Exchange-based Incentive Mechanisms

for Peer-to-Peer File Sharing

Kostas G. Anagnostakis, Michael B. Greenwald

J. Marc Roth

jmroth@mpi-sb.mpg.de

MPI INFORMATIK

Peer-to-Peer Information Systems

January 2005

Overview

- Introduction to Incentive Mechanisms
- Exchange Mechanisms
  Exchange Transfers
  Preventing Cheating
- Simulation & Results
- Measurements, Discussion & Related Work
- Summary & Conclusion

Introduction

- powerful infrastructure
- performance depends on level of cooperation
- non-cooperation may have severe impacts
- **solution: incentive mechanisms**
  = stimulate cooperation
  (reward people who contribute to the system)

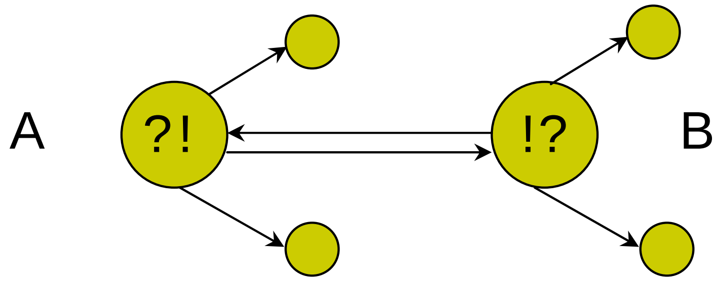

Incentive Mechanisms

- participation level (e.g. KaZaA)
- credit system (monetary economy)
  centralized vs. decentralized
- **proposal:**
  exchange system (barter economy)
  - users trade resources between themselves
  - high priority for contributing users
  - not only 2-way but N-way exchanges

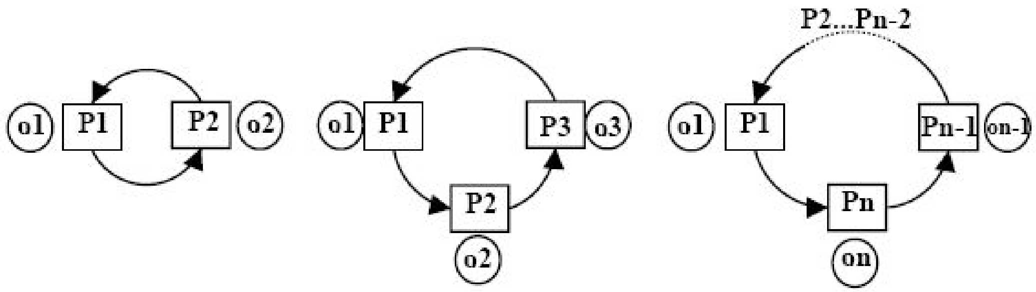

2-way 3-way N-way

Exchange Mechanisms (I)
Basics

- fixed upload and download capacity
- partial transfers
- we ignore the object lookup ☺
- each peer has an IRQ (incoming request queue) = upload queue

Exchange Mechanisms (II)

Rules

- transfer is initiated if
  1. local peer has free upload capacity (slot)
  2. transfer is an exchange transfer (ET)
     OR no ETs in IRQ (incoming request queue)
- upload slots are preemptively reclaimed by ETs!
- fixed-size block transfers
- termination of transfer if:
  - a peer disconnects
  - source deletes object
  - 1st transfer completed

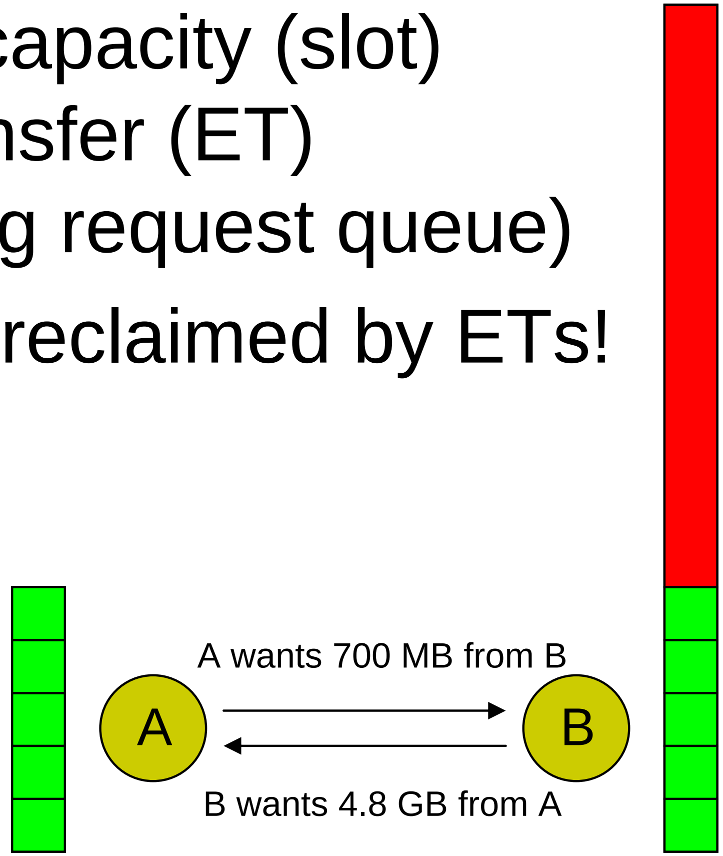

Exchange Transfers (I)

- identify feasible exchanges by looking at IRQ
- 2-way exchanges are simple but frequently do not resolve into convenient pairs
- compute feasible N-way exchanges
  e.g. cycles in (potentially enormous) graph
  G = (nodes, edges) = (peers, requests)

N <= 5 is sufficient

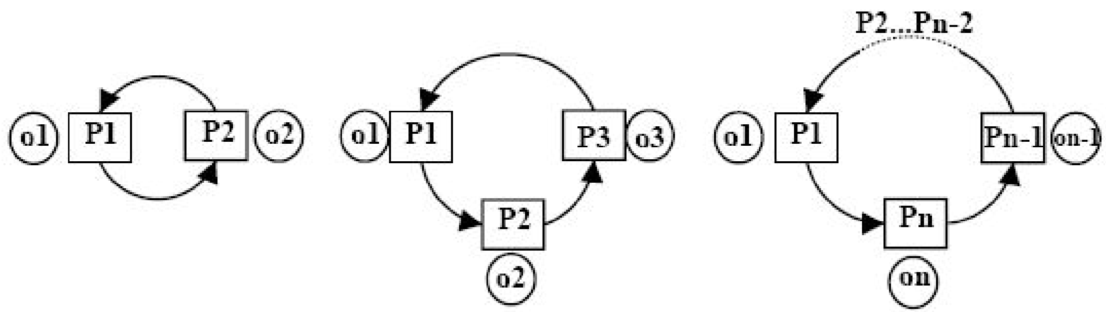

2-way 3-way exchanges N-way

Exchange Transfers (II)

- each peer maintains request tree (RT)
  - empty IRQ → empty RT
  - non-empty IRQ → includes other peer's RTs in their incoming requests
- each peer
  - provides object to predecessor
  - gets object from successor
- A inspects RT
  - before transmitting
  - after receiving

  any request

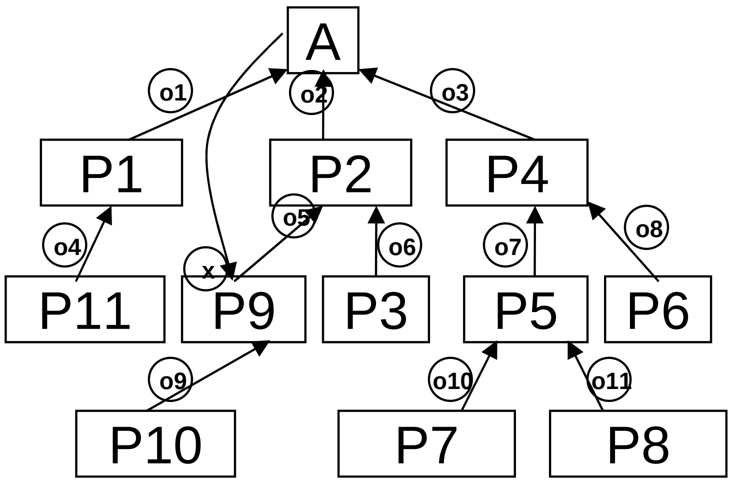

***P9* has object *x* available for *A*.**
The entire request tree is shown.
The cycle for the 3-way exchange that *A* tries to initiate is shown in red.

Exchange Transfers (III)

in practice

- circulate token
- invalidation of the ring if
  - peers offline or crashed
  - member peers have created own rings
- token negotiates transfer rate
- least transfer rate is used and excess capacity is transferred to other exchanges

Exchange Transfers (IV)

how to choose

- larger rings 👎
  - more peers are served
  - high probability for loss of peer
- smaller rings 👍
  - lower search cost
  - higher expected exchange volume
- peers usually care less about global performance than about their own benefit

Preventing Cheating (I)

- claim to be exchange but serve junk
  - local blacklists (bad)
  - cooperative blacklists (medium)
- cheap pseudonyms
- limiting the damage:
  - exchange blocks synchronously with checksum
  - bad performance: $b_{exchange} / t_{rtt}$ bytes/sec
  - use window protocol
  - increase window size if good peer

  → positive effect

Preventing Cheating (II)

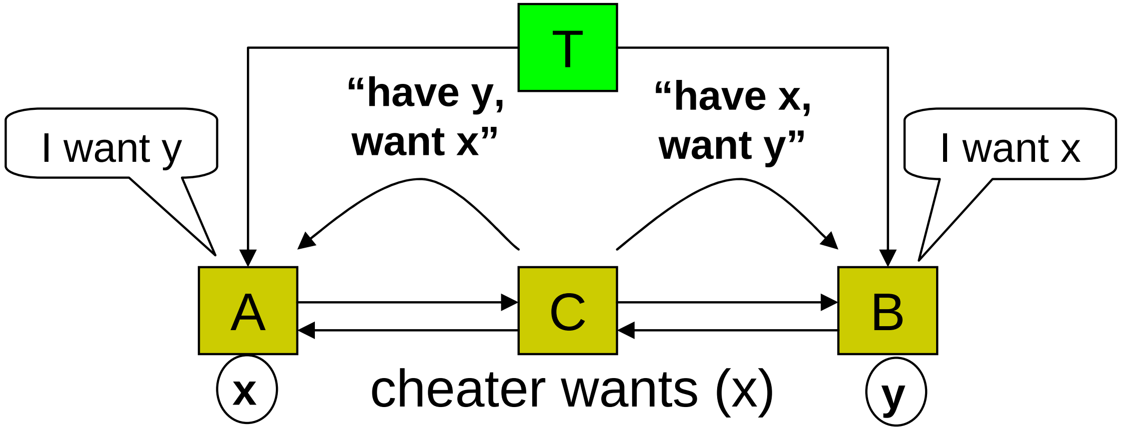

peer	upload	has	wants
A	5	x	y
B	5	y	x
C	10	-	x

- man-in-the-middle attack
- C gets high-priority service but does not contribute to the system ☹
- bidirectional encryption of transfer using secret key
- trusted peer is mediator and verifies data

Preventing Cheating (III)

- self-interest vs. maliciousness
  - solution with better performance at a lower cost
  - useful for system
  - respects desire not to participate
-

peer	upload	has	wants
A	10	-	x
B	5	x	y
C	10	y	x
D	10	y	x

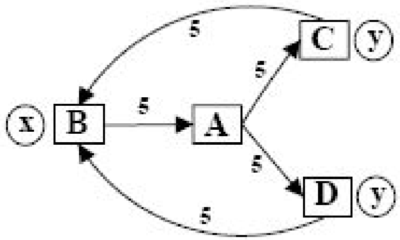

Figure 3. Example middleman scenario resulting in non-ring exchange

- generalization to non-ring topologies: not here!

Simulation

- 200 node file sharing system
- 50% freeriders (freeloaders)
- fixed + asymmetric down-/upload capacity
- neglect delay and loss ☺

Simulation

Object Popularity Model

- uniformly assign subset *m* of total categories to each peer
- (global) popularity rank for each category *c* of rank *i*

$$F_c^i = \frac{1}{1 + i \cdot f_c}, i \in \{0, \ldots, m-1\}$$

probability of request
for object in category *c*

$$p_c^i = \frac{F_c^i}{\sum_{k=0}^{m-1} F_c^k}$$

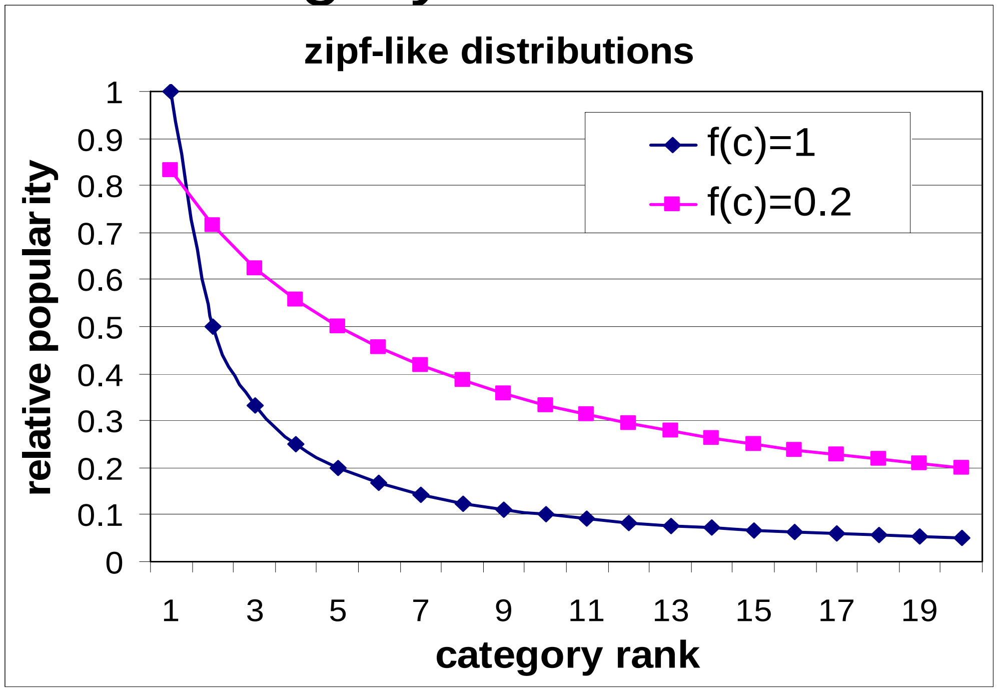

- distribute objects in categories like categories at peers
- uniformly random local preference
  for each category (independent of its global popularity)

Simulation

Setup

- maximum number of pending requests
- request rate is reached and held
- maximum # of objects + cleanup

number of peers	200
download capacity	800 kbit/s
upload capacity	80 kbit/s
ul/dl slot size	10 kbit/s
content categories	300
objects per category	uniform(1,300)
categories/peer	uniform(1,8)
category popularity	f=0.2
object popularity	f=0.2
object size	20 MB (all objects)
storage capacity per peer (nr. of objects)	uniform(5,40)
queue for incoming requests	1000
max pending objects	6
fraction of freeloaders in system	50%

Simulation Results (I)

key metric = download time!

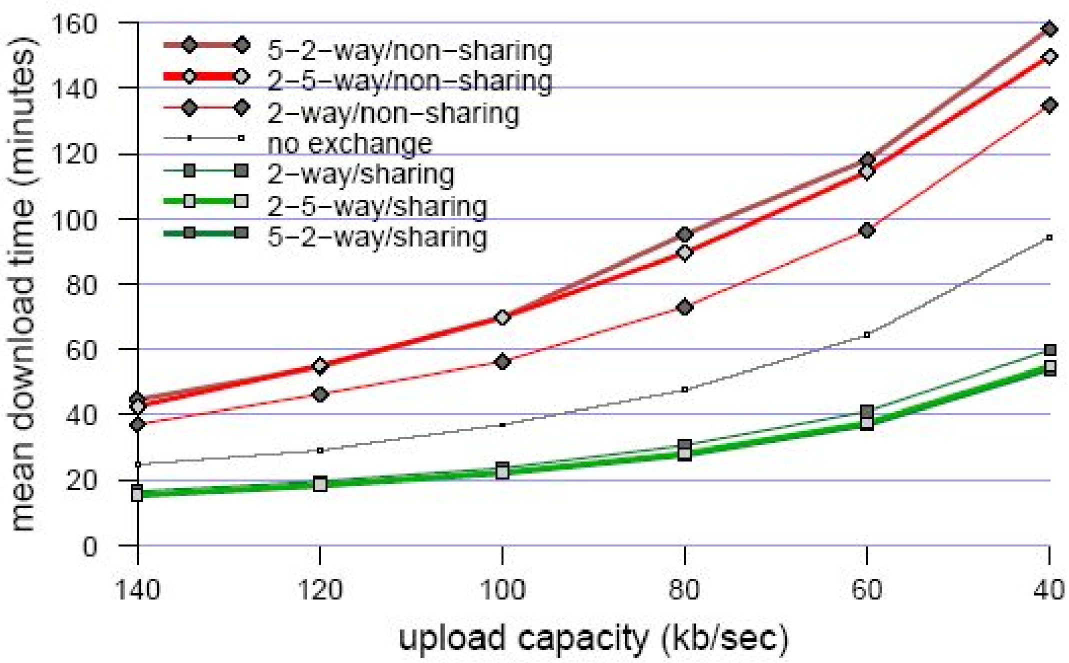

Figure 4. Mean download time vs. upload capacity and exchange policy.

- reduced upload capacity → longer download time
- time to completion increases faster for non-sharing users because exchanges are prioritized

➔ good incentive to deploy the proposed exchange mechanism

Simulation Results (II)

higher order rings + network size

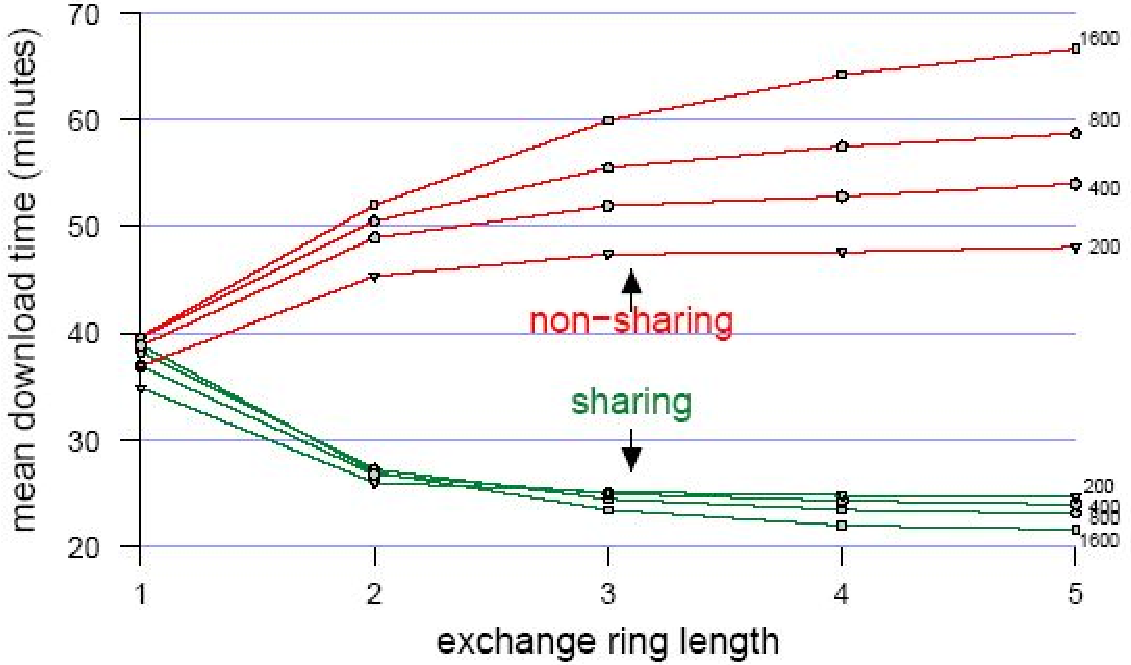

Figure 5. Mean download times vs. maximum exchange size and the number of peers in the network.

N = exchange ring length

- N=5 better than N=2
- as the network grows, difference in performance increases (sharers vs. non-sharers)
- N>5: no real improvement

Simulation Results (III)

object popularity distribution and performance

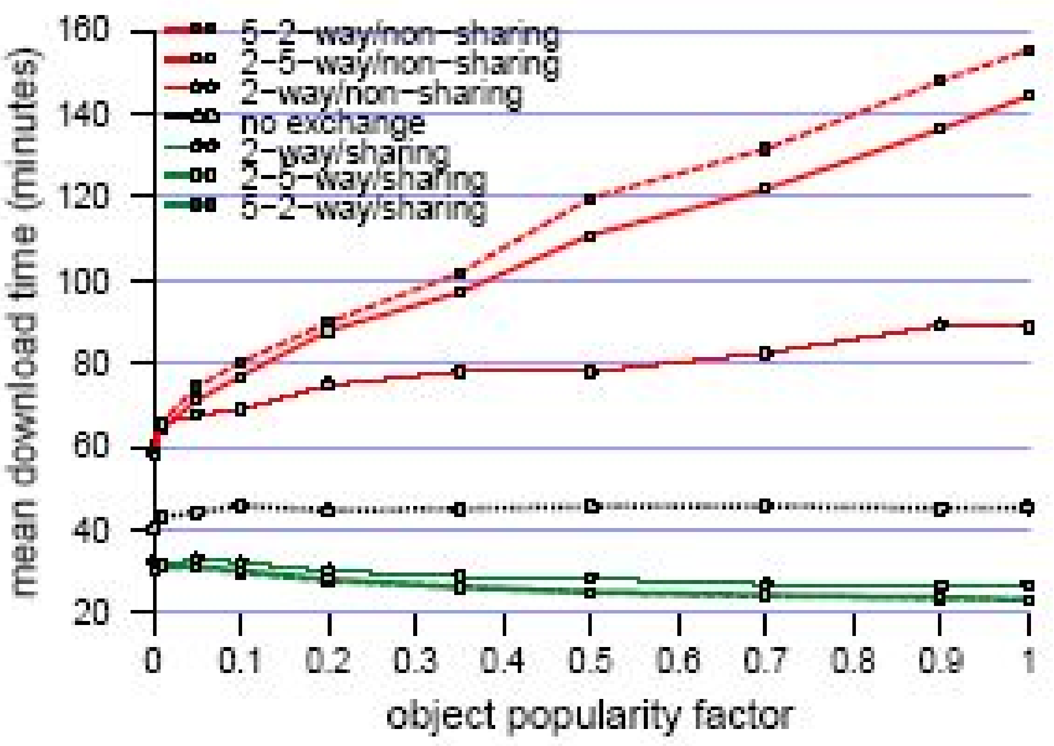

- difference increases as f approaches 1 (zipf)
- 2-5 way slightly better than 5-2 way because performance for non-sharers is reduced (longer lived on average)

Simulation Results (IV)
mean download time vs. (non-)sharing peers

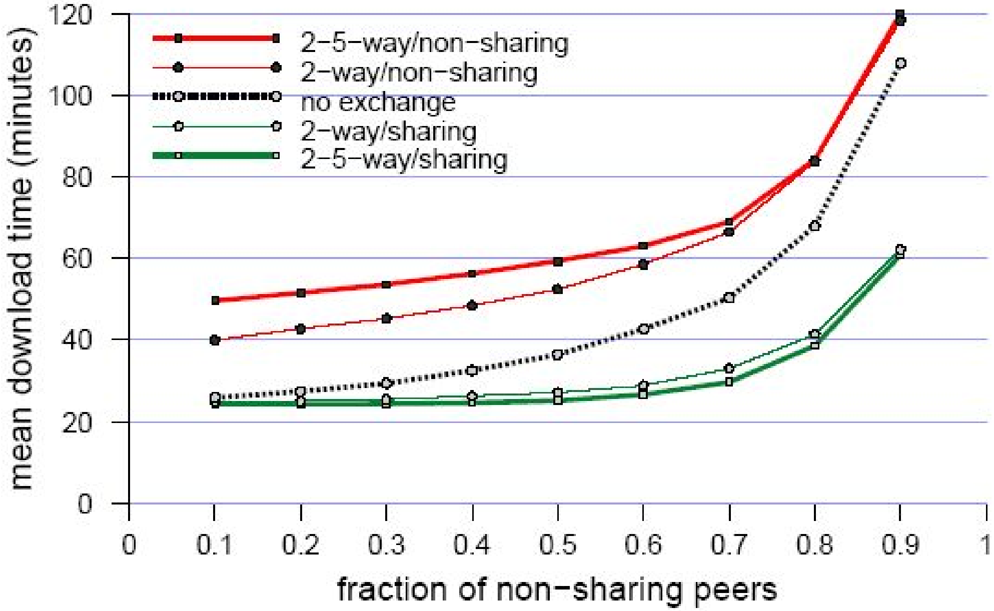

Figure 9. Mean download times vs. fraction of non-sharing peers.

- until now: 50% freeriders
- do the incentives to share always persist?
- yes!
- non-sharers get a large penalty when almost everyone is sharing
- non-sharers tend towards “no-exchange” when no one is sharing
- infrequent sharers get big reward

Simulation Results (V)
waiting time

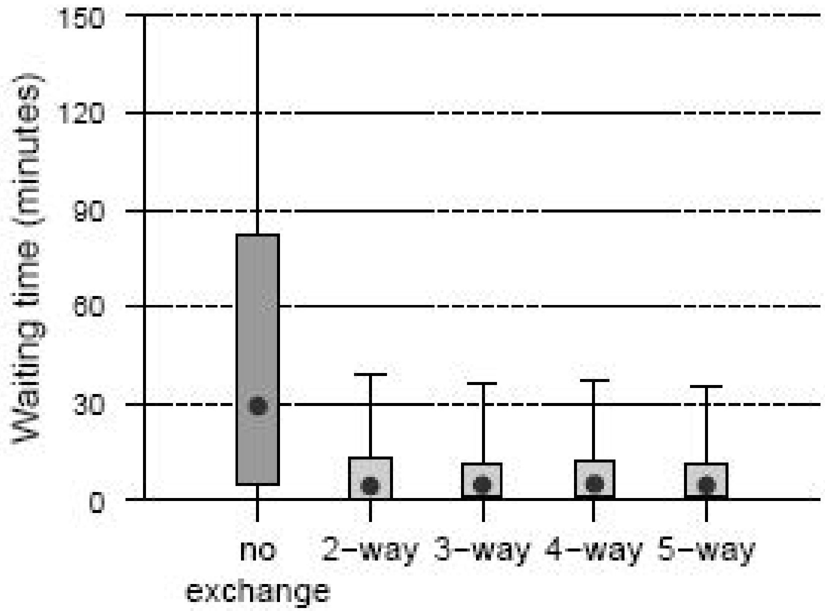

- absolute priority for exchanges = key reason for performance

Real-Life Measurements

The eMule network

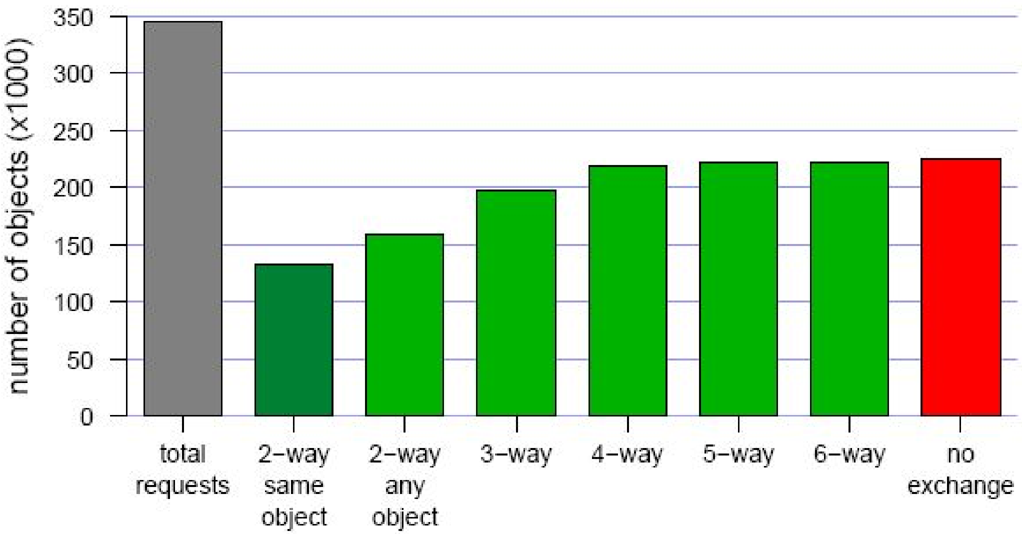

Figure 10. Fraction of requests that can be served in an exchange ring with other nodes in the dataset

- 75% of peers share more than 7 (complete) files
- many users refuse uploads even though data is available
- most peers however had outgoing requests, i.e. were participating in exchanges

Discussion

- simplistic simulation scenario
- limitations & improvements

- real exchanges do serve chunks of incomplete objects → probability for exchanges increases
- heterogeneity of real-world systems
- complexity issues with RT communications
- effect on peer behavior (replication of popular objects = $$$ in exchange economy)

Related Work

- MojoNation (centralized payment-based)
- karma (distributed cash-based system)
  - *bank-set* located via DHT lookup
  - auction mechanism, limitation of new identities
  - simulates full-fledged economic system
- better performance? (no “double coincidence of wants”)
- limitations
  - high cost in terms of user attention
  - cash ⇔ CPU cycles
- lightweight 2-way credit system: eMule
- closely related to this proposal: BitTorrent

Summary & Conclusion

- exchange-based approach provides incentives
- decentralized
- simpler than credit or cash
- higher service priority to peers providing simultaneous and symmetric service in return
- N-way exchanges
- methods for regulating transfers
- protection against malicious users
- simulations show significant performance advantage to cooperating users, especially in a loaded system
- higher-order exchanges offer improvement, if used together with 2-way exchanges

Thank you for your attention!

Any questions?