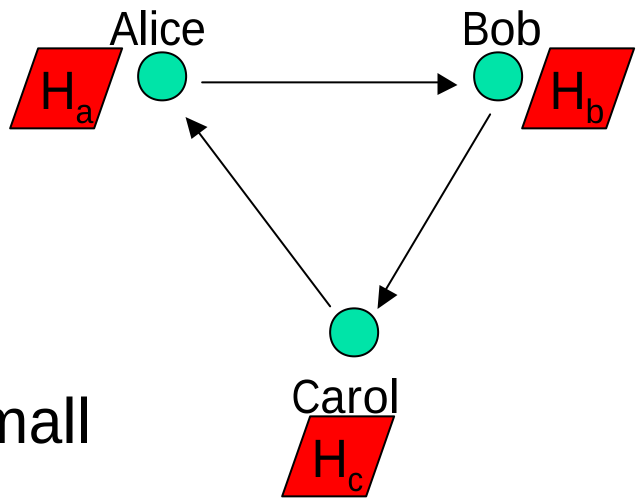

Shared History

- Corresponds to indirect reciprocity
- Advantages
  - Tolerates few repeat transactions (large populations, high turnover)
  - Tolerates asymmetry of interest
- Disadvantages
  - Susceptible to collusion
    - Subjective shared history via max-flow algorithm [Feldman04a]
  - Implementation overhead

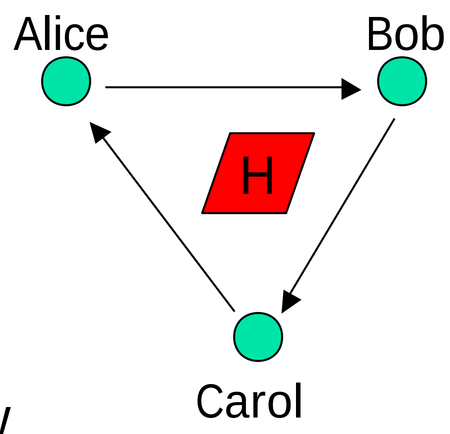

To cooperate or not to cooperate?

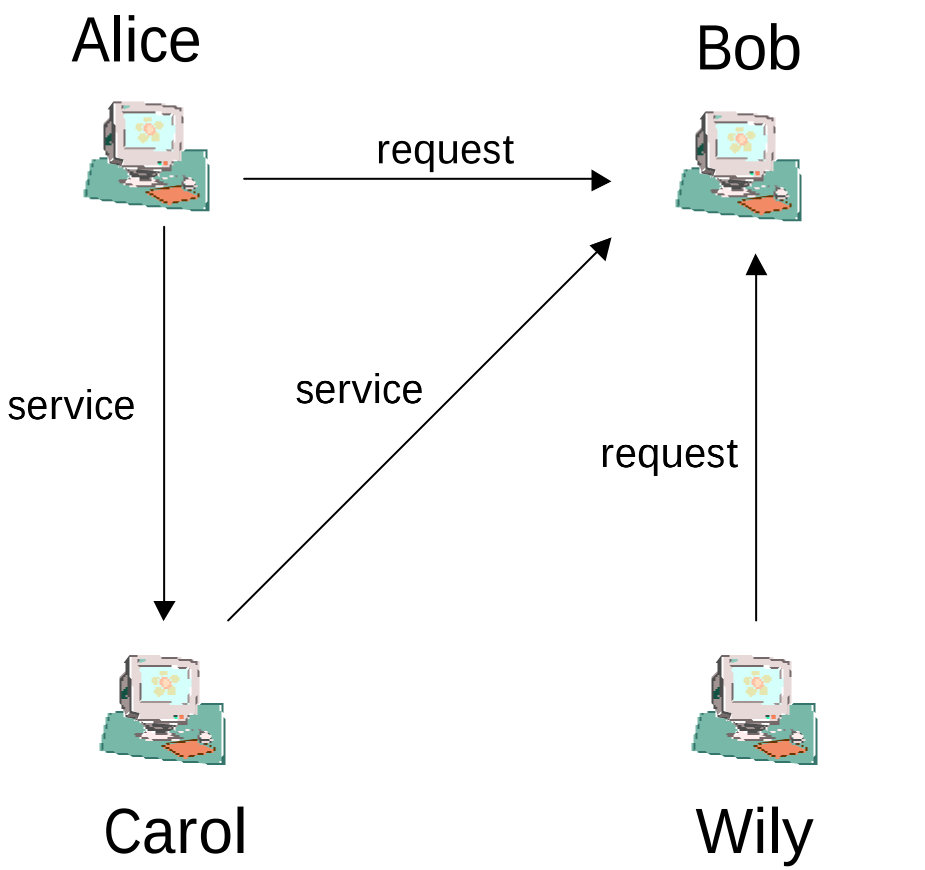

cooperate
or defect?

Cooperate: $P_{Alice} = 7$, $P_{Bob} = -1$
Defect: $P_{Alice} = P_{Bob} = 0$
but the defection is invisible to Alice

private history — *Carol: 1 cooperate*

shared history — *Alice: 1 cooperate w/Carol*
Carol: 1 cooperate w/Bob

stranger policy — *Cooperate with stranger?*

Simulation Framework

- Initial population mixture
  - 1/3 cooperators
  - 1/3 defectors
  - 1/3 reciprocators
- Game composed of rounds in which players are randomly matched, one as client, the other as server
- Learning: players probabilistically switch to strategies with higher payoffs
- Defectors can engage in collusion or whitewashing attacks
- Reciprocators can choose shared vs. private history, and different stranger policies
- Additional simulation parameters
  - Population size
  - Turnover rate
  - Hit rate
  - ...

Dealing with Invisible Defections

- Decision function based only on cooperation, not defection
- Reciprocative decision function: cooperate with probability $g_j(i)$
  - Generosity: $g_i = p_i / c_i$
    - p_i: service i has provided
    - c_i: service i has consumed
  - Normalized generosity: $g_j(i) = g(i) / g(j)$
    - Entity i 's generosity relative to entity j 's generosity

Private vs. Shared History

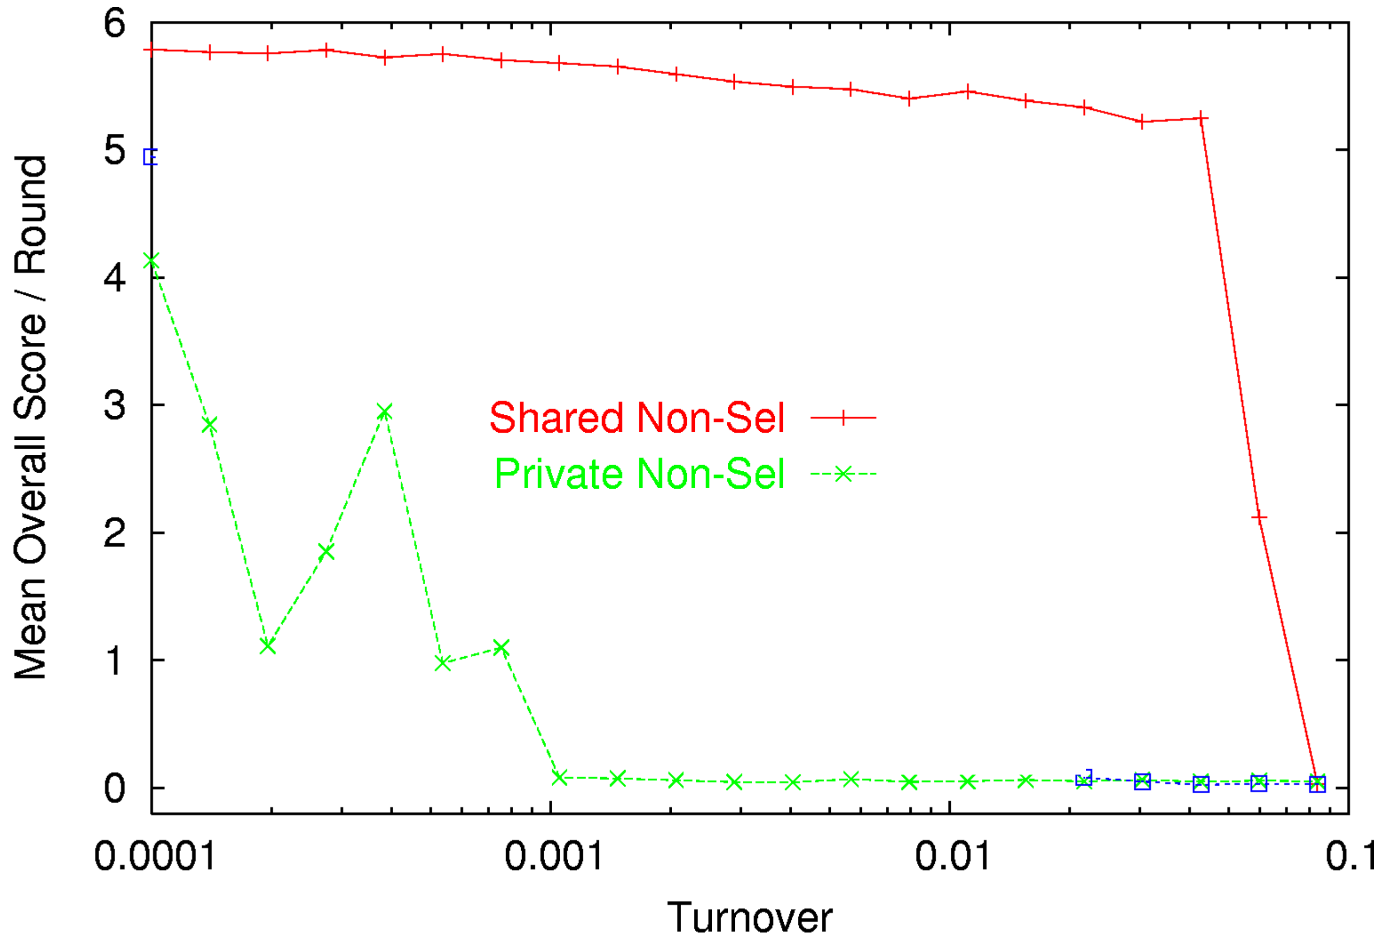

- Shared history scales to larger populations and higher turnover rates

Server Selection

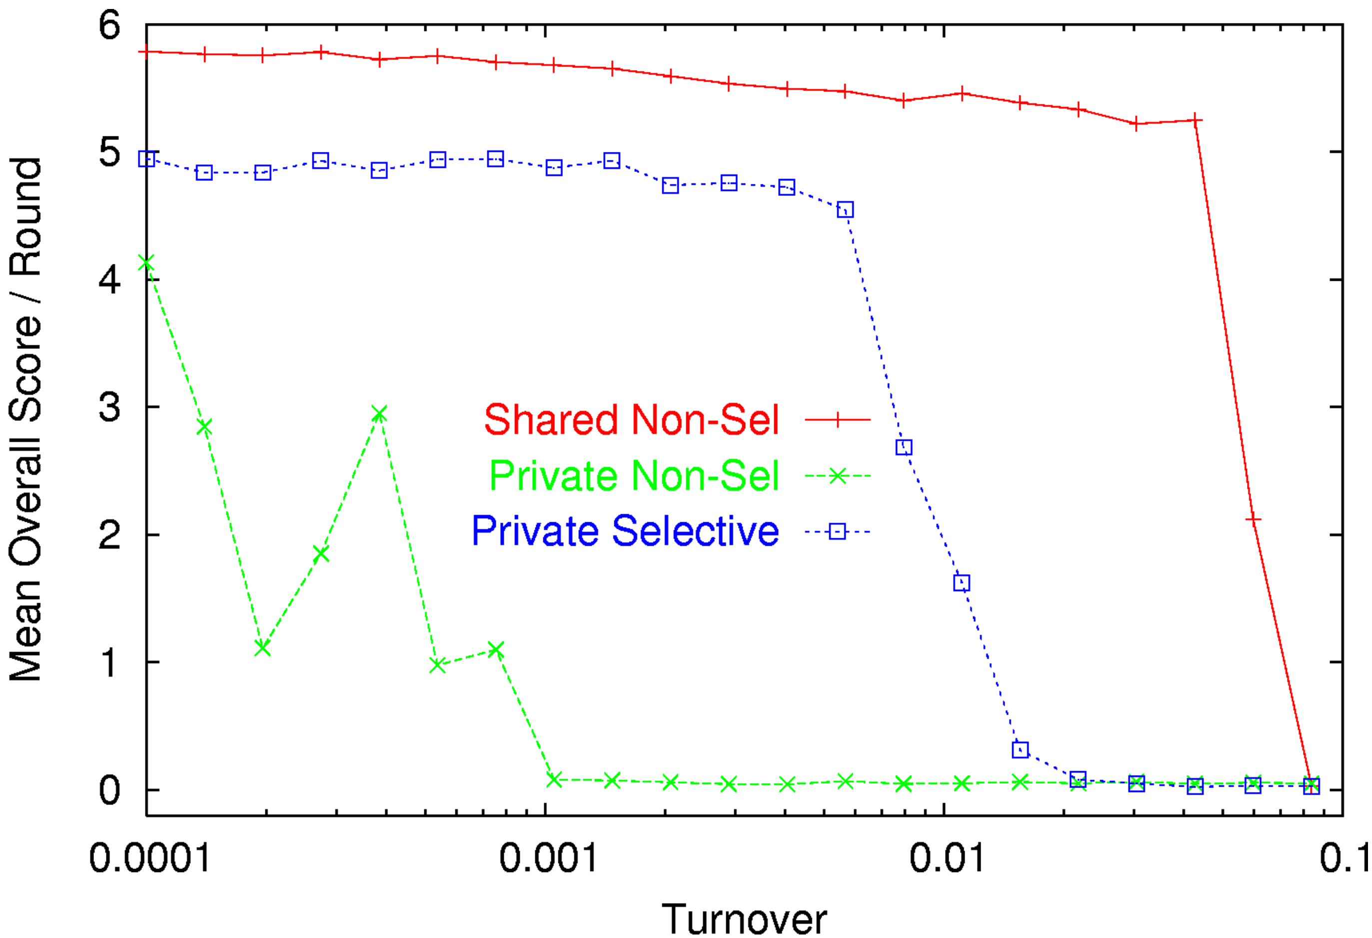

- Server selection improves scalability of private history approach

Collusion

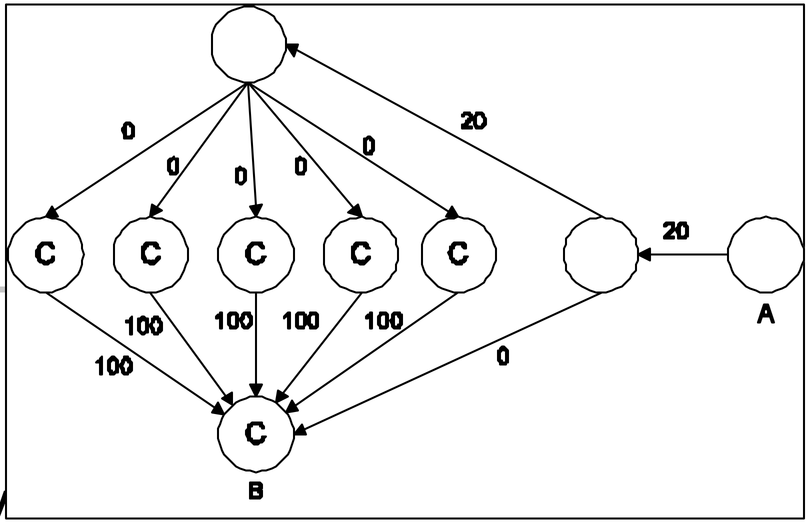

- Shared history susceptible to collusion
- Many forms of collusion may be possible
  - False praise: falsely claiming defectors have cooperated
  - False accusation: falsely claiming cooperators have defected
- *Colluder* strategy: claiming to have received service from other colluders
- Subverts *objective* reputation systems
- Negative effect is magnified when combined with zero-cost identities
- Mitigated by *subjective* reciprocity
  - e.g., leveraging pre-trusted peers [Kamvar03], social links [Marti04], maxflow algorithm

Subjective Reciprocity: Maxflow

- Compute the maximum "reputation capacity" from source to sink
- Proven to be attack resistant for authentication [Levien98][Reiter99]
- Does not require centralized trust
- Mitigate false praise, but not false accusation
- Cost: long running time $O(V^3)$
- Solution: bound mean number of nodes examined during maxflow calculation
  - Bound overhead
  - Bound efficiency

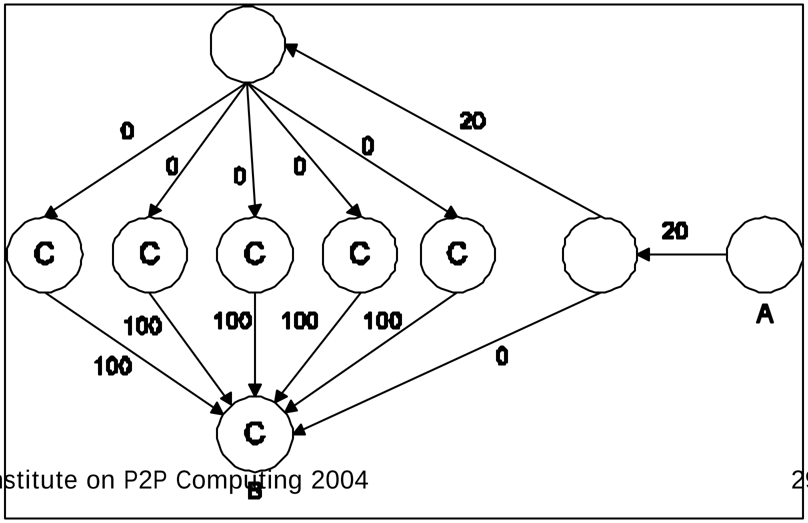

Subjective Reciprocity: Maxflow

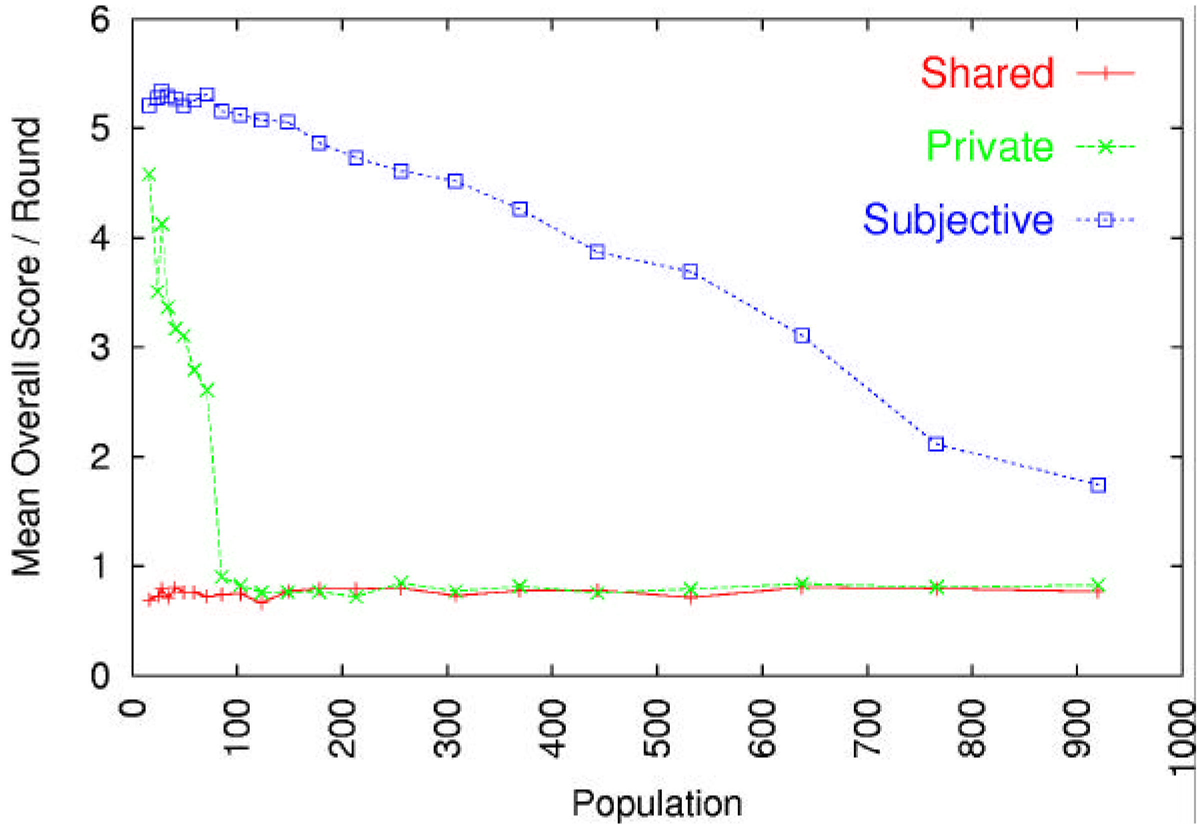

- All defectors are colluders

$$\Pr(i\_cooperate\_with\_j) = \min\left(\frac{MAXFLOW\_j\_to\_i}{MAXFLOW\_i\_to\_j}, 1\right)$$

Whitewashing Attack

- The use of history (or reputation) assumes that entities maintain persistent identities
- Problem: many online systems have zero-cost identities
  - Encourages newcomers to join
  - Circumvents history-based strategies that always cooperate with strangers
- *Whitewash* strategy: always defect, and continuously change identity
- Whitewashers indistinguishable from legitimate newcomers

Stranger Policies

- Always cooperate (e.g., Axelrod's TFT)
  - Fully exploited by whitewashers
- Always defect
  - Provides immunity against whitewashers
  - Incurs "social cost of cheap pseudonyms" [Friedman98]
    - Raises bar to entry (discourage newcomers)
    - May initiate undesirable cycles of defections
- Randomly cooperate
  - Allows exploitation by whitewashers

Stranger Policies

- Adaptively cooperate
  - Cooperate with strangers based on "friendliness" of strangers in system: p_s / c_s
    - P_s: number of services **strangers** have provided
    - C_s: number of services **strangers** have consumed
  - Only taxes newcomers when necessary

Stranger Adaptive

- In the presence of whitewashers:
- SA scales to higher turnover rates with private history
- SA performs as well as SD with shared history

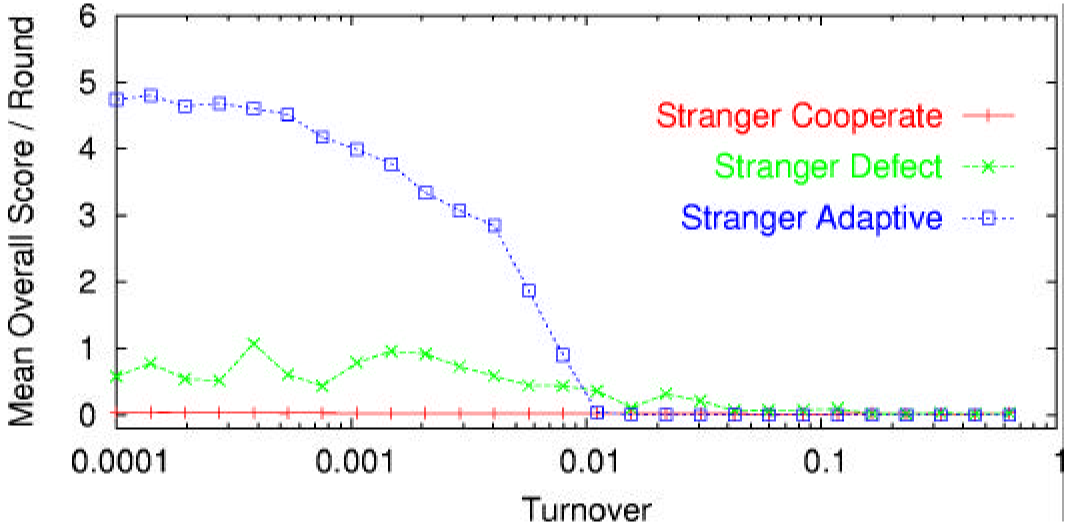

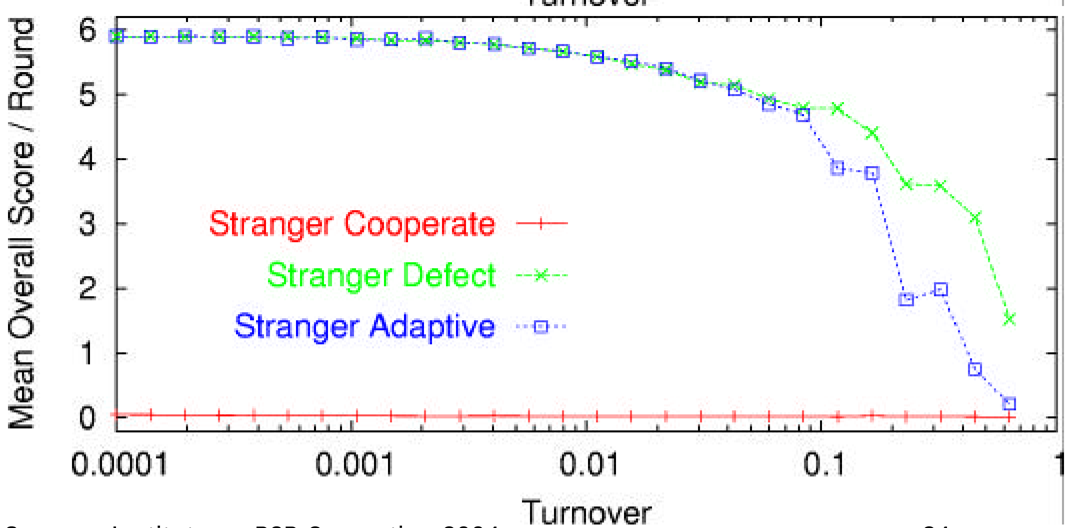

Outline

- P2P system characteristics
  - Disincentives in sharing → free-riding
- Incentive mechanisms
  - Tokens, reputation, taxation, contracts, ...
  - Challenges: whitewashing, collusion, etc.
- Case study:
  - On-demand P2P streaming
  - Live event P2P streaming
- Information Asymmetry
  - Hidden action in multi-hop routing
- Conclusions

Case Studies: P2P Streaming

- Peers contribute forwarding/uploading BW
- On-demand P2P streaming [Habib04]:
  - Many-to-one: each peer can stream from multiple peers
  - Asynchronous consumption & contribution
- Live-event P2P streaming [Chu04]:
  - One-to-many: single publisher, multiple receivers
  - Simultaneous consumption & contribution
- Different incentive mechanisms
  - Implemented for PROMISE and ESM systems, respectively

On-Demand P2P Streaming

- Observation: session quality dictated by peer selection
  - Number, capacity, and location of supplying peers

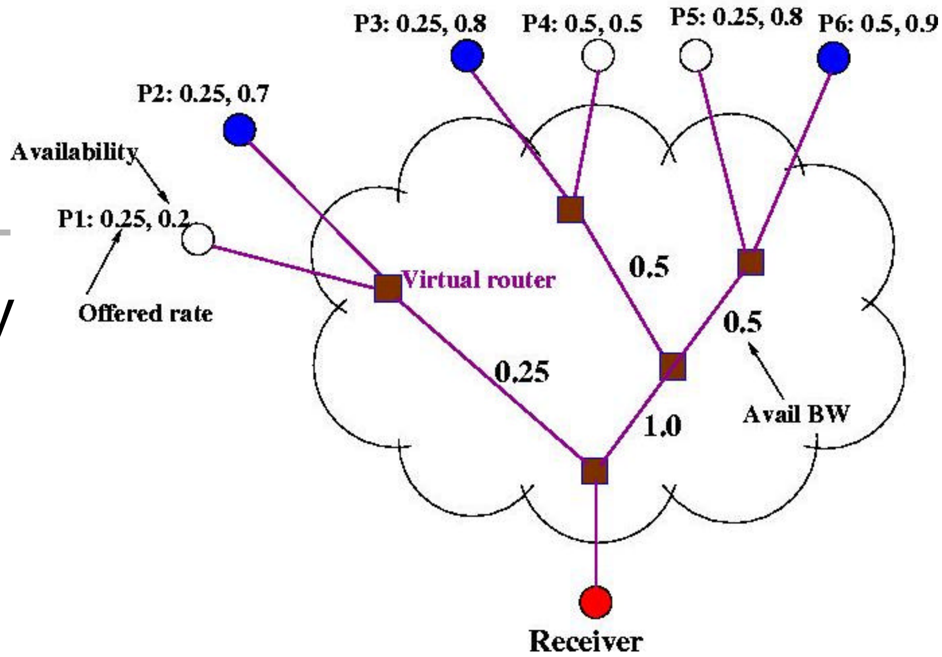

random peer selection

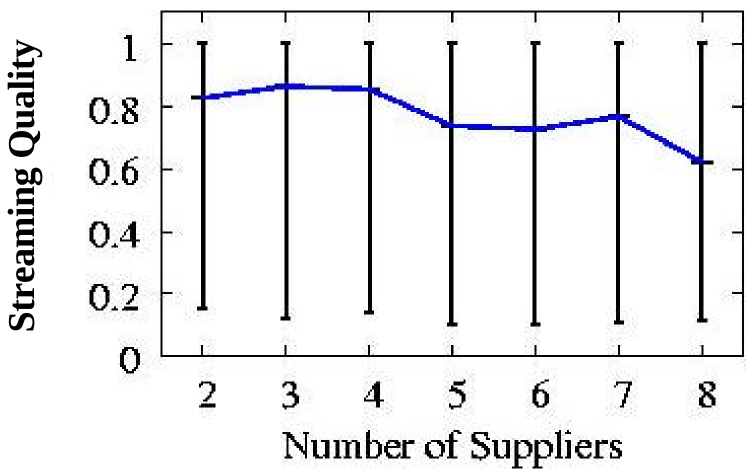

topology aware peer selection

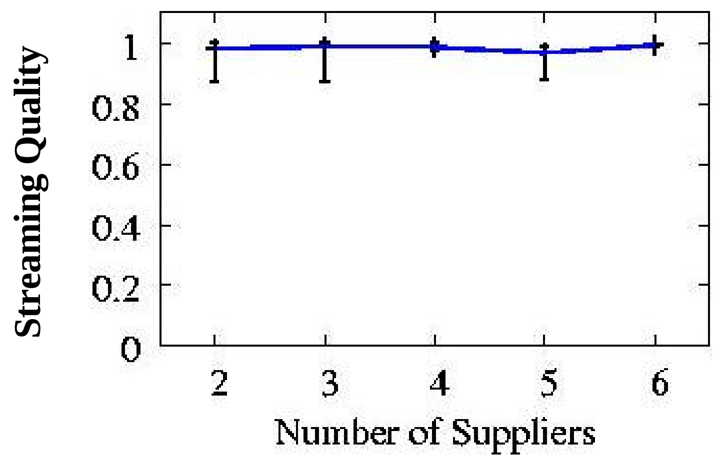

On-Demand P2P Streaming

- Incentive technique: service-differentiated peer selection
  - Contributors get to select the best available peers

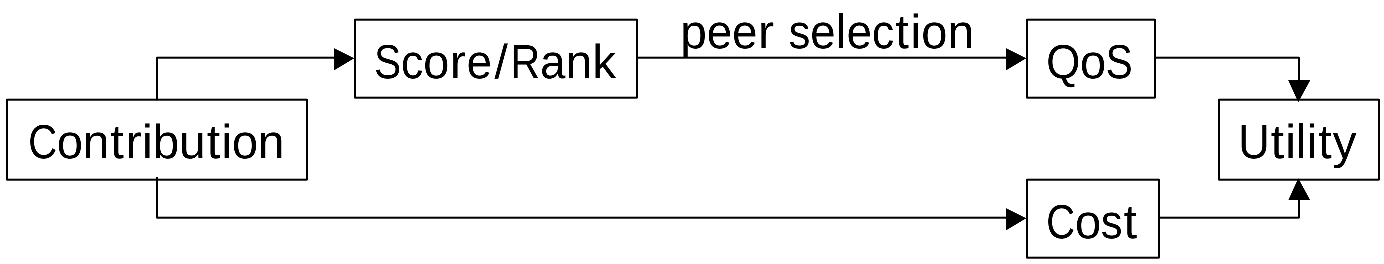

- Since consumption and contribution are independent, need to keep history
- Rational user determines optimal contribution level to maximize utility

On-Demand P2P Streaming

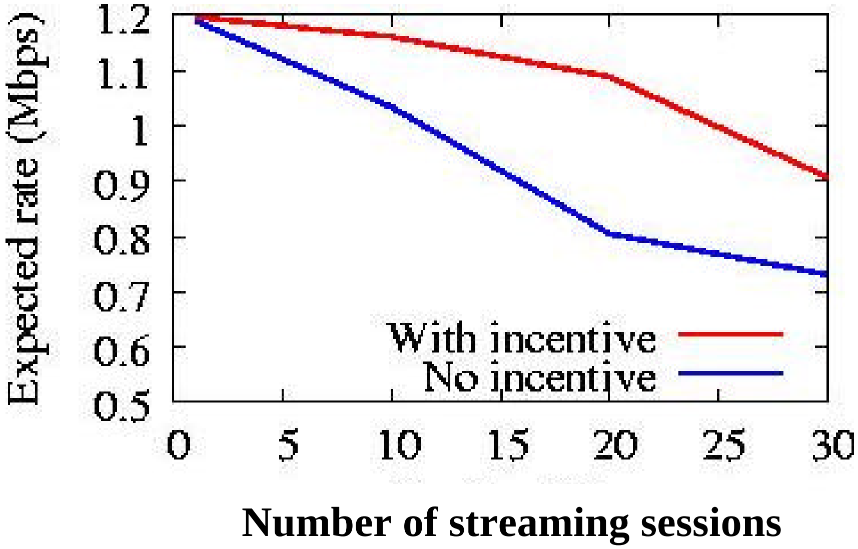

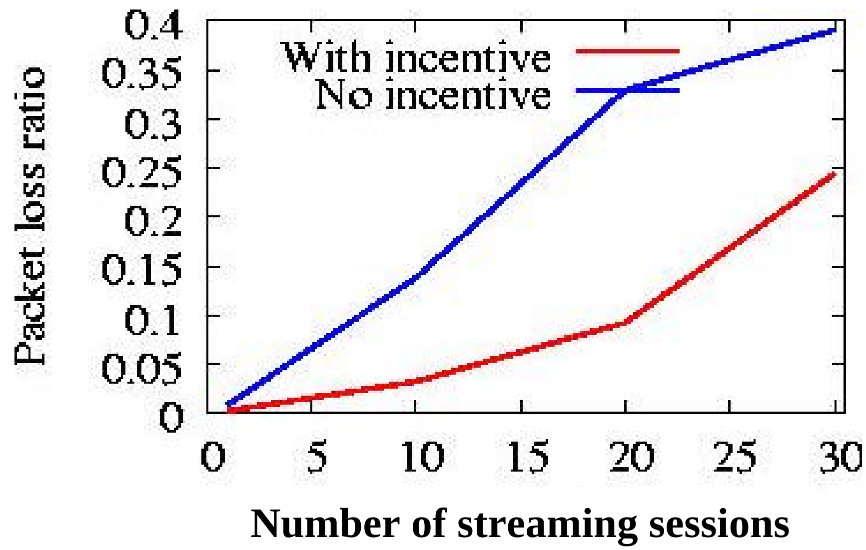

- Use of incentive mechanism improves system performance
  - Except when system load is low, or when network is congested

Live-Event P2P Streaming

- Video stream split into multiple stripes
- Peers form multiple disjoint tree structure
- Simultaneous consumption and contribution
  - No need to maintain history

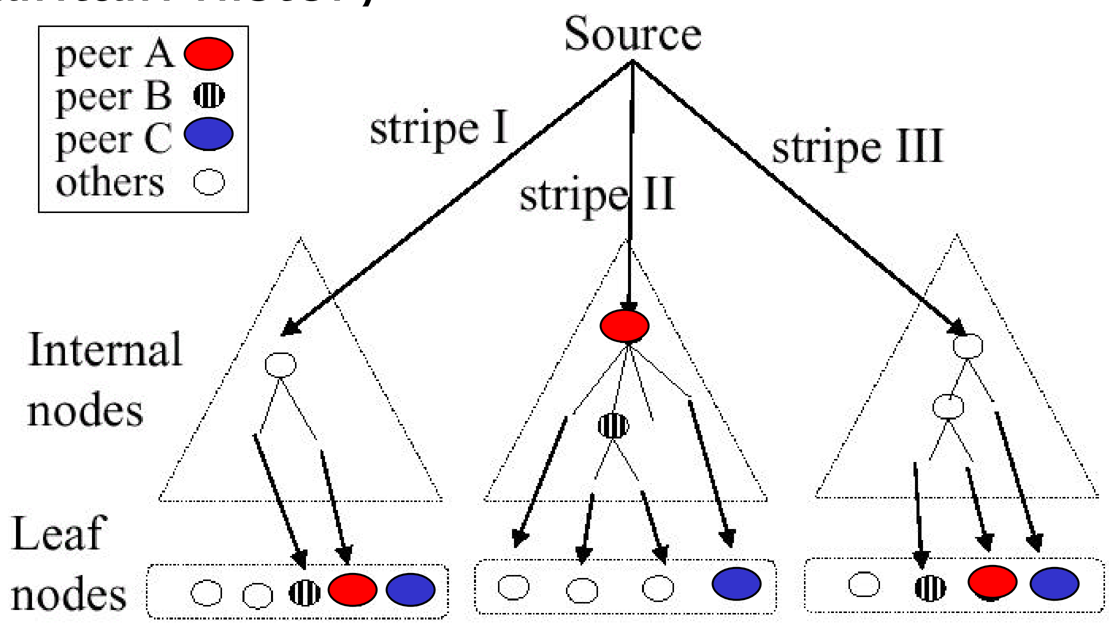

Node Heterogeneity

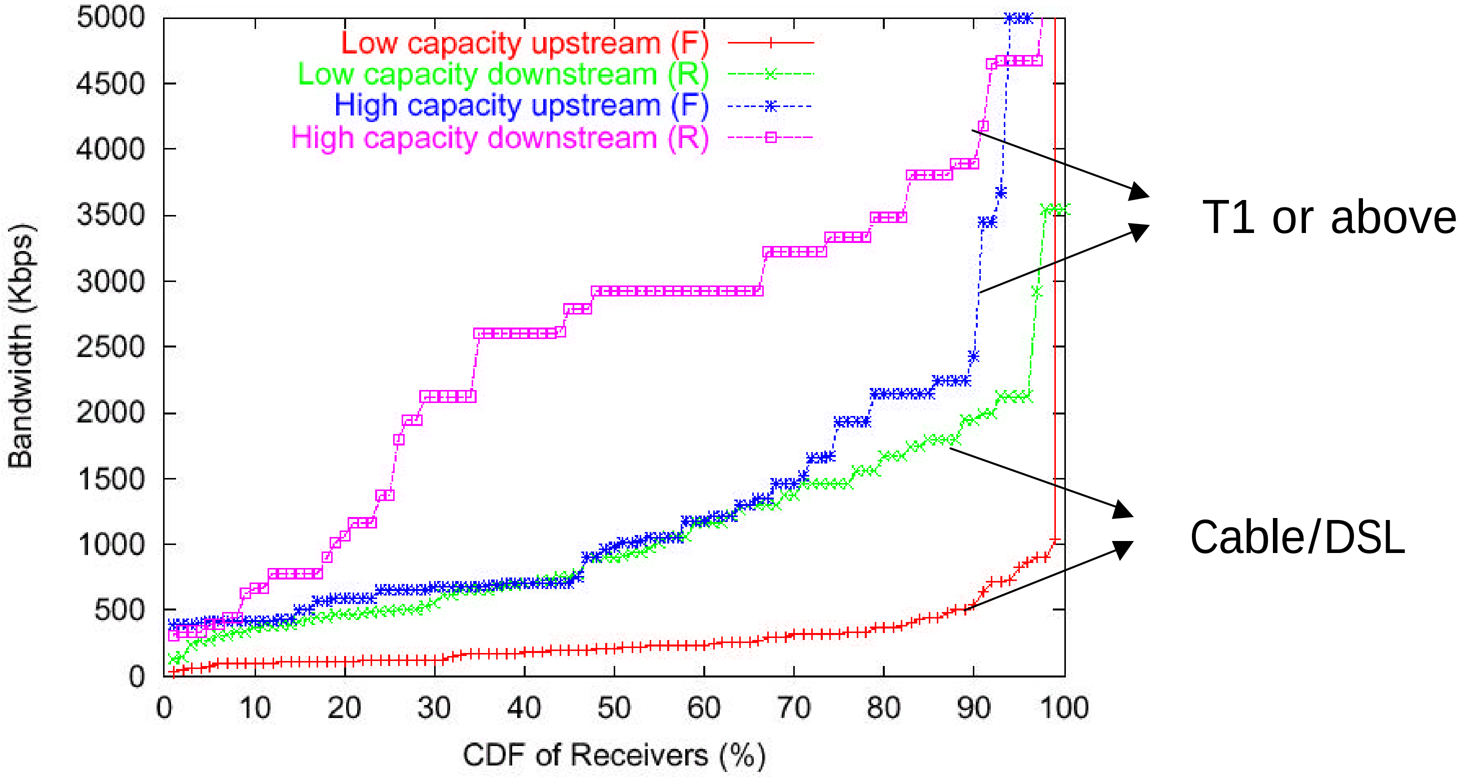

- Measured TCP throughput for slashdot trace
- Not all peers could (should) consume and contribute the same amount of bandwidth

Taxation

- Publisher sets and enforces tax schedule to achieve resource re-distribution
  - Subsidization of resource-poor nodes by resource-rich nodes
- Rich literature in public finance
  - Optimal income taxation

Linear taxation

- Contribution according to tax schedule

$$f = \max[t*(r - G), 0]$$

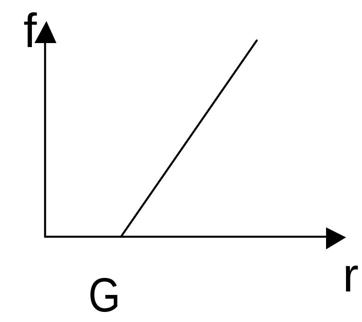

- where
  - f = forwarding bandwidth
  - r = received bandwidth
  - t = marginal tax rate
  - G = demogrant
- Publisher sets t and G, peers choose f and r
- Every peer receives at least a demogrant G
- Note: “tit-for-tat” scheme of Bittorrent [Cohen03] is special case with t=1 and G=0

Evaluation: Social Welfare

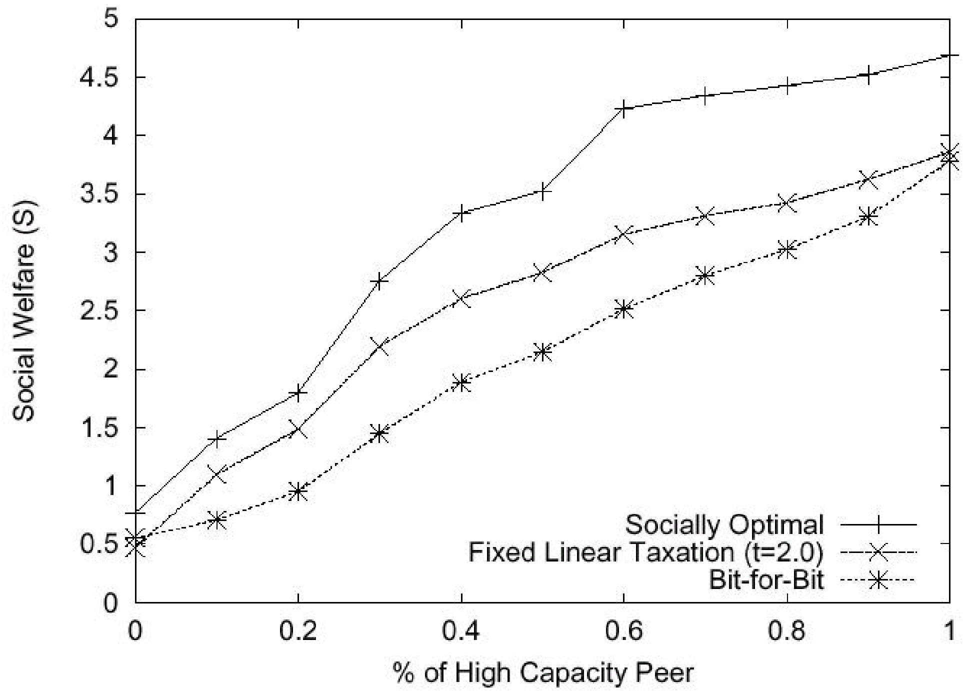

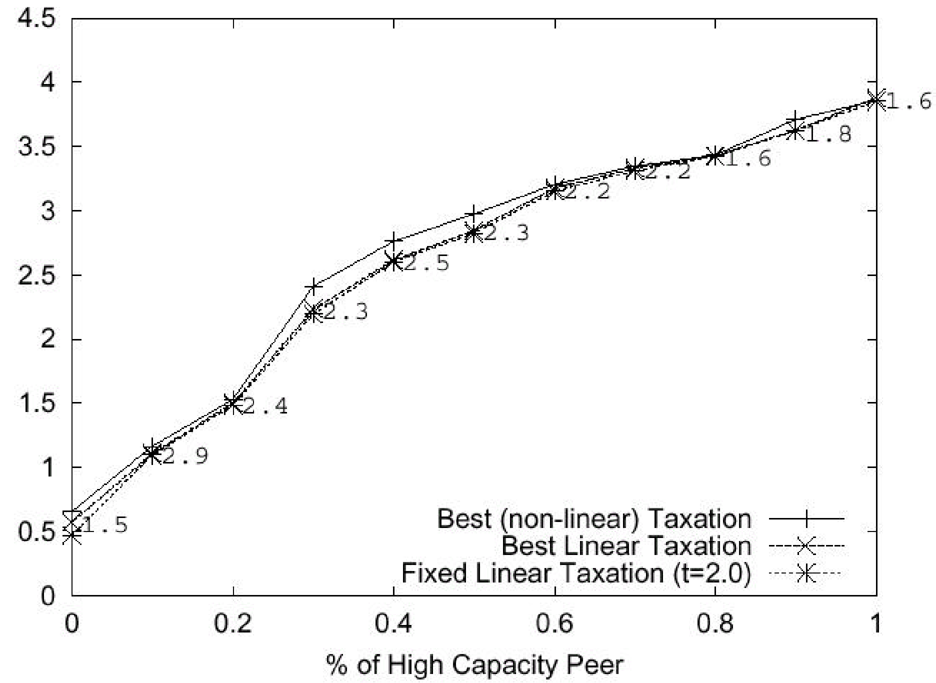

- Simple linear taxation scheme with fixed tax rate and dynamically adjusted demogrant is robust for different peer compositions

Outline

- P2P system characteristics
  - Disincentives in sharing → free-riding
- Incentive mechanisms
  - Tokens, reputation, taxation, contracts, ...
  - Challenges: whitewashing, collusion, etc.
- Case study:
  - On-demand P2P streaming
  - Live event P2P streaming
- Information Asymmetry
  - Hidden action in multi-hop routing
- Conclusions

Information Asymmetry

- Condition in which some relevant information is known to some but not all of the parties involved
  - Hidden information
  - Hidden action

Hidden Information

- Agents possess private information (e.g., individual preferences, costs)
- How to induce truthful revelation to compute allocation outcome?
  - e.g., auction: agents submit truthful bids; auctioneer receives all bids and determine winner and price
- Mechanism design
  - Sometimes referred to as inverse game theory

DAMD

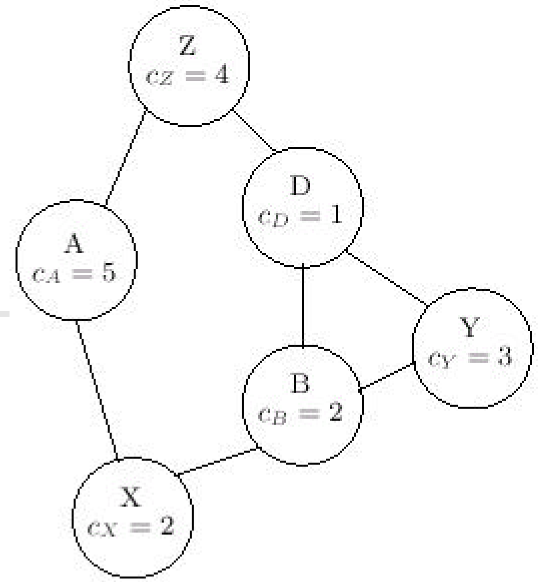

- Mechanism design (MD)
  - Centralized computation
- Distributed algorithmic mechanism design (DAMD)
  - Distributed computation
  - Computation and communication complexity
  - Internet applications [Feigenbaum02a]:
    - BGP routing [Feigenbaum02b] and
      Multicast cost sharing [Feigenbaum01]
    - P2P & overlay networks, web caching, distributed task allocation

Hidden Action

- Agents’ actions may be unobservable by principal
- Objective: the principal designs contract to induce desired action/behavior by the agents
- Also known in economics literature as the “moral hazard” problem

Hidden Action in Multi-hop Routing [Feldman04c]

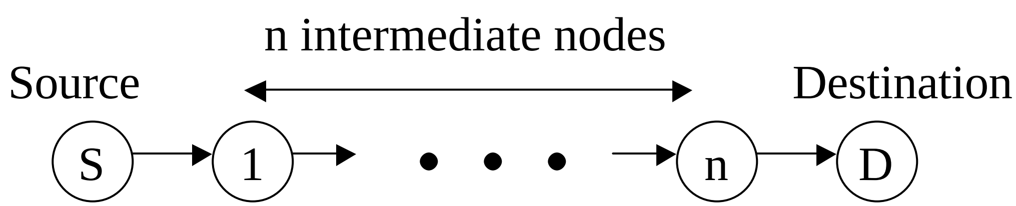

- Multi-hop routing requires cooperation by intermediate nodes
  - P2P overlay networks (e.g., DHT)
  - Wireless ad hoc networks
  - Inter-domain routing
- Intermediate nodes have disincentives to cooperate [Christin04]

Hidden Action in Multi-hop Routing

- Actions of intermediate nodes are hidden from the sender and receiver
  - Multi-hop:
    cannot attribute failure to a specific node
  - Stochastic outcome:
    external factors beyond the node's control
- Rational intermediate nodes may choose to forward packets at a low priority or not forward at all

Research Questions

- Is it possible to design contracts to induce cooperative behavior of intermediate nodes despite hidden-action?
- Under what circumstance, if any, might monitoring mechanisms be useful?
- What are the implications to network design?

Model

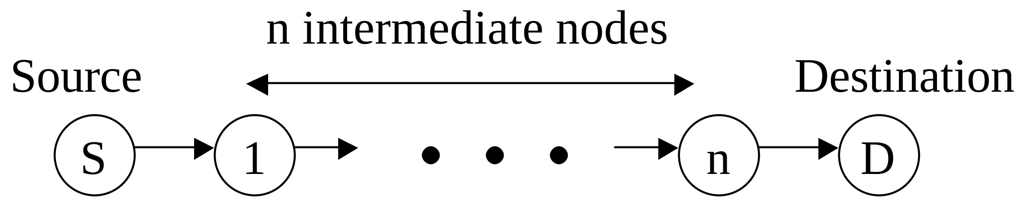

- Principal-agent model with multiple agents performing sequential hidden action
- Agents choose between high and low effort actions
  - Drop vs. forward
  - Best-effort vs. priority forwarding
- Principal can observe
  - Final outcome only (without monitoring)
  - Per-hop outcome (with monitoring)
- Principal signs contract with each agent; payment based on final outcome (without monitoring) or per-hop outcome (with monitoring)

Actions, Costs and Outcomes

- **Actions** $a_i \in \{0,1\}$:
  - Low-effort: $a_i=0$
  - High-effort: $a_i=1$
- **Costs** associated with actions:
  - $C(a_i=0) = 0$
  - $C(a_i=1) = c$
- **Outcomes** $X(a, k)= x \in \{x^L, x^H\}$
  - x^L: packet doesn't reach destination
  - x^H: packet reaches destination

Payments and Utilities

- Individual payments, s_i, depend on outcome
- Utility of participants:
  - Agent i: $U_i(s_i, c_i, a_i) = s_i - a_i c_i$
  - Principal: $W(x, S) = b(x) - S$, *where:* $S = \sum_{i=1}^{n} s_i$
- Principal needs to satisfy two constraints for each agent:
  - IR: individual rationality (participation constraint)
  - IC: incentive compatibility

Assumptions

- Transit cost, *c*, is common knowledge
- Topology is common knowledge
- Nodes are risk-neutral
- *(n+1)* per-hop transmission events are i.i.d.

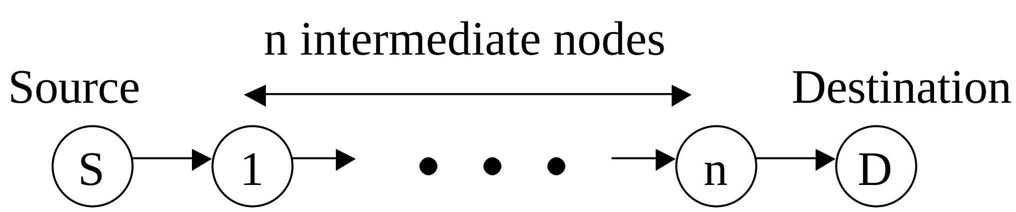

Results

- Scenario 1: drop vs. forward without monitoring
- Scenario 2: drop vs. forward with monitoring
- Scenario 3: best-effort vs. priority forwarding
- Scenario 4: multiple disjoint paths

Scenario 1: Drop Versus Forward without Monitoring

- Probability of a one-hop success:

$$\Pr(x^H_{i\to i+1} \mid a_i) = (1-k)a_i$$

- Principal observes only the final outcome
- Payment schedule to agent *i:* $s_i = (s_i^H, s_i^L)$ where:

$s_i^H = s_i(x = x^H)$ If packet reaches **destination**

$s_i^L = s_i(x = x^L)$ If packet does not reach **destination**

Scenario 1: Drop Versus Forward without Monitoring

Result: Under the best contract that induces high-effort behavior from all agents in a **Nash equilibrium:**

- Agent's expected payment = Agent's expected cost
- Principal achieves the first-best utility
- Payment schedule:

$$s_i^L = 0$$

$$s_i^H = \frac{c}{(1-k)^{n-i+1}}$$

Scenario 1: Drop Versus Forward without Monitoring

Proof sketch:

- IC constraint:

$$\Pr(x^H \mid a_{j\geq i}=1)s_i^H + \Pr(x^L \mid a_{j\geq i}=1)s_i^L$$

$$E[s]_{a_{j\geq i}=1} - c \geq E[s]_{a_i=0, a_{j>i}=1}$$

$$\Pr(x^H \mid a_i=0, a_{j>i}=1)s_i^H + \Pr(x^L \mid a_i=0, a_{j>i}=1)s_i^L$$

- IR constraint:

$$\Pr(x^H \mid a_{j\geq i}=1)s_i^H + \Pr(x^L \mid a_{j\geq i}=1)s_i^L \qquad s_i^L$$

$$\Pr(x^H_{S\to i} \mid a_{j<i}=1)(E[s]_{a_{j\geq i}=1} - c) + \Pr(x^L_{S\to i} \mid a_{j<i}=1)E[s]_{a_{i=0}, a_{j>i}=1} \geq 0$$

Scenario 1: Drop Versus Forward without Monitoring

Proof sketch (continued):

- IC and IR bind at the optimal contract

$$\boxed{\Pr(x^H)s_i^H + \Pr(x^L)s_i^L}$$

- Expected payment to node i: $E[s]_{a_{j=1}\forall j} = \boxed{(1-k)^i c}$
- Expected cost to node i: $\Pr(x^H_{S\to i})c = \boxed{(1-k)^i c}$

Scenario 2: Drop Versus Forward with Monitoring

- With per-hop monitoring, sender knows outcome of each per-hop transmission
- Scenario reduces to n instances of single principal – single agent problem
  - IC: $\boxed{E[s]_{a_i=1} - c \geq E[s]_{a_i=0}}$ → $\boxed{(1-k)s_i^H + ks_i^L - c \geq s_i^L}$
  - IR: $\boxed{E[s]_{a_i=1} - c \geq 0}$ → $\boxed{(1-k)s_i^1 + ks_i^0 - c \geq 0}$
- Principal obtains same utility as first-best contract
- n identical payment schedules:

$$s_i^L = 0$$

$$s_i^H = \frac{c}{1-k}$$

The Value of Per-Hop Monitoring

- The sender derives the same expected utility whether it obtains per-hop monitoring or not
- Yet, several differences

	Solution concept	**Location effect**	**Vulnerability to collusion**
Without monitoring	Nash equilibrium	Location dependent contracts	Not vulnerable
With monitoring	(Weak) dominant strategy	Location independent contracts	Vulnerable

Scenario 3: Best-Effort versus Priority Forwarding

- Priority forwarding reduces the loss rate
- Probability of a one-hop success:

$$\Pr(x_{i \to i+1}^{H} \mid a_i) = 1 - (k - q a_i)$$

where:

$$q \in (0,1] \quad \text{and} \quad k \in [q,1]$$

- Packet may reach the destination under low-effort actions, but with lower probability

Scenario 3: Best-Effort versus Priority Forwarding

- Result: sender derives same expected utility with or without monitoring
- At the optimal contract, the payment upon a failure is negative (transfer from agent to principal)
  - If limited liability constraint is imposed ($s \geq 0$), first-best cannot be achieved
- The sender may maximize its utility by signing a contract with only m out of the n nodes
  - Without monitoring: contract with nodes closest to destination, since expected cost decreases in i

Scenario 4: Multiple Disjoint Paths

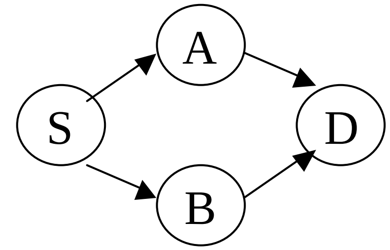

- Multiple disjoint paths exist from source to destination
- Sender elects to send multiple copies of the packets to maximize likelihood of delivery
- Two scenarios:
  - Per-path monitoring: has *a specific copy* of the packet reached destination?
  - No per-path monitoring: has *at least one copy* of the packet reached destination?
- Result: sender derives same expected utility whether it obtains *per-path* monitoring information or not

Discussion

- Appropriate design of contracts achieves cooperative behavior despite hidden-action
- Sender achieves first-best utility in Nash equilibrium in the absence of monitoring under several assumptions
- Per-hop or per-path monitoring:
  - Does not reduce implementation cost to sender under these assumptions
  - Achieves cooperative behavior in dominant strategy
  - Vulnerable to various forms of collusion
  - May yield some benefit under different assumptions, which may or may not justify its cost
- Implications to system design
  - Monitoring vs. contracting

Ongoing and Future Work

- Uniqueness of equilibrium
- Recursive contracts
- Relax assumptions:
  - Correlated transmission events (not i.i.d.)
  - Risk-averse agents
  - Topology and/or transit costs are not common knowledge
- More realistic monitoring mechanisms
- Collusive behavior
- Uncertainty with respect to choice and observability

Outline

- P2P system characteristics
  - Disincentives in sharing → free-riding
- Incentive mechanisms
  - Tokens, reputation, taxation, contracts, ...
  - Challenges: whitewashing, collusion, etc.
- Case study:
  - On-demand P2P streaming
  - Live event P2P streaming
- Information Asymmetry
  - Hidden action in multi-hop routing
- Conclusions

Conclusions

- Inherent decentralization of P2P systems brings incentives to the forefront
  - Peers not just obedient or malicious, but strategic
  - Collective welfare often misaligned with individual rationality
  - Significant challenges and opportunities in designing incentive mechanisms for diversity of P2P systems

Conclusions

- Economics-informed P2P system design
  - Game theory (mechanism design, evolution and learning, network formation)
  - Economics of asymmetric information (incentive and contract theory, agency theory)
  - Public finance
  - Theory on public goods and club goods
  - Social network theory
- Generalizable to various distributed and networked systems, including the Internet

Economics-Informed System Design

- Emerging multidisciplinary research communities
  - p2pecon
    - p2pecon’03: http://www.sims.berkeley.edu/p2pecon/
    - p2pecon’04: http://www.eecs.harvard.edu/p2pecon/
  - PINS
    - Practice and Theory of Incentives and Game Theory in Networked Systems
    - http://www.acm.org/sigs/sigcomm/sigcomm2004/pins.html
  - WEIS
    - Workshop on Economics and Information Security
    - WEIS’04: http://www.dtc.umn.edu/weis2004/

Bibliography

[Adar00] E. Adar and B. Huberman, Free Riding on Gnutella. First Monday 5(10), October 2000.

[Andreoni90] J. Andreoni, Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving.'' Economic Journal, v.100, June 1990, 464-477.

[Asvanund03] A. Asvanund, S. Bagla, M.H. Kapadia, R. Krishnan, M.D. Smith and R. Telang, Intelligent Club Management in Peer-to-Peer Networks. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

[Axelrod84] R. Axelrod, Evolution of Cooperation. Basic Books, 1984.

[Buchegger02] S. Buchegger, J.Y. Le Boudec, Performance Analysis of the CONFIDANT Protocol (Cooperation Of Nodes - Fairness In Dynamic Ad-hoc NeTworks). Proceedings of MobiHoc 2002, Lausanne, June 2002.

[Christin04] N. Christin and J. Chuang, On the Cost of Participating in a Peer-to-Peer Network, 3rd International Workshop on Peer-to-Peer Systems (IPTPS'04), February 2004.

[Chu04] Y.-H. Chu, J. Chuang, and H. Zhang, A Case for Taxation in Peer-to-Peer Streaming Broadcast. ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS), August 2004.

[Cohen03] B. Cohen, Incentives Build Robustness in BitTorrent. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

Bibliography

[Dingledine00] R. Dingledine, M.J. Freedman, and D. Molnar. The FreeHaven Project: Distributed anonymous storage service. Workshop on Design Issues in Anonymity and Unobservability, July 2000.

[Feigenbaum01] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. J. Computer and System Sciences, 63, 2001.

[Feigenbaum02a] J. Feigenbaum and S. Shenker. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, ACM Press, New York, 2002, pp. 1-13.

[Feigenbaum02b] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In Proc. 21st ACM Symposium on Principles of Distributed Computing, New York, NY, 2002.

[Feldman03] M. Feldman, K. Lai, J. Chuang and I. Stoica, Quantifying Disincentives in Peer-to-Peer Networks. 1st Workshop on Economics of Peer-to-Peer Systems, Berkeley CA, June 5-6 2003.

[Feldman04a] M. Feldman, K. Lai, I. Stoica, and J. Chuang, Robust Incentive Techniques for Peer-to-Peer Networks. ACM E-Commerce Conference (EC'04), May 2004.

Bibliography

[Feldman04b] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica, Free-Riding and Whitewashing in Peer-to-Peer Systems. ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS), August 2004.

[Feldman04c] M. Feldman and J. Chuang, Hidden-Action in Multi-Hop Routing. 2nd Workshop on Economics of Peer-to-Peer Systems, June 2004.

[Friedman01] E.J. Friedman and P. Resnick. The social cost of cheap pseudonyms. Journal of Economics and Management Strategy, 10(2):173-199, 2001.

[Habib04] A. Habib and J. Chuang, Incentive Mechanism for Peer-to-Peer Media Streaming. 12th IEEE International Workshop on Quality of Service (IWQoS'04), June 2004.

[Kamvar03] S. Kamvar, M. Schlosser, and H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. In WWW 2003, 2003.

[Kung03] HT Kung and Chun-Hsin Wu, Differentiated Admission for Peer-to-Peer Systems: Incentivizing Peers to Contribute Their Resources.

[Lai04] K. Lai, B.A. Huberman and L. Fine, Tycoon: A Distributed Market-based Resource Allocation Systems", HP Labs Technical Report cs.DC/0404013, http://arxiv.org/abs/cs.DC/0404013, April 5, 2004.

[Levien98] R. Levien and A. Aiken. Attack-resistant trust metrics for public key certification. 7th USENIX Security Symposium, 1998.

Bibliography

[Marti04] S. Marti, P. Ganesan, H. Garcia-Molina, DHT Routing using Social Links, IPTPS 2004.

[Nowak98] M.A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. Nature, 393:573--577, 1998.

[Padmanabhan02] V. Padmanabhan and K. Sripanidkulchai. The case for cooperative networking. In Proc. of 1st International Workshop on Peer-to-Peer Systems (IPTPS '02), Cambridge, MA, USA, March 2002.

[Reiter99] M. Reiter and S. Stubblebine, Authentication Metric Analysis and Design, ACM Trans. Information and System Security, vol. 2, no. 2, pp. 138-158, 1999.

[Sariou02] S. Saroiu, P.K. Gummadi, and S. Gribble, A measurement study of peer-to-peer file sharing systems. In Proceedings of Multimedia Computing and Networking 2002.

[Vishnumurthy03] V. Vishnumurthy, S. Chandrakumar and E.G. Sirer, KARMA : A Secure Economic Framework for Peer-To-Peer Resource Sharing. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

[Wilcox-O'Hearn02] B. Wilcox-O'Hearn, Experiences Deploying a Large-Scale Emergent Network. IPTPS 2002.