Bilateral and Multilateral Exchanges for Peer-Assisted Content Distribution

Christina Aperjis
Social Computing Group
HP Labs

Joint work with Ramesh Johari (Stanford) and Michael J. Freedman (Princeton)

Peer-assisted content distribution

- Users upload files to each other
- Work well only if users share files and upload capacity
- P2P systems try to incentivize users to share

Bilateral and multilateral exchange

Most prevalent P2P exchange systems are bilateral: downloading is possible in return for uploading to the same user

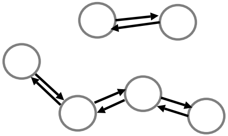

Bilateral and multilateral exchange

Most prevalent P2P exchange systems are bilateral: downloading is possible in return for uploading to the same user

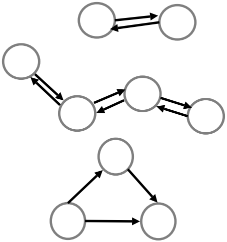

Drawback of Bilateral Exchange:
only works between users that have reciprocally desired files

Bilateral and multilateral exchange

Most prevalent P2P exchange systems are bilateral: downloading is possible in return for uploading to the same user

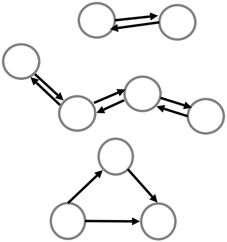

Drawback of Bilateral Exchange:
only works between users that have reciprocally desired files

Multilateral exchange allows users to trade in more general ways
but is more complex to implement
(e.g., virtual currency)

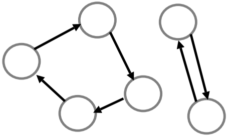

Bilateral and multilateral exchange

Most prevalent P2P exchange systems are bilateral: downloading is possible in return for uploading to the same user

Drawback of Bilateral Exchange: only works between users that have reciprocally desired files

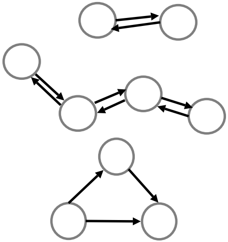

Multilateral exchange allows users to trade in more general ways
but is more complex to implement
(e.g., virtual currency)

Tradeoff: simplicity vs. participation

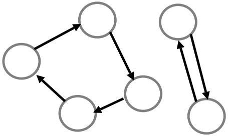

Bilateral vs. multilateral

1. Comparison of equilibria
   What are the efficiency properties
   of the allocations that arise at equilibria?
2. Quantitative comparison
   What proportion of users cannot participate?

Preliminaries

1. View content exchange as an economy:
   Demand = download requests for content
   Supply = scarce system resources
2. What files do peers have?
   We focus on exchange on a timescale
   over which the set of files peers have remains constant.
3. Rates vs. bytes
   We focus on download/upload rates,
   rather than total number of bytes transferred.
4. The network
   In the model we study, the constraint is on upload capacity.
   More generally, a network structure may constrain uploads and downloads.

Notation

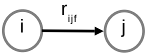

r_{ijf} = upload rate of file f from i to j
$d_{if} = \sum_j r_{jif}$ = download rate of f for peer i
$u_i = \sum_{j,f} r_{ijf}$ = upload rate of peer i
$v_i(\mathbf{d}_i, u_i)$ = utility to peer i from $(\mathbf{d}_i, u_i)$

B_i = bandwidth constraint of user i
$\mathcal{X}$ = set of feasible rate vectors
$\mathcal{X} = \{\mathbf{r} : \mathbf{r} \geq 0; u_i \leq B_i; r_{ijf} = 0$ if i does not have file $f\}$

Bilateral content exchange

- Peers exchange content on a pairwise basis
- Let $R_{ij} = \sum_f r_{ijf}$ = rate of upload from i to j
- Exchange ratio: $\gamma_{ij} = R_{ji}/R_{ij}$
- As if there exist prices p_{ij}, p_{ji},
  and all exchange is settlement-free: $p_{ij}R_{ij} = p_{ji}R_{ji}$
  Thus: $\gamma_{ij} = p_{ij}/p_{ji}$

Bilateral content exchange

- Peers exchange content on a pairwise basis
- Let $R_{ij} = \sum_f r_{ijf}$ = rate of upload from i to j
- Exchange ratio: $\gamma_{ij} = R_{ji}/R_{ij}$
- As if there exist prices p_{ij}, p_{ji},
  and all exchange is settlement-free: $p_{ij}R_{ij} = p_{ji}R_{ji}$
  Thus: $\gamma_{ij} = p_{ij}/p_{ji}$
- Peer i may be effectively price-discriminating (if $p_{ij} \neq p_{ik}$)

Bilateral content exchange

- Peers exchange content on a pairwise basis
- Let $R_{ij} = \sum_f r_{ijf}$ = rate of upload from i to j
- Exchange ratio: $\gamma_{ij} = R_{ji}/R_{ij}$
- As if there exist prices p_{ij}, p_{ji},
  and all exchange is settlement-free: $p_{ij}R_{ij} = p_{ji}R_{ji}$
  Thus: $\gamma_{ij} = p_{ij}/p_{ji}$
- Peer i may be effectively price-discriminating (if $p_{ij} \neq p_{ik}$)

Example: BitTorrent

- Peer j splits upload rate B_j equally among k_j peers with highest rates to j (the "active set")
- For a peer i in the active set: $\gamma_{ij} = \frac{B_j}{k_j R_{ij}}$

Multilateral content exchange

- Users can trade a virtual currency,
  where downloading from peer j costs p_j per unit rate
- Similar to an *exchange economy*

Equilibria

In multilateral exchange,
users optimize given *prices*

Multilateral optimization

$$\begin{array}{ll} \max & v_i(\mathbf{d}_i, u_i) \\ \text{s.t.:} & \sum_j p_j R_{ji} \le p_i \sum_j R_{ij} \\ & \mathbf{r} \in \mathcal{X} \end{array}$$

Equilibria

In multilateral exchange, users optimize given *prices*

Multilateral optimization

$$\begin{array}{ll} \max & v_i(\mathbf{d}_i, u_i) \\ \text{s.t.:} & \sum_j p_j R_{ji} \leq p_i \sum_j R_{ij} \\ & \mathbf{r} \in \mathcal{X} \end{array}$$

In bilateral exchange, users optimize given *exchange ratios*

Bilateral optimization

$$\begin{array}{ll} \max & v_i(\mathbf{d}_i, u_i) \\ \text{s.t.:} & R_{ji} \leq \gamma_{ij} R_{ij} \forall j \\ & \mathbf{r} \in \mathcal{X} \end{array}$$

Equilibria

In multilateral exchange, users optimize given *prices*

Multilateral optimization

$$\begin{array}{ll} \max & v_i(\mathbf{d}_i, u_i) \\ \text{s.t.:} & \sum_j p_j R_{ji} \leq p_i \sum_j R_{ij} \\ & \mathbf{r} \in \mathcal{X} \end{array}$$

In bilateral exchange, users optimize given *exchange ratios*

Bilateral optimization

$$\begin{array}{ll} \max & v_i(\mathbf{d}_i, u_i) \\ \text{s.t.:} & R_{ji} \leq \gamma_{ij} R_{ij} \forall j \\ & \mathbf{r} \in \mathcal{X} \end{array}$$

At an equilibrium all users have optimized, and the market clears

Multilateral equilibrium (ME)
$\mathbf{r}^*$ and prices $\mathbf{p}^*$

Bilateral equilibrium (BE)
$\mathbf{r}^*$ and exchange ratios $\boldsymbol{\gamma}^*$

Under mild conditions, both ME and BE exist

Pareto efficiency

An allocation **r** is Pareto efficient if:
no user's utility can be strictly improved
without strictly reducing another user's utility

- ME are always Pareto efficient
  (First fundamental theorem of welfare economics)
- BE may not be Pareto efficient

Pareto efficiency

When are BE efficient?

Theorem

Assume utility approaches $-\infty$ as upload rate approaches capacity
A BE $(\gamma^*, \mathbf{r}^*)$ is Pareto efficient if and only if there exists a supporting vector of prices $\mathbf{p}^*$ such that $(\mathbf{p}^*, \mathbf{r}^*)$ is a ME

[Hard part to prove is the "only if"]

Pareto efficiency: proof sketch

Given Pareto efficient BE $(\gamma^*, \mathbf{r}^*)$
find price vector $\mathbf{p}^*$ such that $(\mathbf{p}^*, \mathbf{r}^*)$ is a ME

- The proof exploits a connection between equilibria and reversible Markov chains
- Let R^*_{ij} = total rate from i to j at BE and $R^*_{ii} = -\sum_j R^*_{ij}$
- For simplicity, suppose $\mathbf{R}^*$ is an irreducible rate matrix of a continuous time MC (generalizes to nonirreducible case)

Pareto efficiency: proof sketch

Let **p** be the unique invariant distribution of $\mathbf{R}^*$
If $\mathbf{R}^*$ is reversible, then:

$$p_i R^*_{ij} = p_j R^*_{ji} \Rightarrow \gamma^*_{ij} = p_i / p_j \Rightarrow BE \equiv ME$$

What is the intuition for this result?

- The invariant distribution gives a vector of prices at which agents could potentially trade
- When $\mathbf{R}^*$ is reversible, agents' trades balance on a pairwise basis with one vector of prices

Pareto efficiency: proof sketch

What if $\mathbf{R}^*$ is not reversible? $\frac{p_i}{p_j} > \gamma^*_{ij}$ for some $R^*_{ij} > 0$

$\Rightarrow$ i "overpaid" to transact with j at BE

Pareto efficiency: proof sketch

What if $\mathbf{R}^*$ is not reversible? $\frac{p_i}{p_j} > \gamma^*_{ij}$ for some $R^*_{ij} > 0$

$\Rightarrow i$ "overpaid" to transact with j at BE

$\sum_k \gamma_{ki} R^*_{ki} = \sum_k \frac{p_k}{p_i} R^*_{ki} \Rightarrow i$ "underpaid" some j'

Pareto efficiency: proof sketch

What if $\mathbf{R}^*$ is not reversible? $\frac{p_i}{p_j} > \gamma^*_{ij}$ for some $R^*_{ij} > 0$

$\Rightarrow i$ "overpaid" to transact with j at BE

$\sum_k \gamma_{ki} R^*_{ki} = \sum_k \frac{p_k}{p_i} R^*_{ki} \Rightarrow i$ "underpaid" some j'

Can find cycle of users $\{1, ..., K\}$ such that k "overpaid" $k+1\ \forall k$

Pareto efficiency: proof sketch

What if $\mathbf{R}^*$ is not reversible? $\frac{p_i}{p_j} > \gamma^*_{ij}$ for some $R^*_{ij} > 0$

$\Rightarrow$ i "overpaid" to transact with j at BE

$\sum_k \gamma_{ki} R^*_{ki} = \sum_k \frac{p_k}{p_i} R^*_{ki} \Rightarrow i$ "underpaid" some j'

Can find cycle of users $\{1, ..., K\}$ such that k "overpaid" $k+1\ \forall k$

Pareto improvement:
Increase u^*_i and $R^*_{i,i-1}$ by a_i
User i better off if $\frac{a_{i+1}}{a_i} > \gamma^*_{i,i+1}$
Possible to find such a_i's,
because $\prod_i \gamma^*_{i,i+1} < 1$

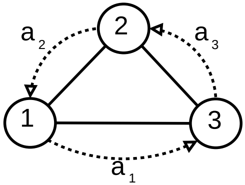

Pareto efficiency: proof sketch

What if $\mathbf{R}^*$ is not reversible? $\frac{p_i}{p_j} > \gamma^*_{ij}$ for some $R^*_{ij} > 0$

$\Rightarrow$ i "overpaid" to transact with j at BE

$\sum_k \gamma_{ki} R^*_{ki} = \sum_k \frac{p_k}{p_i} R^*_{ki} \Rightarrow i$ "underpaid" some j'

Can find cycle of users $\{1, ..., K\}$ such that k "overpaid" $k+1$ $\forall k$

Pareto improvement:
Increase u^*_i and $R^*_{i,i-1}$ by a_i
User i better off if $\frac{a_{i+1}}{a_i} > \gamma^*_{i,i+1}$
Possible to find such a_i's,
because $\prod_i \gamma^*_{i,i+1} < 1$

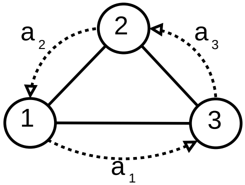

r* Pareto efficient BE $\Rightarrow$ R* reversible

Pareto efficiency: proof sketch

- Of course, in general $\mathbf{R}^*$ may not be irreducible
  Instead, the graph of trades in the BE may have multiple connected components
- To complete the proof, we consider supporting price vectors $\mathbf{p}$ that arise as linear combinations of the unique invariant distributions on each component
- We show that if no supporting prices for the BE exist, then a cycle of agents
  (possibly spanning multiple connected components)
  can be found who have a Pareto improving trade

Bilateral vs. multilateral

1. Pareto efficiency of equilibria
2. Participation
   How many peers are able to trade bilaterally and multilaterally?
   We use a random model to quantify the density of trade produced by the two models

Participation

Two peers are complementary if each has what the other wants
A peer can trade bilaterally if she has a complementary peer

A peer can trade multilaterally if it belongs on a cycle of peers along which peers want to trade

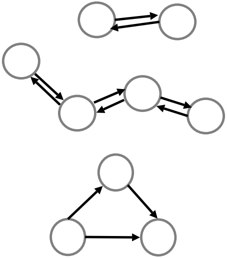

Participation: asymptotic analysis

- N users, K files
- Each user has one file to upload,
  and wants to download one file
- The probability a user wants or has the f-th most popular file is proportional to f^{-s} (Zipf's law)
  - $s = 0$: uniform popularity
  - $s > 1$: popularity concentrated in relatively few files
- Metric: expected proportion of users that cannot participate

Participation: asymptotic analysis for $s < 1$

Let ρ_{ME} (resp., ρ_{BE}) be the expected number of unmatched peers in multilateral (resp., bilateral) exchange

Theorem

When $s \in [0, 1)$:

- If $N > K^2$, then $\rho_{BE} \to 0$
- If $N < K^2$, then $\rho_{BE} \geq (1 - s)^2$
- If $K \log K < N$, then $\rho_{ME} \to 0$

If N scales faster than $K \log K$ but slower than K^2, multilateral is significantly better than bilateral

Participation: asymptotic analysis for $s > 1$

Theorem

If $s > 1$, then $\rho_{BE} \to 0$ for any scaling of K and N

- So in this case, bilateral performs very well
- Intuition:
  high concentration of popularity in a small number of files
- This result also holds:
  - when peers upload and download multiple files
  - for more general random graph models

BitTorrent popularity data

Dataset from [Piatek et al., 2008]
1.4M downloads, 680K peers, 7.3K files

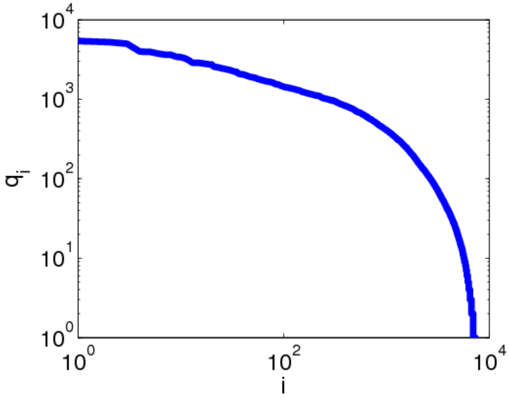

Data-driven comparison

What if we sample a random graph from the BT distribution?

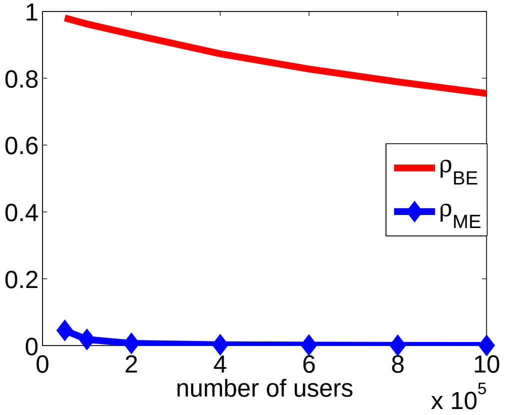

Multilateral exchange matches many more peers than bilateral

Data-driven comparison

What if users can trade in triangles?

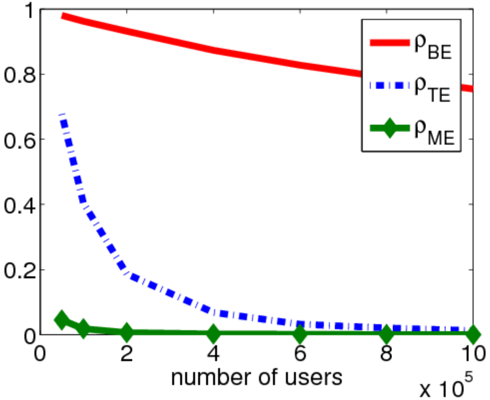

Trilateral exchange converges much faster than bilateral

Data-driven comparison

However, as the number of files a peer has increases, bilateral rapidly approaches multilateral

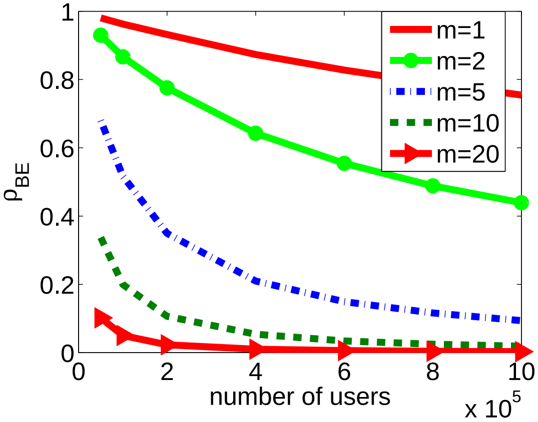

m = # of files a peer has available for uploading

Data-driven comparison

What if the number of files that users possess varies across different users?

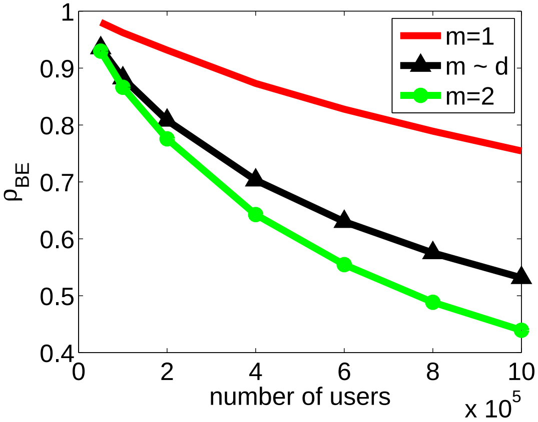

d = distribution from dataset

- mean = 2.0084
- high variance

Conclusions

- A BE is Pareto efficient if and only if it corresponds to a ME
- Bilateral exchange performs very well in expectation if
  the file popularity is very concentrated
  and/or users share a sufficiently large number of files

Conclusions

- A BE is Pareto efficient if and only if it corresponds to a ME
- Bilateral exchange performs very well in expectation if
  the file popularity is very concentrated
  and/or users share a sufficiently large number of files

Open issue: extend comparison to a dynamic setting, where

- downloads complete and preferences change over time
- users join and leave the system