The role of prices in peer-assisted content distribution

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.

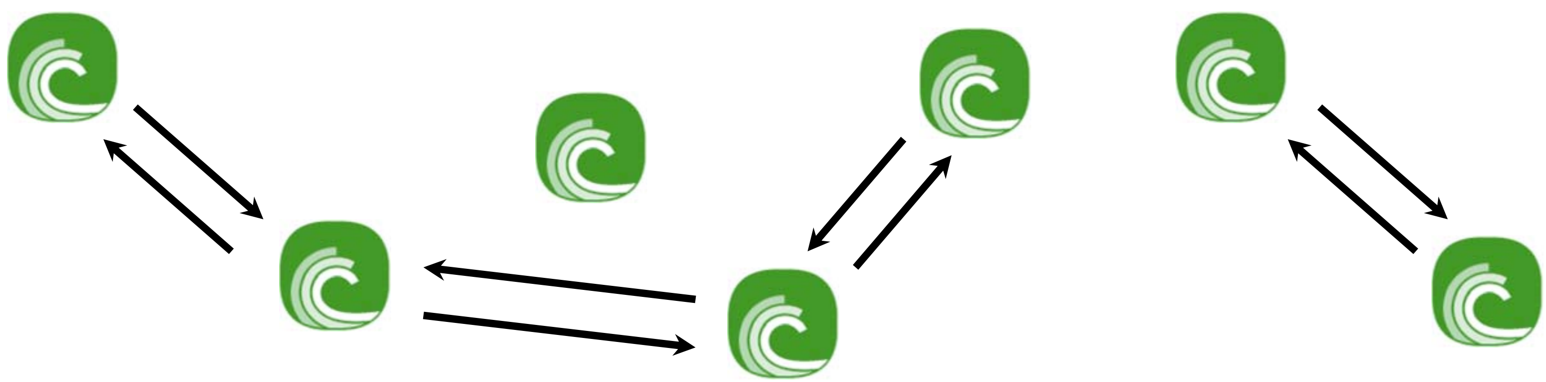

Content exchange mechanisms

- 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

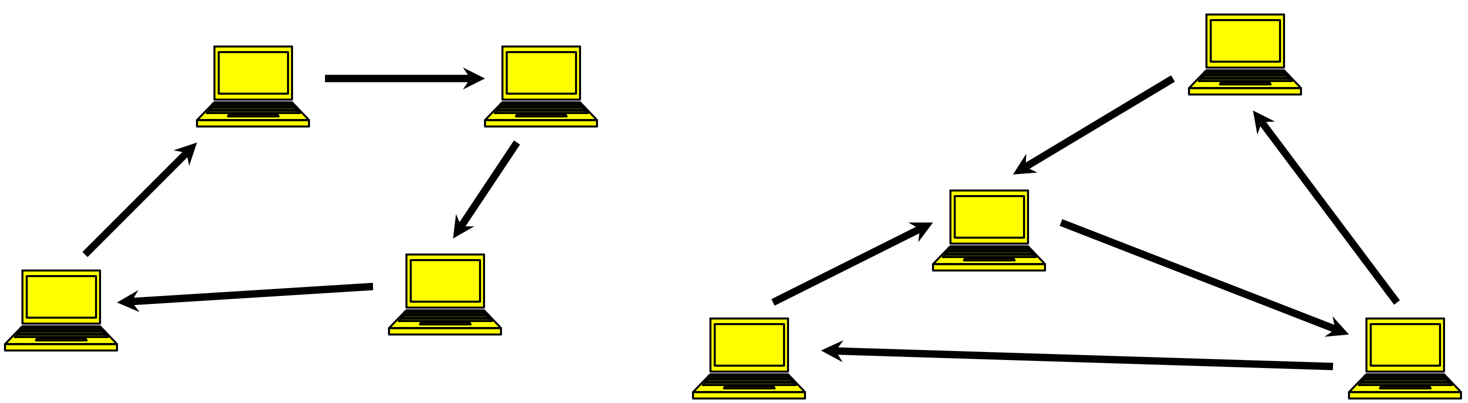

Content exchange mechanisms

- 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 and characterization of bilateral and multilateral content exchange: *How efficient? How robust?*

Outline

- Bilateral content exchange
- Multilateral content exchange
- Bilateral vs. multilateral
- Choosing the right pricing scheme
- Incentives
- Conclusions and future work

Bilateral content exchange

- Peers exchange content on a *pairwise* basis
- Let r_{ij} = 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:

$$\gamma_{ij} = p_{ij} / p_{ji}$$

BitTorrent and BitTyrant

- Canonical 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:

$$\gamma_{ij} = \frac{B_j / k_j}{r_{ij}}$$

- Note:
  This is decreasing in r_{ij}, so there is an incentive for i to make r_{ij} as small as possible while remaining in the active set $\Rightarrow$ **BitTyrant**

Exchange ratios

- The preceding discussion motivates us to model bilateral exchange systems via *exchange ratios.*
- *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

  B_i = bandwidth constraint of peer i

  $V_i(\underline{\mathbf{d}}_i)$ = *utility* to peer i of download rates $\underline{\mathbf{d}}_i$

Bilateral equilibrium

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

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

subject to r_{ijf} = 0, if user i does not have f

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

$\sum_{j,f} r_{ijf} \leq B_i$

- *Bilateral equilibrium* is a vector $\underline{\mathbf{r}}^*$ and exchange ratios $\underline{\gamma}^*$ such that:
All users have simultaneously optimized

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 γ must *coincide* with the optimal r_{ijf} and r_{jif} chosen by **peer j** given γ

Bilateral equilibrium

May not exist:

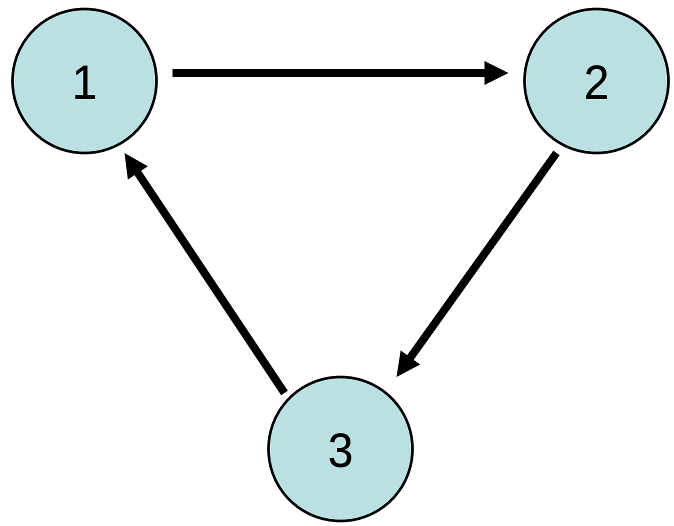

Each has a file that the clockwise neighbor wants.

No bilateral exchange possible!

(May also be inefficient even if it does exist...)

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

subject to $\quad r_{ijf} = 0$, if user i does not have f

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

$$\sum_j \sum_f r_{ijf} \leq B_i$$

Multilateral equilibrium

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

- Multilateral equilibria exist and are efficient (under mild conditions)
- Question:

  What is the "gap" between bilateral and multilateral equilibria?

  (We'll consider two answers.)

Bilateral vs. multilateral: The core

- Suppose exchange ratios γ are fixed.

  The allocation $\underline{\mathbf{r}}$* is in *the core* (w.r.t. γ) if no coalition of peers S can profitably deviate by:

  -trading with those outside S at the given exchange ratios; and

  -trading among themselves however they wish

Example

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

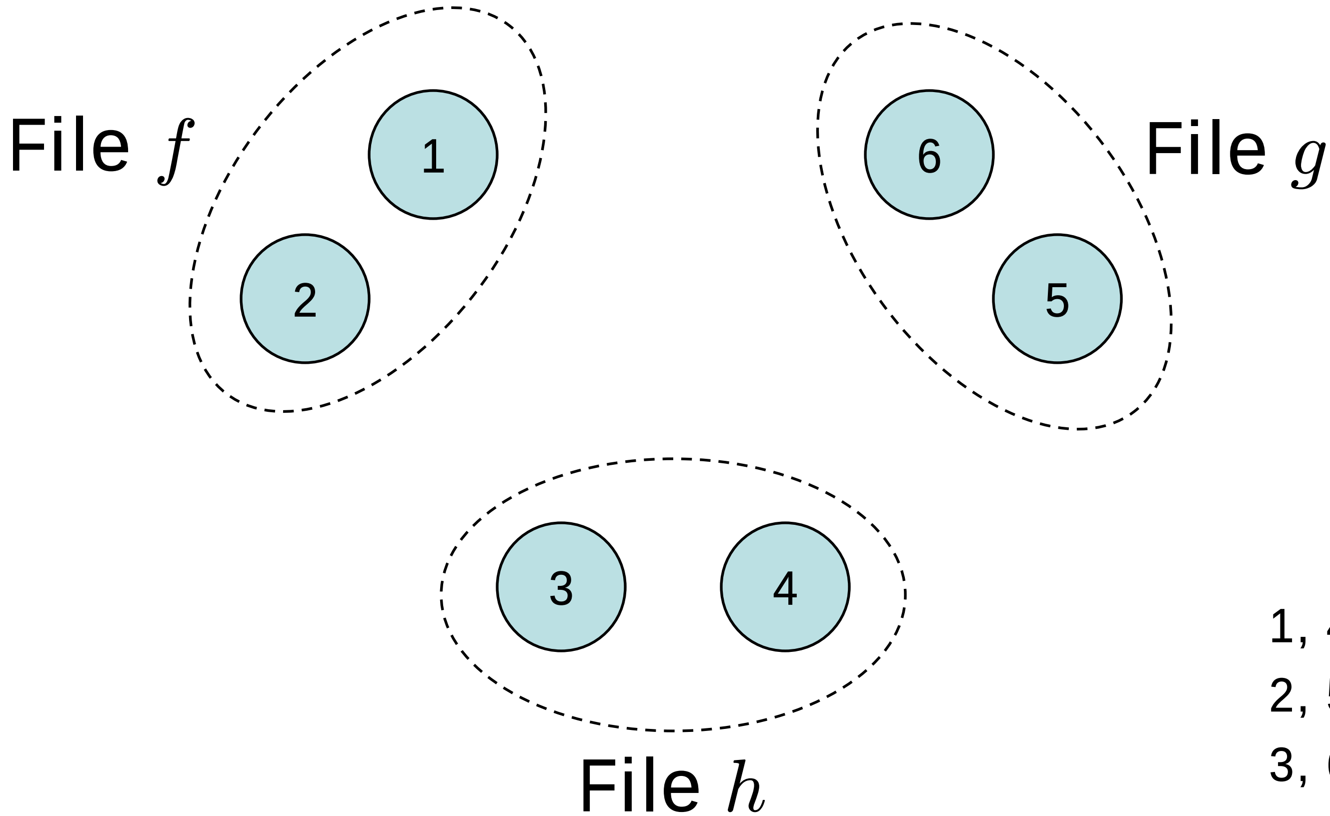

1, 4 want g

2, 5 want h

3, 6 want f

Example

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

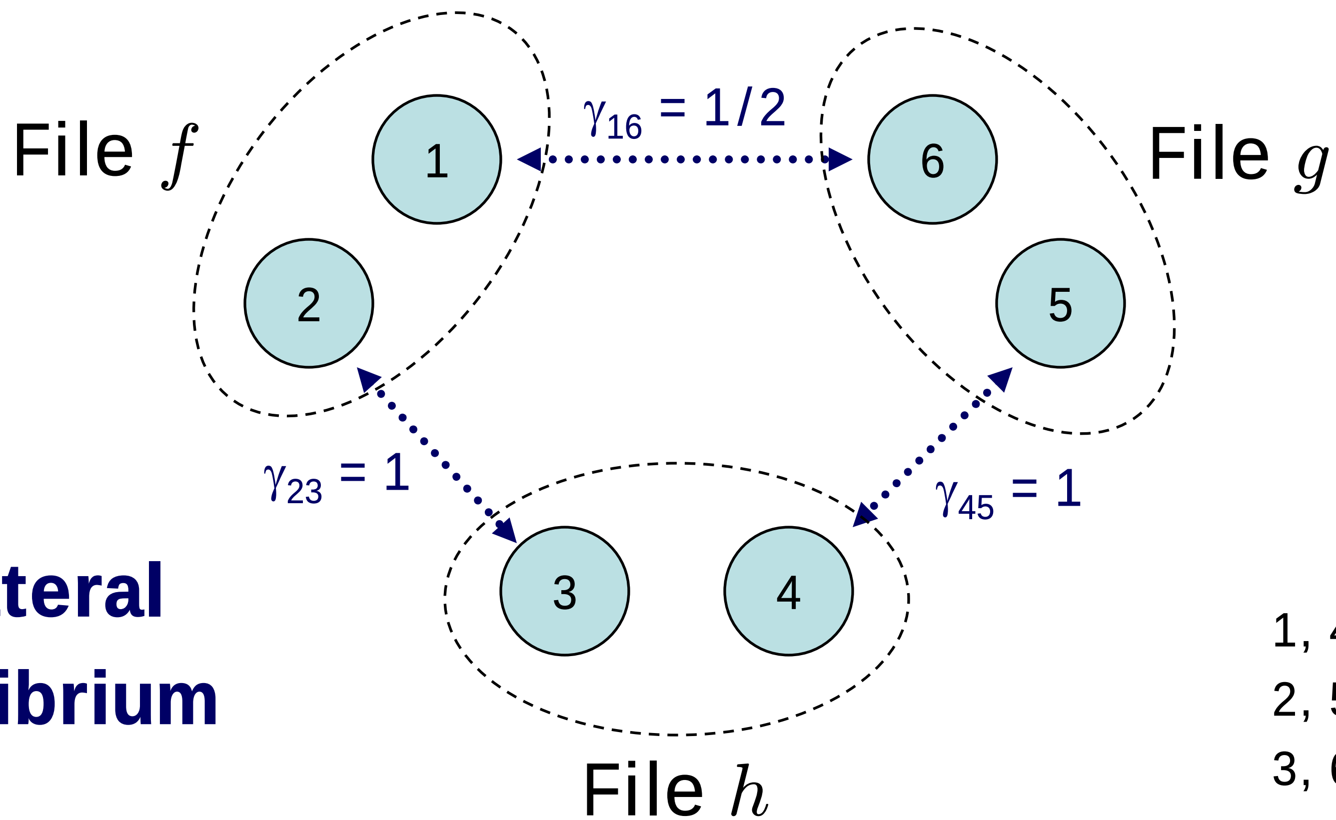

Bilateral Equilibrium

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

Example

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

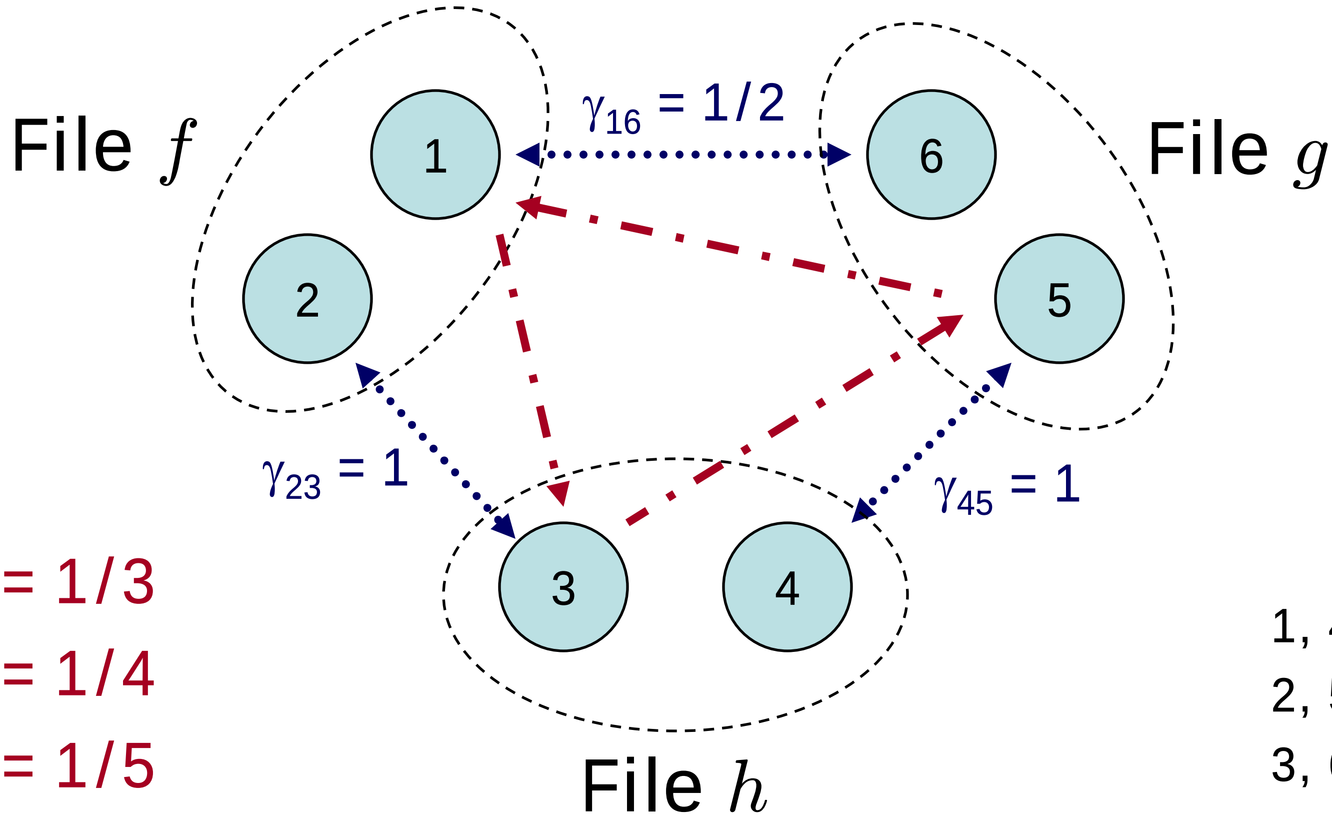

r_{13} = 1/3

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.

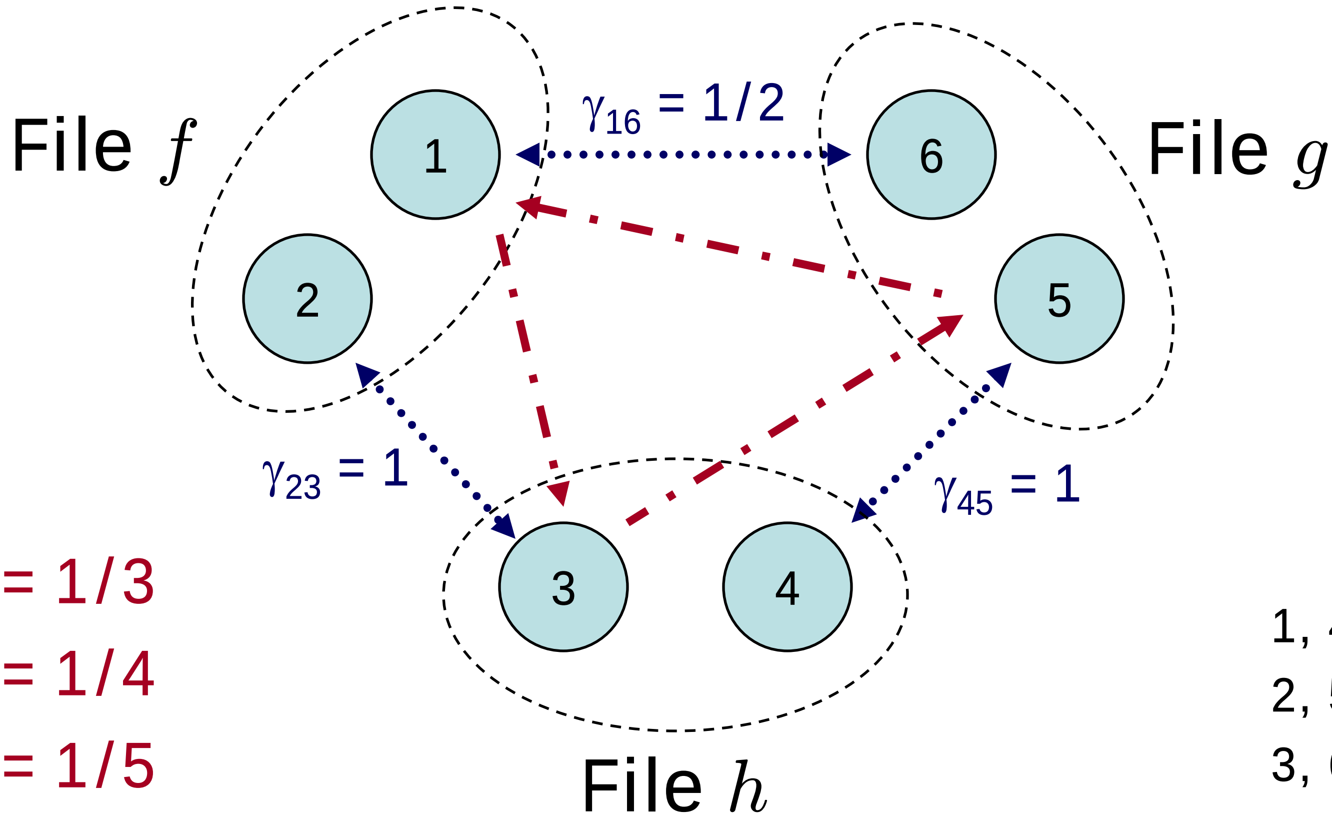

r_{13} = 1/3

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 $\mathbf{p}$ 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

Bilateral vs. multilateral: Existence

Another way to compare bilateral and multilateral equilibria is by determining *how likely they are to exist.*

As previously shown, multilateral equilibria *always exist,* while bilateral eq. may not.

We use a random model to quantify the nonexistence of bilateral equilibrium.

Bilateral vs. multilateral: Existence

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 bilateral equilibrium exists if and only if every peer has a complementary peer.

Bilateral vs. multilateral: Existence

We consider a random model where the probability a peer wants or has file f is proportional to f^{-s}, where $s < 1/2$.

[This corresponds to a light-tailed power law popularity distribution.]

Can show that if $K^{2-2s} > N$, then w.h.p. at least one peer has no complementary peer $\Rightarrow$ no bilateral equilibrium exists.

In particular, K need only be sublinear in N.

Bilateral vs. multilateral: Existence

- Can also show that the *fraction* of peers that must be deleted to ensure that a bilateral equilibrium exists, becomes arbitrarily close to 100% under the same condition

- Currently working to understand existence and nonexistence in heavy-tailed settings,
  and with asymmetric relationships between supply and demand

Choosing the right pricing model

In our model of multilateral exchange,
we set *one price per peer (PP).*

What about other choices?

One price per file (PF)?

Can show that for general networks:
if a PF multilateral eq. exists, then
a PP multilateral eq. exists, but not vice versa.

[Reason: PF pricing can't account for locality of demand in the network.]

Choosing the right pricing model

In our model of multilateral exchange,
we set *one price per peer (PP).*

What about other choices?

~~One price per file (PF)~~

One price per file per peer (PFP)?

Can show:

The set of PFP multilateral eq. is the same as the set of PP multilateral eq.

And PP has less state information...

Choosing the right pricing model

In our model of multilateral exchange, we set *one price per peer (PP).*

What about other choices?

- ~~One price per file (PF)~~
- ~~One price per file per peer (PFP)~~

…so we conclude that PP is the most desirable scheme.

Incentives

- Buyers:
  - prefer local sellers to remote sellers
  - prefer “cheap” wide-area connections
- Sellers:
  - are incented to upload their most valuable content whenever possible
  - will not have an incentive to manipulate the price in a large system (i.e., with *sufficient competition*)
  - create competitors at an exponential rate when the file is uploaded to others

Future directions I: Dynamics

This comparison of bilateral and multilateral equilibria is static.

However, users care about time-dependent statistics: e.g., file completion time

Can local, myopic price update rules provide the right dynamic resource allocation?

We have implemented a simulation of our system design (PACE) that suggests the answer is yes.

PACE vs. BT: Completion time

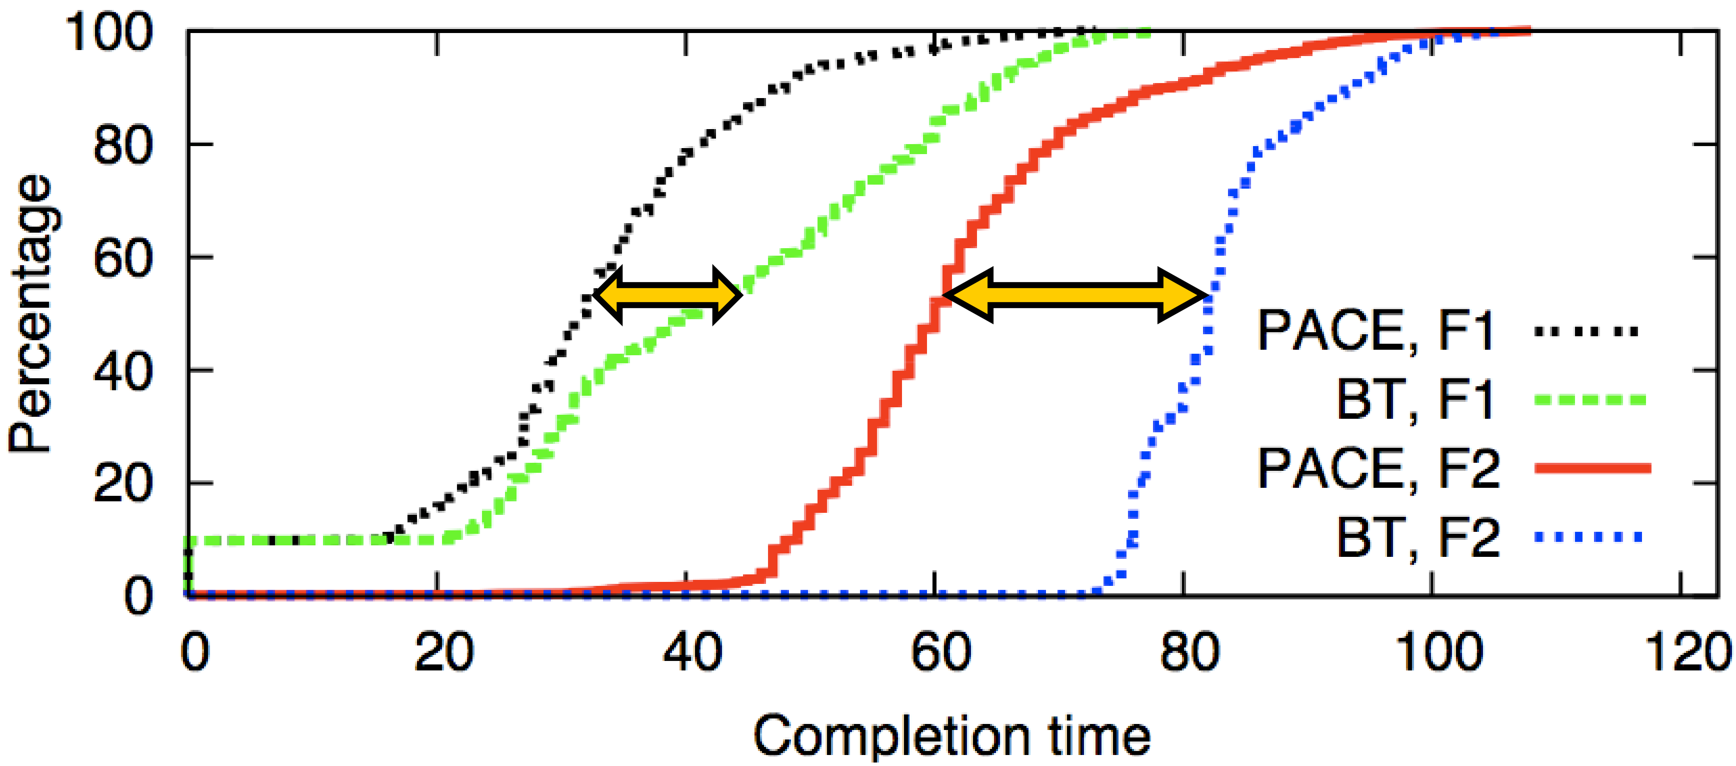

Two files: F1 starts at 10% of nodes;
F2 at a single node

Future directions II: The network

- Our discussion has not included resource constraints at the content distributor, or at network links within ISPs.
- A more general model should harmoniously integrate:
  - User demand for content
  - Load balancing and caching (content provider)
  - Traffic engineering (ISP)
- What are the right “price” signals to align resource allocation across all three?