Bilateral vs. multilateral content exchange

Ramesh Johari

Joint work with Christina Aperjis (Stanford)
and Michael Freedman (Princeton)

Peer-to-peer technology today

- Comprises 35-90% of “all” Internet traffic

- Not just a technology for (illicit) filesharing

Prices and content exchange

- We view content exchange as an *exchange economy:*
  - *Prices* are used to match demand with supply.
- In content exchange:
  - Demand = download requests for content
  - Supply = scarce system resources
- What does a price-based analysis tell us about matching demand with supply?

Content exchange mechanisms

- Most prevalent exchange systems are *bilateral:* downloading possible in return for uploading to the same peer.

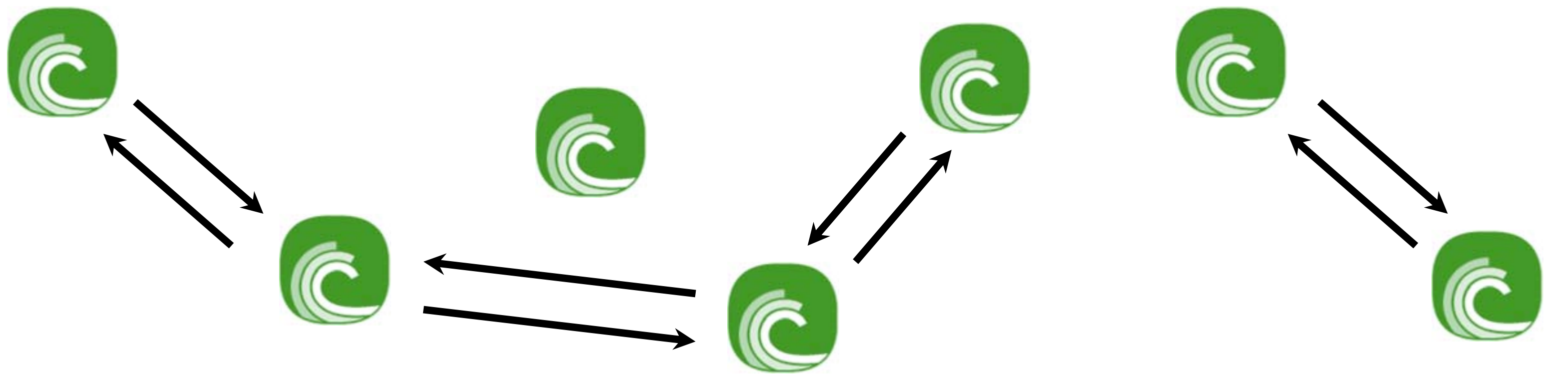

Content exchange mechanisms

- Most prevalent exchange systems are *bilateral:* downloading possible in return for uploading to the same peer.
- In this talk we explore the use of *prices* and a virtual *currency* to enable *multilateral* exchange among peers

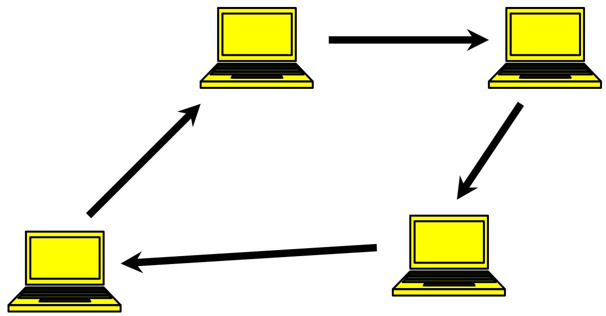

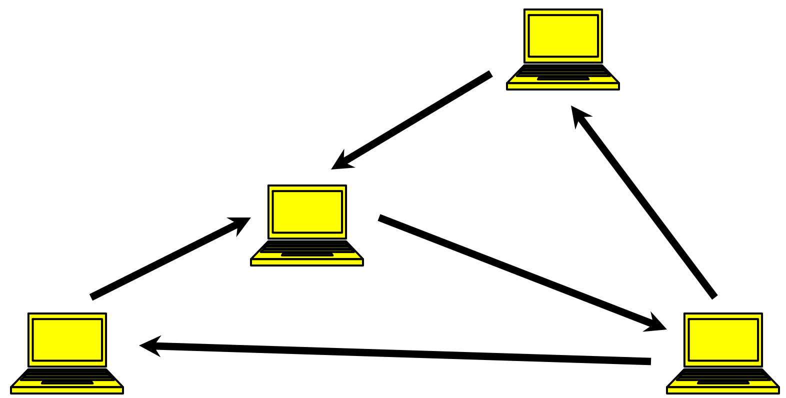

Content exchange mechanisms

- Most prevalent exchange systems are *bilateral:* downloading possible in return for uploading to the same peer.
- In this talk we explore the use of *prices* and a virtual *currency* to enable *multilateral* exchange among peers
- Basic goal:

  Rigorous comparison of *efficiency* of bilateral and multilateral content exchange

Outline

- Bilateral content exchange
- Multilateral content exchange
- Bilateral vs. multilateral: Pareto efficiency
- Bilateral vs. multilateral: Participation
- Conclusions and future work

Preliminaries

- ***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

  $u_i = \sum_{j,f} r_{jif}$ = upload rate of peer i

  B_i = bandwidth constraint of peer i

  $V_i(\underline{\mathbf{d}}_i) - c_i(u_i)$ = ***utility*** to peer i from $(\underline{\mathbf{d}}_i, u_i)$

- Feasible set of rates is:

  X = { $\underline{\mathbf{r}}$: $\underline{\mathbf{r}} \geq 0$; $u_i \leq B_i$ for all i;

  $r_{ijf} = 0$ if user 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 equilibrium

- Bilateral peer optimization *for* i *given* $\underline{\gamma}$:

  maximize $V_i(\underline{\mathbf{d}}_i)$ - $c_i(u_i)$

  subject to $\sum_f r_{jif} = \gamma_{ij} \sum_f r_{ijf}$, for all j

  $\underline{\mathbf{r}} \in X$

- *Bilateral equilibrium (BE)* is a vector $\underline{\mathbf{r}}^*$ and exchange ratios $\underline{\gamma}^*$ such that:
  All users have simultaneously optimized
- We set the following convention:

  $\gamma_{ij}^* = 0 \Leftrightarrow i$ has no file that j wants, or vice versa

Market clearing

Important point:

There is an embedded *market-clearing* operation in the definition of equilibrium.

The optimal r_{ijf} and r_{jif} chosen by **peer** i given $\underline{\gamma}^*$ must *coincide* with the optimal r_{ijf} and r_{jif} chosen by **peer** j given $\underline{\gamma}^*$

Multilateral content exchange

- Suppose instead that users can trade a *virtual currency*, where downloading from peer j costs p_j per unit rate
- Multilateral peer optimization *for* i *given* $\underline{\mathbf{p}}$:

maximize $V_i(\underline{\mathbf{d}}_i)$

subject to $\sum_{j,f} p_j \, r_{jif} = \sum_{j,f} p_i \, r_{ijf}$

$\underline{\mathbf{r}} \in X$

Multilateral equilibrium

- *Multilateral equilibrium (ME)* is a vector $\underline{\mathbf{r}}^*$ and prices $\underline{\mathbf{p}}^*$ such that:
  All users have simultaneously optimized

- Under mild conditions, both BE and ME exist
- We now provide two comparisons of efficiency: one qualitative, one quantitative

Pareto efficiency

- An allocation $\underline{\mathbf{r}}$ is *Pareto efficient* if:
  no user's utility can be strictly improved
  without strictly reducing another user's utility

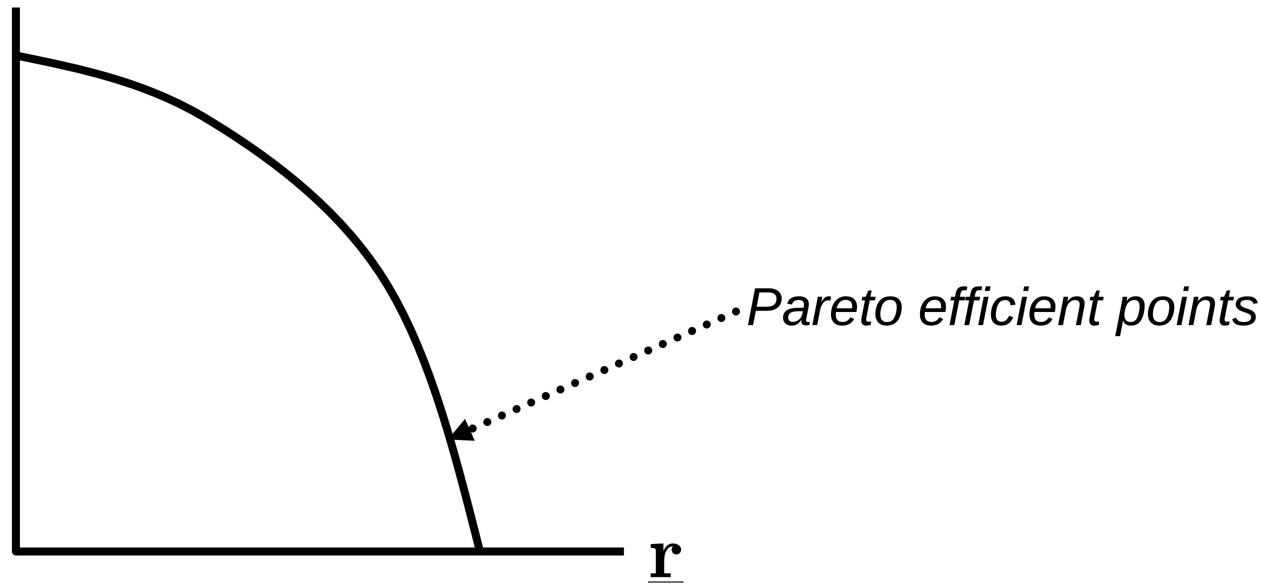

Pareto efficiency

- An allocation $\underline{\mathbf{r}}$ is *Pareto efficient* if:
  no user's utility can be strictly improved
  without strictly reducing another user's utility

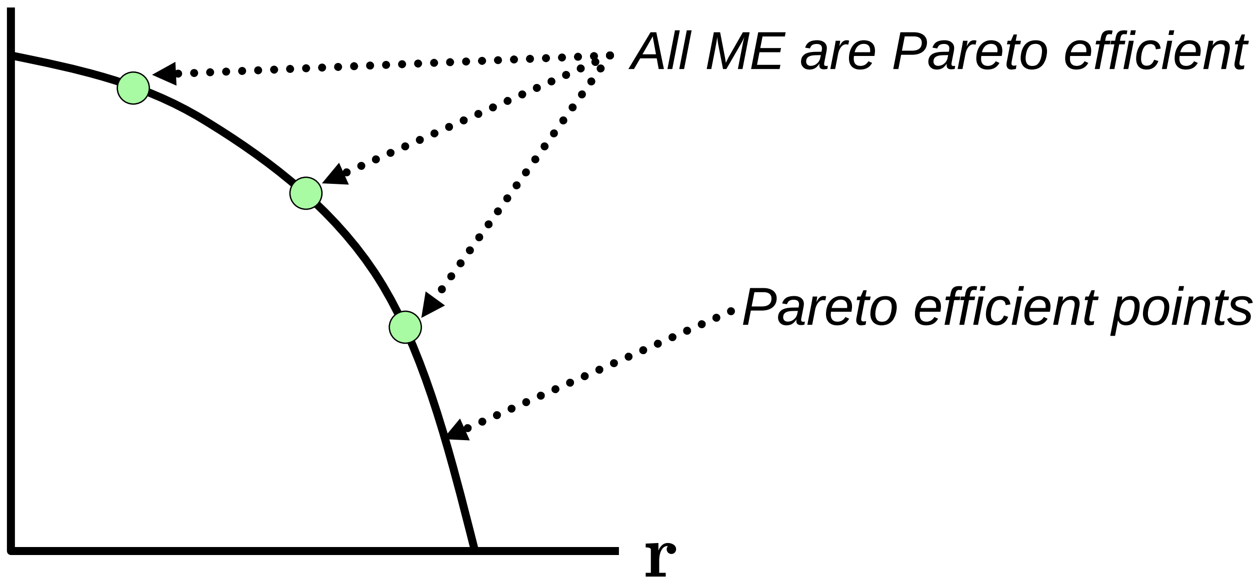

Pareto efficiency

- An allocation $\underline{\mathbf{r}}$ is *Pareto efficient* if:
  no user's utility can be strictly improved
  without strictly reducing another user's utility

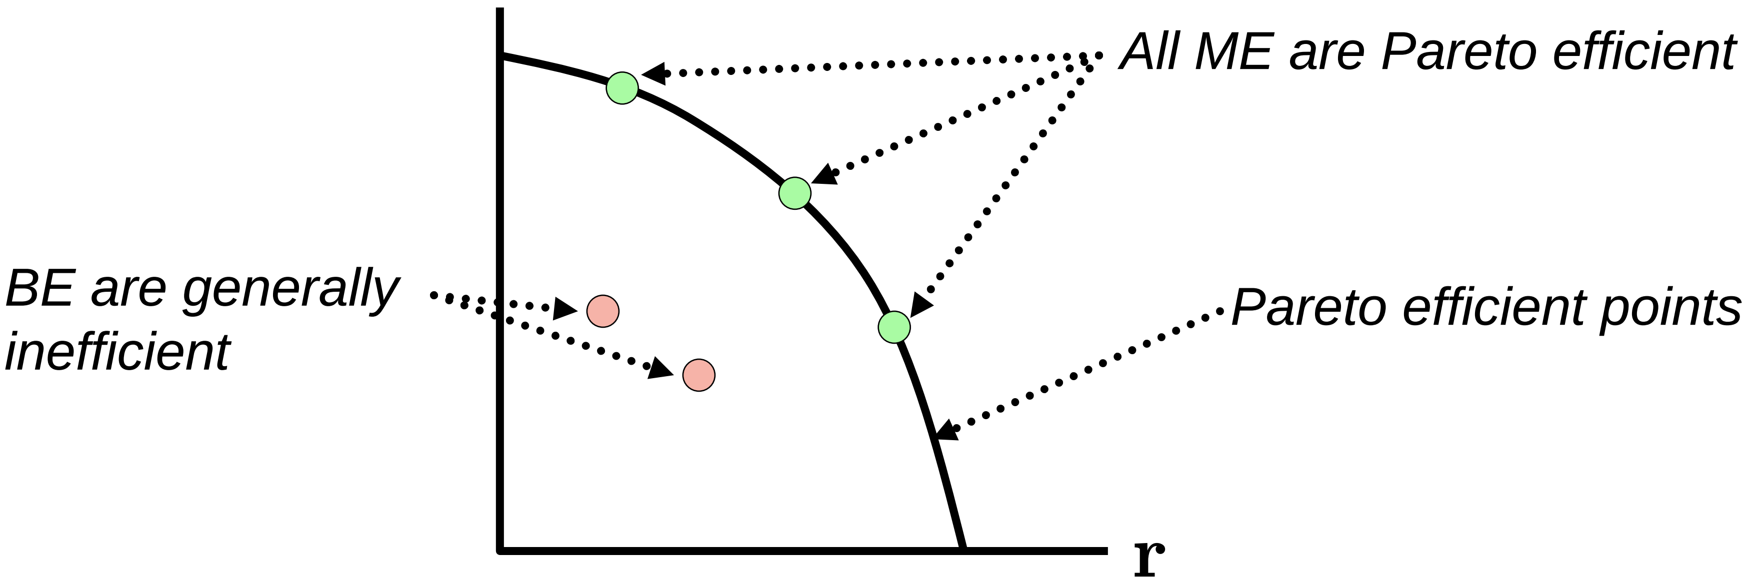

Pareto efficiency

When are BE efficient?

Theorem:

A BE $(\underline{\gamma}^*, \underline{\mathbf{r}}^*)$ *is Pareto efficient*
if and only if there exists a vector of prices $\underline{\mathbf{p}}^*$
such that $(\underline{\mathbf{p}}^*, \underline{\mathbf{r}}^*)$ *is a ME*

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

Pareto efficiency: Proof

- The proof exploits a connection between equilibria and reversible Markov chains
- Let R_{ij}* = $\sum_f r_{ijf}$*, and R_{ii}* = -$\sum_{j \neq i} R_{ij}$*
- For simplicity, suppose $\underline{\mathbf{R}}$* is an *irreducible* rate matrix of a continuous time MC (generalizes to nonirreducible case)
- Let $\underline{\mathbf{p}}$ be the unique invariant distribution of $\underline{\mathbf{R}}$*
- If $\underline{\mathbf{R}}$* is reversible, then:

  $p_i R_{ij}$* = $p_j R_{ji}$* $\Rightarrow \gamma_{ij}$* = $p_i/p_j \Rightarrow$ BE $\equiv$ ME

Pareto efficiency: Proof

What if $\underline{\mathbf{R}}$* is not reversible?

- We construct a sequence of peers **1**, ..., K**+1** with:

  $p_k/p_{k+1} > \gamma_{k,k+1}{}^*$ and **1** $\equiv K$**+1**

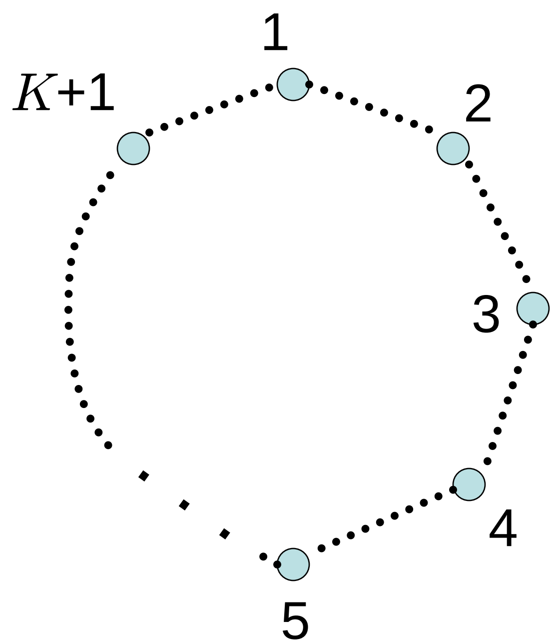

Pareto efficiency: Proof

What if $\underline{\mathbf{R}}$* is not reversible?

- We construct a sequence of peers 1, ..., K+1 with:

  $p_k/p_{k+1} > \gamma_{k,k+1}^*$ and $1 \equiv K+1$
- Consider slightly *increasing* the rates $R_{k+1,k}^*$
- Show that if $\Delta R_{k+1,k} / \Delta R_{k,k-1} > \gamma_{k,k+1}^*$, then peer k is strictly better off
- Since $\prod_k \gamma_{k,k+1}^* < 1$, such a $\Delta\,\underline{\mathbf{R}}$ can be found

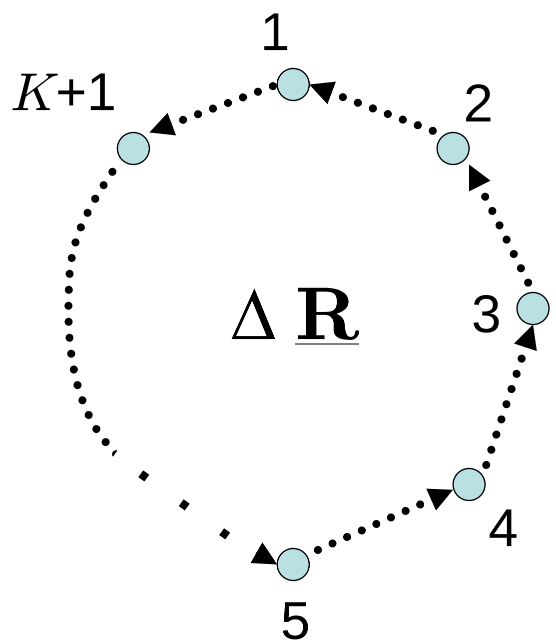

Participation

How many peers are able to trade in equilibrium in BE and ME?

We use a random model to quantify the density of trade produced by the two models.

Participation: Simplified model

Consider a model with N peers and K files.

Each peer has *one* file to upload, and desires *one* file to download.

Two peers are *complementary* if each has what the other wants.

Lemma: A peer participates in a BE if and only if she has a complementary peer.

Participation

We consider a random model where the probability a peer wants or has file f is proportional to f^{-s} (Zipf's law).

We have results on two settings:

$s \to 0$: uniform popularity

$s > 1$: very heavy tailed

Participation

When $s \to 0$:

- If $N^{1-\varepsilon} > K$, then almost all peers trade in ME with high probability
- If $K > \sqrt{N}$, then a constant fraction of peers *do not* trade in BE

- So: If $N^{1-\varepsilon} > K > \sqrt{N}$, then ME has significantly higher participation

Participation

When s > 1:

- High concentration of popularity in a small number of files
- In this case, constant fraction of peers trade in BE with high probability as K, $N \to \infty$ (and same holds for ME as well)

- So in this case, BE performs well

BT popularity data

- 1.4M downloads, 680K peers, 7.3K files

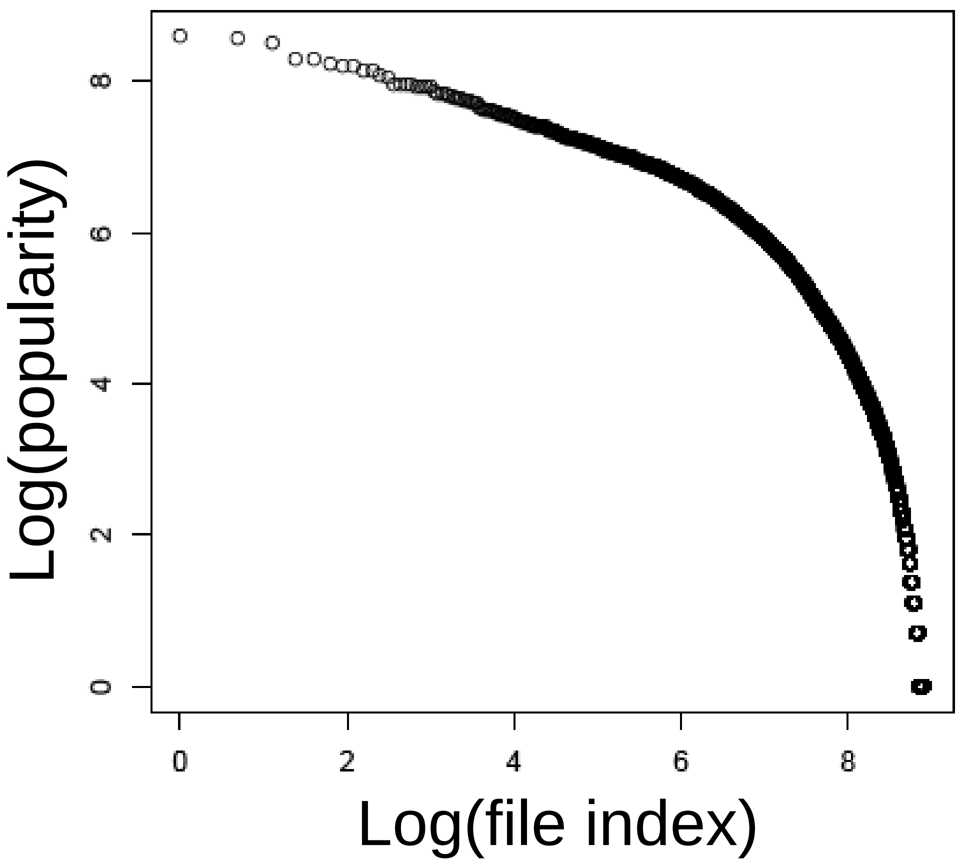

Data-driven comparison

- What if we sample a random graph from this popularity distribution?

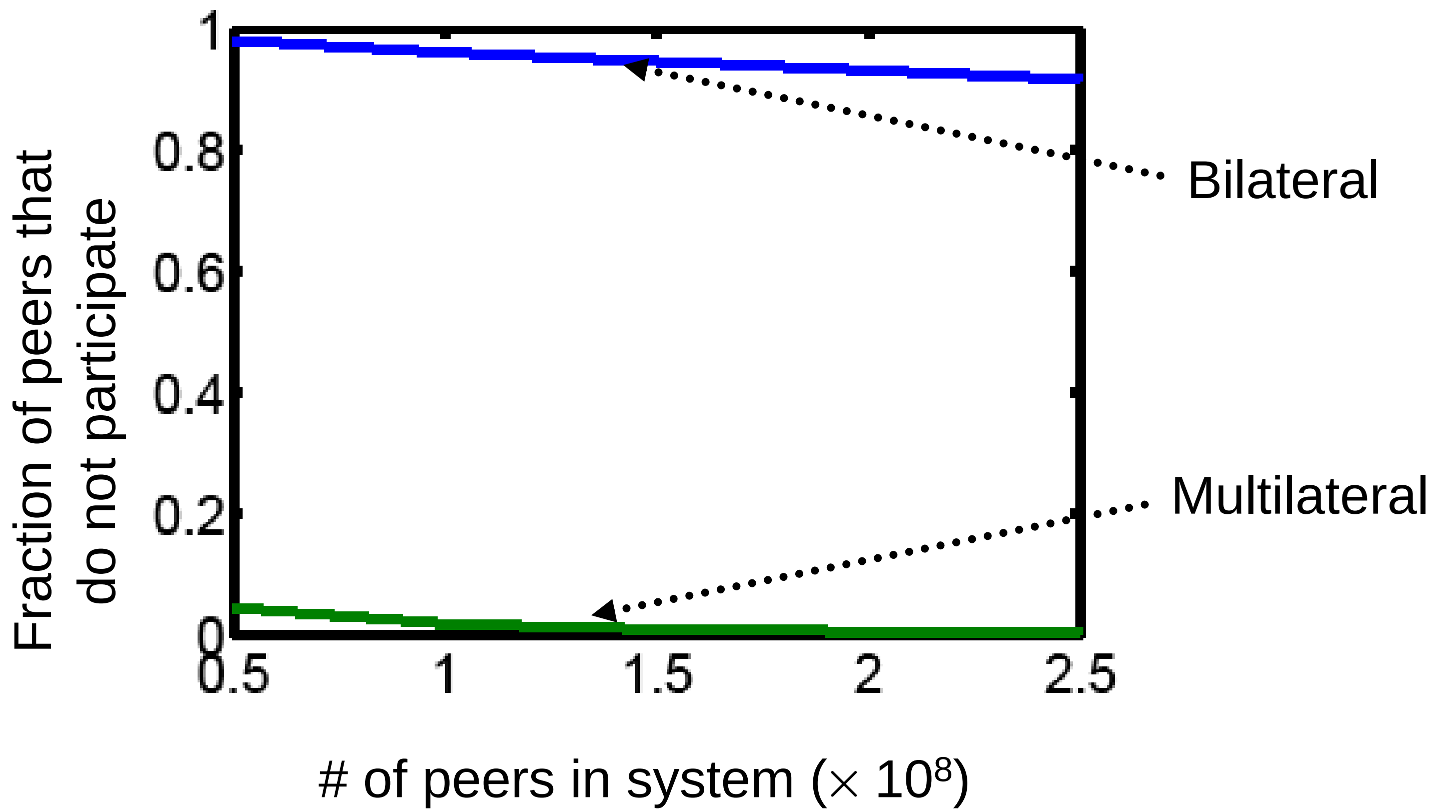

Data-driven comparison

- This comparison suggests that ME matches many more peers than BE
- However, as # of files a peer has *increases*, BE rapidly approaches ME
  - e.g., if all peers have 10 files, # of unmatched peers in BE is <2% in a system of 80K peers

Conclusions

- We have also characterized *why* one price per peer is the best scheme to use.
- We also have simple analysis of peer incentives in a price-based system.

Open issues:

How do we define BE and ME for a system with network constraints?

What is the messaging overhead of a price-based P2P system above a barter P2P system?

Are price-based systems dynamically *efficient?*

Example

$B_1 = 2; B_2 = \cdots = B_6 = 1$

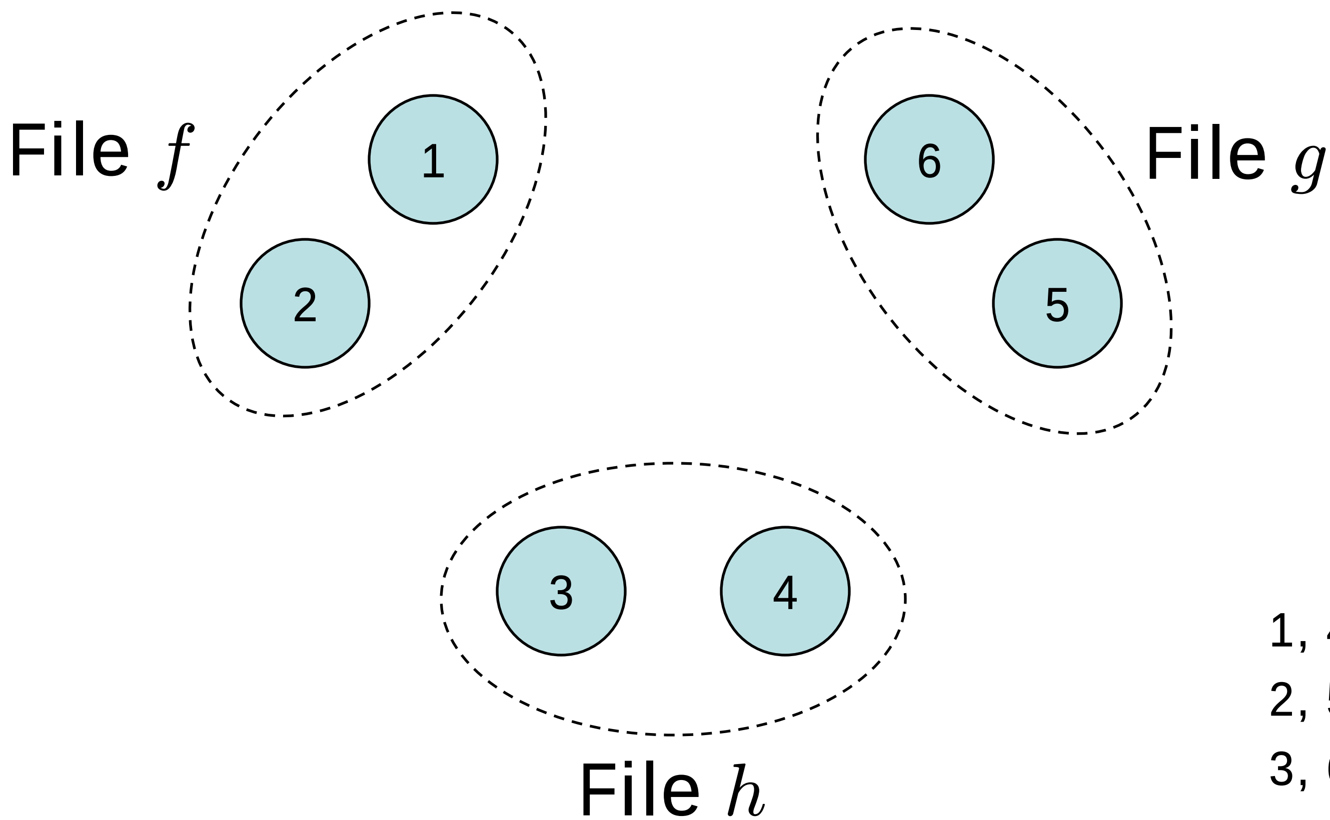

1, 4 want g
2, 5 want h
3, 6 want f

Example

$B_1 = 2; B_2 = \cdots = B_6 = 1$

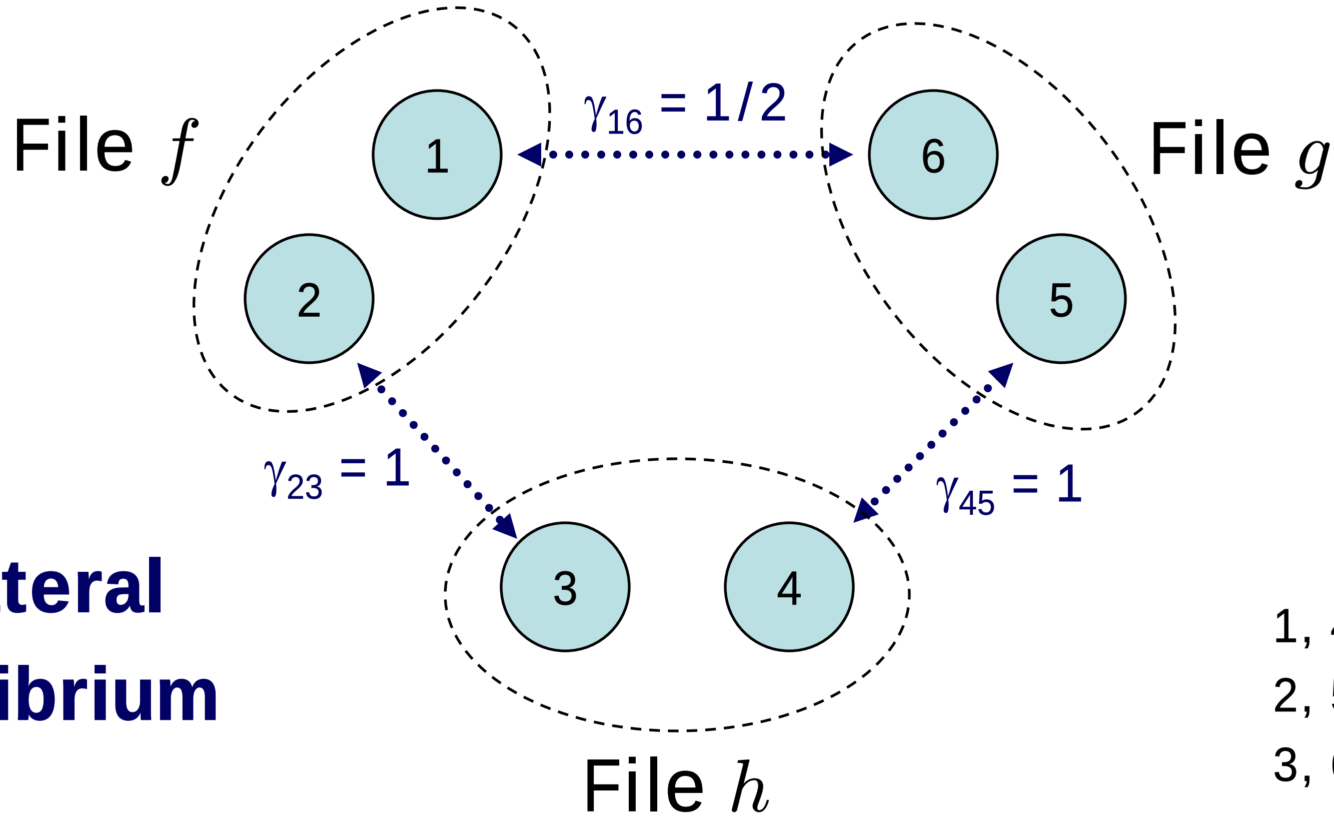

Bilateral Equilibrium

1, 4 want g
2, 5 want h
3, 6 want f

Example

There exists a profitable deviation for {1, 3, 5}:

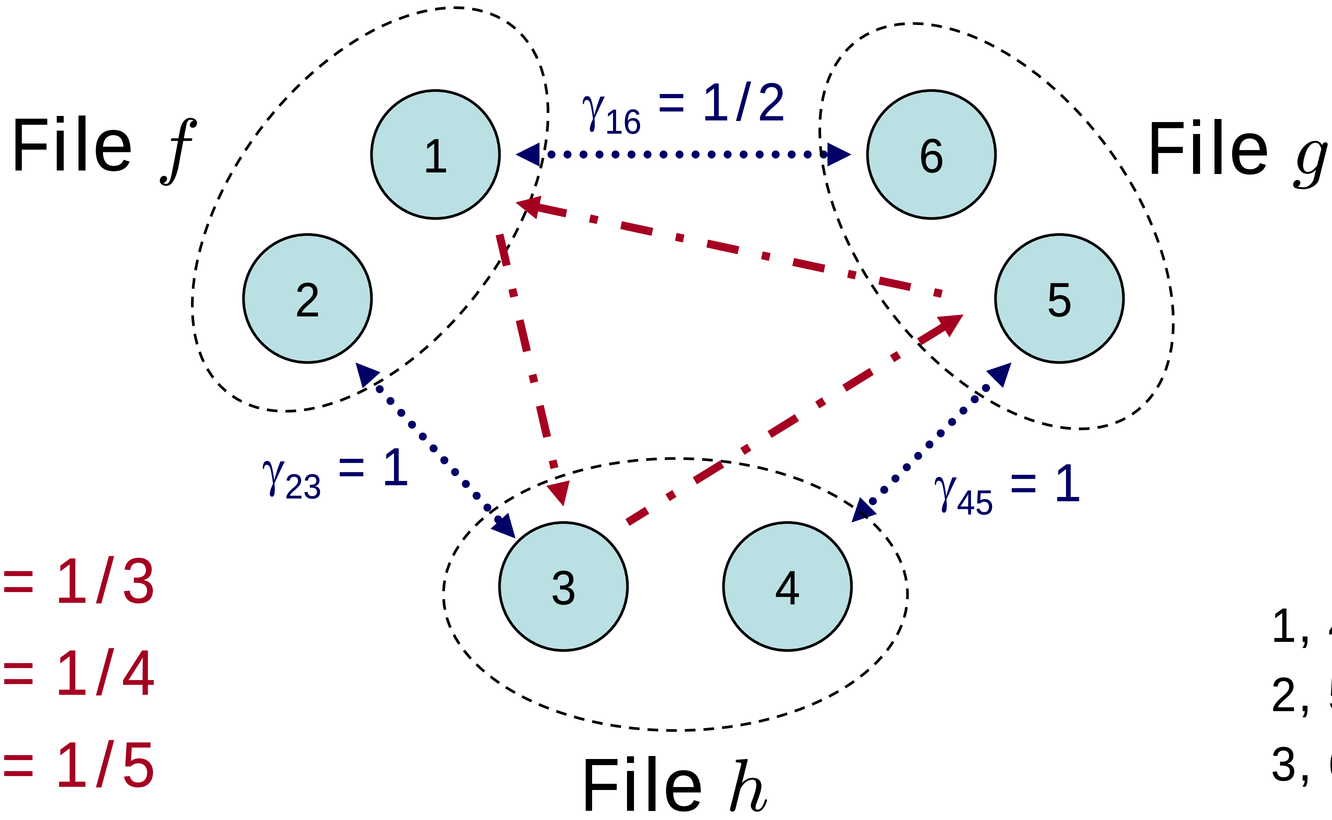

$r_{13} = 1/3$
$r_{35} = 1/4$
$r_{51} = 1/5$

1, 4 want g
2, 5 want h
3, 6 want f

Example

Total rate to 1 = 1/5 + 1/2 × (2 - 1/3) > 1, etc.

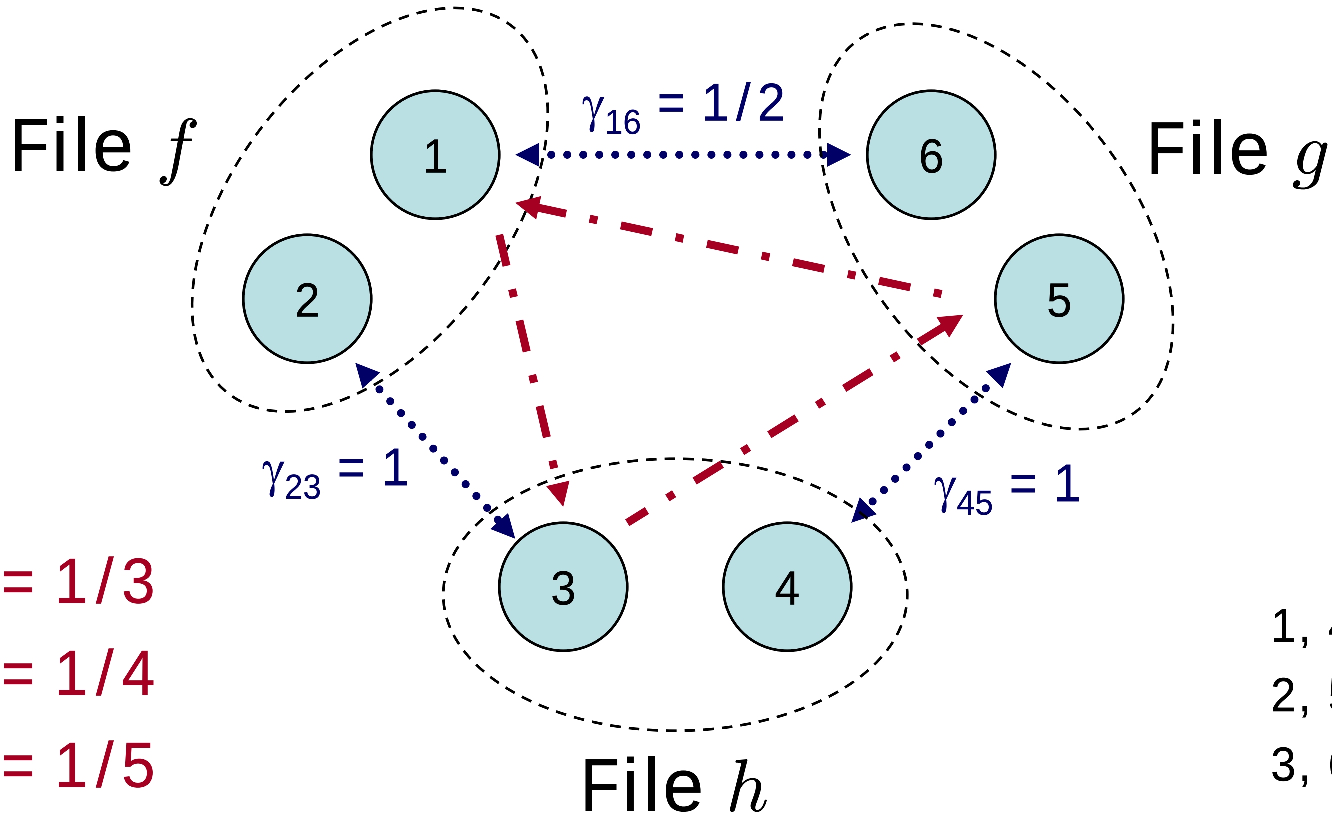

r_{13} = 1/3
r_{35} = 1/4
r_{51} = 1/5

1, 4 want g
2, 5 want h
3, 6 want f

Bilateral vs. multilateral: The core

- Bilateral equilibria are *not generally in the core*
- Key results:

(1) *Multilateral equilibria are* always *in the core*
(w.r.t. $\gamma_{ij} = p_i/p_j$)

(2) Suppose every peer uploads one file.

If $\underline{\mathbf{r}}^*$ is a bilateral equilibrium
with $d_{if} > 0$ for all i and files f that i wants,
and if $\underline{\mathbf{r}}^*$ is in the core,

then $\underline{\mathbf{r}}^*$ *is a multilateral equilibrium*.

Insight into proof of (2)

- Key step in establishing (2):

  Bilateral eq. is a multilateral eq. iff there exists p s.t. $\gamma_{ij} = p_i/p_j$ for all i, j

  [Idea: this ensures the peer optimizations become the same]
- If $\gamma_{ij} = p_i/p_j$,
  then $\Pi\,\gamma_{ij}$ along any cycle must equal 1
- We show that if the product is not equal to 1, then the bilateral eq. is not in the core