

SWIN
BUR
* NE *

SWINBURNE
UNIVERSITY OF
TECHNOLOGY

Fairness of BitTorrent

Gerarde Kelly

Centre for Advanced Internet Architectures

Swinburne University of Technology

gerardekelly@swin.edu.au

2nd August 2007



Outline

- Project Definition
- SCTP Overview
- Experimental Setup
- Problems Encountered
- Conclusion



SWIN
BUR
* NE *

SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 2

The Project (cont...)



- Team
 - Jason But
 - Nigel Williams
 - Myself (for winter break)
- Conduct BitTorrent experiments using TCP and SCTP to determine the following:
 - Is BitTorrent a "fair" way of retrieving content
 - Does localised Content Caching improve performance.
 - Does SCTP offer benefits over TCP
- Vary the bandwidths, connection times and delay of torrent clients and observe "swarm" behaviour



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 3

Objectives



- Different interpretations of fairness:
 - E.g. Peers should download from and upload to other peers with equal measure (equal pieces served/fetched)
- We define fairness as:
 - Each peer in the swarm - under the same conditions (bandwidth) - should receive equal average download and upload rates (ie. No similar capacity peer should be favoured by the swarm over another.
 - This may mean that some individual peers may download more than they have uploaded.



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 4

Objectives (cont...)



- Trials to determine fairness in BitTorrent
 - 82 Jails (1 tracker, 1 seeder, 80 leechers)
 - Base case: 80 peers 512/128
 - Also different access speeds, number of peers, starting times, RTT etc.
 - Some tests with varying RTT ~ simulate peer group of mixed geographic locations
 - How do Australian peers fair with majority US/Europe peers?
 - Could caching at the ISP help?
- Repeat using SCTP



SWINBURNE
UNIVERSITY OF
TECHNOLOGY


CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 5

BitTorrent Overview



BitTorrent tracker identifies the swarm and helps the client software trade pieces of the file you want with other computers.



Computer with BitTorrent client software receives and sends multiple pieces of the file simultaneously.

©2005 HowStuffWorks



SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 6

BitTorrent Choke Algorithm




- Determines which peers to download from
 - Peers can have two states – choked and unchoked
- Helps to ensure peers download/upload in equal measures – reduces “freeloading”
 - Should group together like-bandwidth peers
- Periodically assesses the download rate obtained from connected peers:
 - The slowest peers are ‘choked’ and a new peer is randomly unchoked (“optimistic unchoking”)
 - Optimistic unchoke helps discover faster peers, also distributes pieces to new peers



SCTP Overview




- What is SCTP
 - Stream Control Transmission Protocol
 - Sits on the Transport layer of the TCP/IP Reference Model
 - Can be used in the same way as TCP or UDP
 - Features:
 1. Multi Homing
 2. Message-Based
 3. Multi Streaming
 4. Selective Acknowledgement (SACK)
 5. Better Error Correction



SCTP Overview (cont...)



Multi Homing



Message Based

- TCP uses byte streams
 - Application needs to do the work
- SCTP deals in message streams
 - More logical



CAIA Seminar

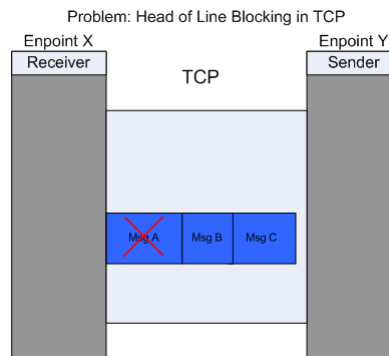
<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 9

SCTP Overview (cont...)



Multi Streaming

- Send 3 independent messages
- Error occurs on Msg A
- B and C in kernel waiting to be processed.
- TCP ensures dependency on the 3 independent messages
- Application needs to wait until A is retransmitted (or all 3) before it can receive B and C



CAIA Seminar

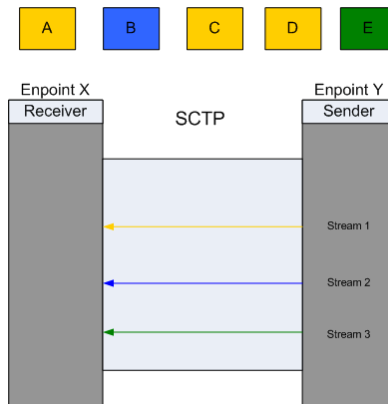
<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 10

SCTP Overview (cont...)



Multi Streaming

- SCTP allows messages to be sent on different streams
 - Stream 1: A, C, D (order preserved)
 - Stream 2: B
 - Stream 3: E
- If one stream has an error, other streams not affected.
- Order affected only within each stream



CAIA Seminar

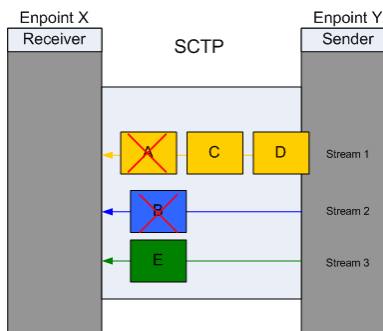
<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 11

SCTP Overview (cont...)



Multi Streaming

- If Msg A and B are dropped
- Stream 3 can still be sent to the application
- Retransmit Msg A and B (SACK)



CAIA Seminar

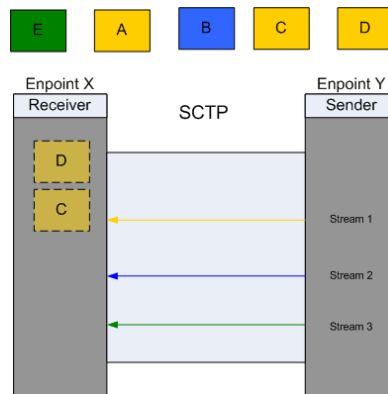
<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 12

SCTP Overview (cont...)



Multi Streaming

- ❑ Msg E can still be passed to the application and msg A and B retransmitted.
- ❑ Msg D and C are buffered in the kernel.
- ❑ Only the order of the individual stream needs to be preserved.



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 13

SCTP Overview (cont...)

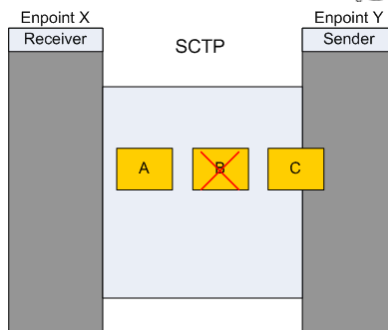


Selective Acknowledgement

- ❑ More descriptive
- ❑ Allows for less retransmissions.
- ❑ An option in TCP, built into SCTP so payload is less.
- ❑ Error occurs in B.
 - ❑ Only resend B, not the entire stream.

Stronger Checksum (CRC32c)

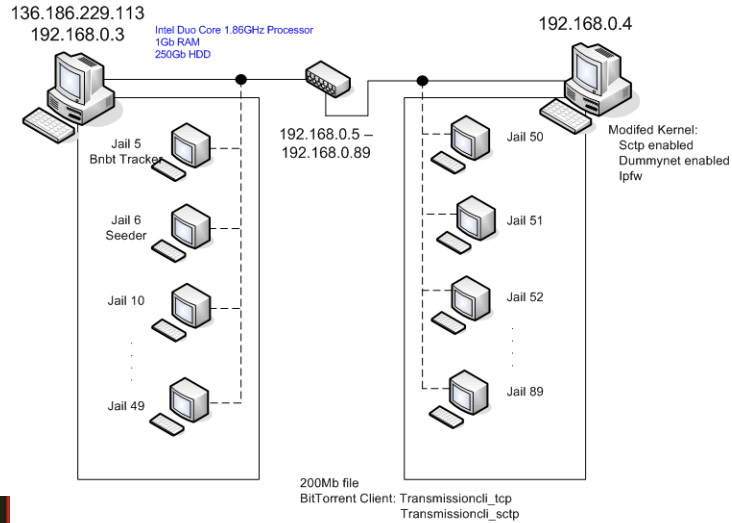
- ❑ TCP uses 16 bit checksum
- ❑ SCTP uses 32 bit so is more reliable



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 14

Experimental Setup



SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 15

Experimental Setup (cont...)



■ Scripts Used

- Jail_setup.sh
 - setups N jails on a system
- Torrent_setup.sh (only script called)
 - Sets up the torrent client
 - Copies the torrent
 - Sets the dummynet limits
 - Sets the cron jobs to automatically start the client
 - Sets up the tcpdump session
 - Sets up the tracker & seeder



SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 16

Experimental Setup (cont...)



- Scripts Used
 - Progress Script (called by torrent_setup.sh)
 - Determines the progress of the experiment
 - If seeding, start torrent_takedown.sh script
 - Torrent_takedown.sh (called by Progress script)
 - Kills tcpdump, transmission and tracker
 - Copies all log files into single directory
- Data Collected
 - Tcpdump logs
 - Bnbt tracker logs
 - Transmission client logs



SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 17

SCTP and BitTorrent



- We are exploring the affect of BitTorrent using SCTP.
- Multi Streaming
 - Send the data on one stream, the protocol chatter on another stream
 - Packet loss on data stream won't interfere with protocol information and vice-versa. (Due to head-of-line blocking)
- Selective Acknowledgment (SACK)
 - Less bandwidth consumed. With our lossy experiments
- Better detection of transmission errors earlier due to SCTP error checking
- Better kernel resources, a native implementation would use less socket and file descriptors



SWINBURNE
UNIVERSITY OF
TECHNOLOGY

CAIA Seminar


<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 18

Problems Encountered



■ Implementation Problems.....

- Sctp sctp_recvmsg
 - Causes system restart on FreeBSD if 'flags' parameter not initialised
- Sctp for jails (work in progress)
 - Binding with ADDR_ANY – listens to IP addresses belonging to other Jail hosts
- Issues with Sctp and SMP kernel
 - cause is under investigation.
 - SMP currently disabled for Sctp tests.



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 19

Problems Encountered (cont...)



- Maximum number of file descriptors concurrently running exceeded.
 - Kern.maxfiles increased to 100000
- Maximum number of mbufs available to the networking stack exceeded.
 - Kern.nmbclusters increased to 100000

Thanks Lawrence



CAIA Seminar

<http://caia.swin.edu.au> gerardekelly@swin.edu.au 1 August 2007 Page 20

Results



- Hope to determine if ISP content caching could benefit torrent users
- Can SCTP enhance the overall experience of BitTorrent.

Conclusion



- Tests are currently in progress
- Many thanks to all who have helped with my 6 weeks at CAIA
 - Jason But
 - Nigel Williams
 - Lawrence Stewart
 - Adam Black
 - Everyone else.