



Free-riding, Fairness and Firewalls

in P2P File-Sharing

name: *Jan David Mol, Johan Pouwelse, Dick Epema, and Henk Sips*
date: September 11, 2008
event: 8th IEEE International Conference on Peer-to-Peer Computing





Common Assumptions in P2P

- *Connectivity*: Any two peers can connect.

- Connectivity is assumed to be high
- Lack of connectivity is a technical detail


- *Load Balancing*: Resource contribution can be balanced.

- Fair (1:1) resource exchange
 - SplitStream, BAR Gossip, Orchard
- Tit-for-tat
 - BitTorrent, ChunkySpread
- DHT load balancing

Connectivity (1)



- Common setups:
 - ADSL modem = NAT box
 - Consumer firewall (Windows/iptables)
 - Corporate firewall
- Firewalls/NATs **block incoming connections**
- Peers behind firewalls/NATs **cannot connect to each other**





Connectivity (2)

- *Solutions exist, but are not trivial:*

- Firewall/NAT configuration

- Manual: too complex
- Automatic: UPnP (1999) support is troubled

- Firewall puncturing techniques

- Not 100% effective, vendor specific
- Rarely used (too complex)

- *Security issues:*

- Firewalls exist to block connections
- Puncturing circumvents security

Connectivity (3)



<i>Study</i>	<i>Measured</i>	<i>Firewall / NAT</i>	<i>Rest (UPnP)</i>
[5] Burch	Internet	93%	7%
[2] Agarwal	P2P Streaming	77%	16% (9%)
[6] Chu	P2P Streaming	76%	24%
[22] Xie	P2P Streaming	70%	30% (19%)
[17] Pouwelse	BitTorrent	40%	60%
[23] Yang	Maze	40%	60%
[21] Wang	Gnutella / eDonkey	24 - 36%	64 - 76%

- None use firewall puncturing.

Load Balancing

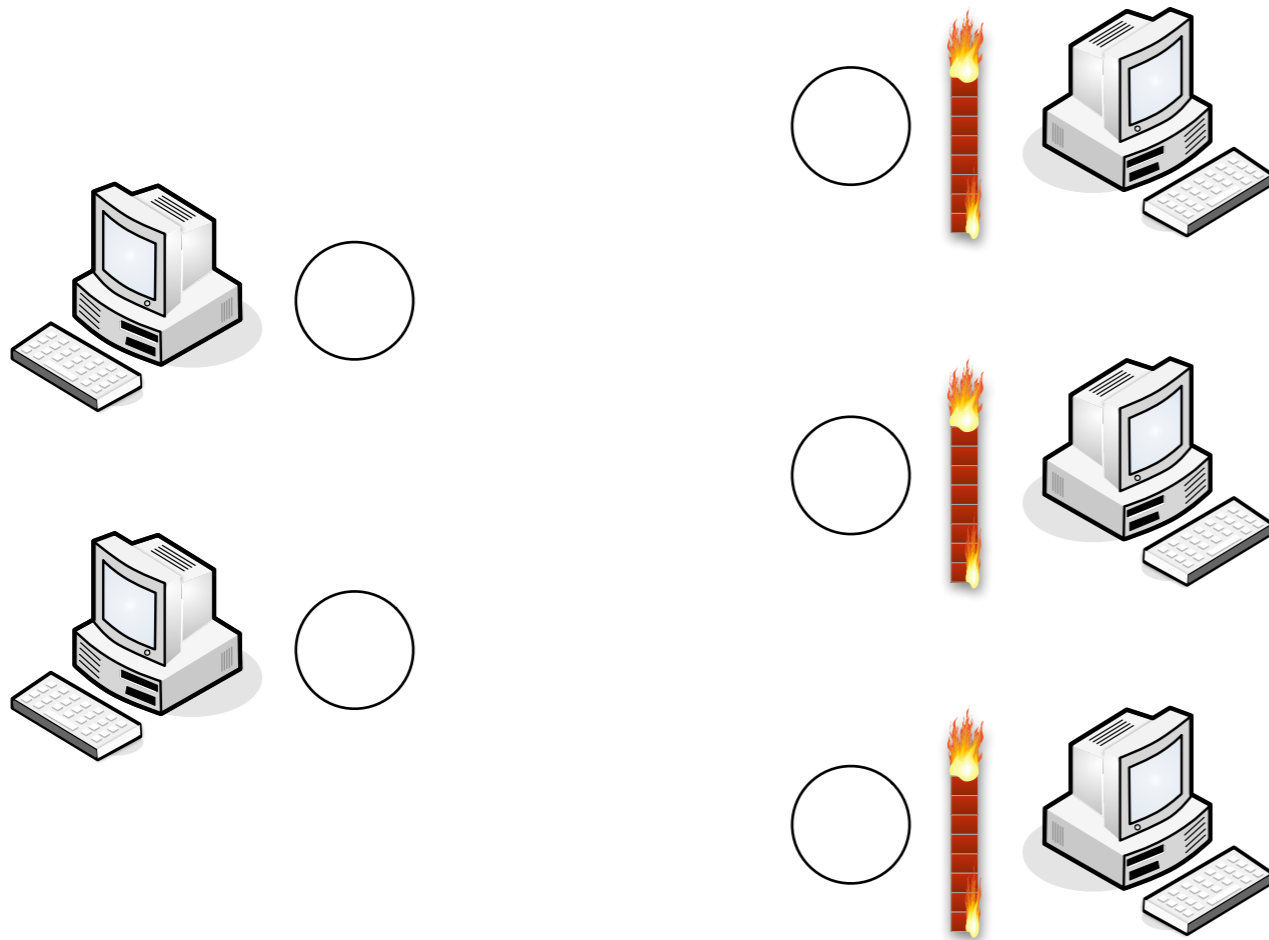
- *Fair* if every peer has contribution \approx consumption
- *Sharing ratio* as a metric of fairness

$$SR = \frac{\# \text{ uploaded bytes}}{\# \text{ downloaded bytes}}$$

- $SR > 1$: Net producer
- $SR < 1$: Net consumer
- Per peer or per group


Conflict Example

- Distributing a file of size L .
- 3 firewalled peers:




Conflict Example

- Distributing a file of size L .
- 3 firewalled peers:



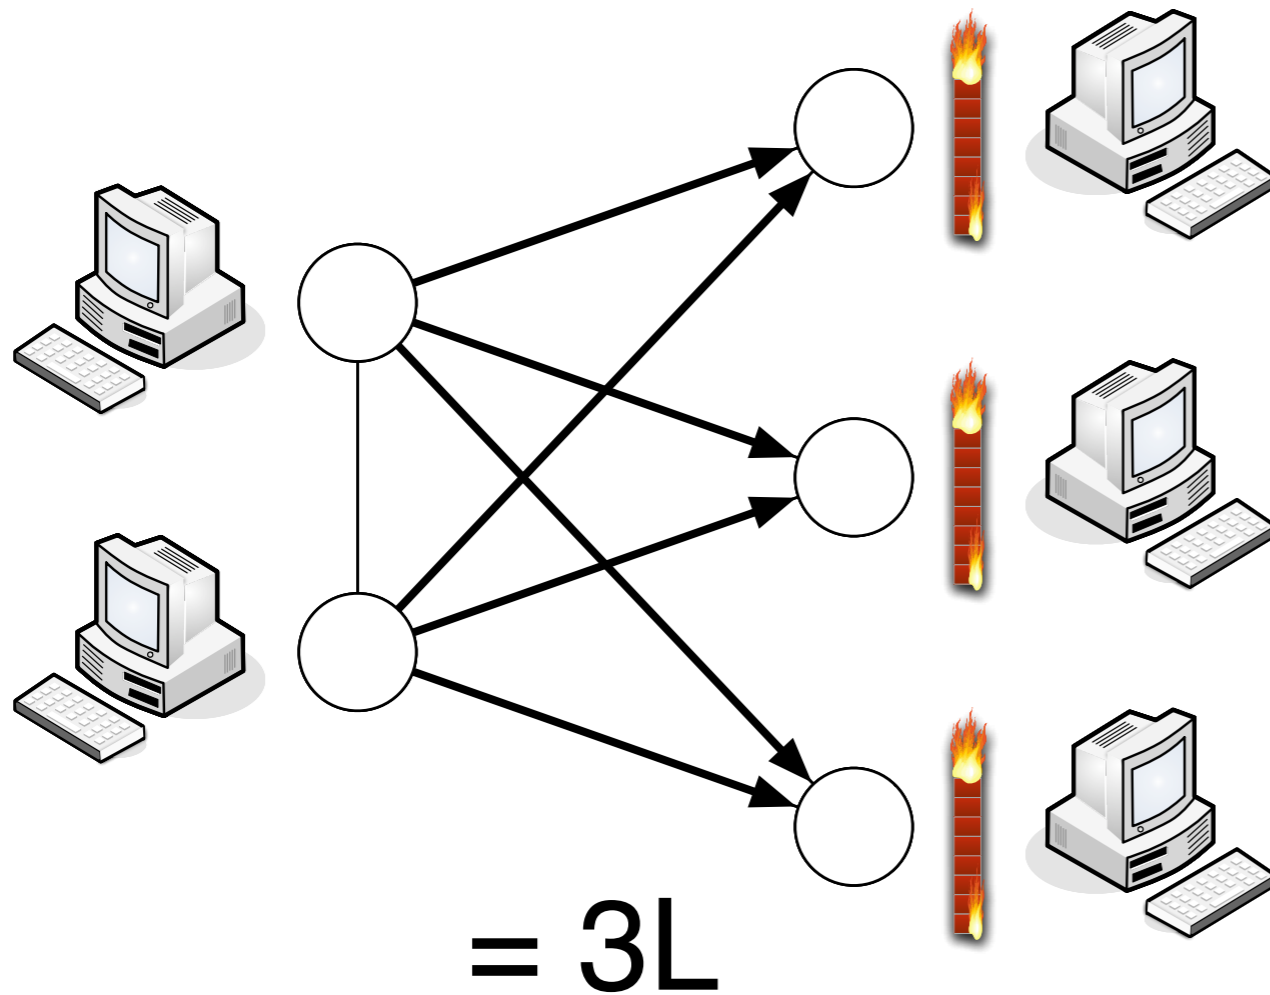
Conflict Example

- Distributing a file of size L .
- 3 firewalled peers:



Conflict Example


- Distributing a file of size L .
- 3 firewalled peers:



down: $3L$

Conflict Example

- Distributing a file of size L .
- 3 firewalled peers:




down: $3L$

up: $\leq 2L$

Conflict Example

- Distributing a file of size L .
- 3 firewalled peers:




down: $3L$

up: $\leq 2L$

$$SR \leq \frac{2}{3}$$

Model



We define:


- *Firewalled peer*: cannot accept incoming connections
- *Connectable peer*: can accept incoming connections
- f = fraction of firewalled peers

We assume:

- A P2P network of N peers sharing data
- If p or q is connectable, p and q can exchange data
- Simplicity: every peer downloads L bytes


General Case

- Distributing a file of size L .
- fN firewalled peers:



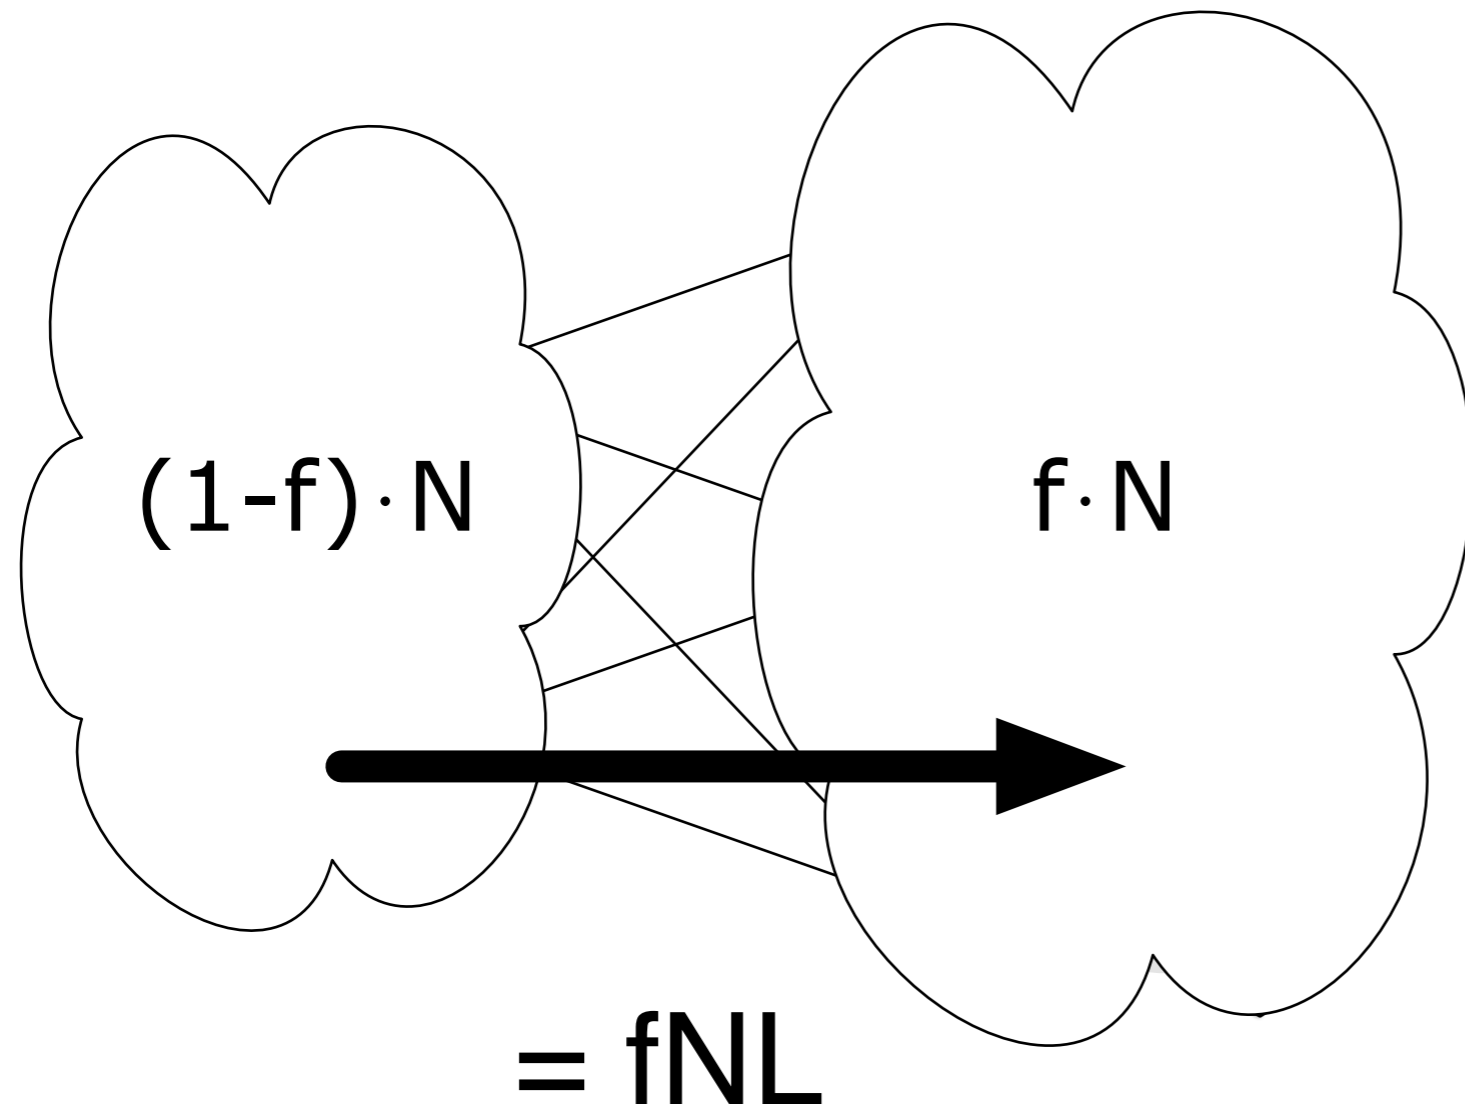
General Case

- Distributing a file of size L .
- fN firewalled peers:



General Case

- Distributing a file of size L .
- fN firewalled peers:

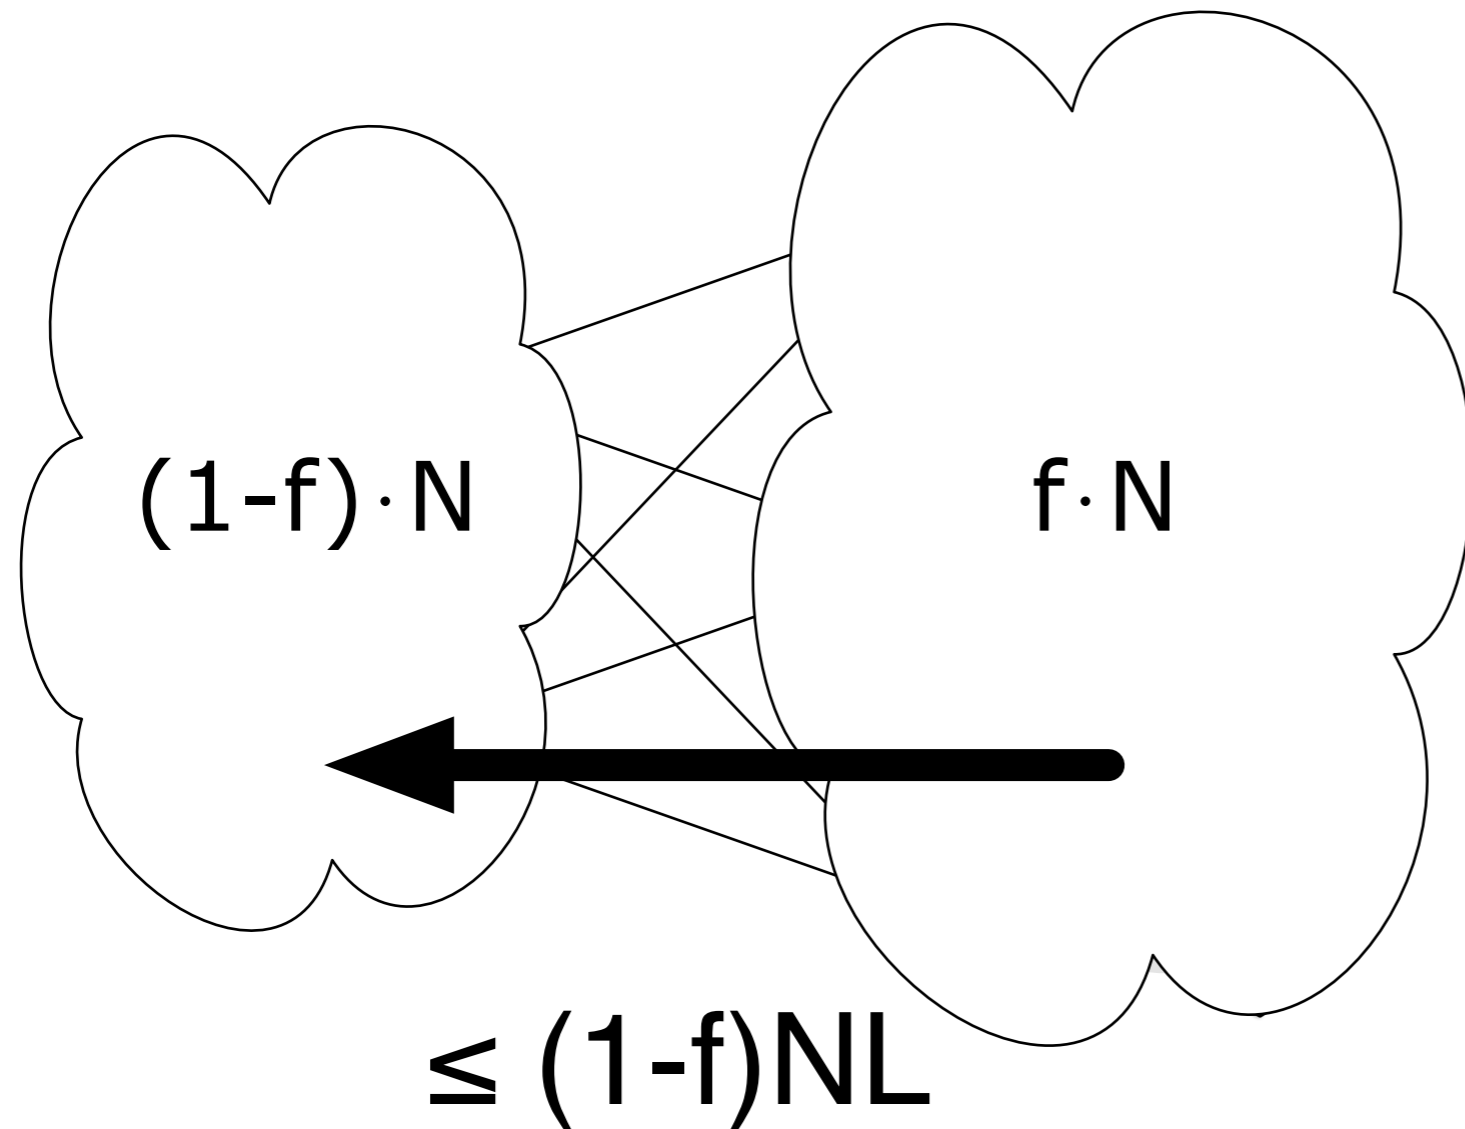


down: fNL



General Case

- Distributing a file of size L .
- fN firewalled peers:




down: fNL

up: $\leq (1-f)NL$



General Case

- Distributing a file of size L .
- fN firewalled peers:



down: fNL

up: $\leq (1-f)NL$

$$SR \leq \frac{1}{f} - 1$$

Bounds on Sharing Ratio



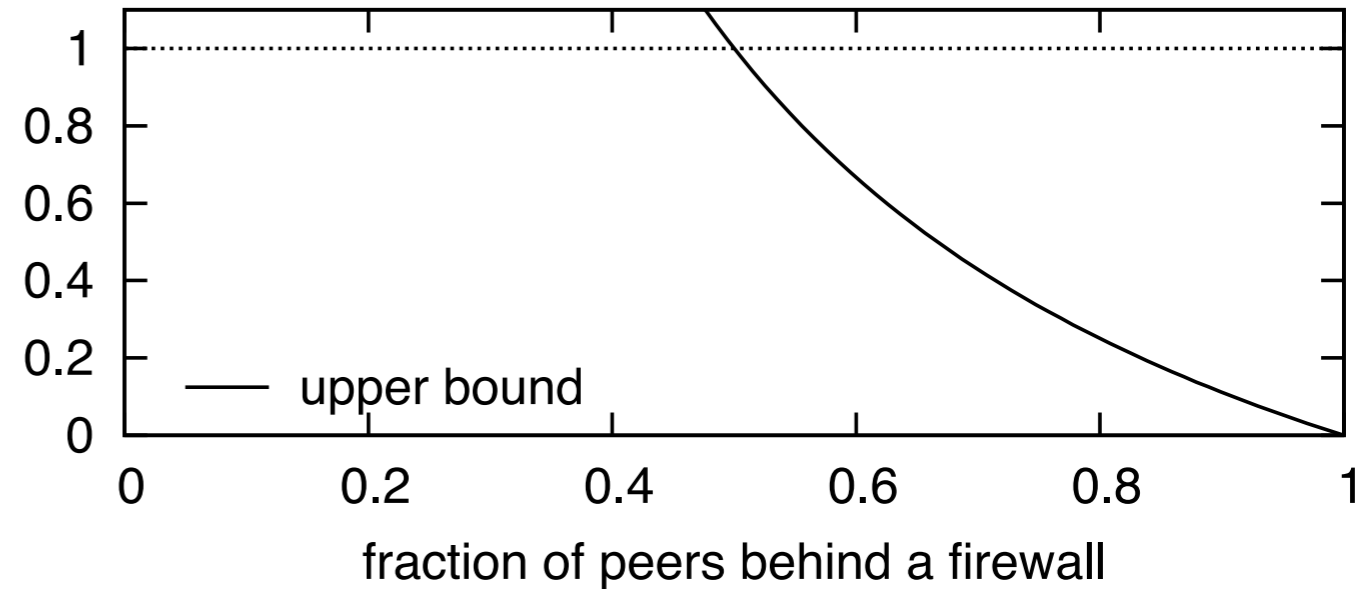
Firewalled peers:

$$SR \leq \frac{1}{f} - 1$$

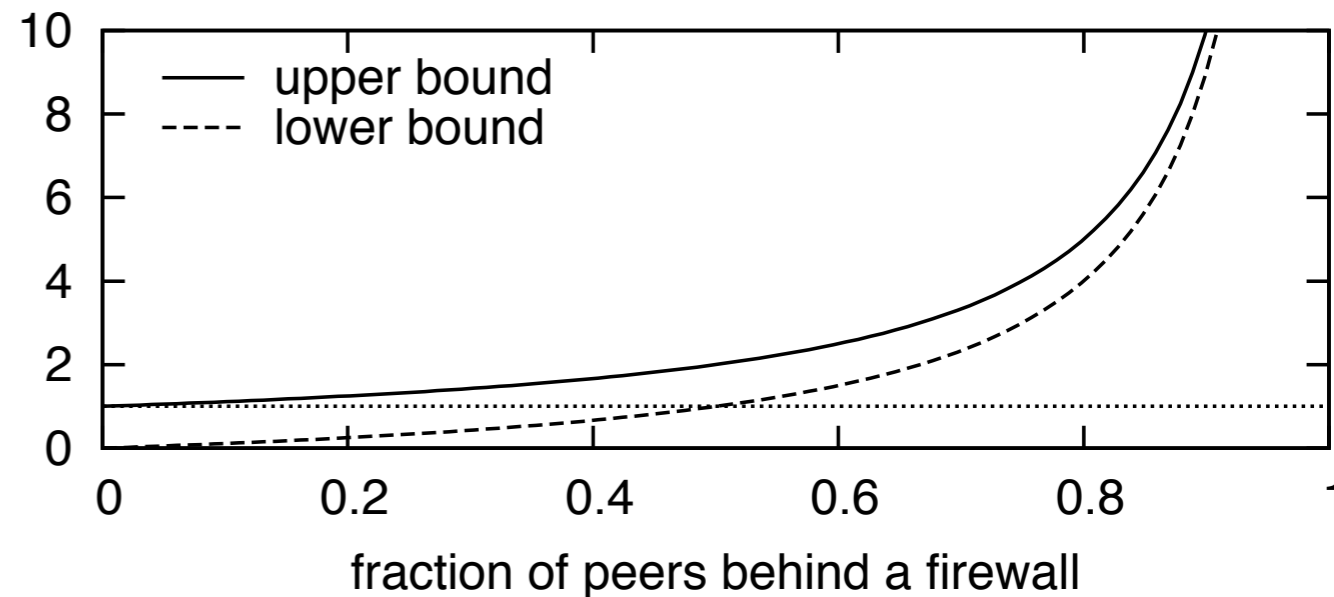
Connectable peers:

$$\frac{1}{1-f} - 1 \leq SR \leq \frac{1}{1-f}$$

expected sharing ratio



expected sharing ratio



Simulations: BitTorrent

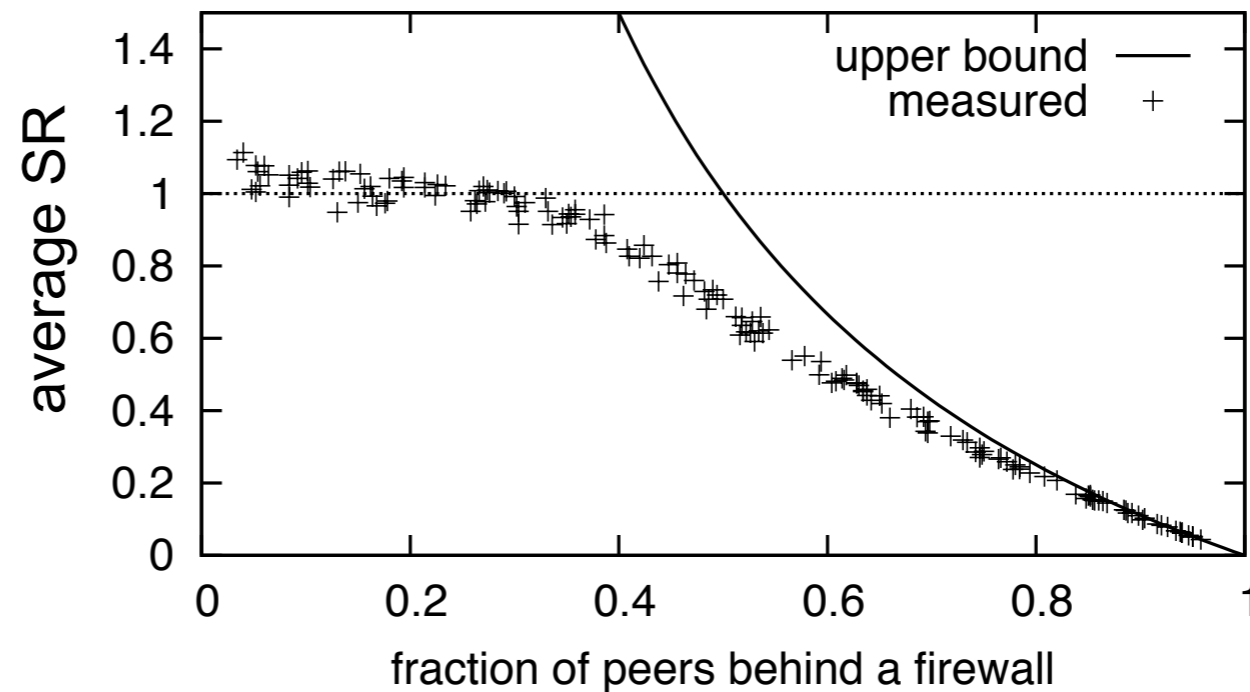


One point =
average SR in
one experiment

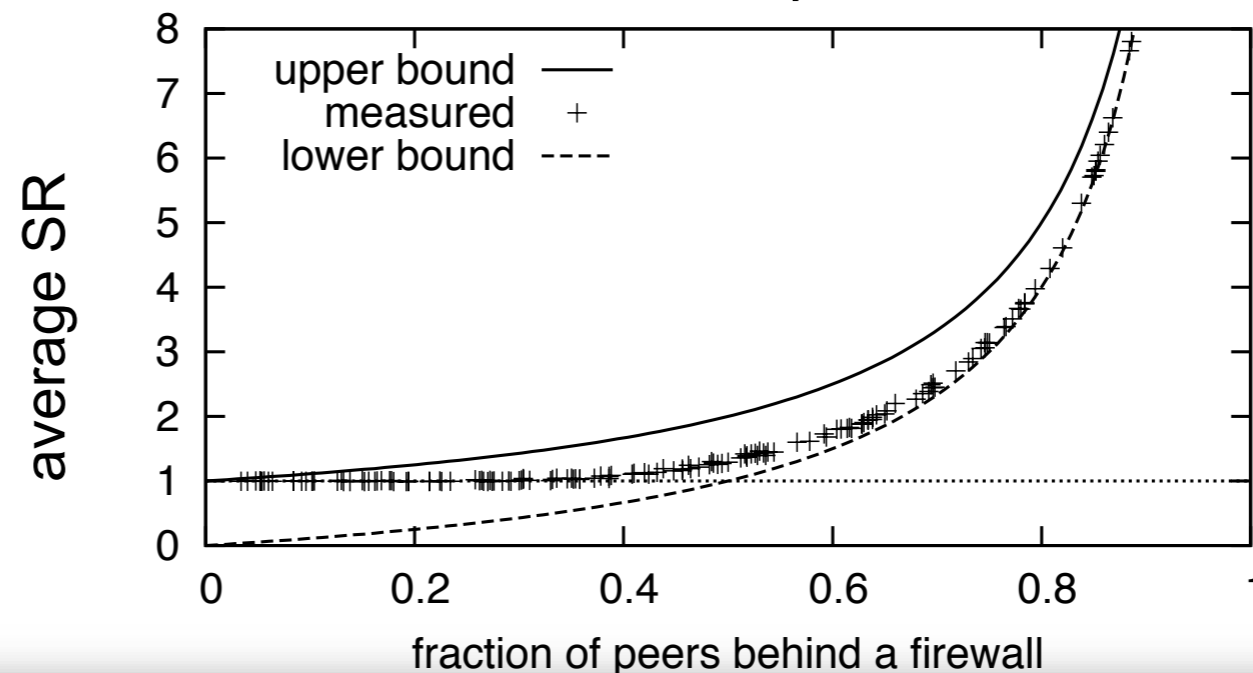
One experiment:

- 500 peers
- share 10MB file
- depart @ $SR=1$

Firewalled peers



Connectable peers





Firewall puncturing

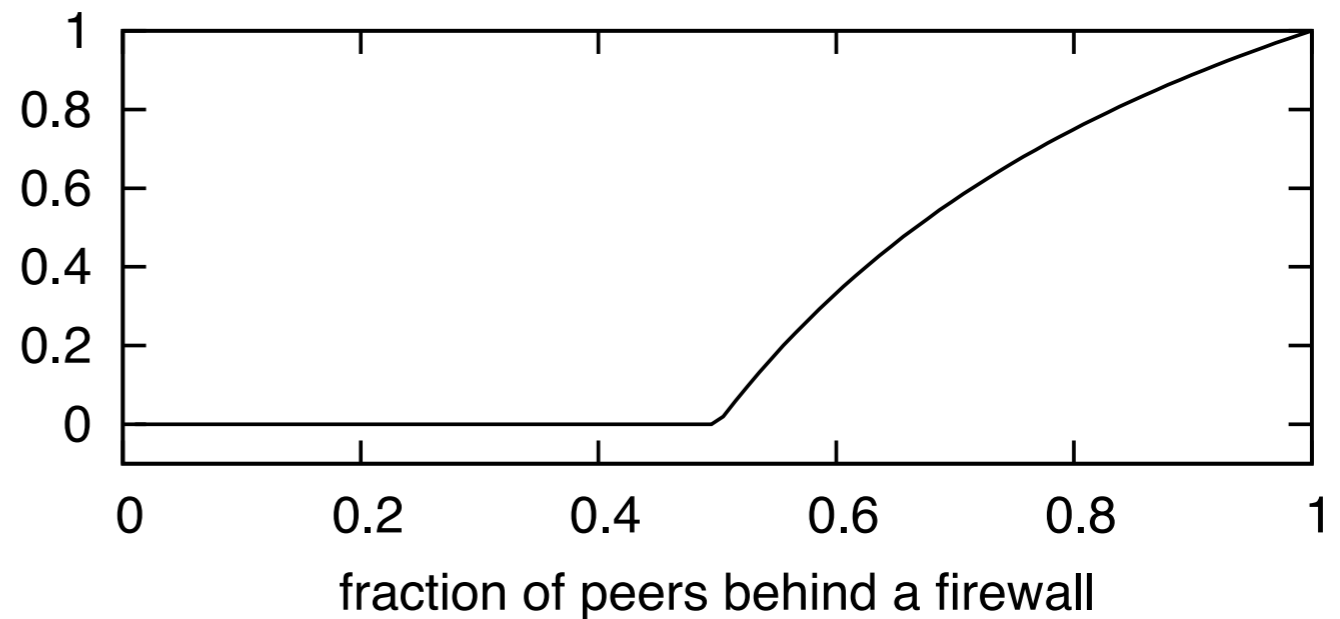
- Firewall puncturing not 100% effective

- Ford et al. [7]: TCP: 64%, UDP: 82%
- What efficiency *is* required?
- To allow fair SRs

- At least:

$$\eta \geq \left(2 - \frac{1}{f}\right) \cdot 100\%$$

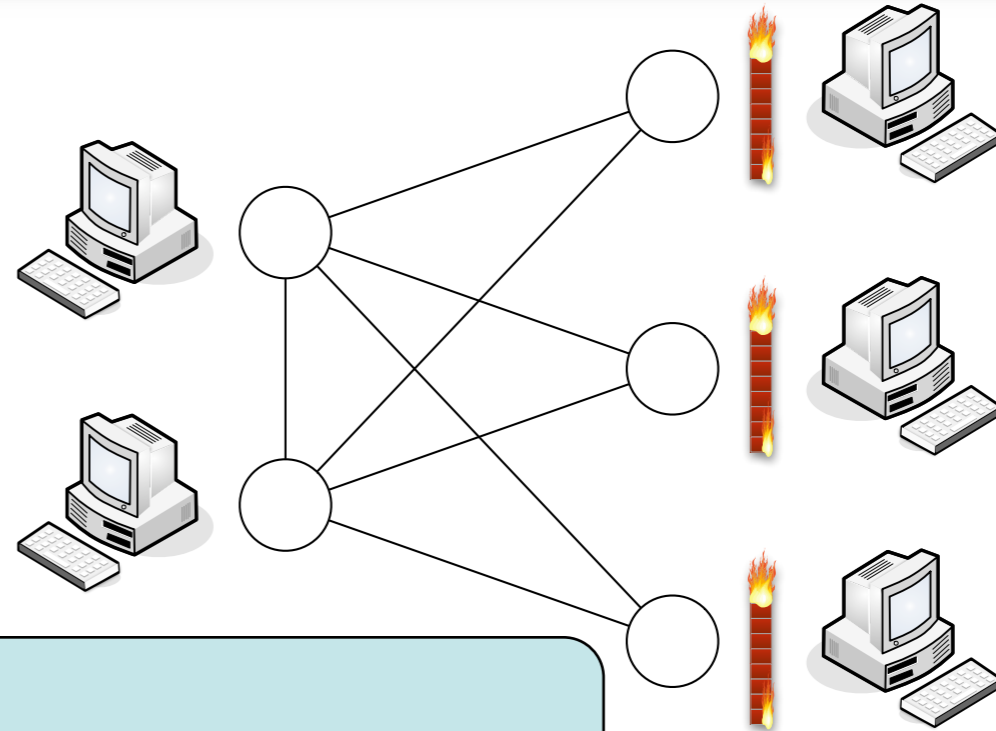
fraction of firewalls
to puncture



- So we can support:

- TCP: $f \leq 74\%$
- UDP: $f \leq 85\%$

Model is Generic



Holds for all P2P networks

Model independent of:

- Connection speeds (assumed infinite)
- Arrival/departure patterns (assume all peers on-line)
- Neighbourhood selection (assumed fully connected)
- Method of distribution (algorithm, downloading/vod/live)



Example Usage

Suppose:

- P2P Live Streaming system
 - 1 Mbit/s stream
- 75% firewalled at some point

Then:

- Firewalled peers:
 - average SR $\leq 1/3$
 - average upload speed ≤ 0.33 Mbit/s
- Connectable peers:
 - average SR ≥ 3
 - average upload speed ≥ 3 Mbit/s
- >66% puncturing efficiency needed for SR=1



Sharing Ratio Enforcement?

Policy: ban peers with low SR

Open community:

- Everyone can join
- No SR enforcement

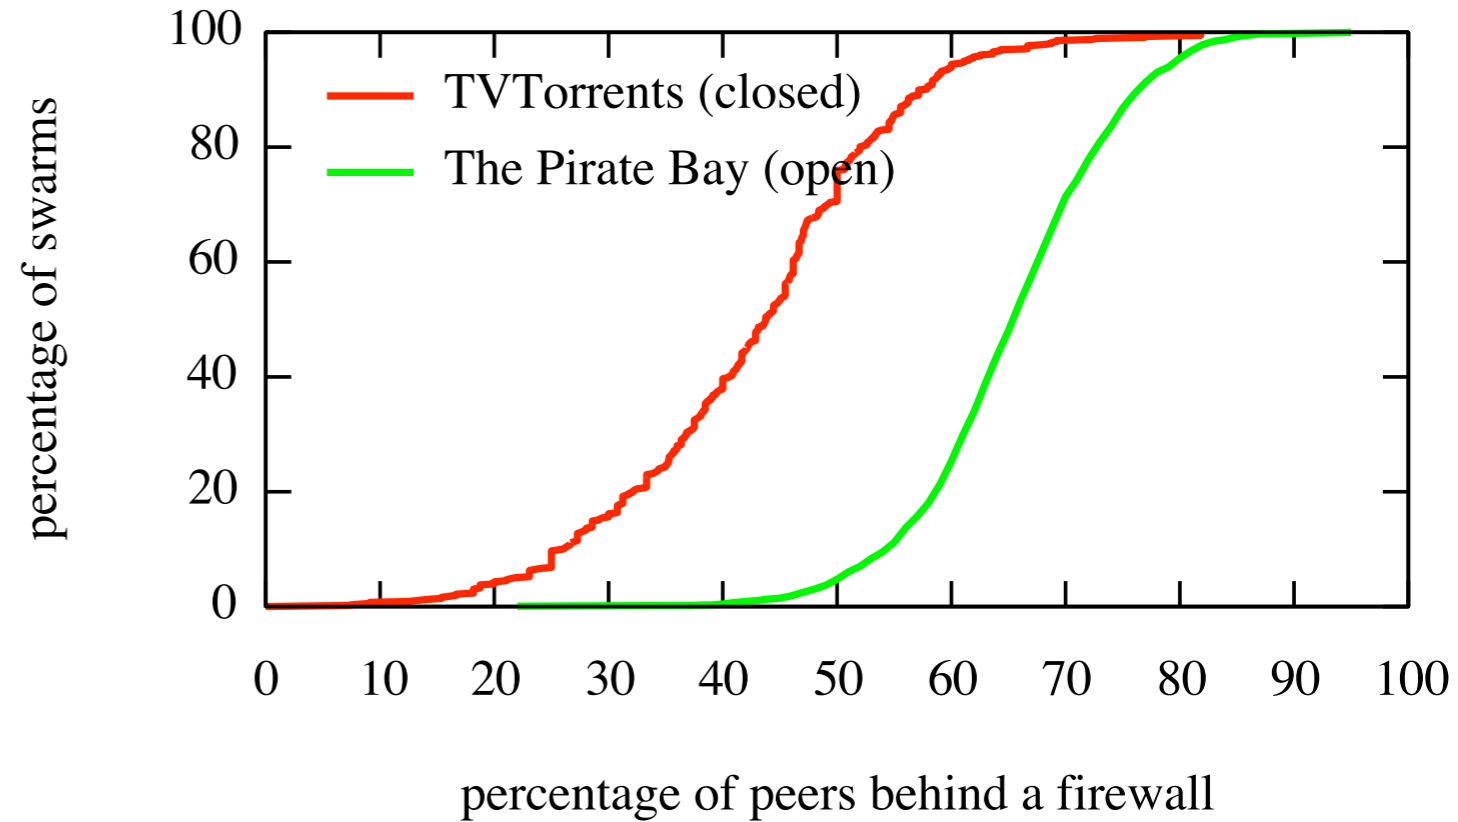
Closed community:

- Membership
- SR Enforcement
- Cannot sustain high **f**



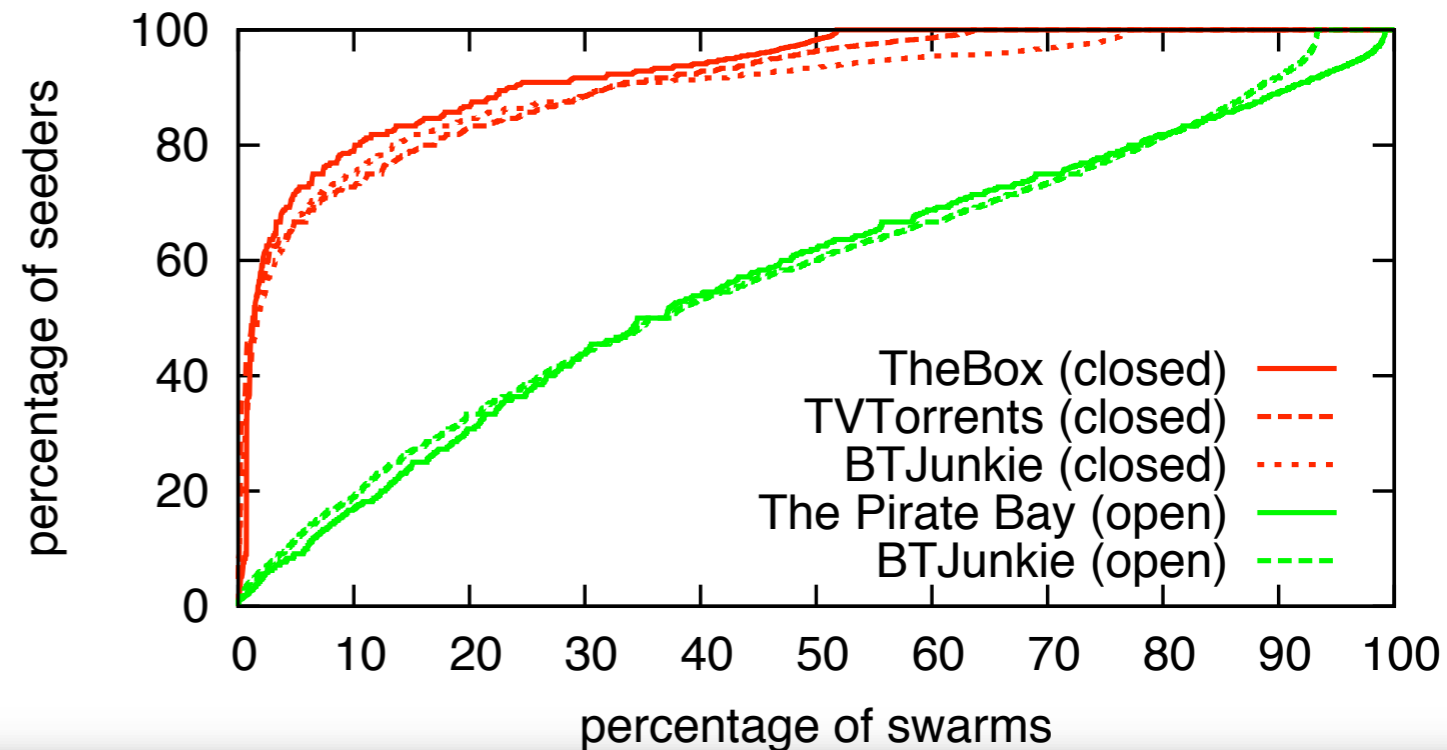
Measurements: 33,823 swarms

Open: 66% firewalled
Closed: 45% firewalled

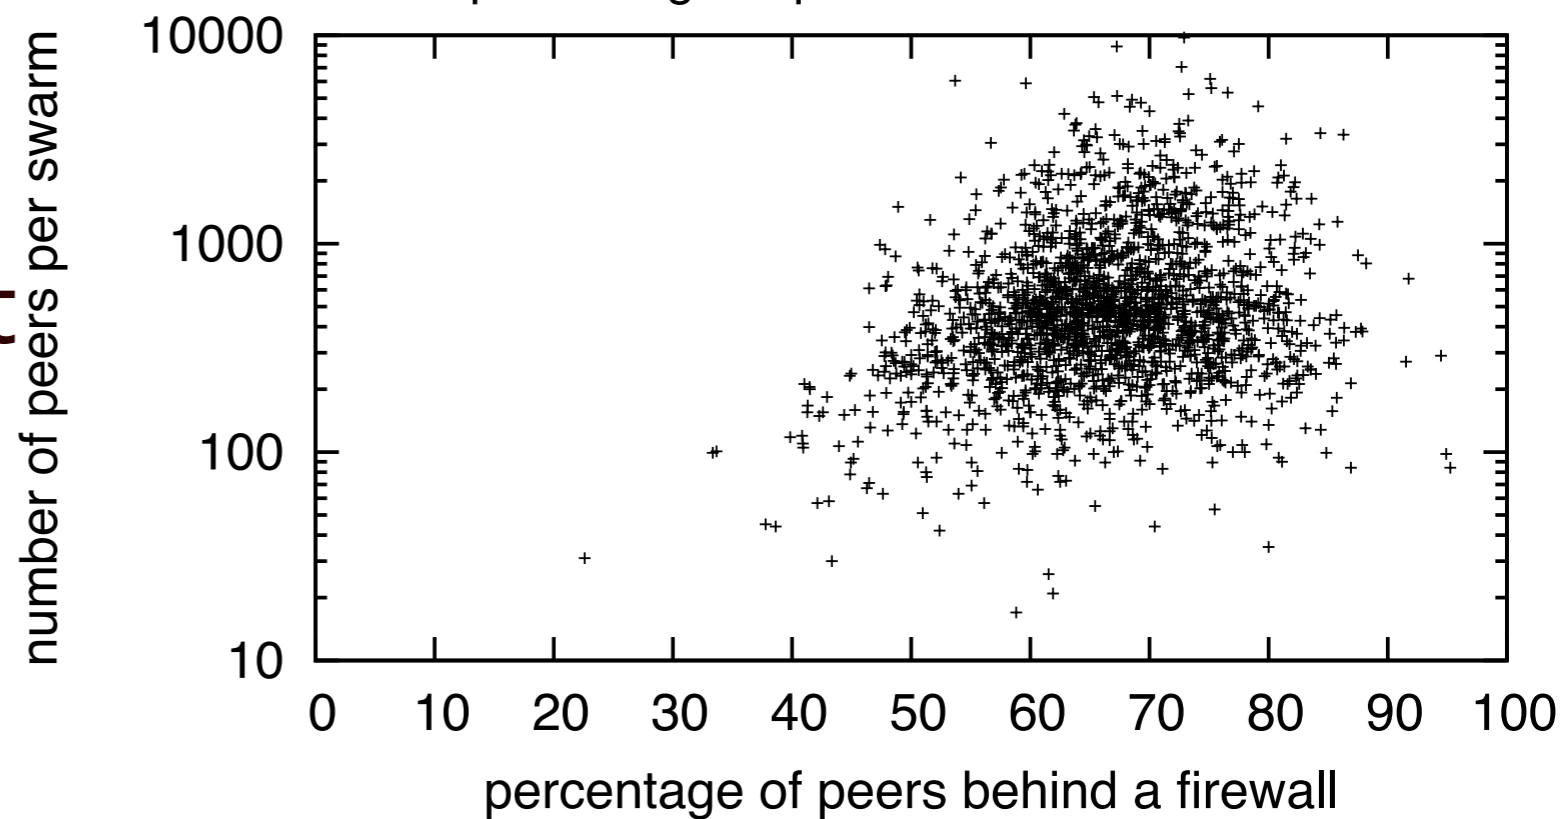
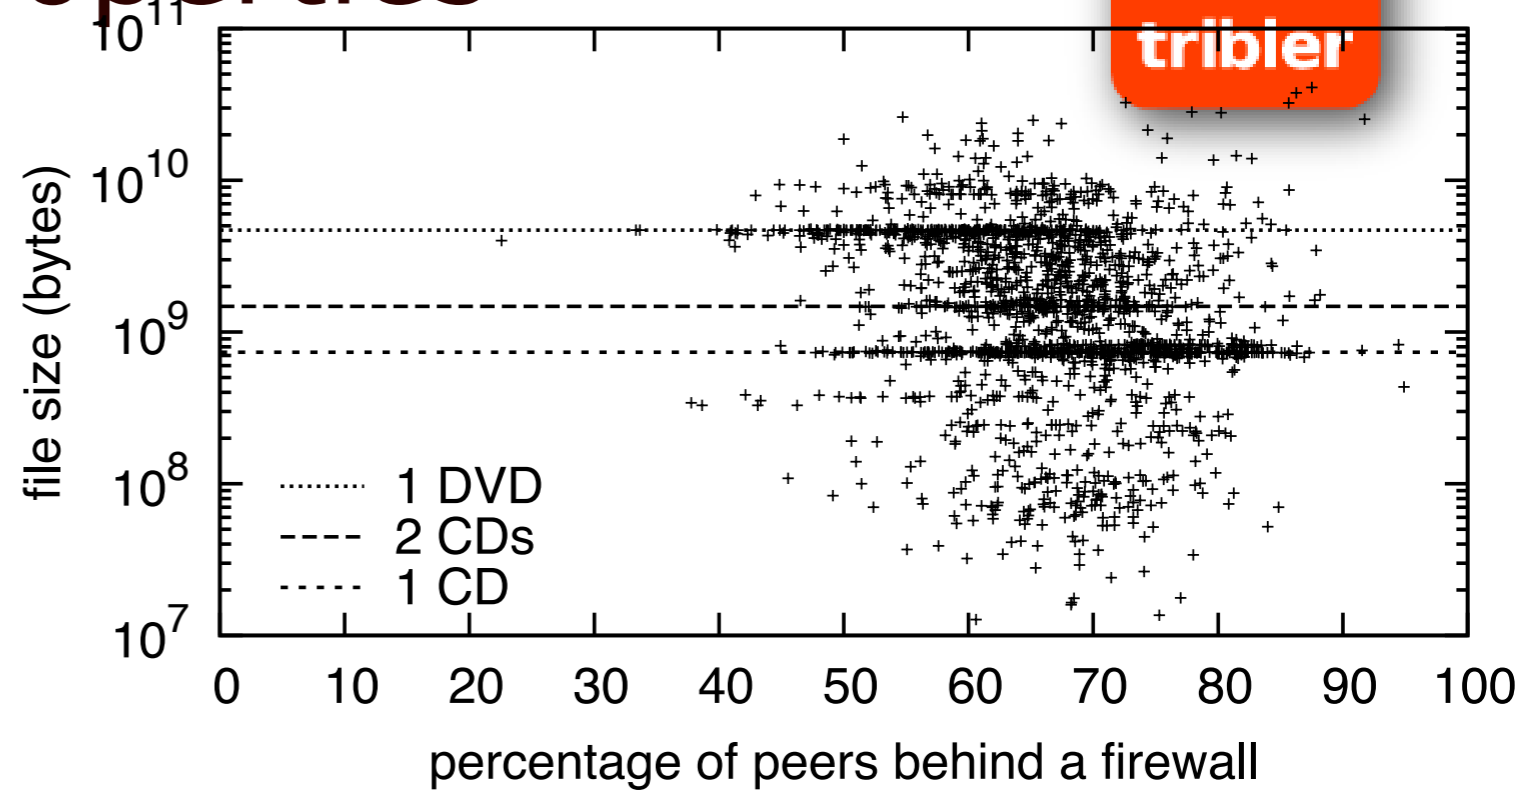


Closed: many seeders

- Trying to increase SR?
- Frustrated by firewalls



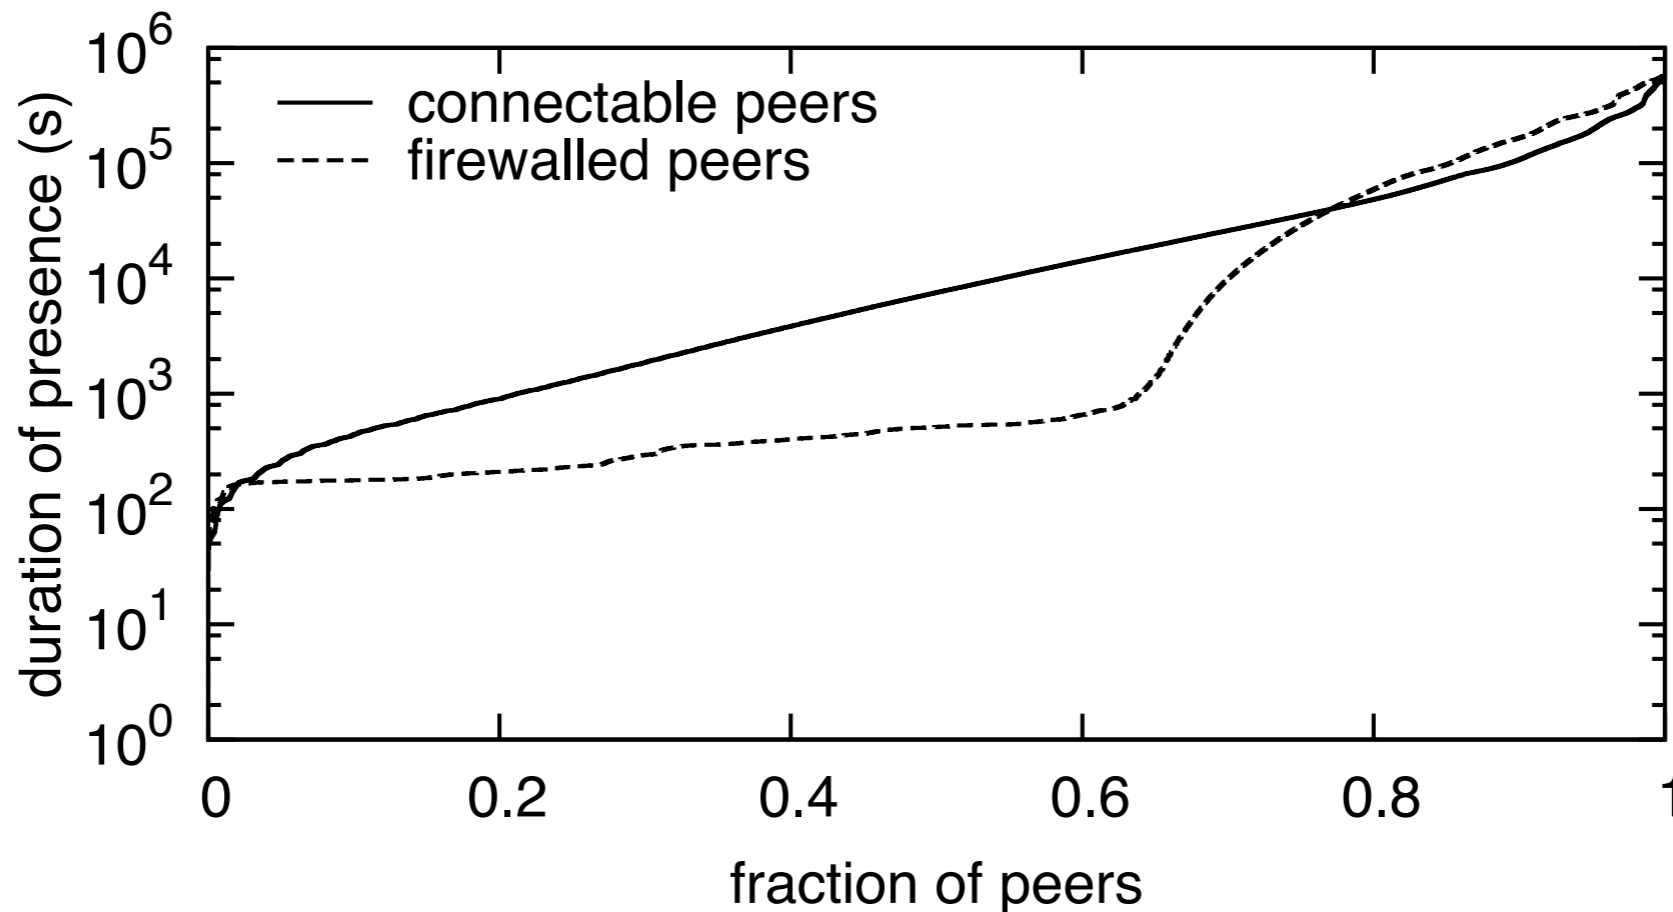
Firewalled peer properties



PirateBay measurement

- 1,995 swarms
- tracked 1 week

Session Duration



- Firewalled peers higher churn, often stay longer
 - Connectivity problems
 - Take longer to download
 - May have less bandwidth (ADSL)

Questions?



Jan David Mol

Delft University of Technology, NL

jjdmol@gmail.com