

Incentives for Cooperation in the Internet


Kevin Lai, Michal Feldman,
John Chuang, Ion Stoica

Computer Science and SIMS

U.C. Berkeley

Cooperation and P2P

- Higher cooperation → higher P2P performance
 - file sharing: more sharing → lower latency
- But, cooperation is not automatic
 - freeriding
- Need an incentive system to increase cooperation

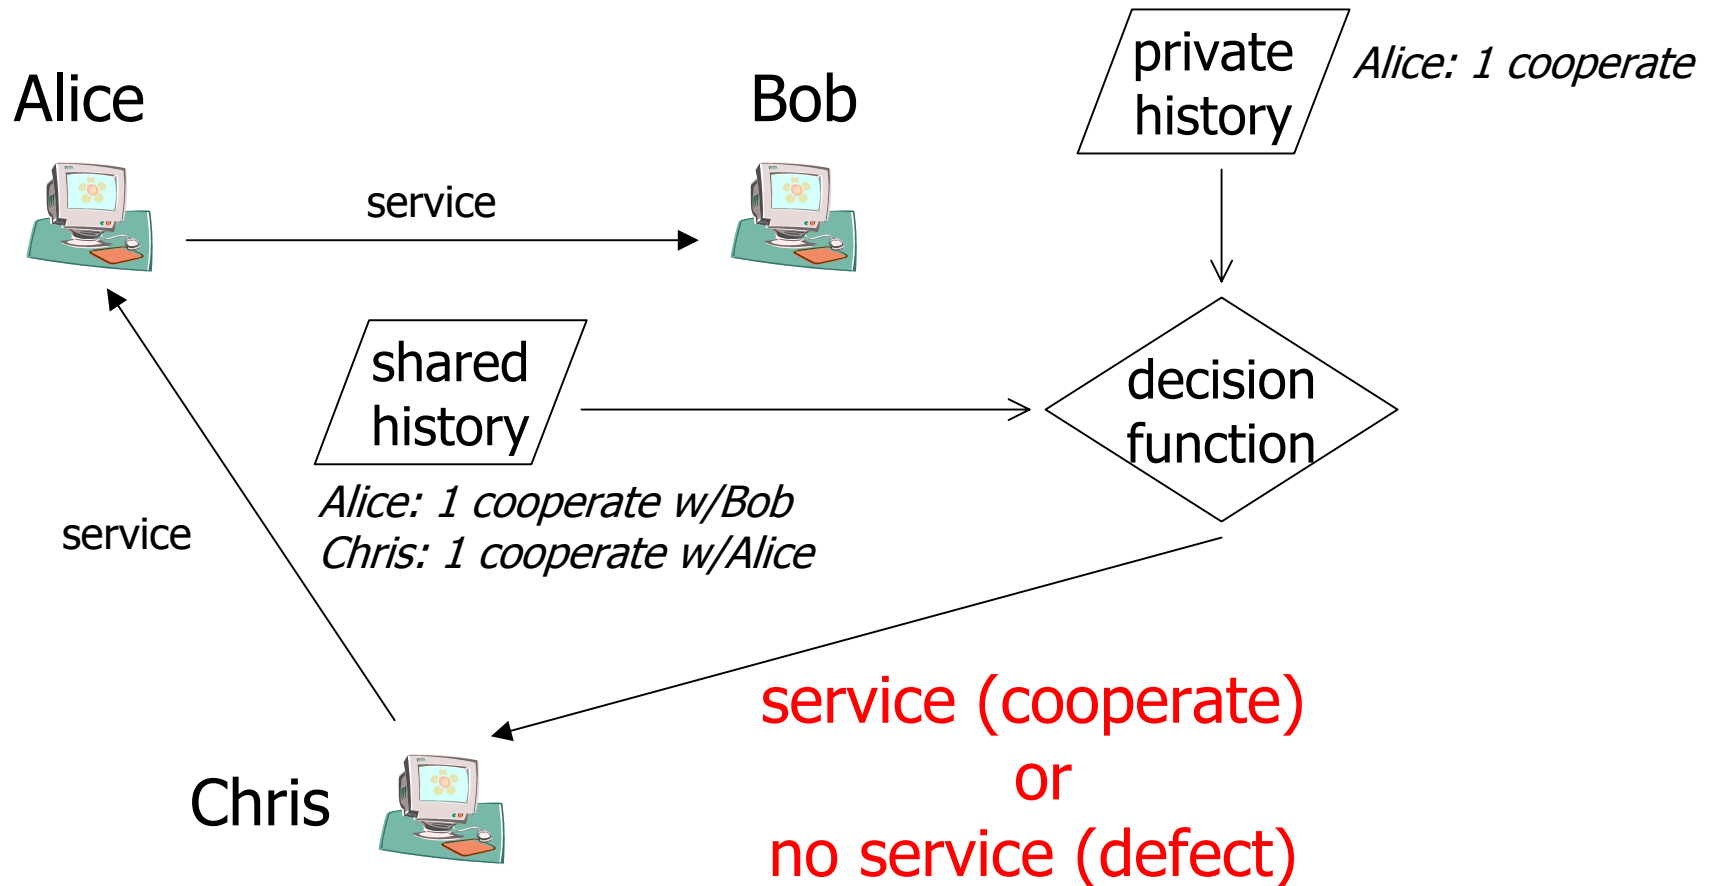


Oath

- Goal: design an incentive system applicable to many applications
 - P2P storage
 - wireless adhoc forwarding
 - grid/utility computing
 - online auctions
 - discussion forums
- First step (of many): understand the tradeoffs in the design space of solutions

Oath Architecture

- Keep track of nodes' actions and give them what they themselves have given



Internet Incentive Challenges

- **scalability**

- file sharing networks have >100,000 participants

- **zero-cost identities**

- being able to easily change identities subverts history

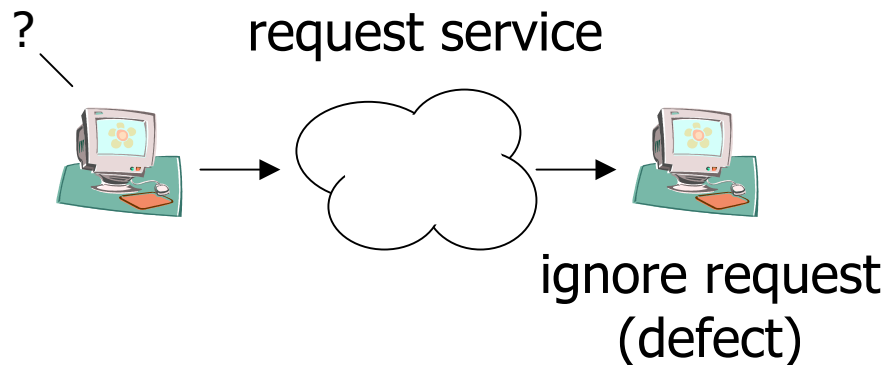

- **collusion**

- cannot verify records in shared history

Outline

- Model
 - Generalized Evolutionary Prisoner's Dilemma
- Decision Function
 - Reciprocative
- Scalability
 - benefit of shared history
- Zero-cost identities
 - adapt to friendliness of strangers
- Collusion
 - subjective reputation

Generalized Evolutionary Prisoner's Dilemma (GEPD)



- Client requests service
- Server chooses to serve or not
 - based on its strategy and history
- Client benefits from service
- Server pays service cost
- Client cannot trace defections
- Peers change to higher scoring strategies in proportion to the difference in scores

GEPD Properties

- Defection is dominant action for 1-shot game
- Universal defection ruins overall score
- Captures essential tension of cooperative applications
- Flexible
 - assignment of payoff matrix
 - definition of cooperation and defection
 - behavior of strategies

Outline

- Model
 - Generalized Evolutionary Prisoner's Dilemma
- Decision Function
 - Reciprocative
- Scalability
 - benefit of shared history
- Zero-cost identities
 - adapt to friendliness of strangers
- Collusion
 - subjective reputation

Decision Function

- Require: converge to cooperation, robust against defection strategies
- Tit-for-Tat: *do to the peer what he last did to me*
 - not effective with shared history, requires tracing of server defections
- Reciprocative: *Cooperate with entity X with probability*

$$\min\left(\frac{\textit{cooperation_X_has_given}}{\textit{cooperation_X_has_received}}, 1\right)$$


- can use shared history, does not require tracing of server defections

Outline

- Model
 - Generalized Evolutionary Prisoner's Dilemma
- Decision Function
 - Reciprocative
- Scalability
 - benefit of shared history
- Zero-cost identities
 - adapt to friendliness of strangers
- Collusion
 - subjective reputation


Private History

- 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

- advantages
 - tolerates few repeat transactions (large populations, high turnover)
 - tolerates asymmetry of interest
- disadvantages
 - susceptible to collusion
 - implementing write-once abstraction requires overhead or centralization:
 - e.g., DHT-based storage w/replication



Outline

- Model
 - Generalized Evolutionary Prisoner's Dilemma
- Decision Function
 - Reciprocative
- Scalability
 - benefit of shared history
- Zero-cost identities
 - adapt to friendliness of strangers
- Collusion
 - subjective reputation

Zero-Cost Identities

- History assumes that entities maintain persistent identities
- Problem: most online systems have zero-cost identities
 - lowers bar to entry
 - allows pseudonymity through multiple identities
 - circumvents history-based strategies that always cooperate with strangers
- Whitewash: *100% defection, continuously changes identity*

Stranger Policies

- Always defect
 - forces newcomers to allow exploitation by existing players
 - raises bar for entry
- Adaptively cooperate
 - separately estimate stranger friendliness

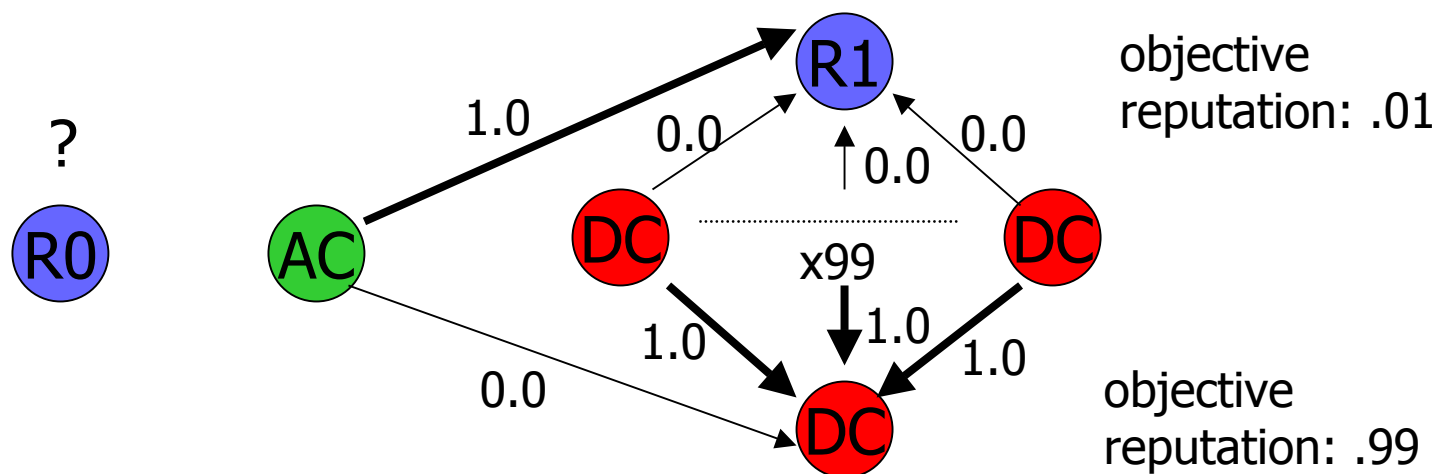
$$\min\left(\frac{\textit{cooperation_strangers_have_given}}{\textit{cooperation_strangers_have_received}}, 1\right)$$

- only taxes newcomers when necessary
- achieves highest level of cooperation

Outline

- Model
 - Evolutionary Prisoner's Dilemma (EPD)
- Decision Function
 - Reciprocative
- Scalability
 - benefit of shared history
- Zero-cost identities
 - adapt to friendliness of strangers
- Collusion
 - subjective reputation

Collusion



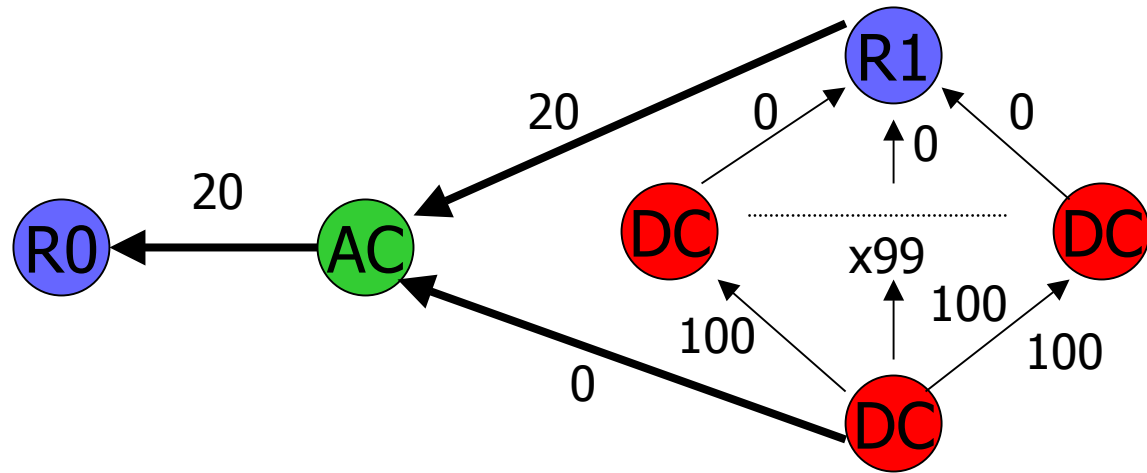
- Secure shared history can still be subverted
 - positive collusion
 - Defecting Colluder: *100% defect and claim other colluder gave +100 cooperation*
 - combine with zero-cost identities = Sybil attack [Douceur 2002]
- Most existing reputation systems are vulnerable or depend on trusted infrastructure

Subjective Reputation

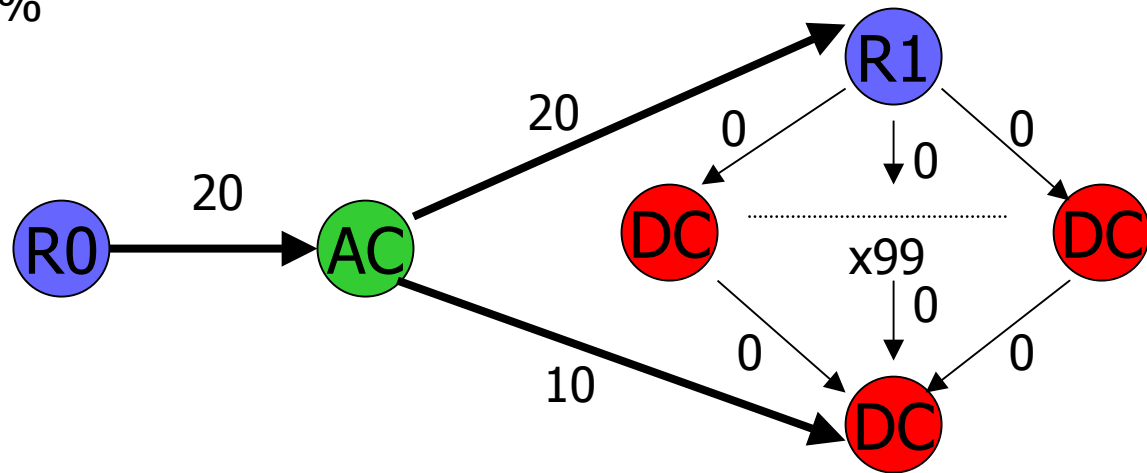
- Objective reputation is meaningless
- Need to account for who is reporting history
 - weigh nodes by how much they have contributed to source
- Calculate how much sink has benefited source, however indirectly
- Compute *max flow* from source to sink
 - *max flow*: using any number of paths, compute the maximum capacity from the source to the sink

$$\min\left(\frac{\text{MAXFLOW}_{\text{from } X \text{ to } Y \text{ of cooperation } X \text{ has given}}}{\text{MAXFLOW}_{\text{from } Y \text{ to } X \text{ of cooperation } X \text{ has received}}, 1\right)$$

Subjective Reputation Example



max flow DC→R0: 0
max flow R0→DC: 10
R0 cooperates w/DC: 0%



Subjective Properties

- resists any number of colluders
- fully decentralized
 - no trusted peers, infrastructure, etc.
- running time
 - worst case: $O(VE)$
 - incrementally: $O(V)$
 - bounded accuracy: $O(1)$

Recent Related Work

- Prisoner's Dilemma
 - Ranganathan, Ripeanu, Sarin, Foster
- Shared History / Distributed Reputation
 - Dingleline, Mathewson, Syverson
 - Dutta, Goel, Govindan, Zhang
 - Vishnumurthy, Chandrakumar, Sirer
- Stranger Policy
 - Rosenthal, Roussopoulos, Maniatis, Baker
- Intelligent Selection
 - Asvanund, Bagla, Kapadia, Krishnan, Smith, Telang

Conclusion

- Gain the benefits of shared history
 - scalability, tolerates high turnover
- Without the vulnerabilities
 - adaptive stranger policy
 - subjective reputation