

BitTorrent Transport: Mechanisms and Impact

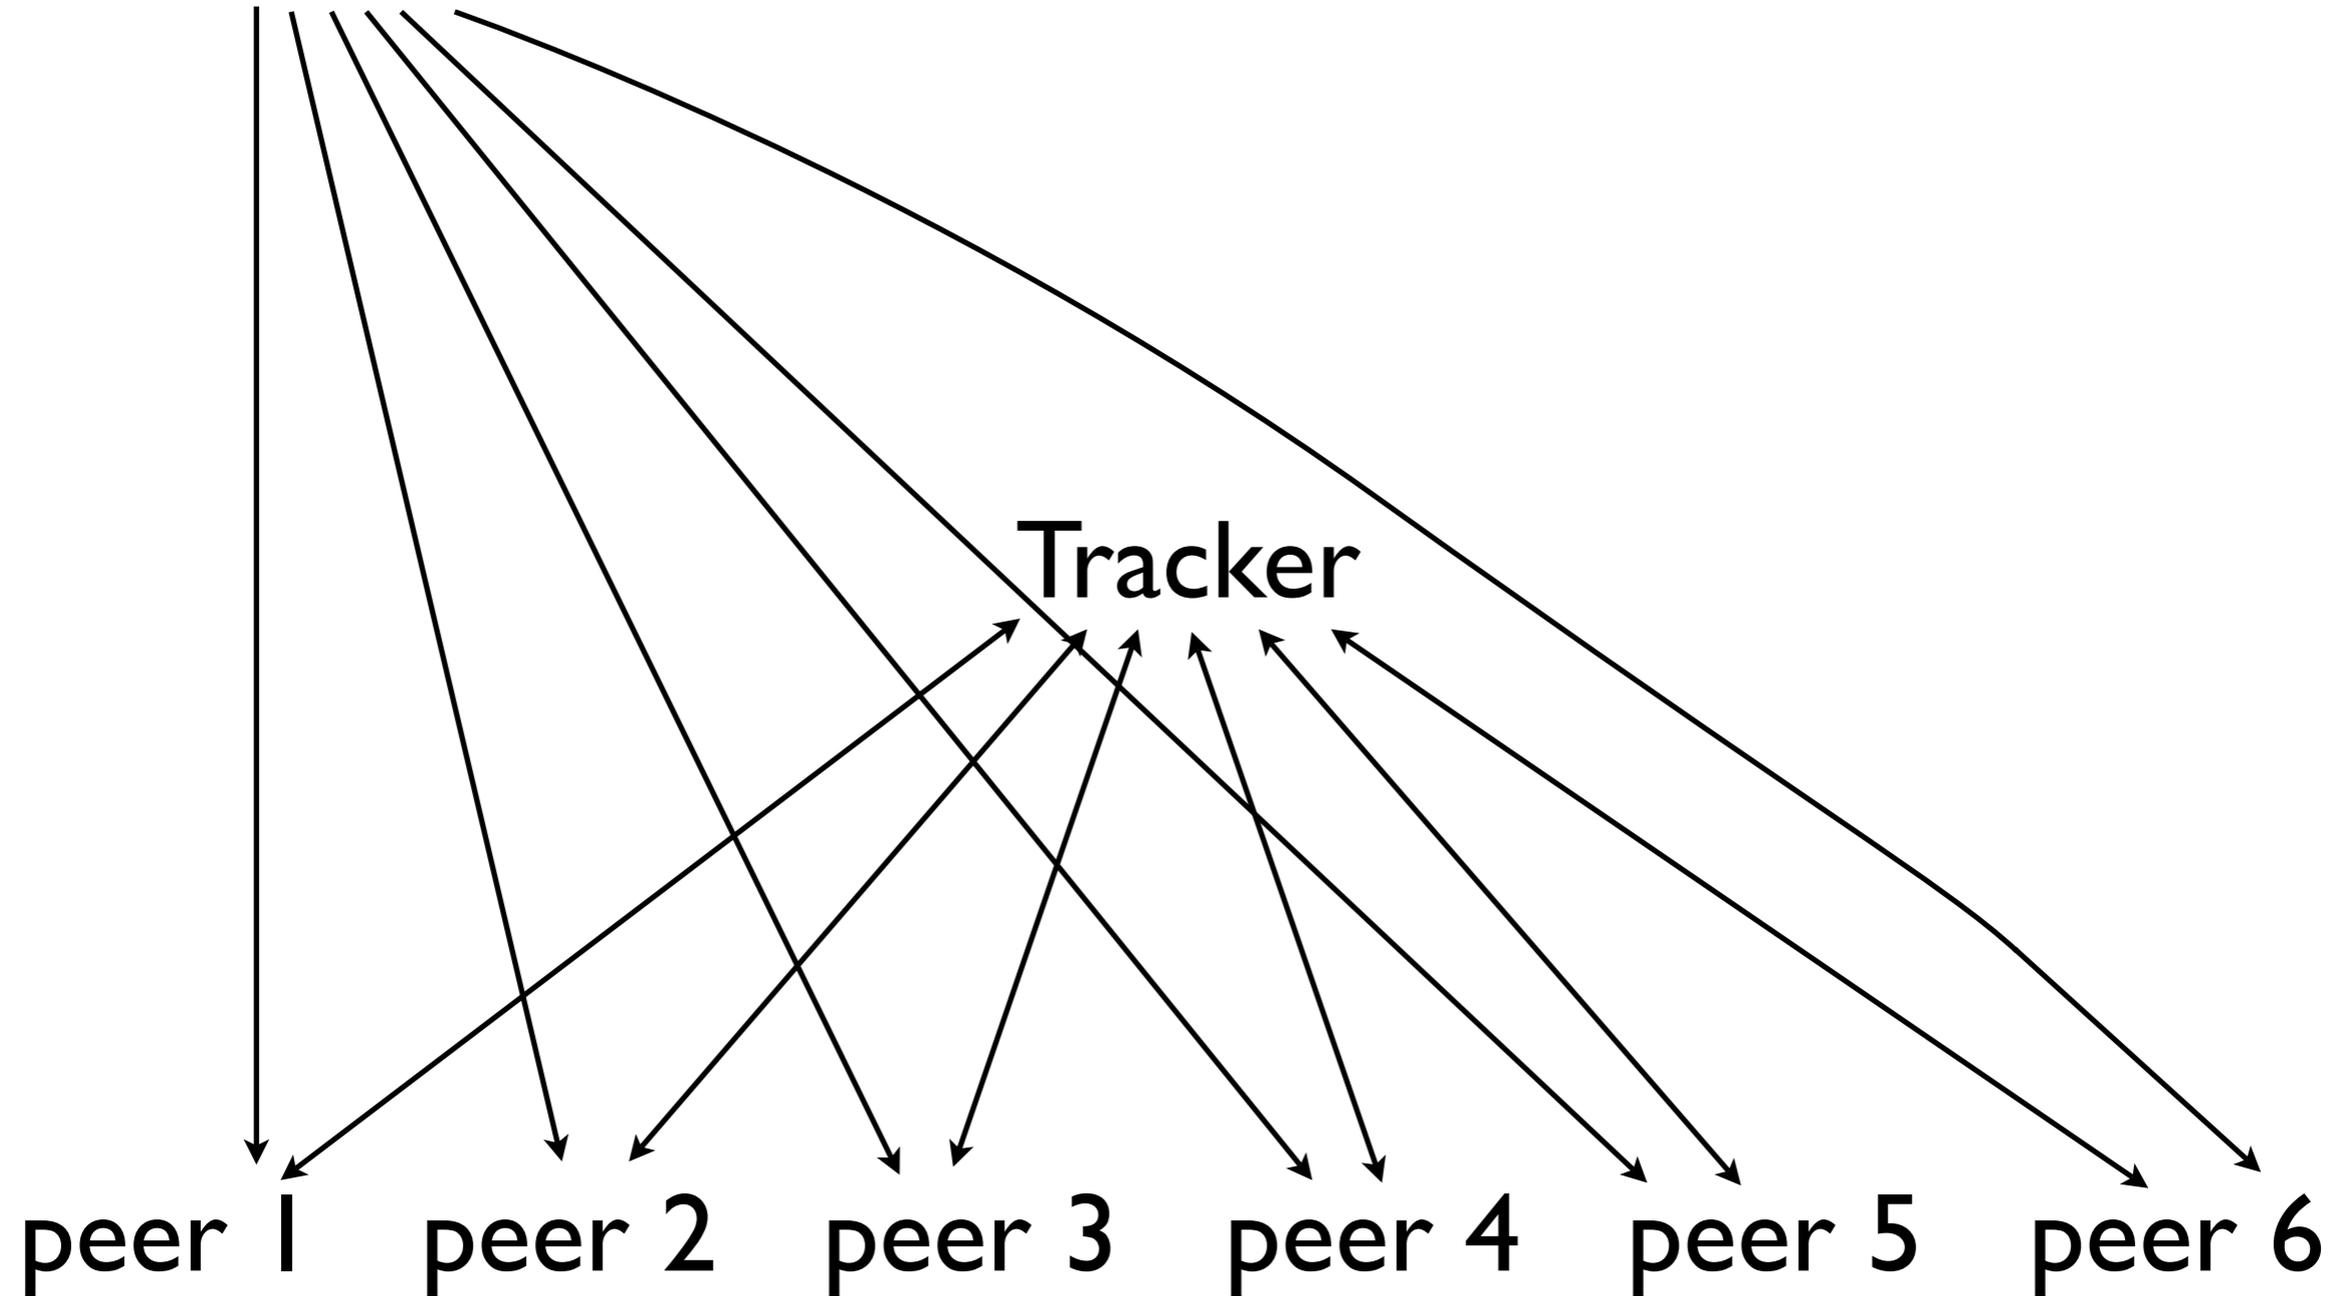
Stanislav Shalunov <shalunov@bittorrent.com>
IETF 72, Transport Area Meeting, Dublin, July 29, 2008

BitTorrent: purpose and brief overview

- Deliver the same file to many places
- Use upload capacity of the peers
- Upload while downloading and beyond
- Reduce load on the origin

General architecture

torrent file source



Pieces and chunks

- Files are broken up into pieces
- SHA1 of each piece is in the torrent file
- Pieces are broken up into chunks
- Chunks are requested and sent
- Pieces are verified and reported

Peers: connected graph

- Tracker gives each peer ~50 peers
- New peer quickly learns of more through PEX and DHT
- Usually end up knowing all peers within a minute
- Maintain n connections ($n = 30$ or 50)

Most P2P connections are nearly idle

- ~90% of connections only send metadata
- Space-efficient protocol lets peers know who has what
- More connections send metadata than data so rarest-first works better

Choke/unchoke

- Choking is the mechanism used to ensure only metadata is sent
- “You’re choked”
- Send to interested peers who unchoked me
- Get from peers I unchoked
- The sets are driven to be the same because of tit for tat

How many to unchoke?

- if 1, pairs: really bad
- if 2, chains: nearly as bad
- if 3, a graph that begins to be useful
- if 4, resilience is added, diameter reduced
- +1 for optimistic unchoke (search for new)

Tit for tat

- Purpose: reduce leeching
- Only give to those from whom you receive
- If below threshold, then forget it
- Otherwise, sort by rate, unchoke top 4
- Does not apply to optimistically unchoked peer

Rarest first

- First, download rarest pieces
- Reduces probability that some pieces disappear from swarm
- Increases probability that I'll have something interesting for others so they'll give me something
- Local knowledge (only limited to 50 peers)

Rate limiting

- Off by default
- Users can choose a max upload rate
- Almost always used by advanced users
- Without it, your connection is TCPed
- Form of congestion control: user enters a rate based on pain (delay)
- Incredibly crude, but still better experience than with straight TCP

Incentives to use few unchoked connections

- Upload rate is scarce
- Min useful threshold puts upper bound
- Rate limiting works better with few
- Get whole pieces faster

Multiple torrents

- RSS feeds
- Concurrent or serial downloads
- Multiple torrents are treated entirely separately
- Concurrent downloads mean more unchoked connections
- Which is why it's often faster to download serially

P2P apps: congestion as a norm

- Most Internet traffic does not experience congestion
- When it does, congestion is mostly transient
- No congestion means no congestion control
- P2P experiences sustained congestion

Effect of TCP uploads

- TCP fills the buffer
- Buffer in DSL or cable modem can be in seconds
- Interactive apps don't work
- Does not significantly depend on number of connections

TANA BoF

- Transport for Advanced Networking Apps
- End-to-end scavenger congestion control
- Discussion of multiple connections
- **Thursday, 1-3PM**, Convention 1