Incentives for Sharing in P2P Networks

P. Golle, K. Leyton-Brown
I. Mironov, M. Lillibridge

Speaker: Georgiana Ifrim
Advisor: Jens Graupmann

Outline

- Introduction
- A game theoretic model
- Payment schemes
- Experiments
- Conclusions

Introduction

- The *free-rider* problem
  - taking advantage of the network without contributing to it
    - Napster: 60% peers share only 20% files
    - Gnutella: 70% do not share any

Motivation

- Providing incentives for peers to make active contributions to the network
- If the individual components are selfish can we somehow get good aggregate behavior?
- A need and an opportunity to improve the P2P file sharing systems

Model

- The model proposed addresses file sharing systems that make use of centralized servers
  - maintain a database of the files currently available on the network
  - connect dowload requests with available clients

Defining a 'Game' for P2P Sharing

- non-cooperative game among rational and strategic players
- n ‘agents’ (peers): $a_1,..,a_n$
- each agent has a number of possible ‘strategies’
  - agent a_i has strategy $S_i = (\sigma, \delta)$; 2 'actions':
    - σ = sharing
    - δ = downloading
- the strategies chosen determine the ‘outcome’
- associated with each outcome is a collection of ‘payoffs’, one to each agent

Game Setup

- **Sharing** : Agents select what proportion of files to share in three levels: σ_0 (none), σ_1 (moderate) and σ_2 (heavy)

- **Downloading** : Each agent determines how much to download from the network in three levels: δ_0 (none), δ_1 (moderate) and δ_2 (heavy)

- **Agent Utility** : Agents' utility functions describe their preferences for outcomes.

Game Setup

- **Factors** :
  - Positive: Amount Downloaded (AD), Network Variety (NV), Altruism (AL)
  - Negative: Disk Space Used (DS), Bandwidth Used (BW)
  - Financial Transfer (FT)

Game Setup

- **Agent ai's utility function** :
- $U_i=[f_i^{AD}(AD)+f_i^{NV}(NV)+f_i^{AL}(AL)]-[f_i^{DS}(DS)+f_i^{BW}(BW)]-FT$
- f-functions
  - associated with:
    - an agent
    - a particular variable
  - describe that agent's preference for different values of the variable, in money

Game Setup

- **Assumptions:**
  - **agents' relative preferences for outcomes:**
  - $f^{AD}(k) > k*\beta$
    - the utility agents gain from downloading k files is more than what they pay; β = cost per file
  - $f^{DS}(k)+f^{BW}(k) < k*\beta$
    - the cost to agents of sharing and uploading k files is less than what they are paid; β = reward per file

Equilibria

- **Assumptions:**
  - **agents**
    - have the same type (same f-functions)
      - it is enough to analyze the choice made by a single agent
    - economically rational
    - act to maximize expected utility w.r.t knowledge about other agents’ actions and their own payoffs

Equilibria

- **Weak Equilibrium**
  - No agent can ‘gain’ by changing his strategy
- **Strict Equilibrium**
  - Every agent is strictly worst off if he changes strategy
- **Dominant Strategy** (of an agent)
  - the agent's best action does not depend on the action of any other agent

Micro-Payment Mechanism

- **Scheme:**
  - charge downloads, reward uploads
  - central server tracks the number (per user)
    - d = downloads
    - u = uploads (downloads by other agents)
    - for a given period of time
  - after each period, users are charged
    - $C = g(d - u)$
    - linear with coefficient β (cost/reward per file)

Micro-Payment Mechanism

- In a time period, let
  - σ^{-i} = total number of files shared by others
  - δ^{-i} = total number downloaded by others
  - agent a_i chooses (σ_s, δ_d); s = # units shared; d = # units downloaded; n agents; β = cost per unit downloaded
- a_i's expected payment to the system

$$E[FT] = \beta * \left(d - \delta^{-i} * \frac{s}{\frac{n-2}{n-1} * \sigma^{-i} + s} \right)$$

  - server matches downloaders uniformly at random with shared units; no agent will download from himself

Micro-Payment Mechanism

- Analysis

$$E[FT]=\beta*\left(d-\delta^{-i}*\frac{s}{\frac{n-2}{n-1}*\sigma^{-i}+s}\right)$$

  - $f^{AD}(1) > \beta$
    - utility gained from downloading one file exceeds the cost (incentive for downloading)
  - $f^{DS}(1) + f^{BW}(1) < \beta$
    - cost incurred from sharing and uploading less than the gain (incentive for sharing)
- Results in strict and unique equilibria
  - $\Sigma=((\sigma_2, \delta_2),\ldots,(\sigma_2,\delta_2))$

Micro-Payment Mechanism

- Advantages:
  - unique strict equilibrium:
    - share and download maximally

- Disadvantages:
  - equilibrium doesn't hold for risk averse agents
  - users can make a profit
  - users dislike micro-payments

Quantized Micro-Payment Mechanism

- **Scheme:**
  - charge a fixed price for each block of *b* files downloaded
  - reward uploads as before
  - round up number of files downloaded after each period to next multiple of *b*
- Advantages:
  - may be preferable to users (flat pricing)
  - unique strict equilibrium as before

Quantized Micro-Payment Mechanism

- Disadvantages:
  - users can redirect their zero-marginal cost download to credit their friends with uploads
  - Proposals:
    - hide identities of users
    - reply to searches with random subsets

Points-Based Mechanism

- **Scheme:**
  - 'points' currency: points can be bought (with money or contribution), but not sold
  - penalize downloads, pay agents for size of material shared

Rewarding Sharing

- Agents' payment for sharing

$$\int M(t)\,dt$$

  - M(t) the amount of data in megabytes available for download at time t
- Downloading a file costs c*m points
  - m = file's size in megabytes
  - c = system constant
    - How long a new file must be shared to waive its download cost

Rewarding Sharing

- Analysis
  - Assume each file is exactly 1MB
  - Each agent shares for 1 period
  - Each level of sharing earns 1 point per period
    - e.g. σ_2=2 points
  - Each level of downloading costs 1 point (c=1); one point costs β
  - Downloaders are matched uniformly at random with shared units; no agent may download from himself

Rewarding Sharing

- Analysis
  - expected number of uploads: $E[u]=\delta^{-i}*\frac{s}{\frac{n-1}{n-2}*\sigma^{-i}+s}$
  - n-1 agents play $S=(\sigma_2, \delta_2)$
  - agent a_i's strategy:
    - $f^{AD}(k) > k*\beta$
      - δ_2 dominates $\delta 1$ and $\delta 0$
    - $f^{DS}(k) + f^{BW}(k) < k*\beta$
      - agents prefer to share and upload at level k, than to pay the system for k points

Rewarding Sharing

- Advantages
  - no agent makes a profit
  - maximal sharing, downloading is a strict equilibrium
- Disadvantages
  - no sharing, maximal downloading is also a strict equilibrium
  - agents don't want their shared files to be downloaded (BW – negative utility)

Rewarding Sharing

- share at off-peek time, share unpopular files
  - solution:

$$\int M(t)\lambda(t)dt$$

    - $\lambda(t)$ scaling factor proportional to expected demand

Experiments

- Validate and enrich the theoretic model
  - levels of risk-aversion
  - different utility functions (characterize agents)
  - different types of files
- Experimental results
  - strategy convergence in this richer setting
  - interesting effects

Experimental Setup

- Types of agents
  - Altruism
    - Uniformly random from $[AL_{min}, AL_{max}]$
  - Disk space
    - Uniformly random from $[DS_{min}, DS_{max}]$
  - File type preference
    - Weighted combination of file types
  - Other parameters: fixed and equal for all agents

Experimental Setup

- Simulation:
  - multi-agent reinforcement learning model
  - TD Q-learning algorithm
    - agents learn the expected utilities of (state, action)-pairs
    - strategy convergence corresponds to a Nash equilibrium

Strategy Convergence

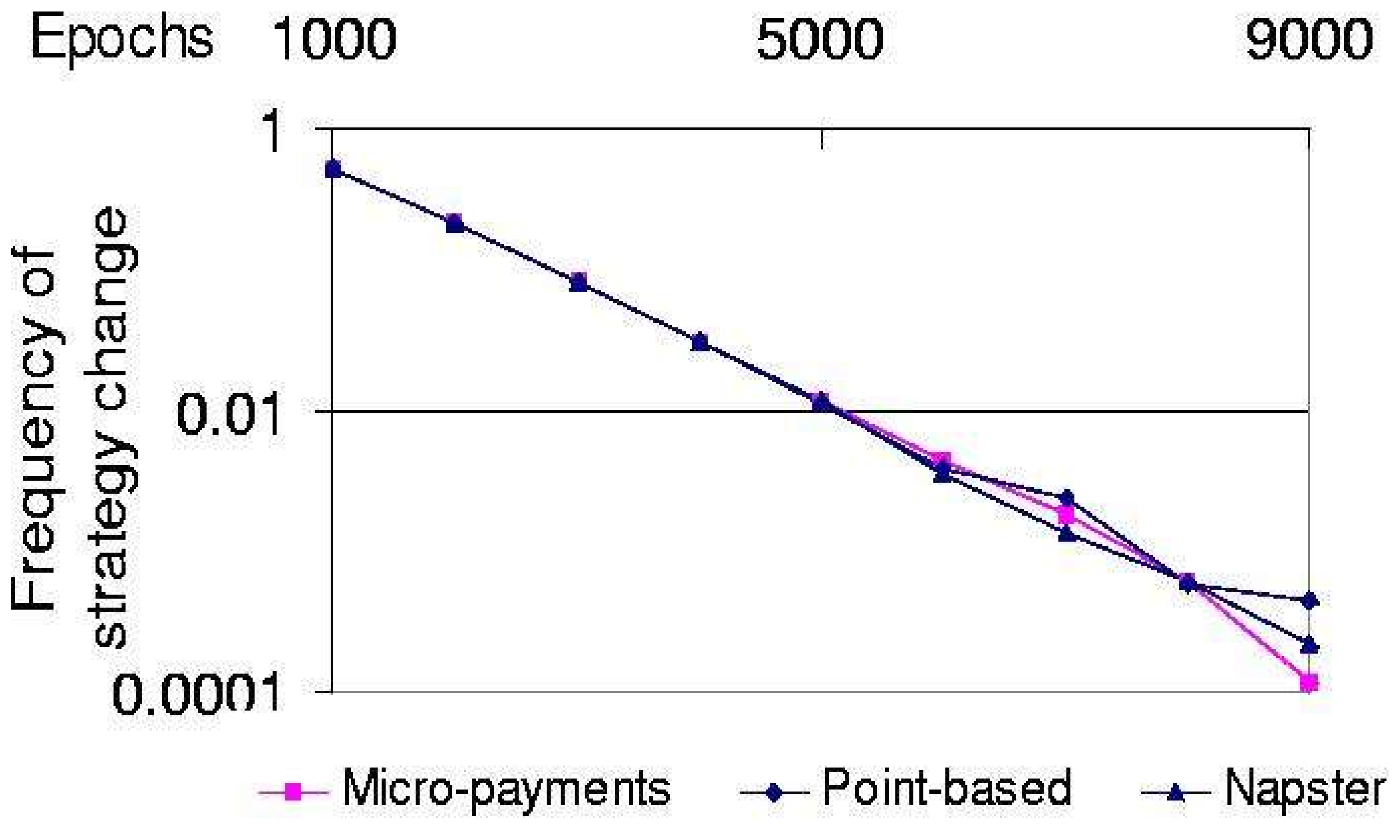

Points:
Effect of Altruism on Sharing

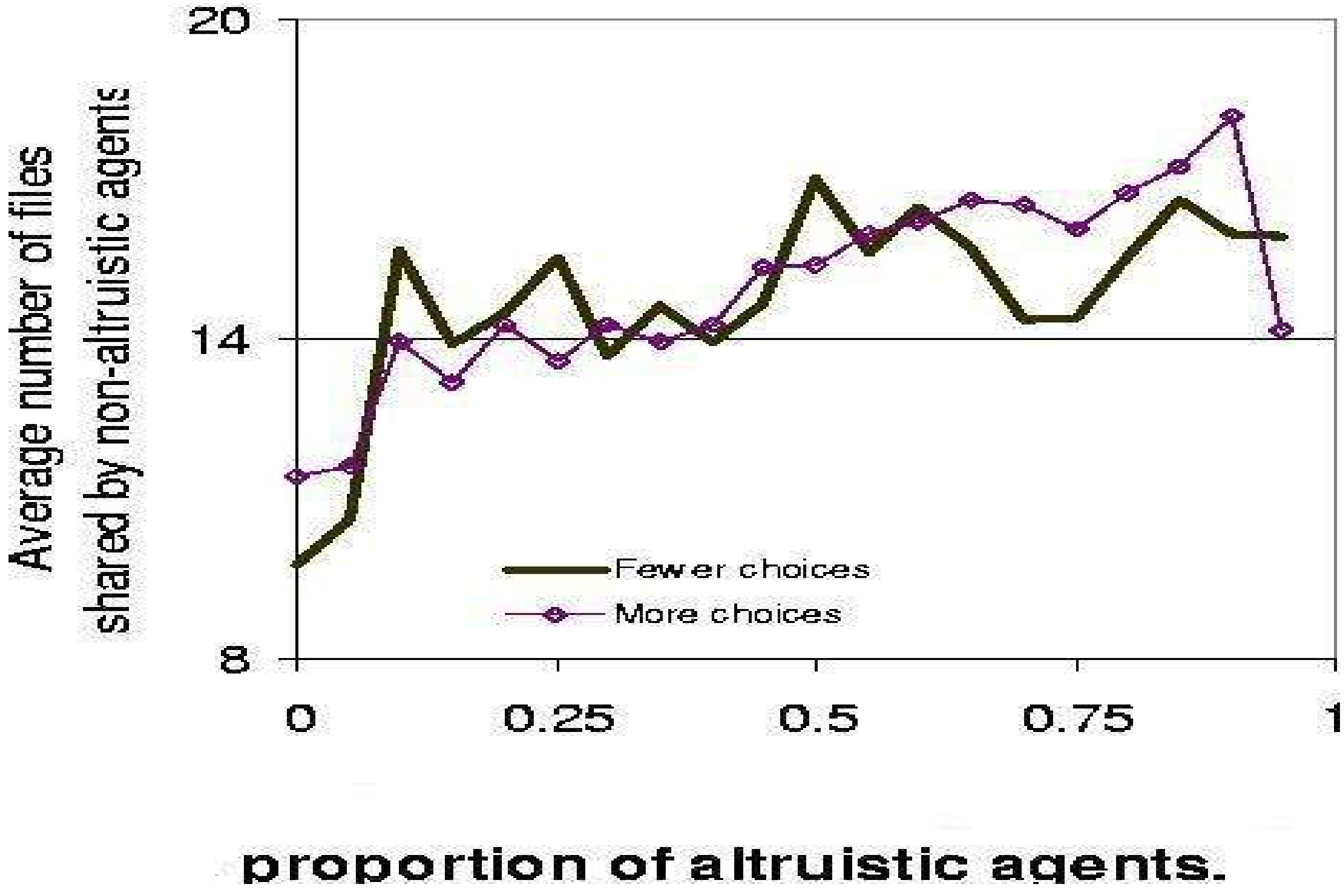

Micro-Payments: Effect of Risk Aversion on Sharing

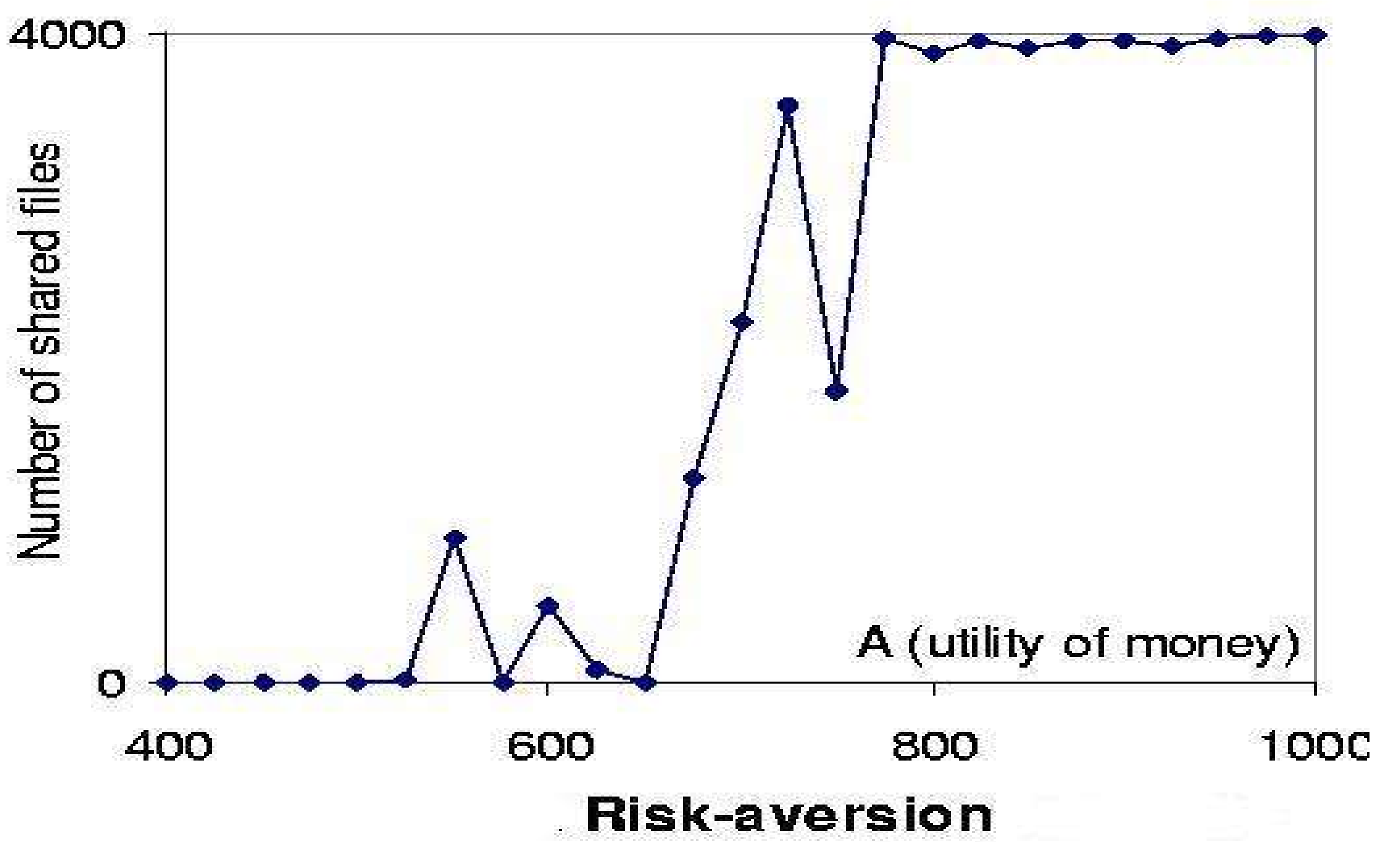

smaller values of A = greater risk aversion

Conclusions

- Model:
  - a game-theoretic model for centralized P2P file sharing systems
- Theory:
  - three payment schemes that give rise to equilibria in which free-riding does not occur, pros & cons
- Experiments:
  - showed convergence to the same equilibria in an enriched model; also some non-trivial behaviors

Thank you!

Questions?