BitTorrent

Kevin Kaminski
CISC856
11/27/2012

P2P In General

- Various protocols(Gnutella, eDonkey, Fastrack)
  - BitTorrent most popular
    - Next is Gnutella, though declining due to LimeWire takedown and FrostWire drop
- Various clients
  - uTorrent seems to be best/most popular
  - Vuze, eMule, Frostwire other alternatives

Worldwide Traffic Share

Asia-Pacific

Aggregate	
Application	Share
BitTorrent	27.19%
YouTube	14.94%
HTTP	10.44%
PPStream	6.36%
Thunder	4.62%
Flash Video	3.36%
QVoD	3.44%
Facebook	2.08%
iTunes	1.86%
MPEG Streaming	1.61%
Top 10	75.08%

Europe

Aggregate	
Application	Share
BitTorrent	20.32%
HTTP	17.70%
YouTube	15.25%
eDonkey	9.39%
Flash Video	4.70%
RTMP	2.47%
Facebook	2.43%
SSL	1.74%
MPEG Streaming	1.66%
iTunes	1.53%
Top 10	77.19%

Latin America

Aggregate	
Application	Share
YouTube	26.61%
HTTP	19.98%
BitTorrent	9.24%
Facebook	6.97%
Flash Video	5.83%
Skype	3.93%
Ares	2.77%
MPEG Streaming	2.04%
Windows Update	1.85%
SSL	1.45%
Top 10	80.67%

- 2012 statistics courtesy of Sandvine

NA Traffic Statistics

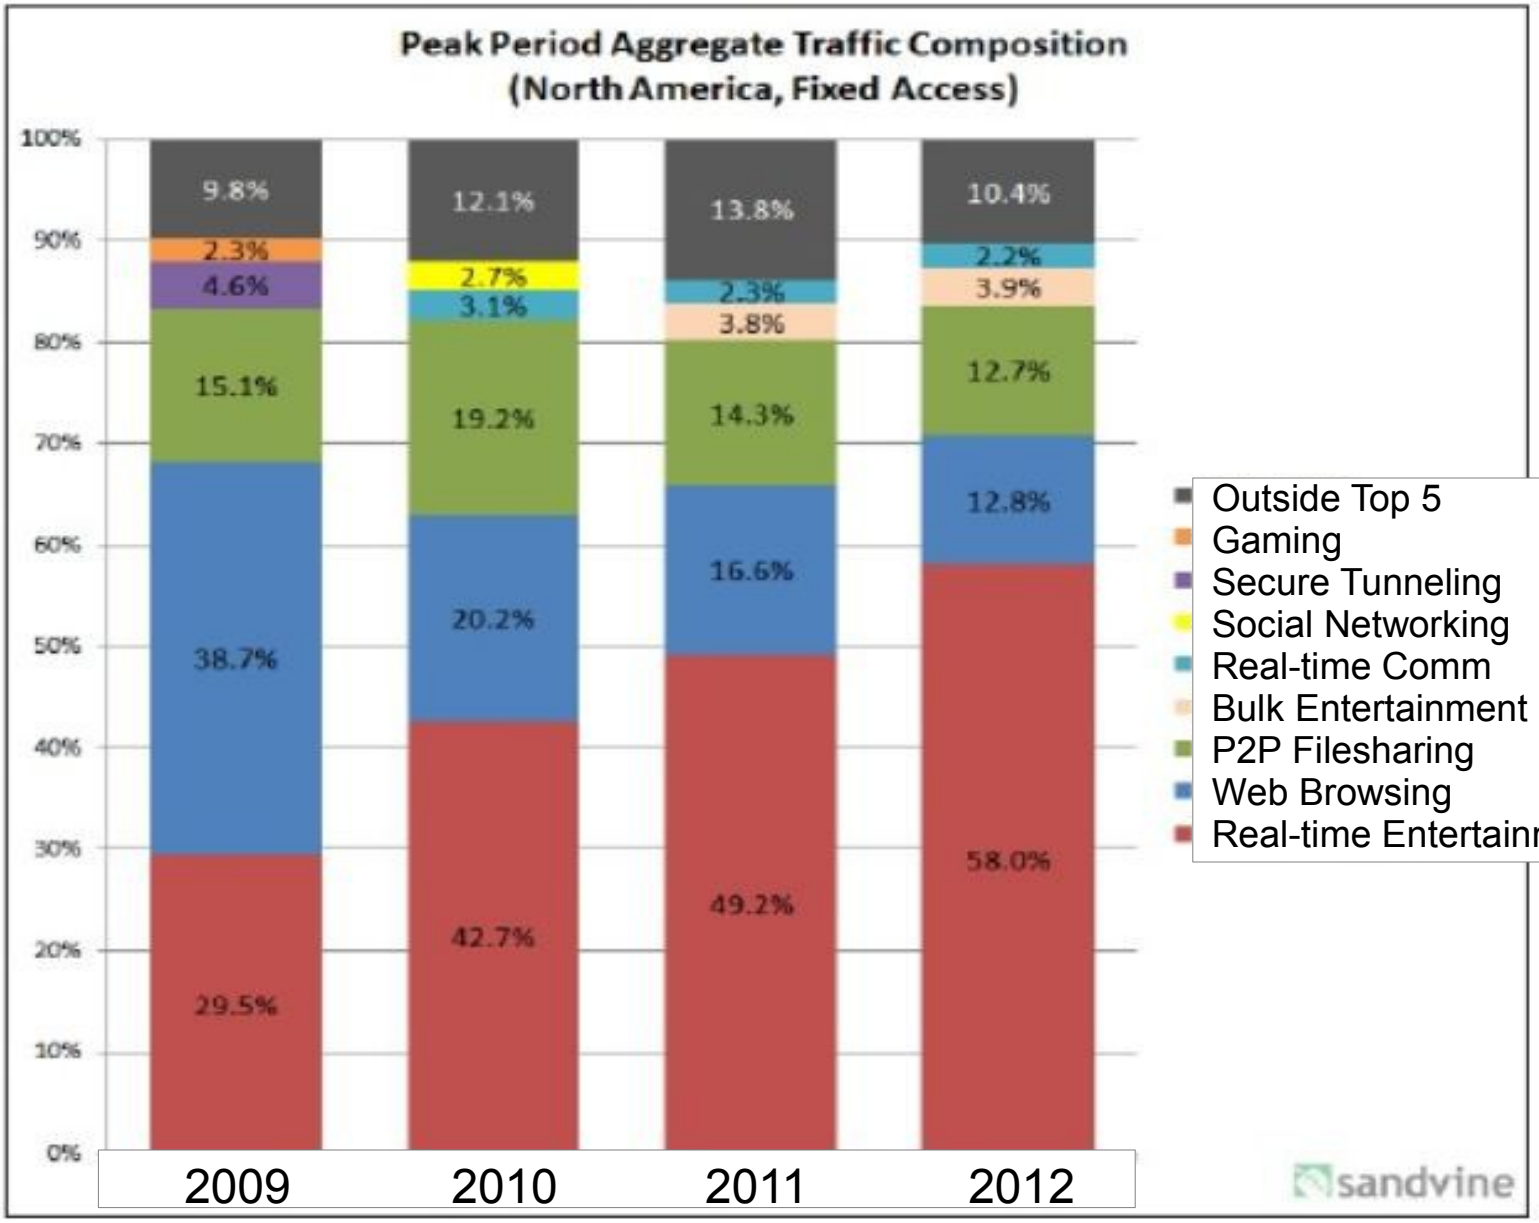

Server Client Model

- Centralized
- Fast speeds
- Stable

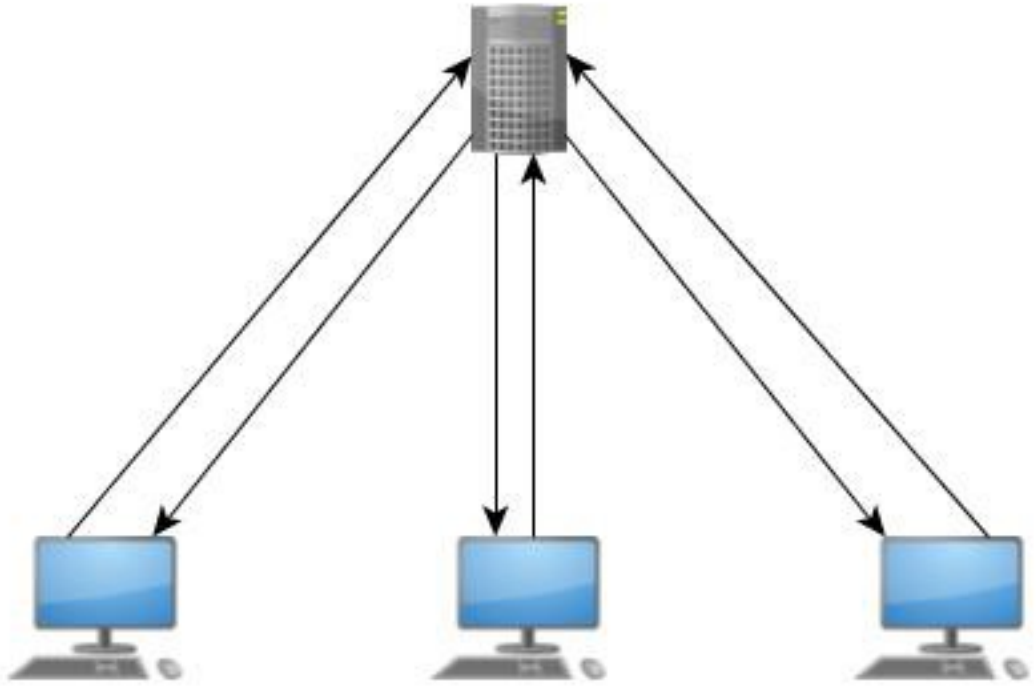

P2P Model

- Each node is a "peer"
- Peers both clients and servers
- Decentralized

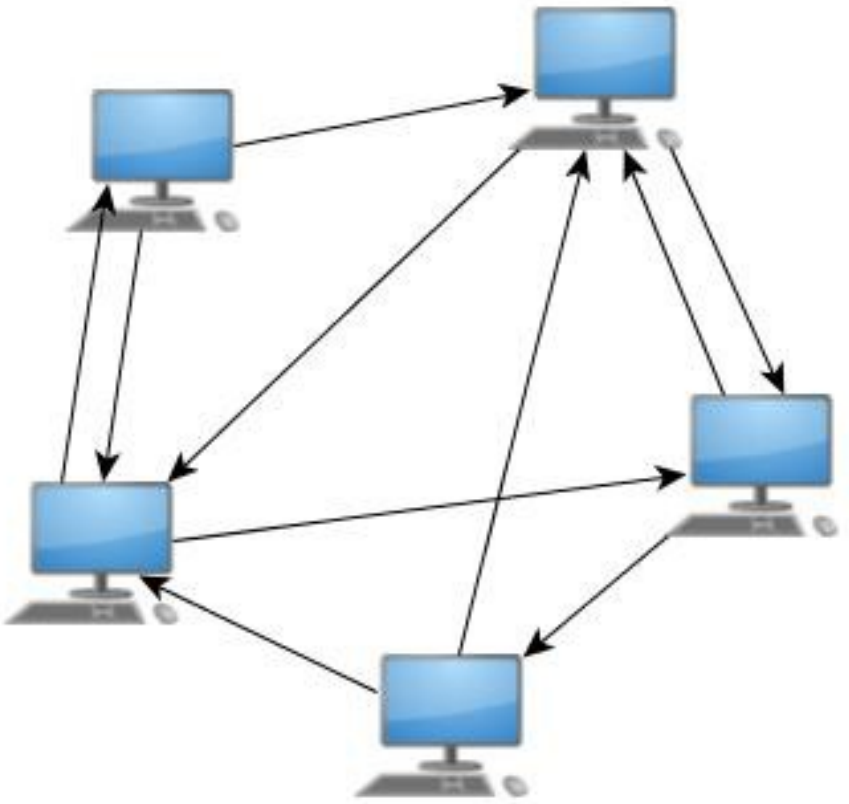

Why P2P?

- Advantages
  - Servers require expensive hardware, P2P can be used on common desktops
  - Redundancy - Resources located in multiple locations, no single point of failure
- Disadvantages
  - Slower speeds near endpoints
  - Peers constantly connecting/disconnecting, unstable
  - If no one sharing the resource, cannot be downloaded

BitTorrent Component Overview

- Target file
  - Broken into indexed “pieces”
- Metainfo file
  - Hosted on BitTorrent index, contains meta information about the torrent file
- Tracker
  - Centralized server that coordinates downloads
- Peer Wire Protocol
  - Used to transfer pieces between peers

General Process

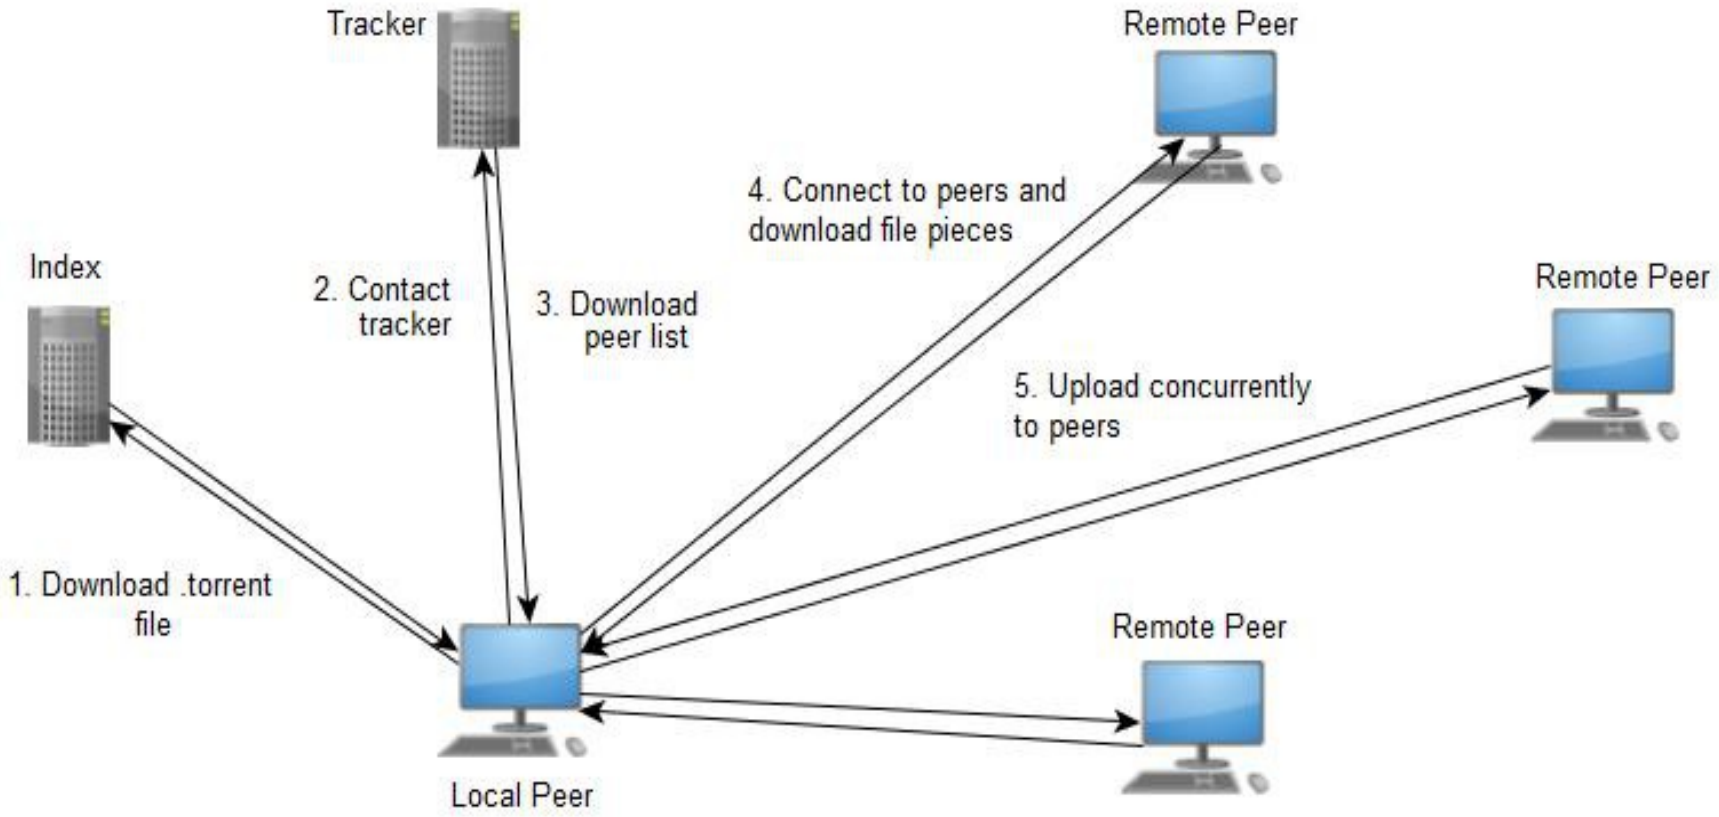

Metainfo File

- Hosted on web server (torrent index) and downloaded out-of-band
- Contains tracker and file information in “dictionary” of keys
- Dictionary maps one value to another
- File is encoded using “bencoding”

Metainfo File Contents


```
2f 74 6f 72 72 65 6e 74 2e 75 62 75 6e 74 75 2e 63 6f 6d 3a  d8:announce39:http://torrent.ubuntu.com:
75 6e 63 65 2d 6c 69 73 74 6c 6c 33 39 3a 68 74 74 70 3a 2f  6969/announce13:announce-listll39:http:/
36 39 36 39 2f 61 6e 6e 6f 75 6e 63 65 65 6c 34 34 3a 68 74  /torrent.ubuntu.com:6969/announceel44:ht
75 6e 74 75 2e 63 6f 6d 3a 36 39 36 39 2f 61 6e 6e 6f 75 6e  tp://ipv6.torrent.ubuntu.com:6969/announ
65 6e 74 20 64 6f 77 6e 6c 6f 61 64 65 64 20 66 72 6f 6d 20  ceee7:comment33:Torrent downloaded from
62 79 38 3a 61 74 69 6b 30 37 38 36 31 33 3a 63 72 65 61 74  1337x.org10:created by8:atik078613:creat
34 3a 69 6e 66 6f 64 36 3a 6c 65 6e 67 74 68 69 37 33 35 33  ion datei1335433885e4:infod6:lengthi7353
75 2d 31 32 2e 30 34 2d 64 65 73 6b 74 6f 70 2d 69 33 38 36  58976e4:name29:ubuntu-12.04-desktop-i386
35 32 34 32 38 38 65 36 3a 70 69 65 63 65 73 32 38 30 36 30  .iso12:piece lengthi524288e6:pieces28060
71 53 f6 7d c3 db 47 60 73 ca f3 e5 0d c8 35 69 b2 39 0b 6f  :qwFuAáÅ¶Ç.'àÐ..ßÇ1qSö}ÅÛG`sÊóå.È5i²9.o
9d 0a 34 a0 11 49 f1 ac bd 60 2a e8 8a 7a 53 e0 4a 12 5f 92  ..ãKJx\z¯¦ èl u8..)$..4 .Iñ¬½`*è.zSàJ._'
c9 78 e4 75 6d 47 72 b8 e3 5e 1a 26 e7 02 38 65 62 d2 89 a4  ð.Mò.¯RÀ®Ü...å.öè.'.ÉxämGr¸ã^.&ç.8ebÒ.¤
d8 fe 59 27 9f d6 a8 36 d5 c6 eb f1 ca 80 ea 34 3c 38 bf 75  .M.ÁøáÎ..Q..Àú/½µÇ}.ØþY'.Ö¨6ÕÆëñÊ.ê4<8¿u
81 ec 81 80 08 d1 ee 98 a5 b9 50 1a 04 42 1c 4a 5c 7b 1f f7  [á.WÞml̦ç.ñ*Dw.=cä+&".ì...Ñî.¥¹P..B.J\{.÷
7a 53 a2 a7 f9 3b bd 5a 08 7d 4c 9f b7 83 d5 17 47 a8 5c 39  9Ú4{AÆº.ç.µ..ÜrÔÿ.Ó.zS¢§ù;½Z.}L..Õ.G¨\9
ee dd 30 49 91 b2 41 18 77 44 73 e4 dc 54 50 e8 2f d7 5e 70  ?Ð(¶-rÊU.È}.m¦¦%)..¿îÝ0I'²A.wDsäÜTPè/×^p
b9 3e 62 85 12 de ea fa c4 87 02 f4 b4 7d bd 94 27 80 a7 b9  .§w..{O$t.lù.Ì>\ruó.¹>b..ÞêúÄ..ô´}½.'.§¹
a9 fd cb 13 55 5d d6 6a c1 86 30 b9 6e 44 4c ba 1c 8f 9a f8  ð.t.ñ¤ø....ï.wYe.mìµ©ýË.U]ÖjÁ.0¹nDLº...ø
30 ee 15 35 07 d0 71 98 d2 dd 1d b3 21 98 2f 87 b8 c2 89 7a  cQ}.aÁ7.ÒÅþ#.ò;ýø¬y'0î.5.Ðq.ÒÝ.³!./.¸Â.z
9c 8e 6f 34 33 ba 7e 7c 7c 20 35 d5 ca cf 0e c8 07 a0 3d 9a  <;.V<.S.fbæç'º.%R..µ..o43º~|| 5ÕÊÏ.È. =.
90 41 fc a0 10 d4 ee d0 bc 2f a7 d5 2a 90 e7 43 e9 c4 4e 96  F±.z':pG£-"¯Áns¿Ék}O.Aü .ÔîÐ¼/§Õ*.çCéÄN.
8b 2b fd 5c 1c 98 a0 f2 55 50 0d c4 32 0e 8e f0 3b ed 25 43  .c..£ä5»ê¬U.E.ômøGu³.+ý\.. òUP.Ä2..ð;í%C
48 fb 65 66 7d 02 71 32 e0 42 ee b9 ed a9 1d d4 6f 7b e1 5c  >lbq^~¸.6)'6/.x_á èÞHûef}.q2àBî¹í©.Ôo{á\
24 a3 39 2c 21 74 72 c1 0e cd b7 c0 bc e3 99 a9 d8 9a e6 1b  õñL¾.Ü;Ê4.Y.P.øâ[.·Þ$£9,!trÁ.Í·À¼ã.©Ø.æ.
e1 10 8c 44 52 e1 fd cb ac 1f 80 d7 4d 2d 0d eb 86 bd 51 d6  µ@§j...bp®ì±.Y.uÙ\ªvá..DRáýË¬..×M-.ë.½QÖ
21 96 ad 29 9a ca b5 c8 c8 98 03 e9 a9 1f ed c1 75 6c 5b 01  ãÁp÷.&¢vò¿.|.XNóö .7!.-).ÊµÈÈ..é©.íÁul[.
19 c8 35 06 a4 81 ca 01 9f c1 57 74 ae 48 a4 73 3f c5 a4 ad  l.ª.7¨..Eð.º¬u¿C.Ìþd.È5.¤.Ê..ÁWt®H¤s?Å¤-
86 2b fe 53 18 f4 c1 40 28 0d 4d ea 8a 6f 0c 00 3f a0 8f 3b  V°Ç×...ã¹@.Ü¤»Ä.í.I¿.+þS.ôÁ@(.Mê.o..? .;
b7 cd ed ef 02 f9 f1 6a 80 e6 2b 70 ef 58 80 d5 c8 9e d1 64  ¨ ¥a$.4'Íî¸~Ø.qZé,@¬·Ííï.ùñj.æ+pïX.ÕÈ.Ñd
7a 8a 92 f0 0e 3c ec 07 62 7a 06 d5 d6 a9 51 bf e7 45 25 9c  ów<.'.@ªLs Åêð&nò.¯Oz.'ð.<ì.bz.ÕÖ©Q¿çE%.
05 ce 33 2e 2a bd 69 21 03 28 27 14 88 82 e4 45 0d 0e 8e 3d  lÎ~B.K.ä.ÜJ¦.$Àè..6..Î3.*½i!.('...äE...=
94 8f 2c a1 74 c2 0d 8b d8 51 ef f2 45 47 7d f1 7c dd 77 3a  Î`¹Áf..ý.`Ïú<¦ÙywÇw:..,¡tÂ..ØQïòEG}ñ|Ýw:
99 29 96 17 fb f6 4f 32 6c 0d c5 d2 cc c0 d8 89 74 49 9d 63  ea.Ç..3Ñ^Á.½.ã.ý þ:6.)..ûöO2l.ÅÒÌÀØ.tI.c
1e c5 22 70 f0 7b 68 65 13 5b 6a 67 ac 55 c0 29 c3 03 0d 74  ÒaÉ¼.'ûe_Í..ý9Q.Yæ.Á.Á"pð{he.[jg¬UÀ)Ã..t
```

Metainfo File Contents (cont'd)

d8:announce39:http://torrent.ubuntu.com:6969/announce
13:announce-listll39:http://torrent.ubuntu.com:6969/announcee
l44:http://ipv6.torrent.ubuntu.com:6969/announceeee
7:comment33:Torrent downloaded from 1337x.org
10:created by8:atik0786
13:creation datei1335433885e
4:info
d6:lengthi735358976e
4:name29:ubuntu-12.04-desktop-i386.iso
12:piece lengthi524288e
6:pieces28060:.....ee

Bencoding

- Integers: **i** <number> **e**
  - **i**404**e** is the integer 404
- Strings: <string length> **:** <string data>
  - 7**:**example is the string “example”
- Lists: **l** <bencoded values> **e**
  - **l**7:example**i**404**e****e** is [“example”, 404]
- Dictionaries: **d** <string><bencoded value> **e**
  - **d**5:alpha1:a4:beta1:b**e** is [“alpha”:”a”, “beta”:”b”]

Metainfo File Keys

announce http://torrent.ubuntu.com:6969/announce

announce-list http://torrent.ubuntu.com:6969/announce
http://ipv6.torrent.ubuntu.com:6969/announce

comment Torrent downloaded from 1337x.org
created by atik0786
creation date 1335433885

info
- **length** 735358976
- **name** ubuntu-12.04-desktop-i386.iso
- **piece length** 524288
- **pieces** 28060:.....

Metainfo File Translated

[“announce”:“http://torrent.ubuntu.com:6969/announce”,
“announce-list”: { [“http://torrent.ubuntu.com:6969/announce”],
[“http://ipv6.torrent.ubuntu.com:6969/announce”] },
“comment”:”Torrent downloaded from 1337x.org”,
“created by”:”atik0786”,
“creation date”:1335433885,
“info”: [“length”:735358976,
“name”:”ubuntu-12.04-desktop-i386.iso”,
“piece”:524288,
“pieces”: <28060 byte string of concatenated 20 byte SHA-1 hashes of each piece>]]

Tracker

- Coordinates the communication between peers
- Tracks statistics of torrents
- Typically a public server
  - Not all torrent indexes have a tracker
  - Public vs private trackers

#	Tracker	Torrents	Peers	Software
1	PublicBitTorrent	2,484,145	21,694,091	Opentracker
2	OpenBitTorrent	2,388,738	21,186,589	Opentracker
3	Denis.Stalker	1,614,356	10,527,993	Opentracker
4	1337x.org	302,799	4,889,991	XBT
5	Torrent.to	326,467	3,205,170	Opentracker

torrentfreak.com

Tracker (cont'd)

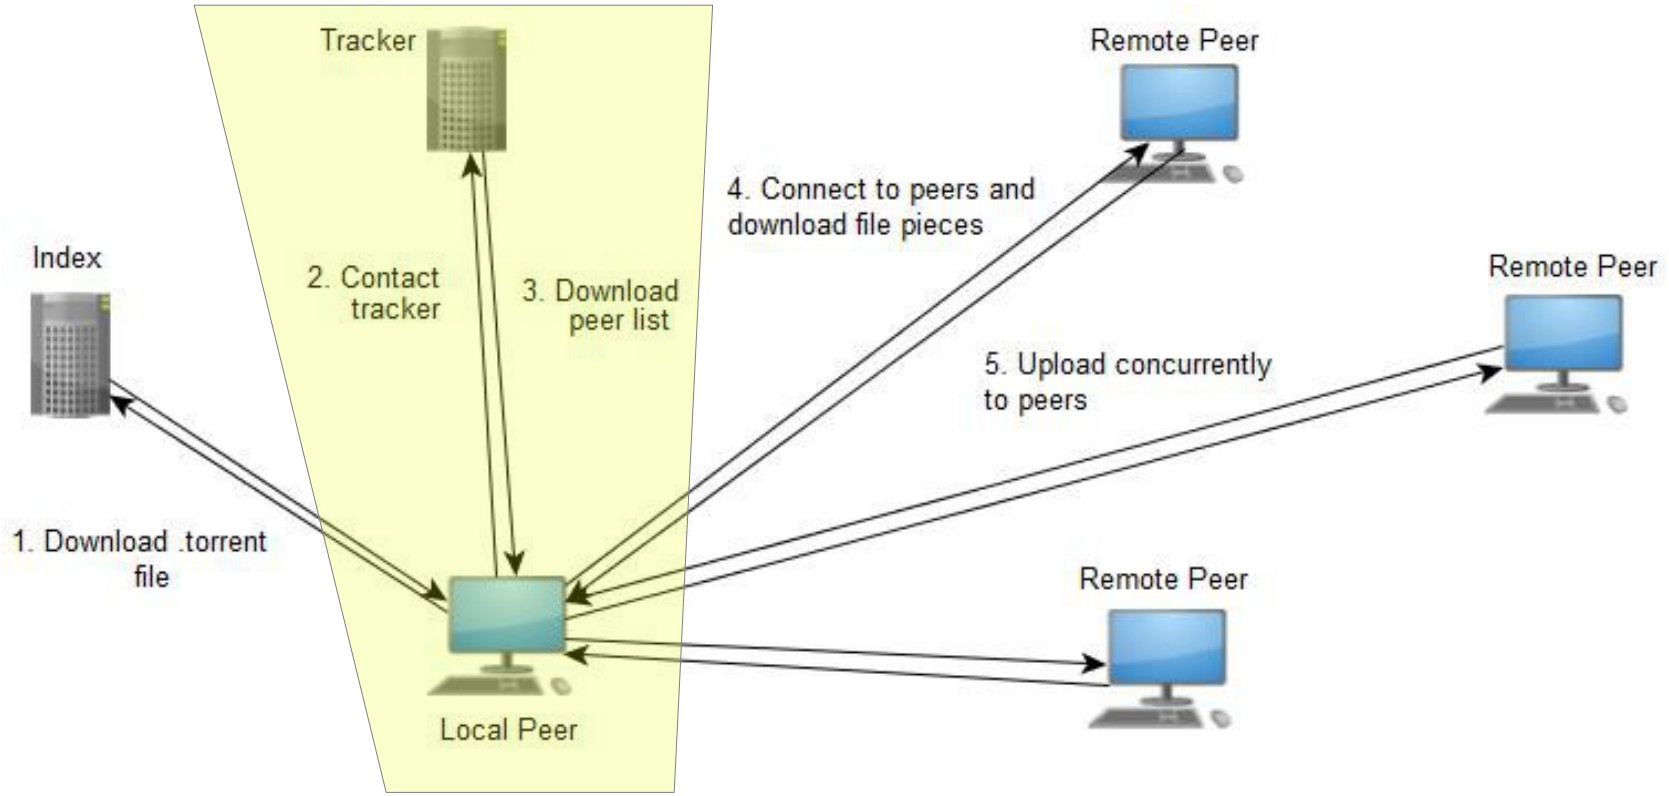

Tracker Communication

GET /announce?**info_hash**=%bb%b6%dbi%96Z%f7i%f6d%b6cny%14%f8sQA%b3&
peer_id=-UT3210- %b6m%10%ea%bb%01%1eDkV%bc %e5
&**port**=42176&**uploaded**=0&**downloaded**=0&**left**=735358976&**corrupt**=0&**key**=7F5DA749&
event=started&**numwant**=200&**compact**=1&**no_peer_id**=1 HTTP/1.1
Host: torrent.ubuntu.com:6969
User-Agent: uTorrent/3210(28086)
Accept-Encoding: gzip
Connection: Close

Tracker Response

HTTP/1.1 200 OK
Content-Type: text/plain
Content-Length: 313

d8:**complete**i33e10:**downloaded**i401e10:**incomplete**i3e8:**interval**i1924e12:**min**
intervali962e5:**peers**216:p.NVNcp.N.....'..4.F.%....r.....o.....*.!?.M...-.r."T.iu.
+.l.u.63l.u..#.Yp.....b.&.R._g.\9E^..j..Y.@7..Y.5?..We.i..TJ.
$V.SwV..3R.;A..Q.iZh.O....QLo....L.....D.H"6.>-.i=.;.@b.<.c,.IB)...Hf)...gF..Z.5o....$.e

[“**peer_id**”:“UT3210-%b6m%10%ea%bb%01%1eDkV%bc%e5”,
“**ip**”:”192.168.1.8”,
“**port**”:”42176”]

Peer Wire Protocol (PWP)

- Communication and data transfer between peers
- Requests pieces of target file in 16KB “blocks”
  - Pieces generally 512KB-2MB in size
- States
  - **Choked**: remote peer does not respond to requests for blocks from peer
  - **Unchoked**: remote peer will upload data to the peer
  - **Interested**: peer expresses want to request blocks from remote peer
  - **Uninterested**: peer will not request blocks from remote peer

Establishing Connections

- Client peer opens PWP connections with all peers in the list from the tracker
- Handshake PDU
  - **Pstrlen**: 1 byte, length of pstr
    - 19 for Bittorrent
  - **pstr**: variable length, protocol ID string
    - “BitTorrent protocol”
  - **reserved**: 8 bytes for extensions
  - **info_hash**: 20 byte SHA-1 of entire info key in metainfo file
  - **peer_id**: 20 byte ID, same as transmitted in tracker request

Handshaking

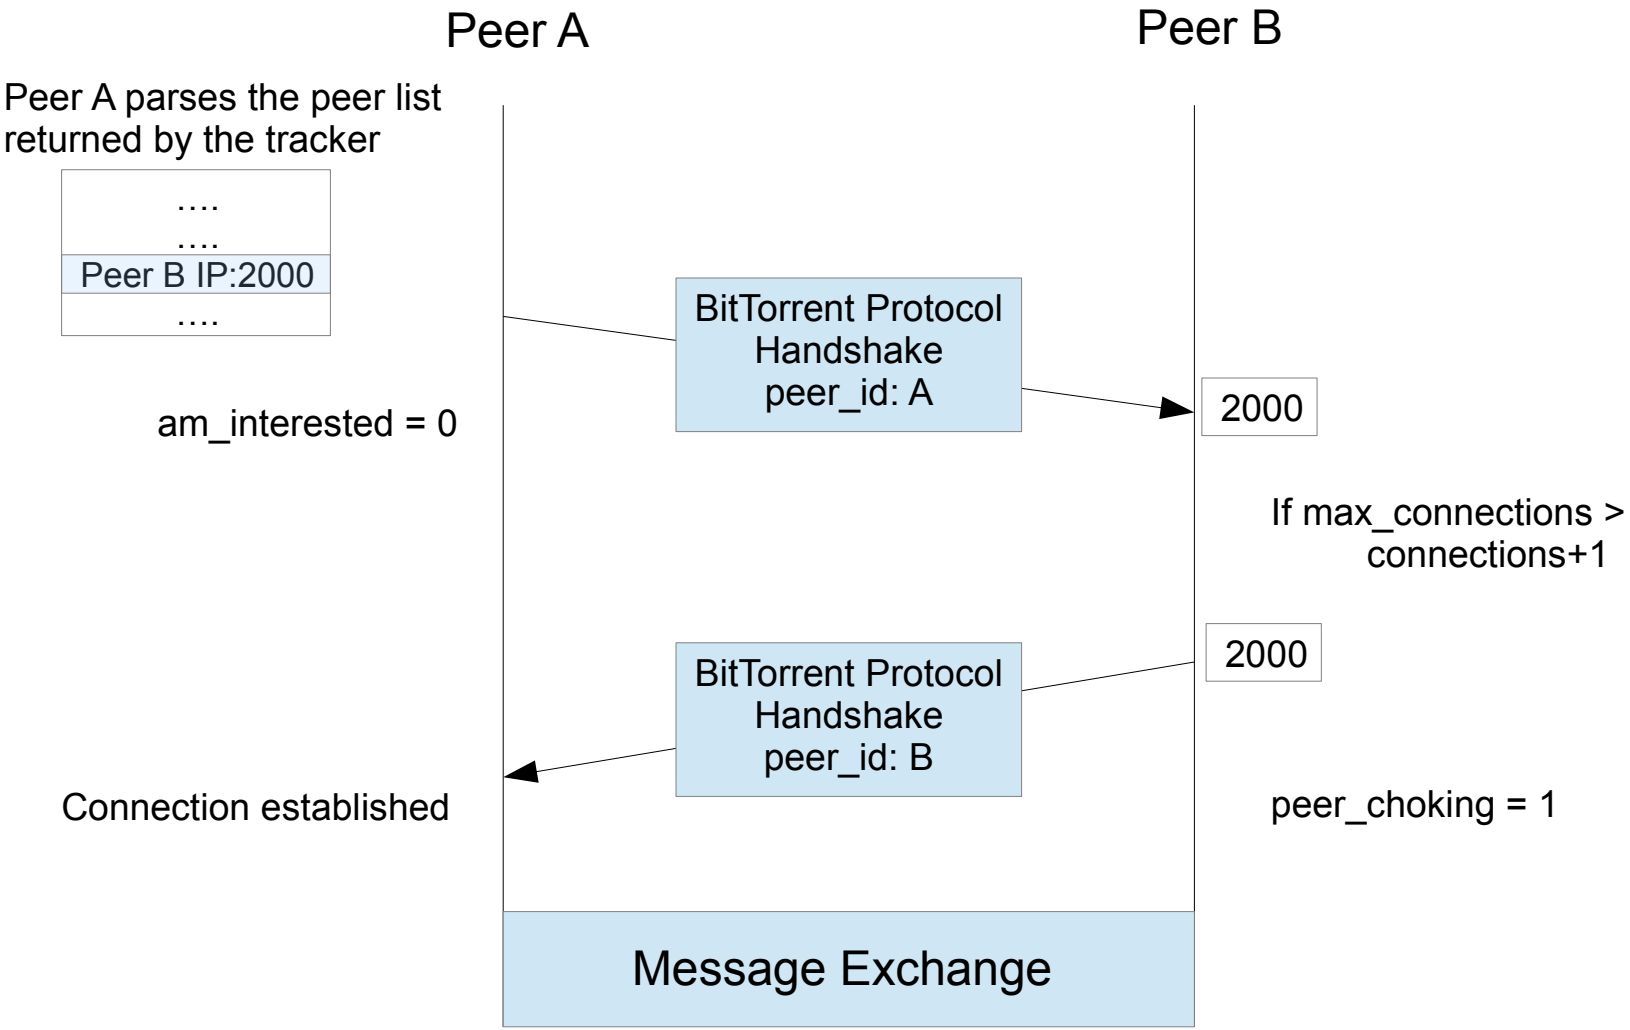

BitTorrent Message PDU

- Format
  - 4 byte **length** value, 1 byte **ID** value, variable **payload**
- Multiple PDU's can be sent in a single TCP PDU
- Types
  - **keep-alive** – sent if no commands sent to keep connection alive
    - Connection timeout = ~2 minutes
  - **choke** (0) – local peer choking remote peer
  - **unchoke** (1) – local peer unchoking remote peer
  - **interested** (2) – local peer is interested
  - **uninterested** (3) – local peer is uninterested

BitTorrent Message PDU (cont'd)

- **Have** (4) – payload is piece index
  - Sent after downloaded and hash verified
- **Bitfield** (5) – payload is bitfield of piece indexes
  - Bit set to 1 if peer has the piece at that index
  - First message sent after handshake, optional if peer has no pieces
- **request** (6)
    - **index** – piece index
    - **begin** – byte offset of block within piece
    - **length** – length of requested block
  - Default block length is 16KB, peer drops connections for any request lengths over 32KB

BitTorrent Message PDU (cont'd)

- **piece** (7)
  - **index** – piece index
  - **begin** – byte offset of block
  - **block** – the requested block data
- **cancel** (8) - **index**, **begin**, **length**
  - Cancels a block request
- **port** (9) - 2 byte port number
  - Used for DHT

Message Flow

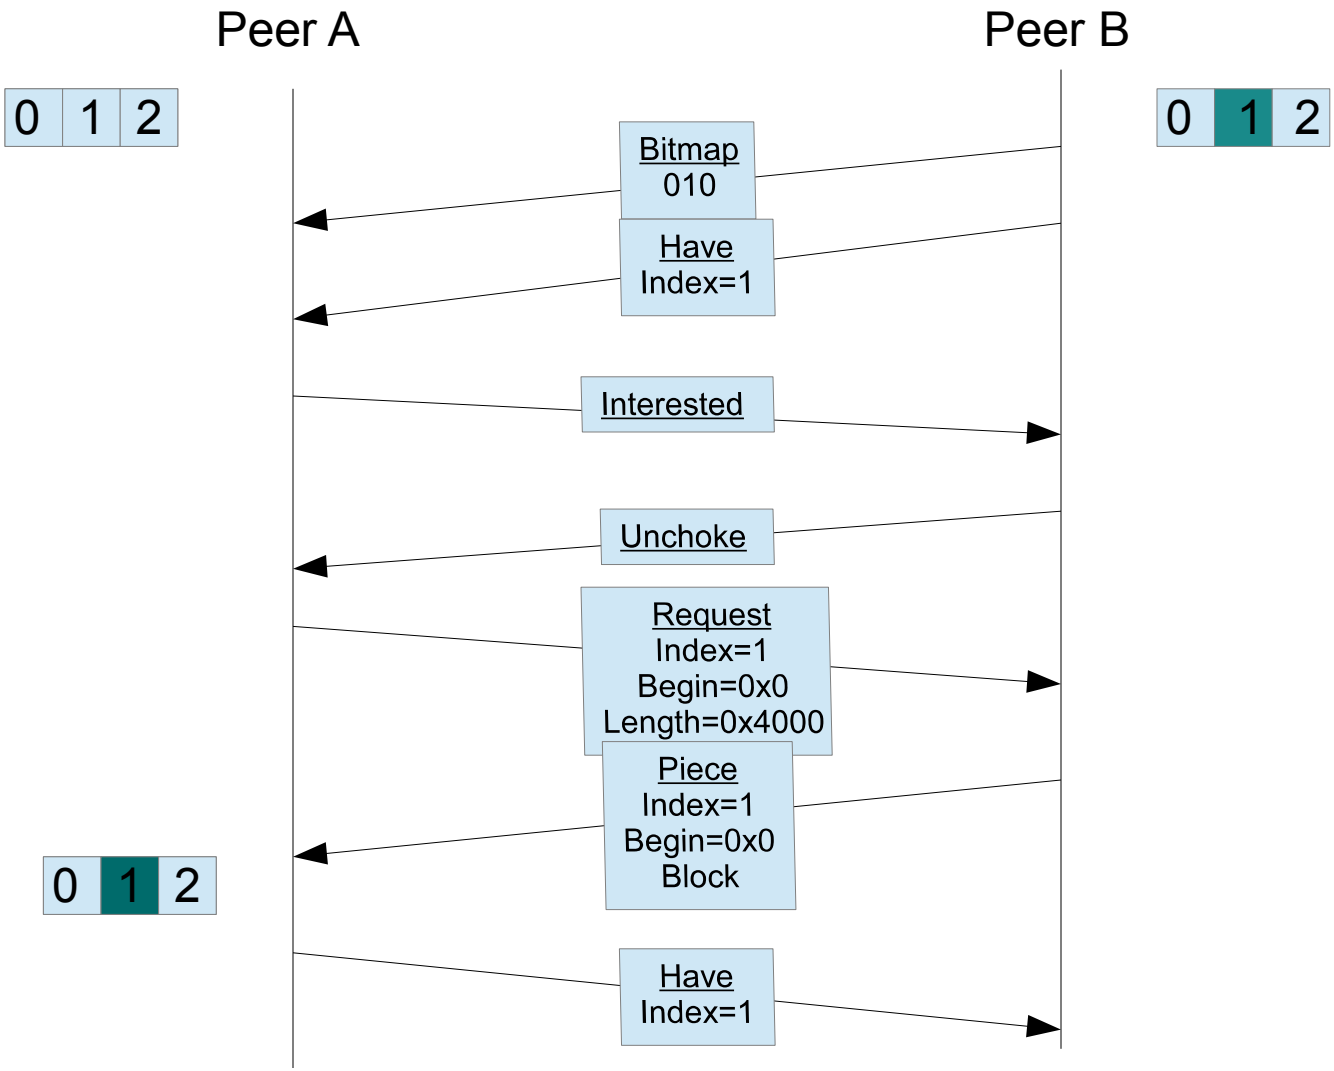

Choke Algorithm

- Local peer keeps a list of interested and uninterested remote peers
- Interested peers are ranked by their upload rate to the local peer
- The top four peers are unchoked - downloaders
  - Every 10 seconds, rates and downloaders are recalculated
    - Time cycle reduces fibrillation – rapid choking and unchoking
  - If a peer has a better upload rate than the downloaders but is uninterested, unchoke
    - If it becomes interested in the future, it replaces the downloader with the lowest upload rate
- Reduces leechers
  - in order to become unchoked, must upload to peer

Optimistic Unchoke

- Every 30 seconds, one random and choked peer is unchoked regardless of its upload rate
  - If interested, counts as one of the four downloaders
    - If uninterested, unchoke and randomly select a new choked peer
- Advantages
  - Allows faster connections to be discovered
  - Selected peer may be new and have no pieces to share
    - Optimistic unchoke will give peer its first piece so it can upload to others and become a downloader

Peer Selection Process

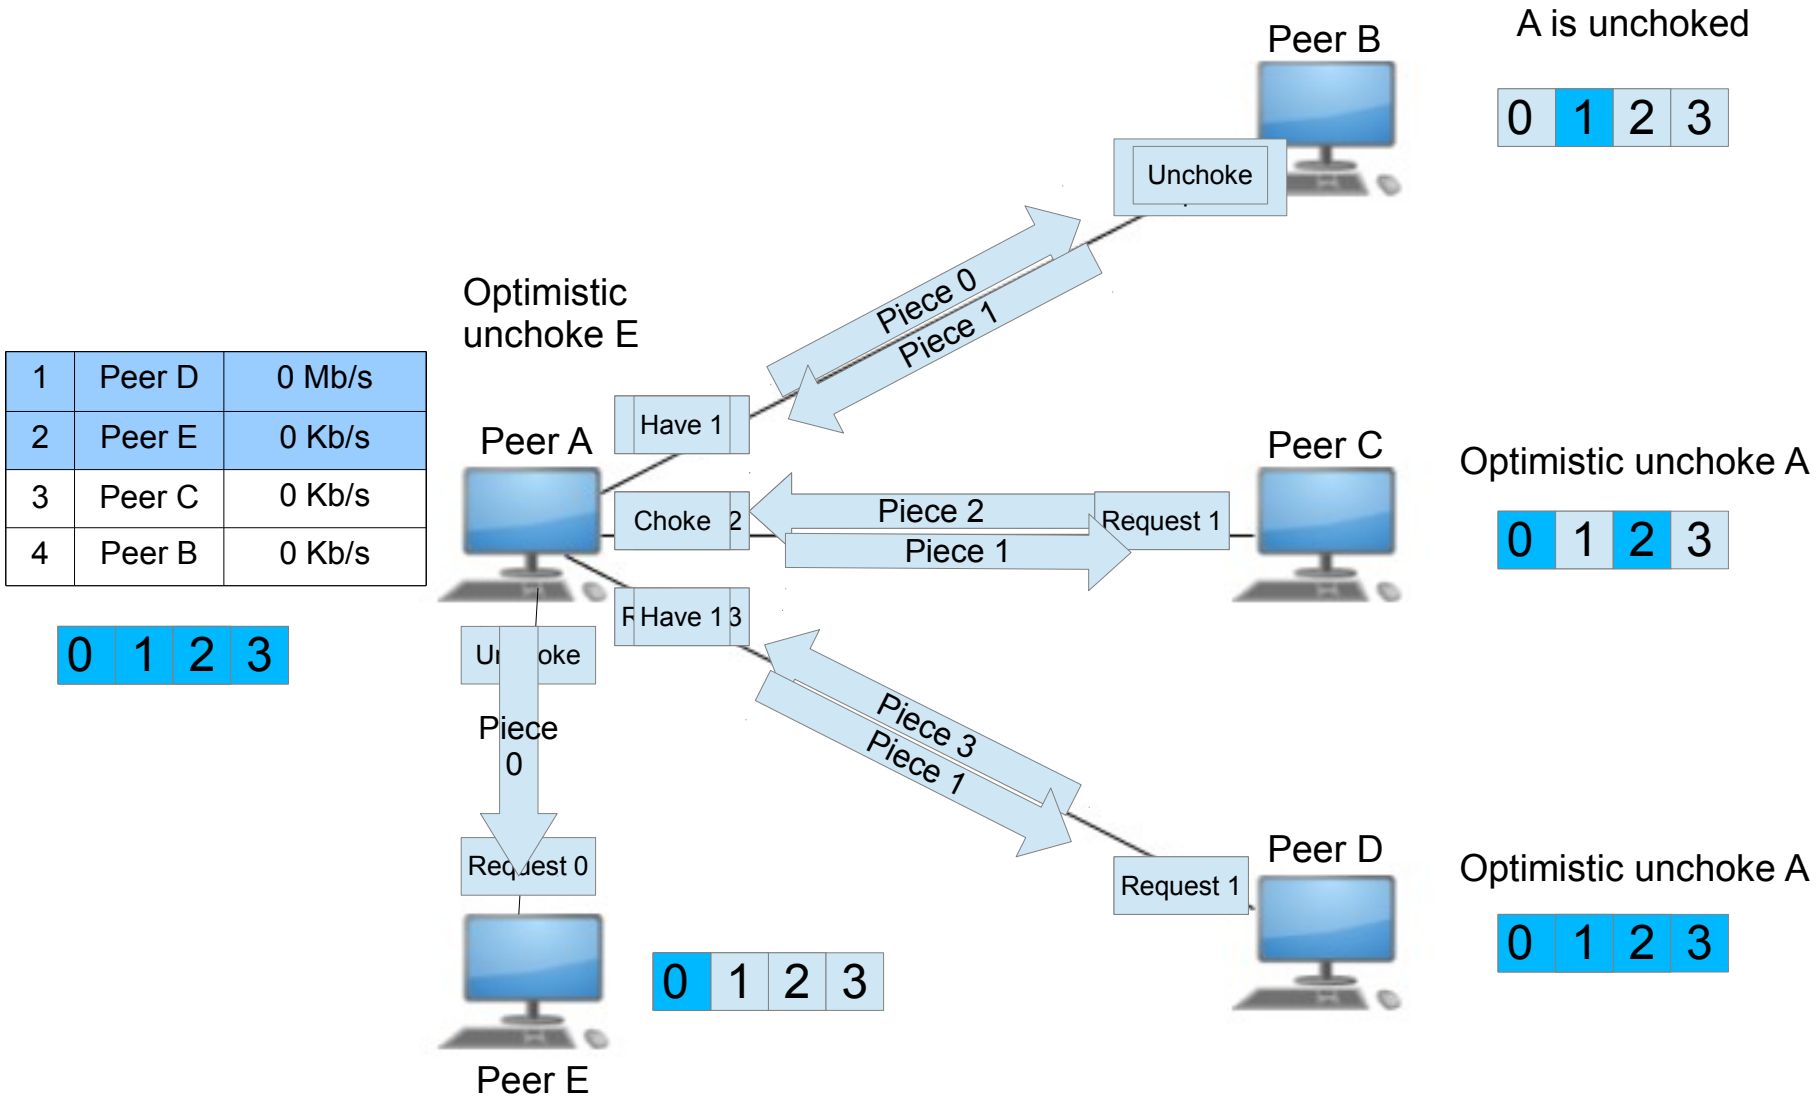

Seeder Mode

- Entered once peer has all pieces
- Uploads to four peers, ranked by upload rate
- Optimistic unchoke

Snubbing

- Peer may become choked by all uploaders
- Anti-snubbing – local peer cannot unchoke a peer if
  - The peer has received blocks from the local peer
  - The peer has not uploaded any blocks to the local peer
- Snubbing peer is not unchoked unless by optimistic unchoke
- Allows snubbed peer to find better peers
  - Can select multiple peers for optimistic unchoke, improves recovery time

Piece Selection Algorithm

- **Random first** – randomly selects a piece to request
  - Occurs when peer has downloaded less than 4 pieces
  - Allows fast acquisition of pieces
- **Rarest first** – requests the rarest pieces first
  - Counter tracks the number of peers who have a specific piece
  - Pieces are ranked by rarity, requested in low to high order
  - Allows pieces to be have a more equal distribution
- **Strict block policy** – when a block is requested, all other requested blocks from the same piece have higher priority
  - Complete pieces faster so they can be uploaded to other peers

End Game Mode

- Last blocks of a download can be delayed
- Send block request to all peers
  - Once received, send cancel to other peers
  - Helps last blocks download faster
  - Not reliant on a single, possibly slow peer
- Usually occurs after all blocks are requested or a download threshold is reached

End Game Mode Process