Introduction to P2P Networks

B. Sc. Florian Adamsky

`florian.adamsky@iem.thm.de`
`http://florian.adamsky.it/`

Internet Protocols and Applications – SS 2013

Outline

1 Introduction
- What is P2P?

2 The first P2P networks
- Napster
- Gnutella

3 BitTorrent
- Introduction
- Choking Algorithm
- Rarest-Piece First Algorithm
- Visualization

Contents

1 Introduction
- What is P2P?

2 The first P2P networks
- Napster
- Gnutella

3 BitTorrent
- Introduction
- Choking Algorithm
- Rarest-Piece First Algorithm
- Visualization

What is an ”Overlay” Network?

- Overlay Network is a network which is build on top of another network
- Mostly over the Internet

What is an ”Overlay” Network?

- Overlay Network is a network which is build on top of another network
-

Yo Dawg!

YO DAWG I HEARD YOU LIKE NETWORKS?
SO WE PUT A NETWORK IN A NETWORK SO YOU CAN COMMUNICATE WHILE YOU COMMUNICATE
memegenerator.net

What does ”Peer” mean?

The term “peer” is defined by Dictionary.com in the following way:

> ***peer*** *(pɪə) 3. a. a person who is an equal in social standing, rank, age, etc (from Old French per, from Latin pār equal)*

What is a P2P network?

Definition (P2P-Network)

Peer-to-Peer (P2P) network is an *overlay-network* where all peers are on a par, without a central entity.

P2P is not dead:

Skype	P2P VoIP application
TOR	The Onion Routing – Anonymity Network
BitCoin	P2P based digital currency
NameCoin	P2P DNS system
BitTorrent	P2P file sharing protocol
BT Live	P2P live video streaming
BT Sync	Dropbox-like synchtool via BitTorrent

What P2P Networks are good for?

- Different resources can be shared within a P2P network:
  - Processing power
    - SETI@home (distributed)
  - Disk storage
    - Tahoe-LAFS, BT Sync
  - Network bandwidth
    - BitTorrent, Gnutella

Difference between P2P and client-Server

Client-Server Paradigma

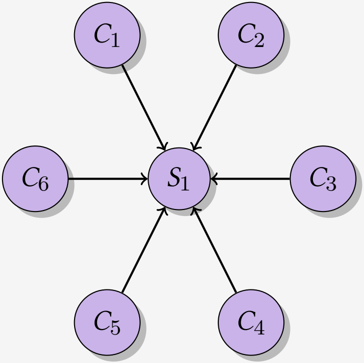

Difference between P2P and client-Server

Client-Server Paradigma

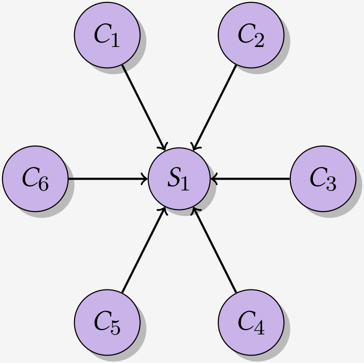

Peer-to-Peer Paradigma

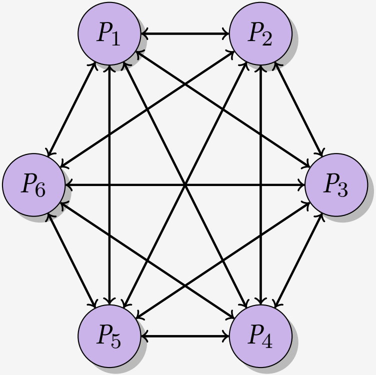

Why is P2P more efficient?

Definition (Flash crowd)

Simultaneous download start of n peers

- Mean download time of a flash crowd scenario of n peers where each peer has the same bandwidth capacity:
  - P2P: $\mathcal{O}(\log(n))$
  - Client-Server: $\mathcal{O}(n)$
- Discussion:
  - Simple model
  - No piece selection algorithm
  - No peer selection algorithm
  - No overhead

Contents

1 Introduction
- What is P2P?

2 The first P2P networks
- Napster
- Gnutella

3 BitTorrent
- Introduction
- Choking Algorithm
- Rarest-Piece First Algorithm
- Visualization

History of Napster

1999

- is known as the first P2P application
- was 1999 download of the year
- is mostly known for its lawsuit against music industry
- however, architecture is rather based on client-server
- only the direct download is P2P

Functionality of Napster

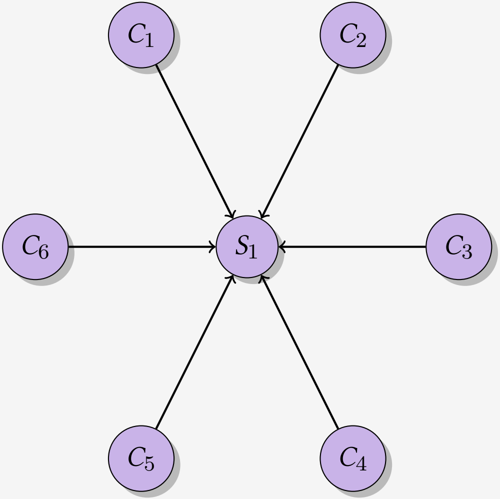

Functionality of Napster

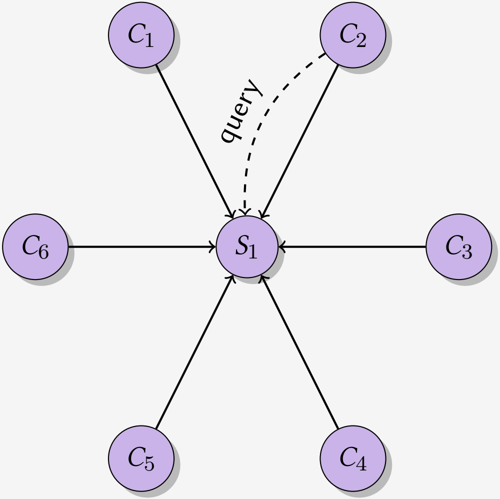

Functionality of Napster

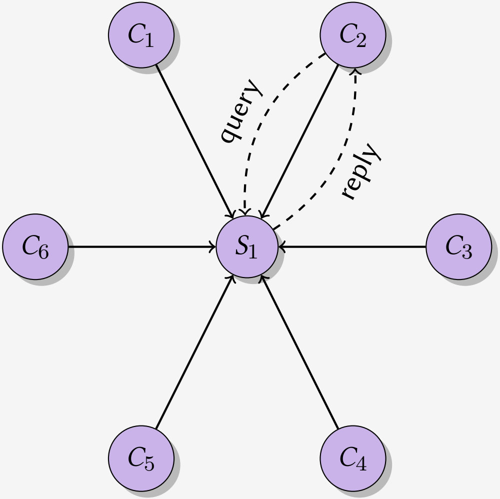

Functionality of Napster

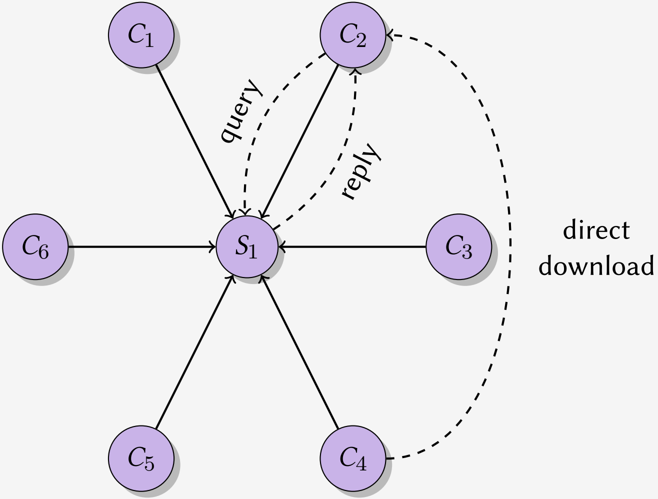

Discussion about Napster?

Advantages

- Revolutionary idea

Disadvantages

- No real P2P and Single Point of Failure (SPOF)

History of Gnutella

March 2000

- Gnutella is a simple P2P protocol
- in contrast to Napster it is a real P2P network without any central point
- developed from *Nullsoft*, a subsidiary of AOL
- after the process of reverse engineering public domain

How does Gnutella work?

- first problem: How to join the network? (Bootstrapping problem)
  - Gnutella software provides a list of peers
  - first active peer returns the requesting peers its k neighbours

How does Gnutella work?

- first problem: How to join the network? (Bootstrapping problem)
  - Gnutella software provides a list of peers
  - first active peer returns the requesting peers its k neighbours

- Gnutella provides 5 message types:

  Ping announce the presence of a node and discover new hosts in the network

  Pong reply message for *ping*

  Query search for a specific file

  QueryHit reply message for a *query* message. It contains IP, port number, transmission speed and a description of the file

  Push is a mechanism to bypass a firewall

How does Gnutella work?

- first problem: How to join the network? (Bootstrapping problem)
  - Gnutella software provides a list of peers
  - first active peer returns the requesting peers its k neighbours

- Gnutella provides 5 message types:

  Ping announce the presence of a node and discover new hosts in the network

  Pong reply message for *ping*

  Query search for a specific file

  QueryHit reply message for a *query* message. It contains IP, port number, transmission speed and a description of the file

  Push is a mechanism to bypass a firewall

- Every message will be forwarded to all of its k neighbours
  - TTL = Time To Live (default setting is 7)

Search Mechanism

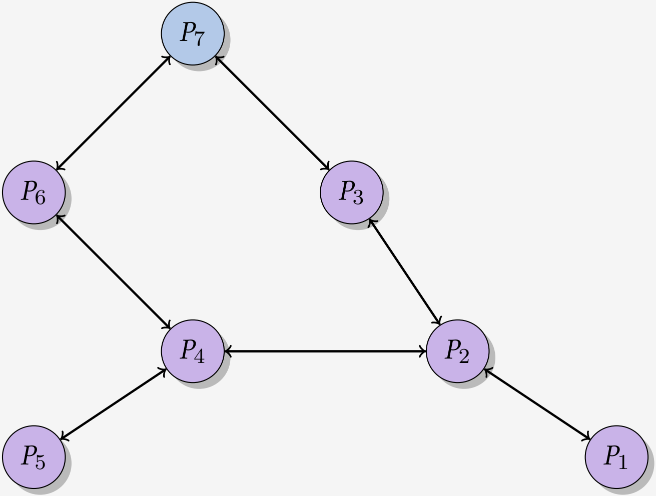

Search Mechanism

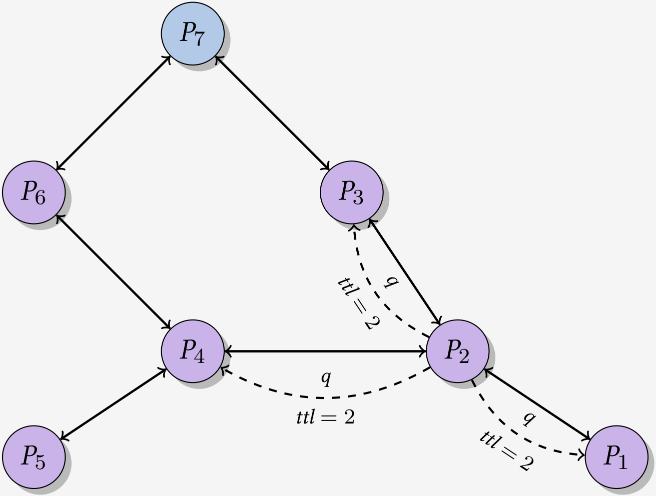

Search Mechanism

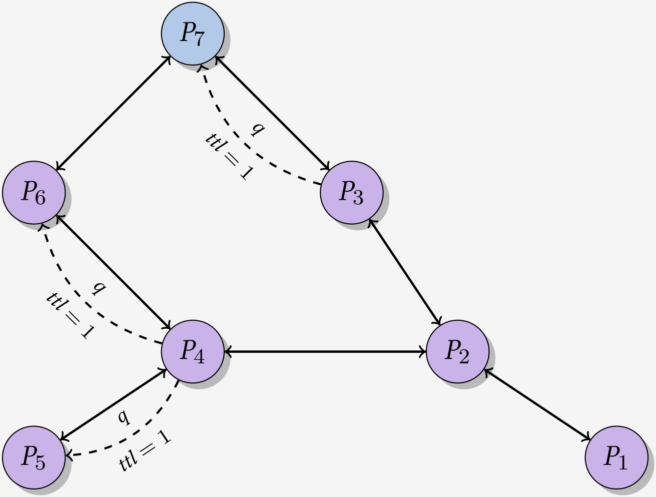

Search Mechanism

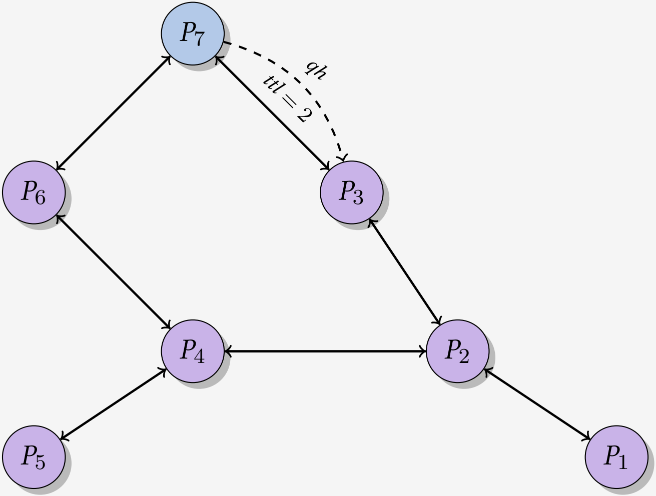

Search Mechanism

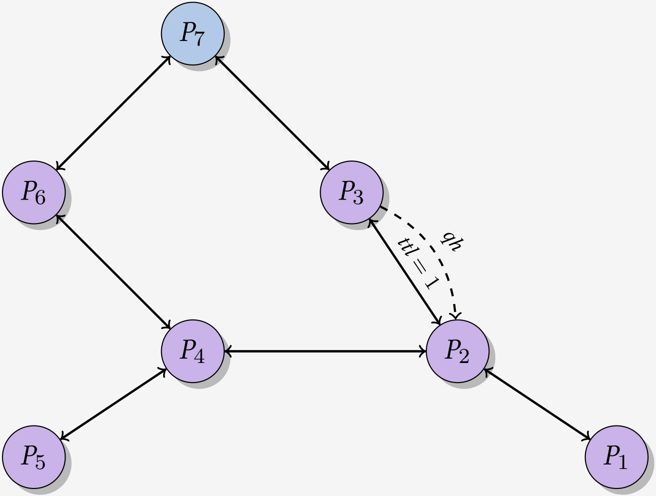

Search Mechanism

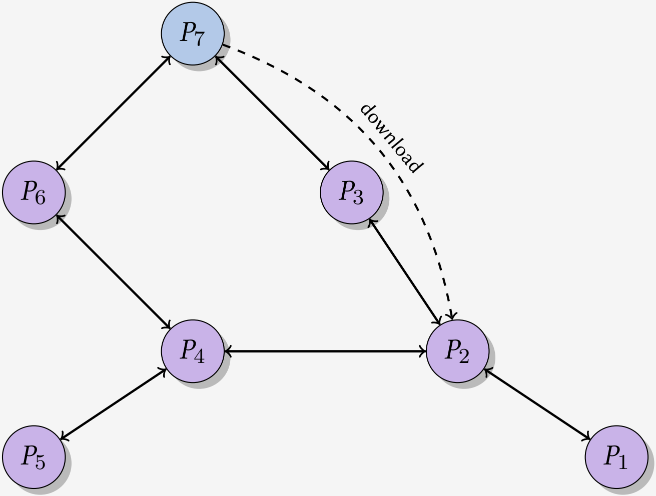

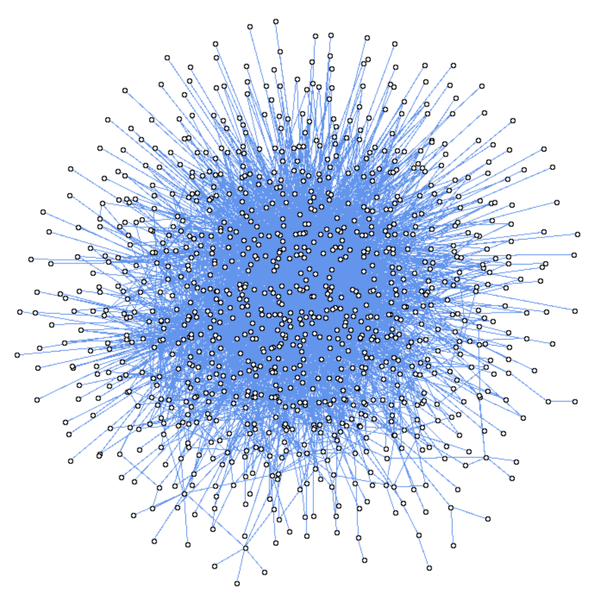

Figure : Snapshot of the Gnutella network [1] (2001)

Discussion about Gnutella?

Discussion about Gnutella?

Advantages

- in contrast to Napster, a real P2P network
- quite robust and hard to attack
- network structure scales very well

Discussion about Gnutella?

Advantages

- in contrast to Napster, a real P2P network
- quite robust and hard to attack
- network structure scales very well

Disadvantages

- rare files are hard to find because of partial search
- message flooding
- no file splitting in the original specification
- free riding problem
    - study [2] showed that 70 % of peers are not sharing any files
    - and 25 % provided 99 % of all query hits

Contents

1 Introduction
- What is P2P?

2 The first P2P networks
- Napster
- Gnutella

3 BitTorrent
- Introduction
- Choking Algorithm
- Rarest-Piece First Algorithm
- Visualization

BitTorrent Protocol

- Invented by Bram Cohen in April 2001
- Is the most used P2P protocol
  - 2011: Greater share of network bandwidth then Netflix and Hulu
- comprises the largest share of P2P traffic on the Internet
- Is heavily under attack from anti-P2P companies (MediaDefender, . . .)

Terminology

Leecher is a peer which hasn't the complete file but shares the pieces which it has

Seeder is a peer which has the complete file and justs uploads it to other peers

Statistics 2011

Top applications in Europe during peak hours (May 2011)

	Upstream		Downstream	
Rank	**Application**	**Share**	**Application**	**Share**
1	BitTorrent	59.68 %	BitTorrent	21.63 %
2	Skype	7.16 %	HTTP	20.47 %
3	HTTP	7.02 %	YouTube	14.13 %
4	PPStream	3.64 %	RTMP	4.58 %
5	Spotify	2.91 %	Flash Video	3.99 %

Table : Source: Torrentfreak.com, URL: `http://goo.gl/UIRmS` (2011)

Statistics 2012

Top applications in Europe during peak hours (May 2012)

	Upstream		Downstream	
Rank	**Application**	**Share**	**Application**	**Share**
1	BitTorrent	31.80 %	HTTP	26.30 %
2	HTTP	11.40 %	YouTube	22.30 %
3	eDonkey	11.20 %	BitTorrent	12.10 %
4	YouTube	6.66 %	Flash Video	3.95 %
5	Skype	6.00 %	Facebook	3.71 %

Table : Source: Torrentfreak.com, URL: `http://goo.gl/BbZts` (2012)

Fields of application

- Facebook and Twitter are using BitTorrent to update their servers
- Twitter published Murder

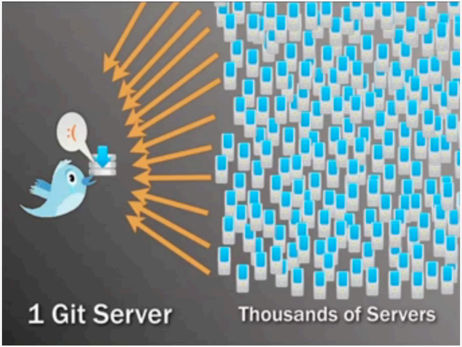

Fields of application

- Facebook and Twitter are using BitTorrent to update their servers
- Twitter published Murder

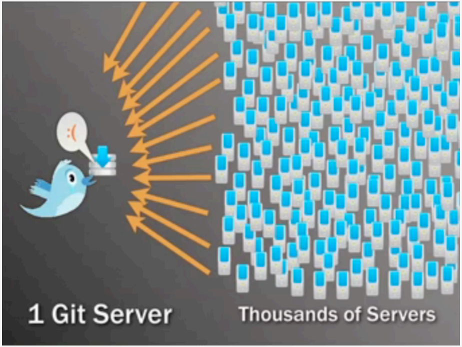

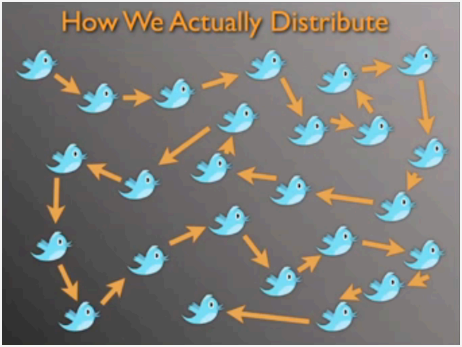

Fields of application

Difference

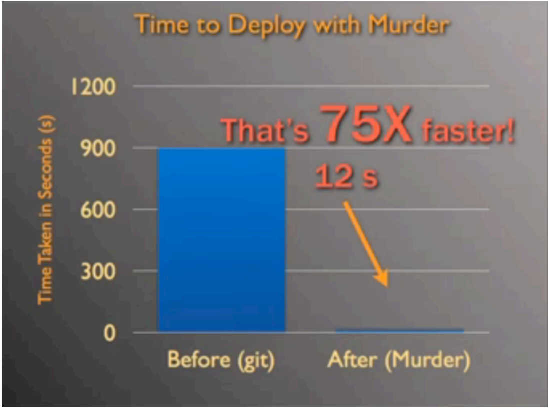

File Splitting

- Split up a file in pieces with the same size
- Size depends on the size of the file, but normally 256–1024 KiB
- Calculate a SHA-1 for each piece

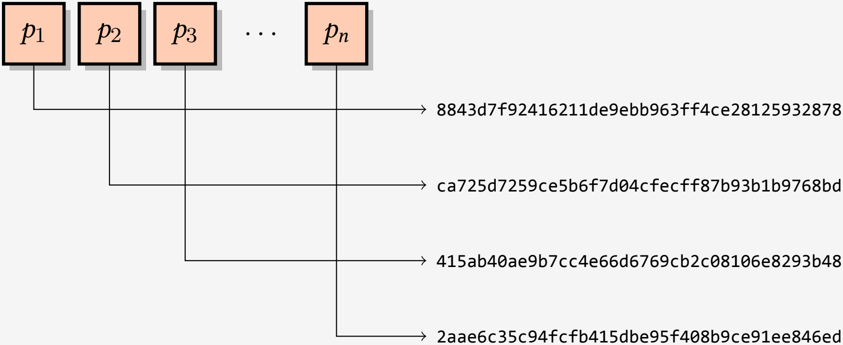

```
8843d7f92416211de9ebb963ff4ce28125932878
ca725d7259ce5b6f7d04cfecff87b93b1b9768bd
415ab40ae9b7cc4e66d6769cb2c08106e8293b48
2aae6c35c94fcfb415dbe95f408b9ce91ee846ed
```

Metadata file: torrent

- Metadata file (`*.torrent`) is distributed to all peers
  - bencoded
  - usually distributed via HTTP
- the `*.torrent` fiel contains
  - SHA-1 hashes of all pieces
  - Mapping from pieces to files
  - list of trackers

Tracker

- Central server which keeps a list of all peers in the swarm
- Tracker-less peer discovery:
    - Distributed Hash Table (DHT)
    - Peer Exchange (PEX)
    - Local Peer Discovery (LDP)
- A peer joins a swarm by asking the tracker for a peer list

Unique Features of BitTorrent

- Separated the download process and the search process
- Rarest first and choke algorithm
- File splitting
- Pipelining
  - keep enough requests pending

Choking Algorithm

- Chocking is a notification that no data will be sent until unchoking happens
- Connections contains two bits of state on either end:
  - choked/unchoked
  - interested/uninterested
- Data will be sent whenever one side is interested and the other is unchoked
- Choking is done for several reasons:
  - Each peer use a tit-for-tat-ish algorithm to ensure a consistent download rate
  - Prevents *free riders*
  - TCP congestion control behaves very poorly when sending over many connections at once

Choking Algorithm: Leecher/Seeder

- Decision which peer will be choked is different for a leecher and a seeder

Choking Algorithm: Leecher/Seeder

- Decision which peer will be choked is different for a leecher and a seeder

Leecher State

- Every 10 sec peers will be ordered according to their upload rate
- 3 fastest and interested peers will be unchoked
- Periodically (30 seconds) select a peer at random and upload to it (Optimistic Unchoking)

Choking Algorithm: Leecher/Seeder

- Decision which peer will be choked is different for a leecher and a seeder

Leecher State

- Every 10 sec peers will be ordered according to their upload rate
- 3 fastest and interested peers will be unchoked
- Periodically (30 seconds) select a peer at random and upload to it (Optimistic Unchoking)

Seeder State

- best downloader will be unchoked

Choking Algorithm: Visualization

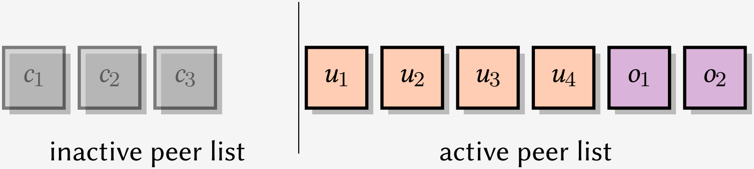

Choking Algorithm: Visualization

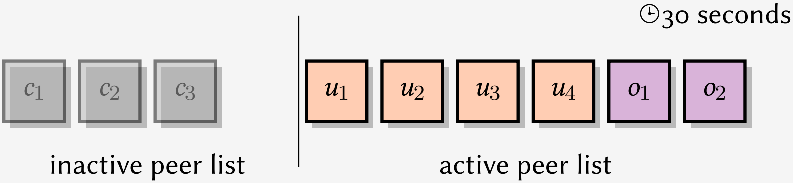

Choking Algorithm: Visualization

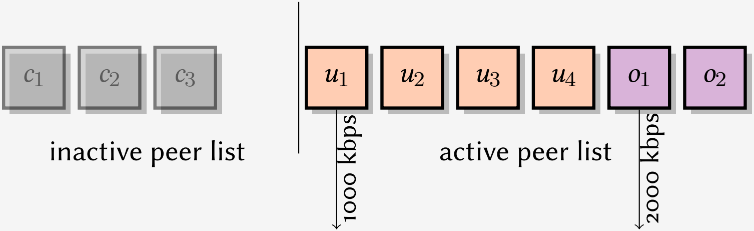

Choking Algorithm: Visualization

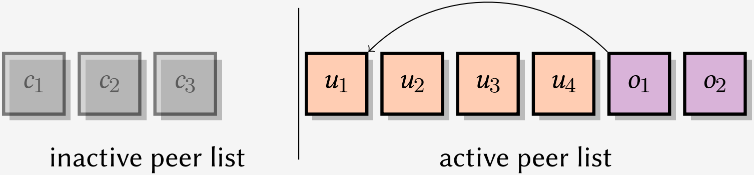

Choking Algorithm: Visualization

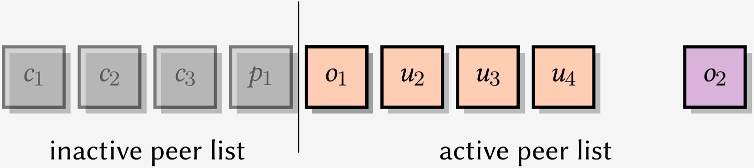

Choking Algorithm: Visualization

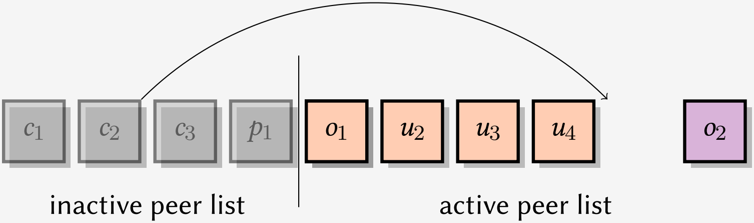

Choking Algorithm: Visualization

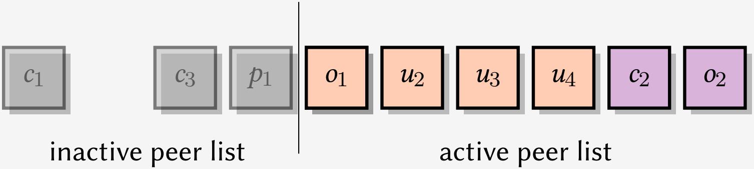

Piece Selection Algorithm

- Piece Selection Algorithm is crucial to the performance of a P2P protocol
- BitTorrent uses a combination of 4 policies: Rarest First Algorithm [3]
  - Strict priority policy
  - Local rarest first policy
  - Random first policy
  - End game mode
- Each peer maintains a list of pieces from its peer set
  - Every peer sends a HAVE message to its peer set when finishing a piece

Strict Priority Policy

- A shared file is divided into pieces and pieces are divided into *sub-pieces*
- Once a single sub-piece has been requested, the remaining sub-pieces are requested first
- This helps to get complete pieces as soon as possible

Rarest First Policy

- peers download pieces first, which the fewest of their own neighbors have first
- this strategies has a lot of advantages:
  - peers have pieces which other peers wants to acquire
  - it reduces the probability that a peer sits on rare pieces and goes offline
  - performance is much better if different peers have different pieces

Rarest First Policy (Example)

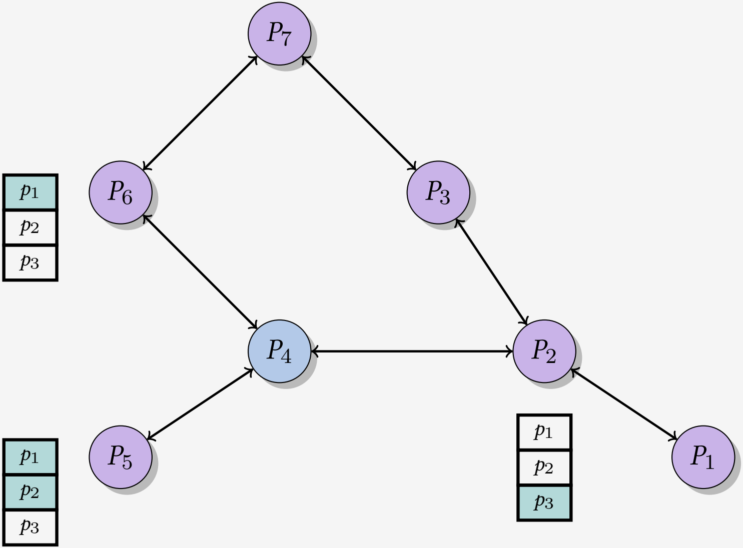

Rarest First Policy (Example)

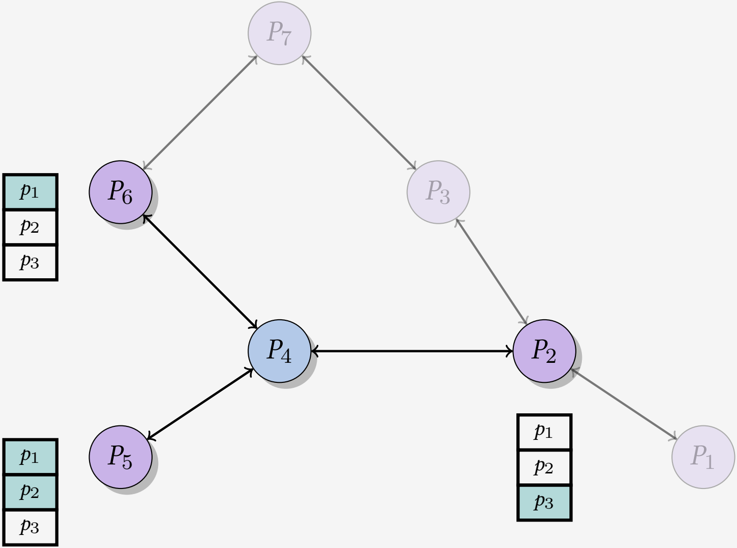

Rarest First Policy (Example)

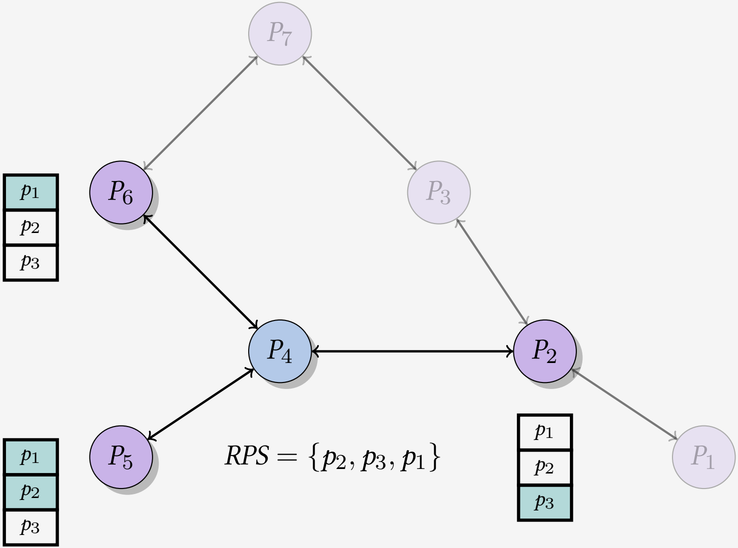

$$RPS = \{p_2, p_3, p_1\}$$

Random First Policy

- if a peer has downloaded less than 4 pieces it uses the *Random First Policy*
- the peer chooses a piece randomly
- the aim is to get complete pieces as soon as possible to engage in tit-for-tat
- after finishing at least 4 pieces the peer switches to the rarest first policy

End Game Mode

- this mode starts at the very end of a download
- more precisely, this mode starts once a peer has sent requests for pieces
- in this mode a peer sends the requests to all peers in the active peer set
- each time the peer got a piece from a peer it will cancel the request for the receiving piece
- this ensures a quick download at the end of a file

Animation

- With Internet: `http://mg8.org/processing/bt.html`
- Without Internet: BitTorrent.ogv

Discussion about BitTorrent?

Discussion about BitTorrent?

Advantages

- Focus only on file sharing and nothing more
- Very efficient

Discussion about BitTorrent?

Advantages

- Focus only on file sharing and nothing more
- Very efficient

Disadvantages

- No incentive to be a seeder

Thank you

Questions?

florian.adamsky@iem.thm.de

http://florian.adamsky.it/

Bibliography

[Jovanovic, 2001] Jovanovic, M. A., Annexstein, F. S., Berman, K. A.
Scalability Issues in Large Peer-to-Peer Networks — a Case Study of Gnutella
Technical Report, University of Cincinnati, 2001.

[Adar, 2000] Adar, Eytan and Huberman, Bernardo A.
Free Riding on Gnutella
Journal Paper, First Monday, 2000.

[Cohen, 2003] Cohen, Bram
Incentives Build Robustness in BitTorrent
Conference Proceedings, Workshop on Economics of Peer-to-Peer systems, 2003.