SWIN BUR NE

CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Synchronisation Analysis over Concurrent TCP Flows

Philip Jay, Qiang Fu
{pjay, qfu}@swin.edu.au

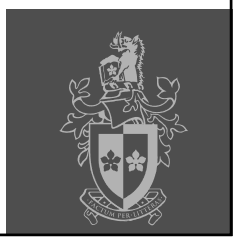

Outline

- Motivation
- Implementation
- Method
- Results
- Conclusion
- Further Work

Motivation

- Flow characteristics focus on single TCP flows
- Characteristics of concurrent TCP flows are unknown
- There may or may not be global synchronisation

Motivation

- TCP flows should share the bandwidth
- Convergence time to fairly share bandwidth
- Synchronised packet losses increase convergence time

Implementation

- Modified FreeBSD 5.4 host at Swinburne
  - TCP stack modified to include current RTT estimates in TCP packets
- Remote hosts in various locations
  - South Australia
  - China
  - Germany
  - USA

SWIN BUR NE — CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Implementation

- Generate tcpdump trace data
- Unix cron utility to schedule data collection
- Script to automatically run various programs
  - tcpdump to capture the TCP/IP headers
  - scp to perform a data transfer

SWIN BUR NE — CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Implementation

- Analysis of the data
  - Six gigabytes of trace files for a 1 week trial
  - Custom Java program

SWIN BUR NE — CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Method

- Loss events vs. lost packets
- Correlate lost packets into loss events
  - Period between sending time and detection time of the first lost packet
  - Any lost packets in that time increase the duration of the loss event
  - Verified

SWIN BUR NE — CENTRE FOR ADVANCED INTERNET ARCHITECTURES

Method

- Compare the flow with the most loss events to all other flows
- Count how many other flows have losses at the same time

Results

- South Australian remote host (12ms round trip time)

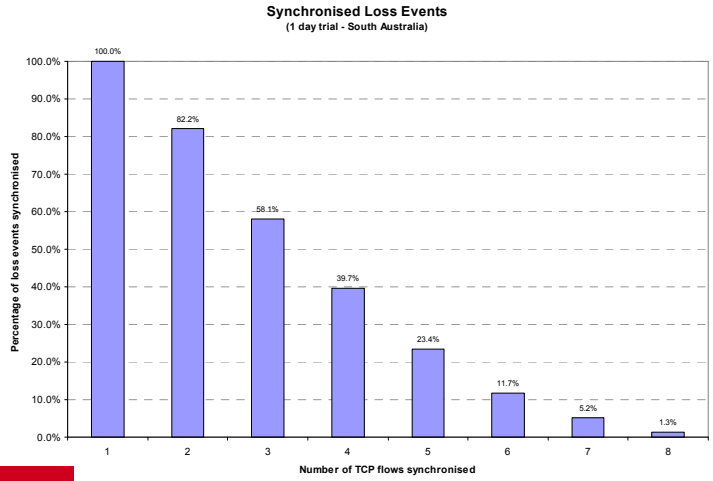

Results

- German remote host (340ms RTT)

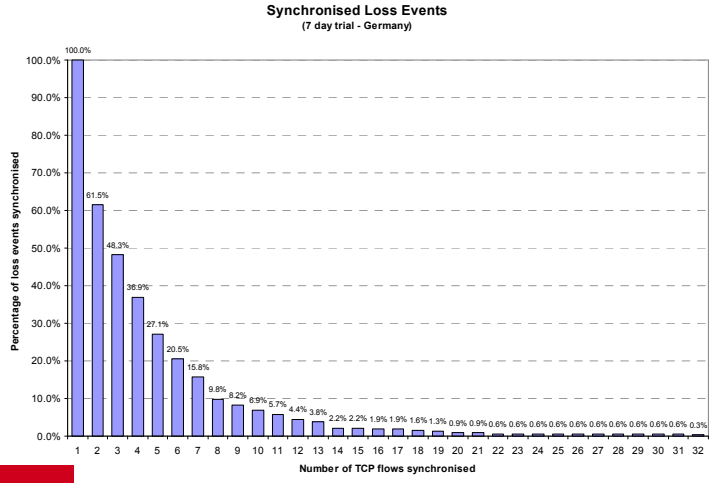

Results

- USA remote host (230ms RTT)

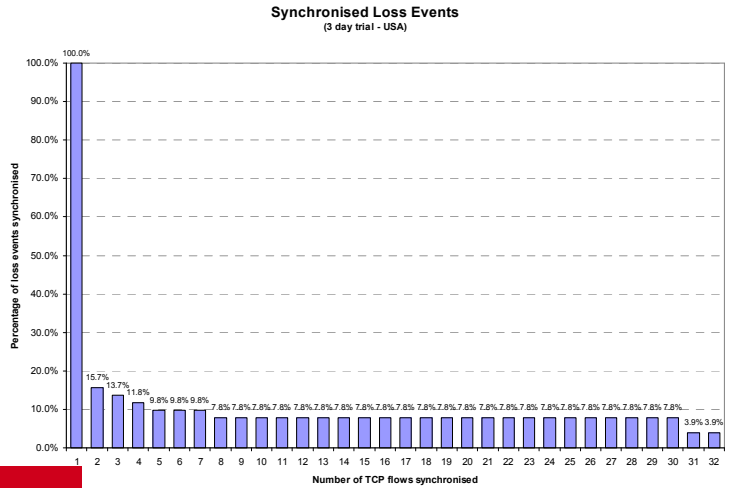

Results

- Chinese remote host (170ms RTT)

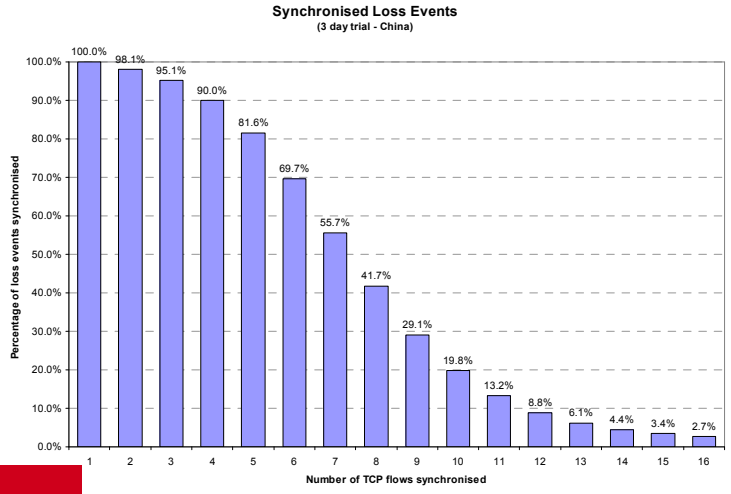

Conclusion

- Some new TCP algorithms assume 100% synchronisation
- Losses are not 100% synchronised
- Serious congestion likely to be synchronised

Further Work

- Analyse the data in more depth
  - Correlate round trip times
  - Group flows together
- Capture the congestion window size
  - Accidentally captured the advertised size
- Use different versions of TCP
- Attempt method again in a controlled environment