




ttprobe v0.1 - Packet-Driven TCP Stack Statistics Gathering for TEACUP

Rasool Al-Saadi
ralsaadi@swin.edu.au

Supervisors

Prof. Grenville Armitage & Dr. Jason But

Centre for Advanced Internet Architectures (CAIA)
Swinburne University of Technology



Outline




- TEACUP
- TCP stack information loggers
- Web10g logger
- ttprobe v0.1
- Linux Jprobe/Kprobe mechanism
- Comparing ttprobe v0.1 and Web10g

TCP/IP performance testing



- Bandwidth utilisation (Highest throughput)
- Network Congestion (Avoid/reduce congestion)
- Link sharing (Fairness)
- Latency (Smallest Latency)

TEACUP: TCP Experiment Automation

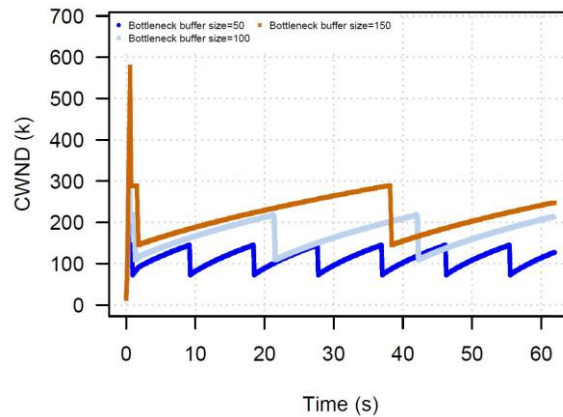


Example of a TEACUP testbed topology

TEACUP (cont.)



- Running and analysing TCP/UDP experiments.




TEACUP experiment to compare three different scenarios



TCP stack data loggers



- Capture state variables from inside TCP/IP stack.



- TEACUP uses (FreeBSD-> SIFTR), (Linux->Web10g)

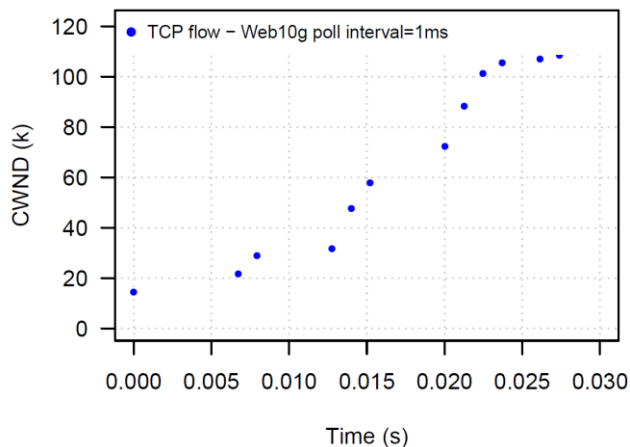


Web10g TCP state logger



- Implements many TCP Extended Statistics MIB (RFC4898).
- Using regular polling per fixed time interval (≥ 1 ms).
- Includes user-land and kernel patch codes.
- Latest Linux kernel official support is 3.17

Challenges of using Web10g



CWND plot of a TCP experiment captured using Web10g

tprobe v0.1: TEACUP TCP Probe




- Event-driven (per incoming and/or outgoing packet)
- Based on TCP probe, utilises JProbes/KProbes architecture
- Output format: tprobe, binary and web10g formats
- Capture TCP states in both directions
- Hooks `tcp_v4_do_rcv`, `tcp_v6_do_rcv` and `tcp_transmit_skb`
- Implements buffer flushing and reader termination

Kprobes mechanism in Linux



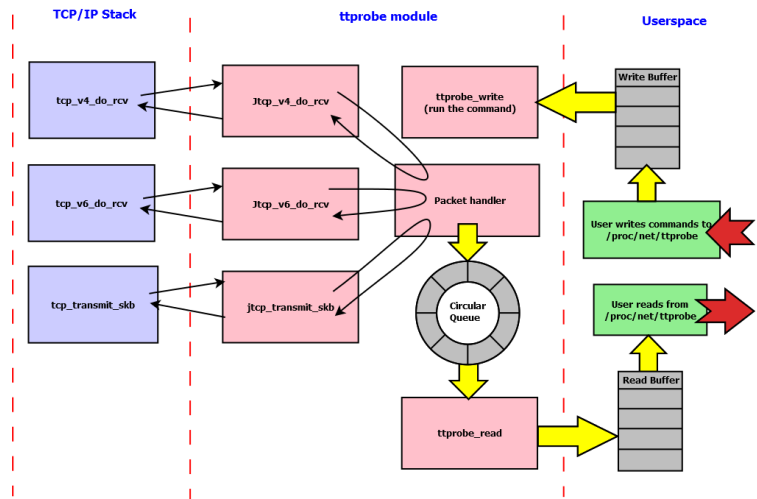
- Kprobes = Kernel Probes
- Used for kernel debugging and information collection
- Dynamically break into any kernel instruction/routine




Execution of a KProbe

Image copyright Sudhanshu Goswami 2005

tprobe v0.1 Design



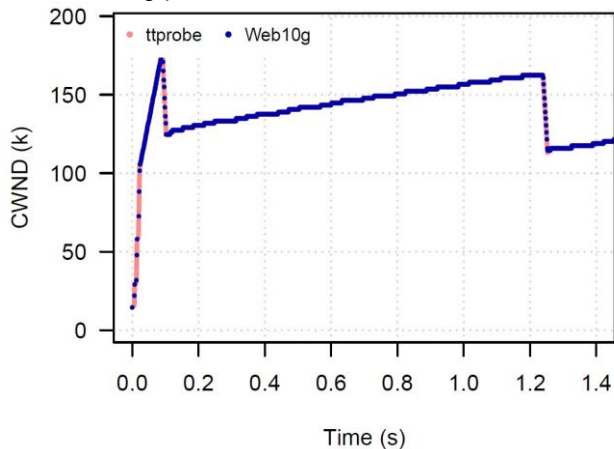
Comparing Web10g with tprobe



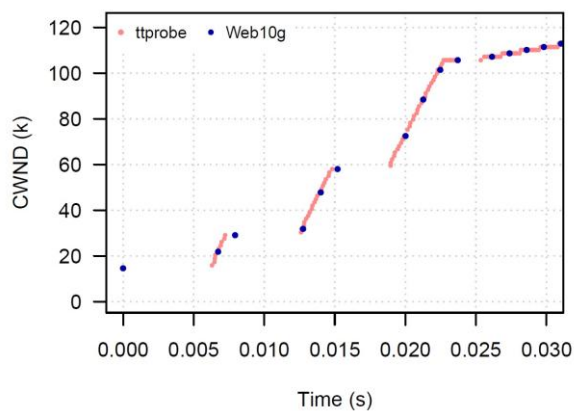
Sampling: ttprobe vs. web10g



- TCP flow (iperf), 100Mbps Up/Down, 6ms RTT, TCP CUBIC
- 1ms Web10g poll interval



Sampling: ttprobe vs. web10g (cont.)



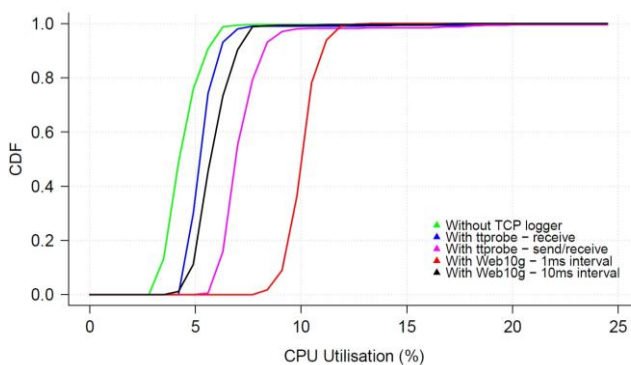
CWND of the first 30ms of the experiment



CPU usage: tprobe vs. web10g (cont.)



- Five individual experiments with ten runs each
- Each run lasted 60 seconds



CDF of CPU usage of iperf TCP traffic over 1Gbps link

Summary



- Event-driven TCP loggers produce more accurate data than TCP loggers that use polling approach.
- CPU overhead of event-driven loggers depends on the traffic.
- tprobe v0.1 captures more samples with reasonable CPU overhead than Web10g.
- Code and technical report
 - <http://caia.swin.edu.au/tools/teacup/>

Questions



Thank you for listening

