

Exchange-based Mechanisms and Cooperative Distributed System Design

An experimental study of P2P file sharing

Kostas Anagnostakis

University of Pennsylvania

(joint work with Michael Greenwald)


Approaches to incentive design

- **Payment**-based systems:
 - Actions coupled with economic transactions (payments)
 - Incentive problem reduced to computing and communicating prices, making price-based decisions
 - Highly flexible but potentially very complex
 - Vulnerable to pseudospoofing, sandbagging, market pathologies
 - Requires centralized trust, centralized service
- **Reputation**-based systems:
 - Maintain shared history of user behavior
 - Make reputation-based decisions
 - Vulnerable to pseudospoofing, sandbagging
 - Perform poorly unless there are trusted nodes

Approaches to incentive design

- **Exchange**-based mechanisms

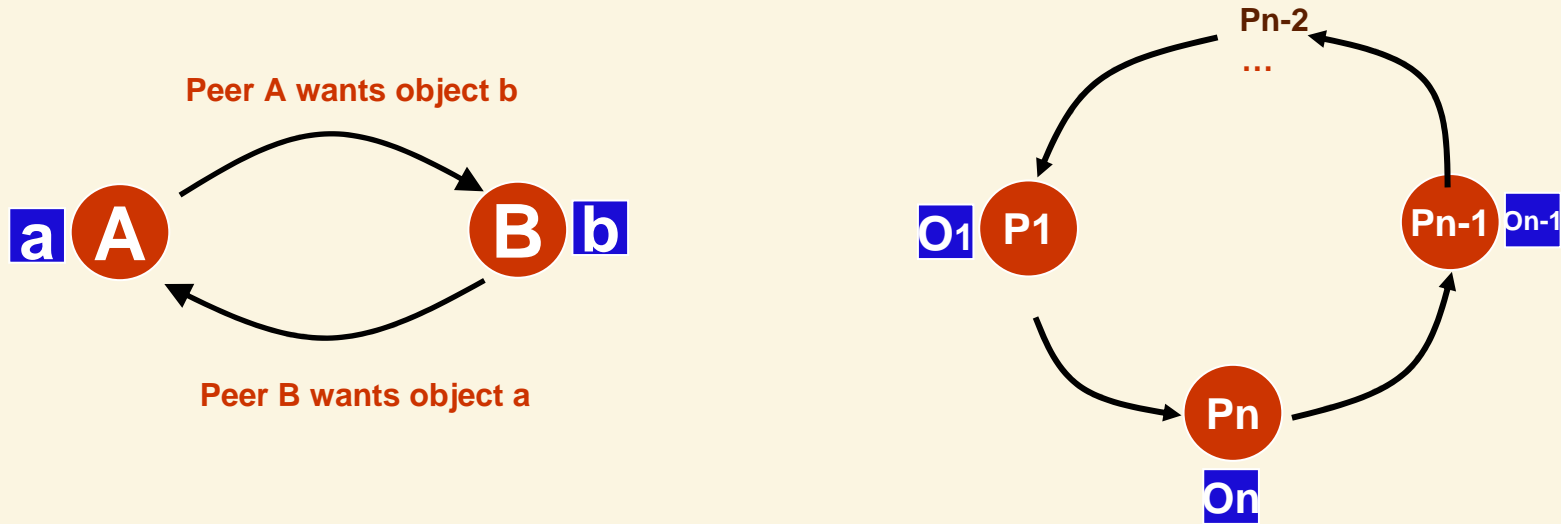
- Actions bound to verifiable reciprocal actions within groups of participants
- Example: simple file exchange



- Approach *seems* very promising:

- Simple, decentralized, memory-less, directly based on evidence of cooperation rather than indirectly through identities
- But "double coincidence of simultaneous wants" limits applicability

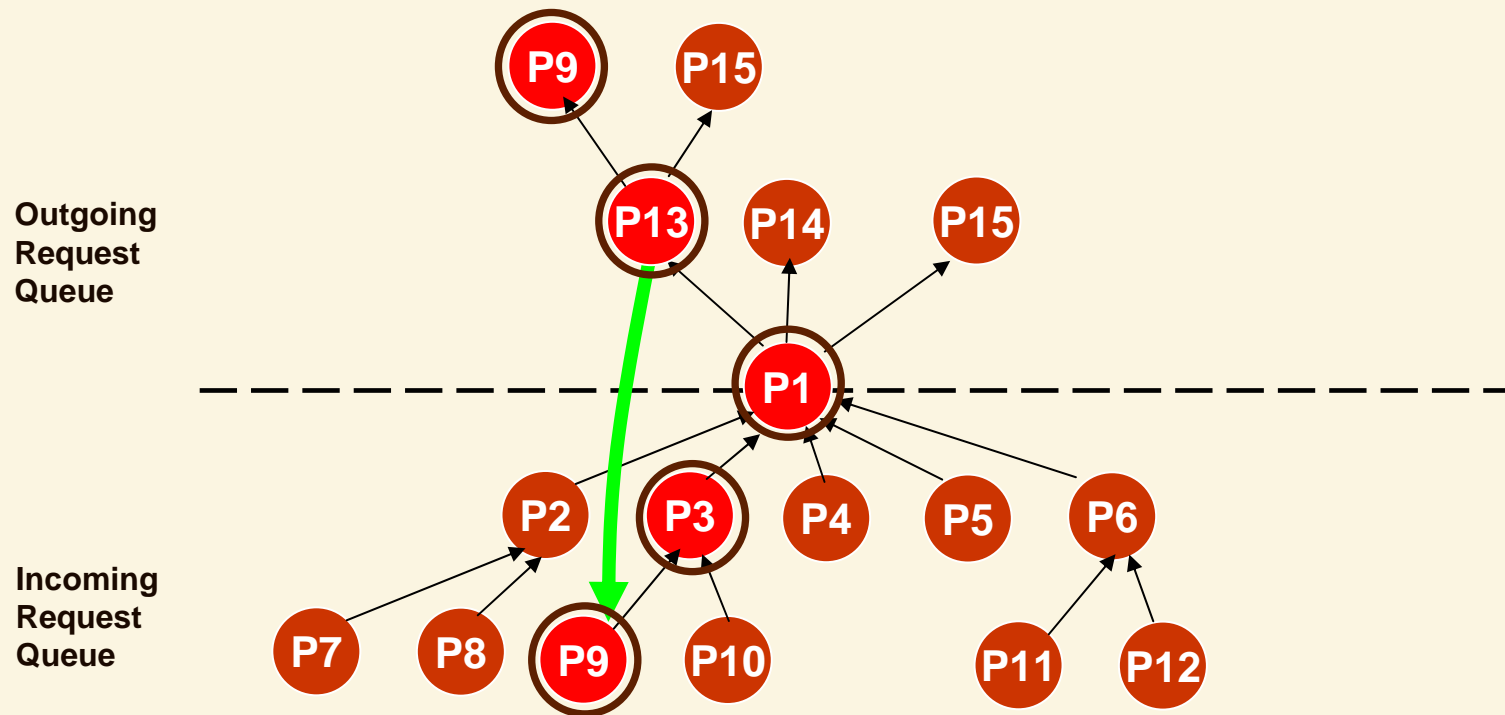
File sharing: an exchange mechanism



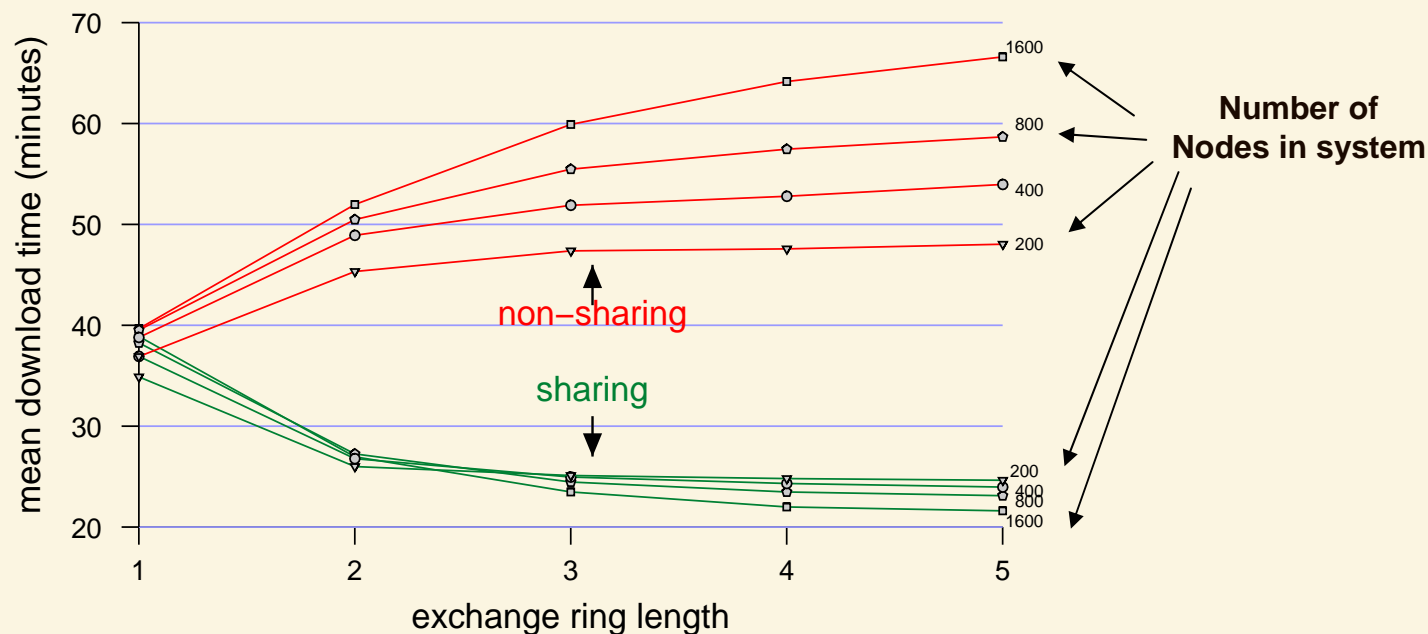
- **Key idea:** give priority to peers that can provide a reciprocal download in return
 - As long as there is at least one exchange ring for each peer, freeloaders will have to wait in low-priority queues
- Two-way, or N-way exchange rings

File sharing: discovering exchange rings

- Maintain Incoming and Outgoing Request Queues: IRQ/ORQ
- Incrementally extend IRQ and ORQ until some node is on both
- Limit each branch to K nodes - complexity $O(K^{[N/2]})$




Simulation: Importance of N-way exchanges



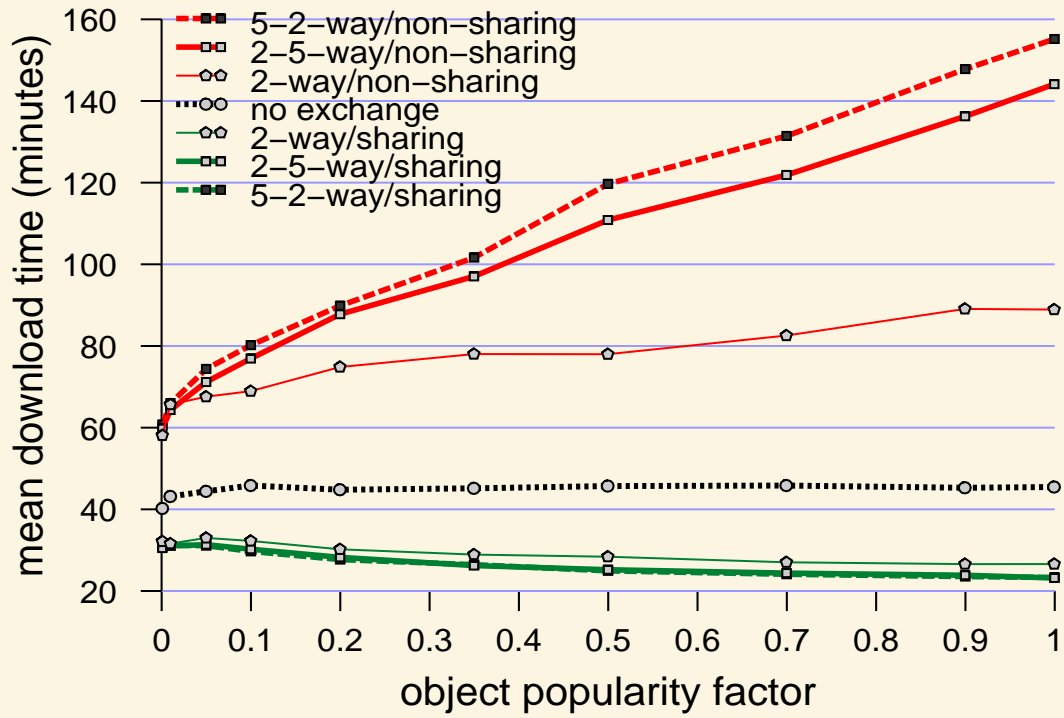
- Strong incentive to share (2x-3x faster downloads)
- Substantial improvement for $2 < N < 5$
- Diminishing returns for $N > 4$
- Incentive improves with system scale

Simulation: Incentive to deploy exchange



- Under current conditions (70% non-sharing), and even worse, sharing peers have strong incentive to deploy exchange
- Freeloaders are generally worse-off compared to sharers regardless of sharing/non-sharing population mix

Simulation: Performance vs. object popularity



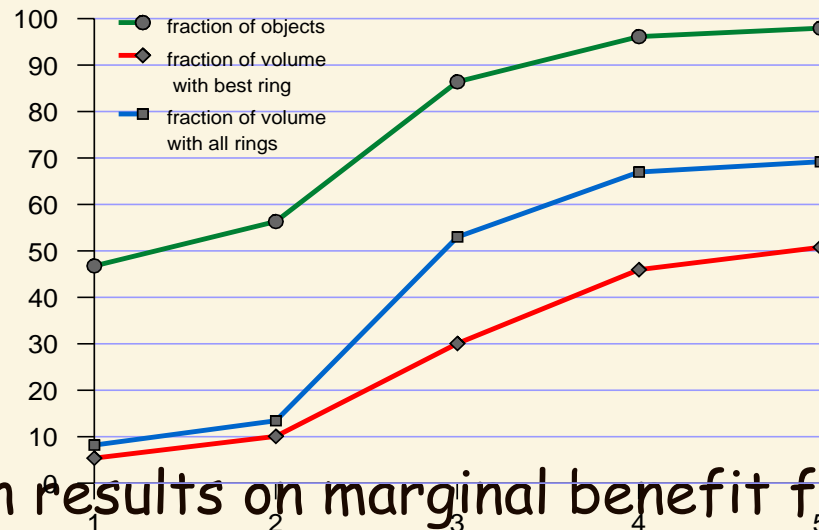
(uniform)

(ZIPF)

- Gap between sharing and non-sharing increases as popularity distribution approaches zipf
- Difference between shorter and longer rings increases
- Small increase in difference between 2-5- and 5-2-way

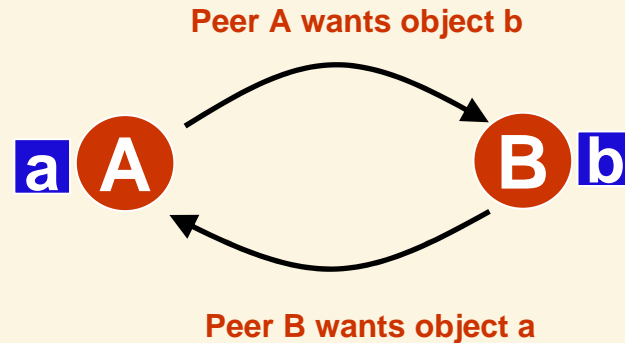
Validation with real-world measurement

- Collected data on objects shared and requested from 30,000 nodes of eMule P2P file sharing system
- Executed ring discovery algorithm for objects requested by nodes in dataset



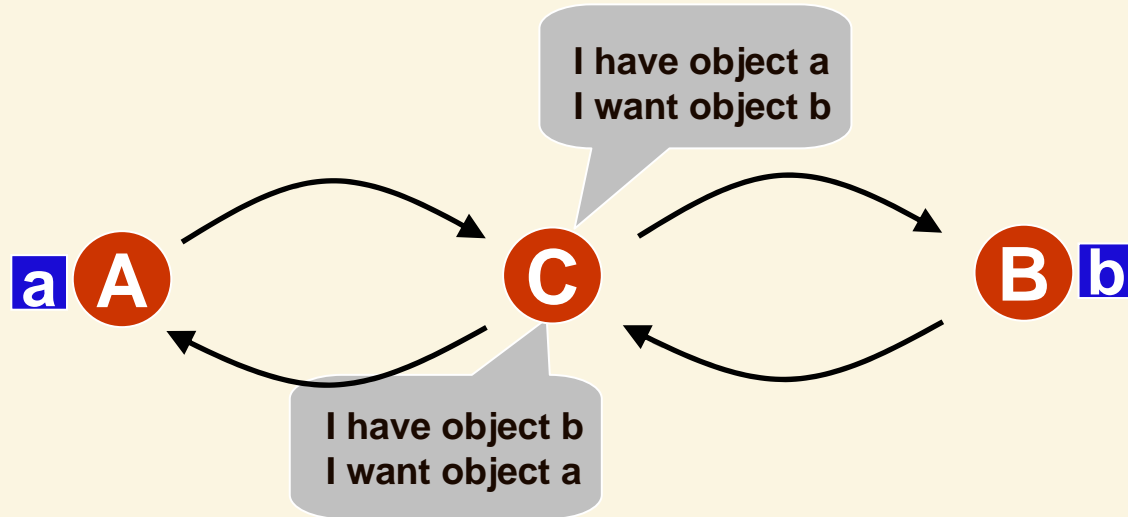
- Confirms simulation results on marginal benefit for $N > 5$
- Shows that two-way-same-object exchange (as in BitTorrent) covers only a small fraction (10%) of transfers

Threats



- **One-sided defection:** cheater leaves ring after acquiring object, without honoring exchange
 - Block-by-block transfers (as in BitTorrent) not good enough: cheater may be happy even with single block exchanges. Pseudo-spoofing makes things worse
 - *Enforce repeated interaction: prioritize long exchanges, encrypt transfers, and swap keys only when done*
- **False content:** cheater joins ring, sends bogus content
 - *Requires content validation service and a mediator to validate transaction and assist in final exchange of keys*

Threats (cont'd)

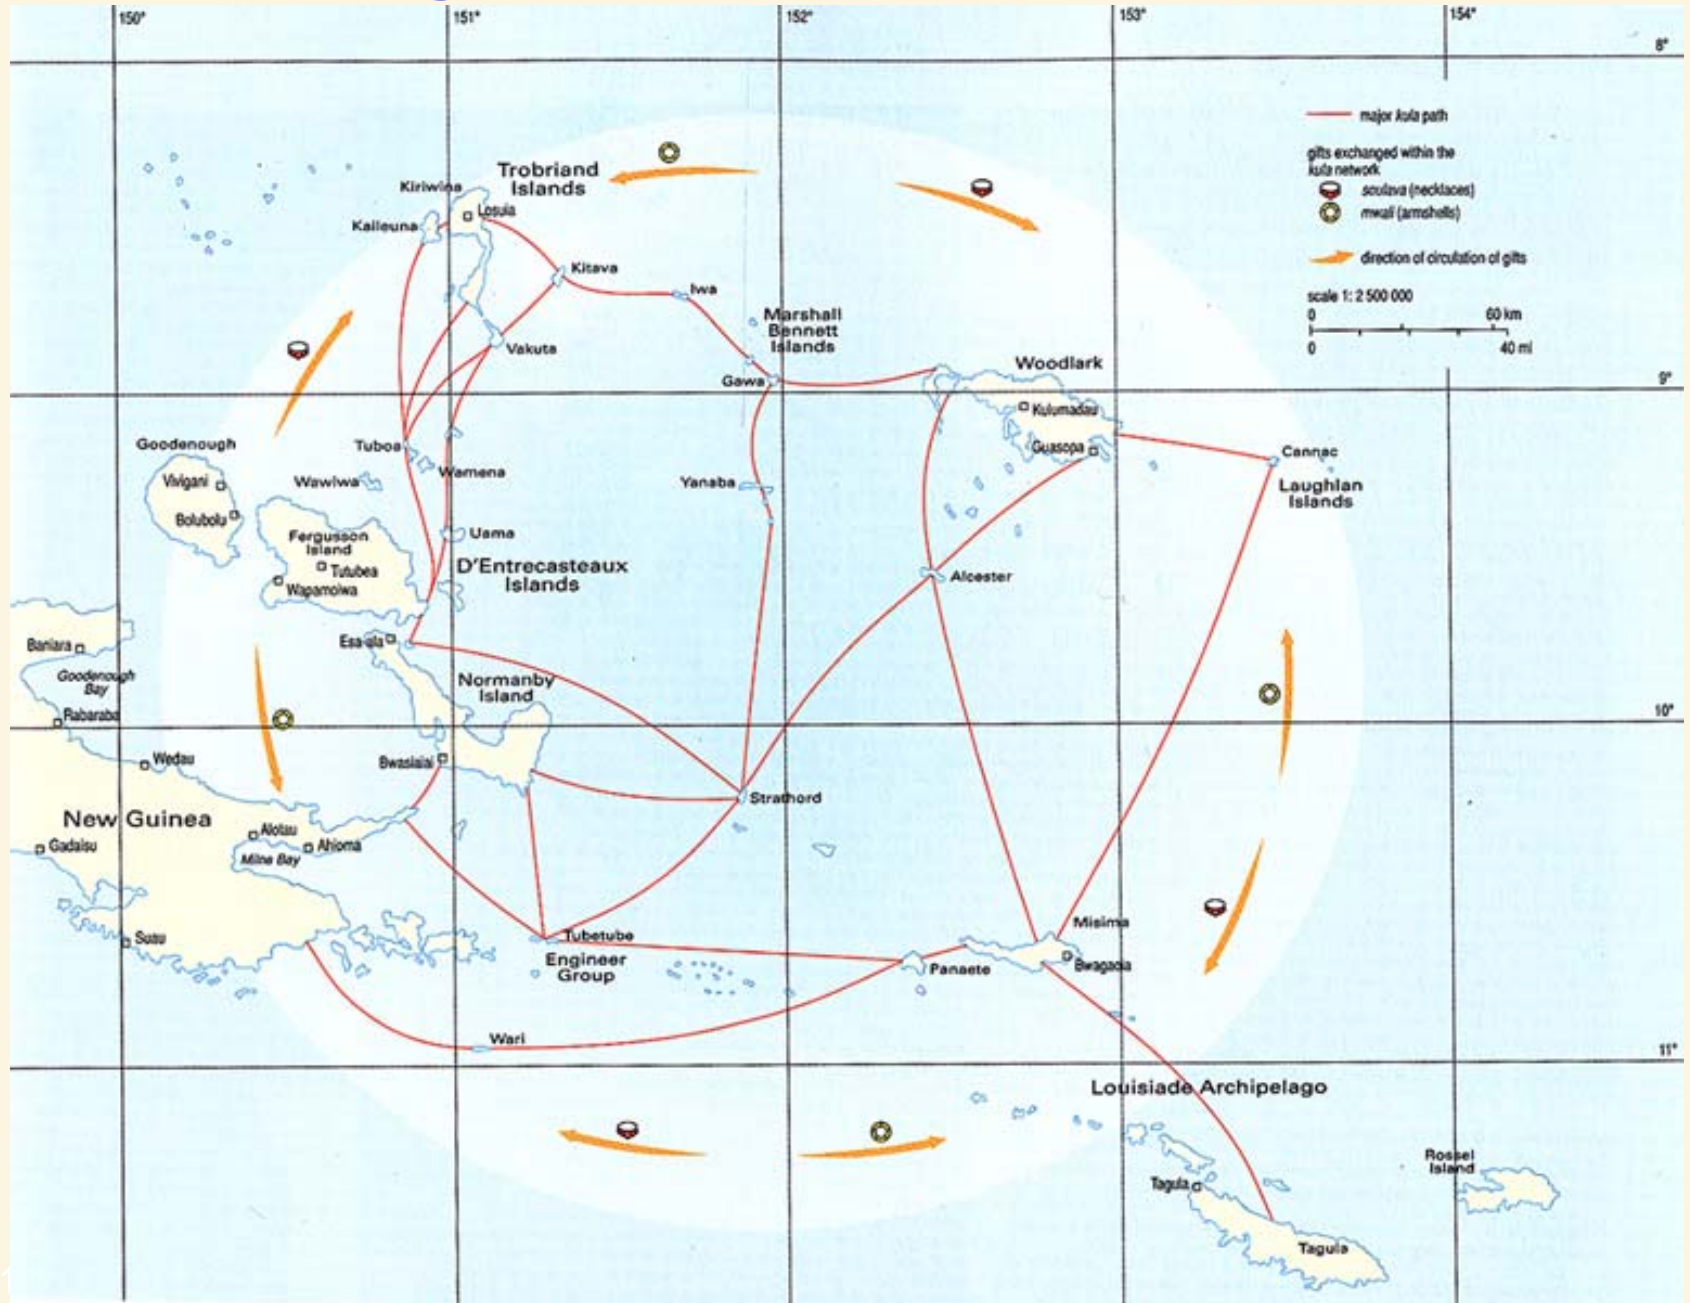


- **Middleman #1:** cheater joins ring, peddles blocks between A,B
 - *Embed **peer-id** in encrypted blocks, use third party to complete exchange of keys (third party sends key to **peer-id**)*
- **Middleman #2:** cheater obtains seed blocks using non-exchange transfers, uses seed blocks for one-block exchanges
 - *Prioritizing longer exchanges helps, but no complete solution*


Ongoing and future work

- Deploy exchange and determine effect on user behavior
- Characterization of freeriders in real P2P systems
- Evidence of cheating and freeloading on BitTorrent
- Extend approach to bw sharing, telephony (skype), etc.
- Look at dual problem of exchange-based collusion

The Kula ring

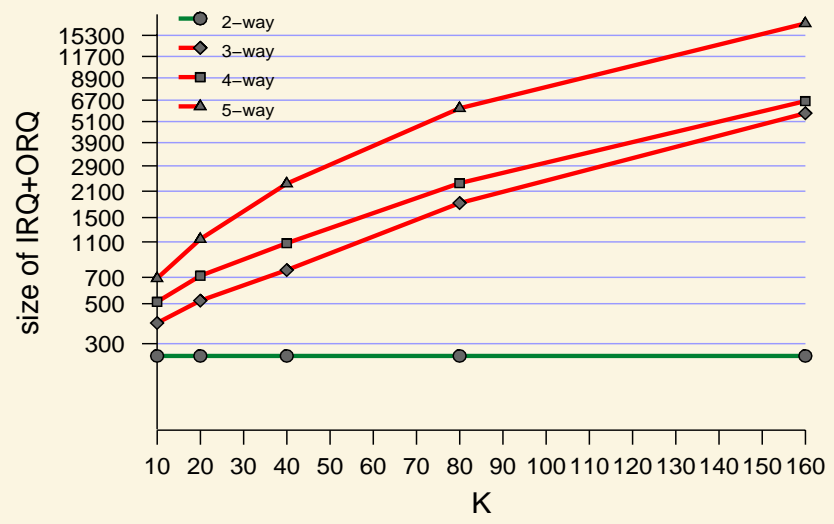



Performance vs. system load



- Sharing peers get 2x-4x faster downloads compared to freeloaders
- Advantage increases with network load
- Notable difference between 2-way and 5-way
- Small difference between 2-5-way and 5-2-way

Search cost



- Can get >95% of all rings with $K < 60$
- Size of IRQ+ORQ reasonable
 - E.g., $3900 \times 16\text{byte ID} = 62\text{KB} \ll \text{object size}$