Timing Precision on Ark - RADclock

Darryl Veitch

dveitch@unimelb.edu.au

Department of Electrical & Electronic Engineering
THE UNIVERSITY OF MELBOURNE

The Big Picture

- **Host has hardware counters (TSC, HPET, ACPI..)**
  - but these drift, need disciplining
  - extra hardware (GPS, atomic clock) expensive
  - convenient to access a reference clock over the network
- **Network timing is hierarchal**
  - *Stratum-1* has access to reference hardware
  - *Stratum-2* references stratum-1, etc.
- **NTP protocol**
  - transports timestamps (of its own pkts!) between server & client
- **Kernel**
  - maintains a system clock, timestamps packets
- **Userland**
  - runs a clock sync daemon, providing key parameters defining a clock
- **Timestamping**
  - needed in kernel and userland, must be fast
  - critical both for clock sync itself, and network measurement

What RADclock Provides

- **Basic**
  - more accurate absolute timestamps (100's of μs rather than ms)
  - much higher robustness to network delays, disruptive events
- **Basic ++**
  - a *difference clock* (specialist clock for `short' time differences)
    - far more accurate (<1μs, even 10's of ns)
    - extraordinarily robust (lose server for hours, no problem)
  - reliable error bounds
- **Advanced**
  - Ability to `replay' raw timing data
    - capture raw timestamps fast, convert to UTC later
    - upgradable final timestamping
  - naturally compatible with virtualized operating systems

Potential of Ark + RADclock

- **Coordinated distributed experiments (trust timing)**
  - eg: Internet coordinate systems, route disambiguation
  - think of Spanner (Google's time aware distributed database)
- **Time Server vetting (can't trust stratum 1's)**
  - to better select Ark's own stratum 1's
  - as a service provided by Ark
- **Waiting in the wings**
  - RADclock servers as well as clients
  - more support to ease advanced use

time.nist.gov

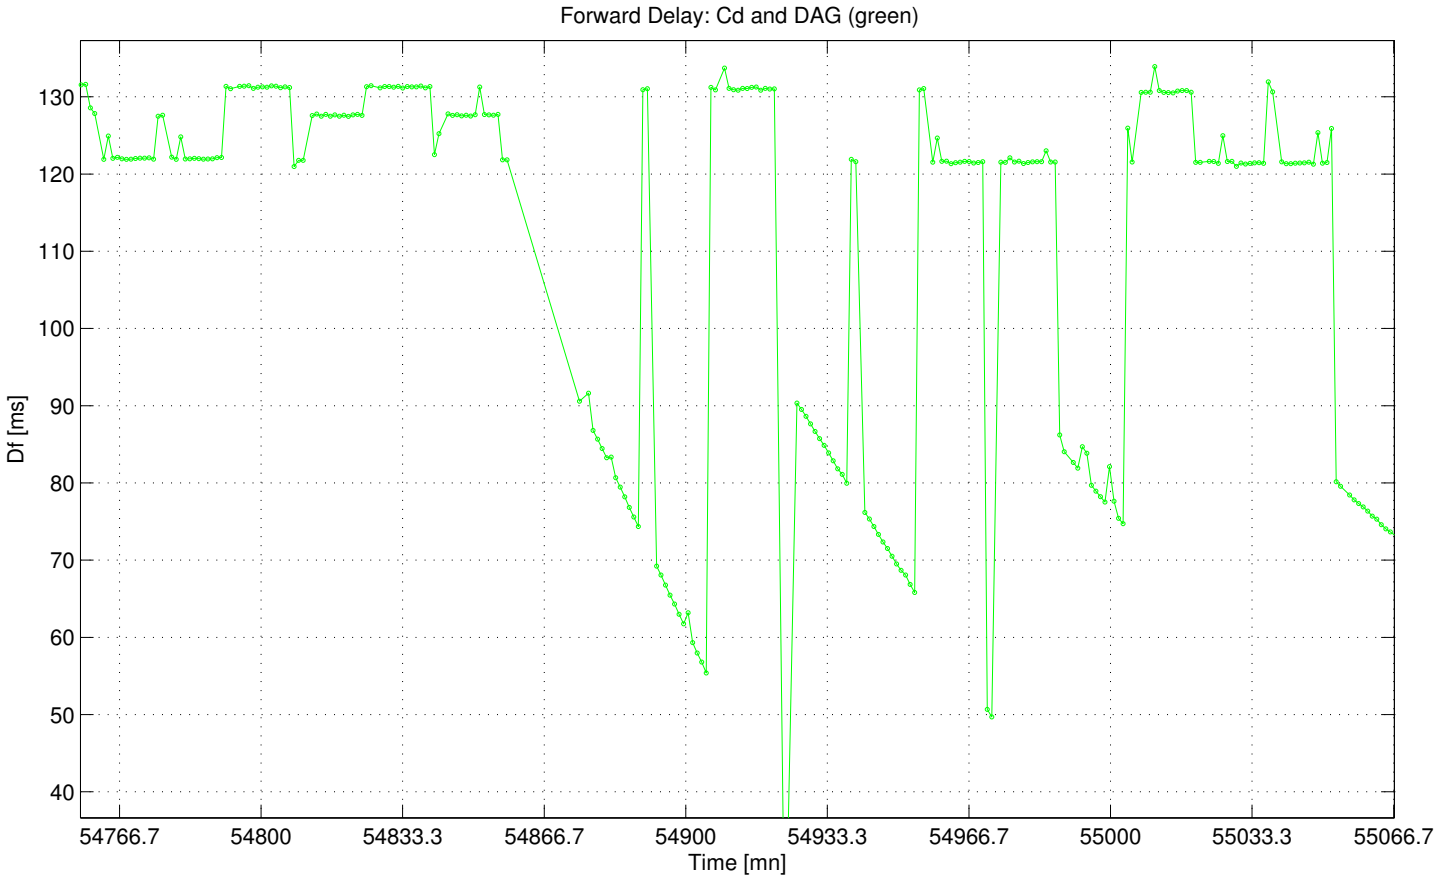

Client-Server OWD: Server shows frequent jumps in the 10-100ms range plus some regions with skew of one PPM. (meanwhile the RTT is very close to a constant.)
min(RTT) = 187 ms

Potential of Ark + RADclock

- **Coordinated distributed experiments (if one could trust timing)**
  - eg: Internet coordinate systems, route disambiguation
  - think of Spanner (Google's time aware distributed database)
- **Time Server vetting (can't trust stratum 1's)**
  - to better select Ark's own stratum 1's
  - as a service provided by Ark
- **Down the track**
  - RADclock servers as well as clients
  - more support to ease advanced use (like replay)

How Ark can Help Timing Research

- **Network Timing System (NTS) to replace NTP**
  - involves inter-linking strata 1 & 2, server recommendation, ..
  - Ark an ideal platform to develop and test NTP
  - requires some Ark monitors to be stratum-1
- **Network Timing Health Monitoring**
  - exploit Ark's vantage point diversity to perform wide scale vetting
  - detailed one-off studies, on-going monitoring of public infrastructure
- **Synergy**
  - each project helps the other
  - even one or two GPS-enabled Ark nodes a big benefit
  - each will improve Ark's timing service further

Support

- **FreeBSD**
  - patches for 8.1 (includes Zen support) + userland code
    - *ntpd* based system clock untouched
    - RADclock + *ntpd* system clock available in parallel
  - more mature version (but incomplete) adopted into 10.1
    - can select RADclock as the system clock transparently, plus extras
  - Expectation of full inclusion in 11
- **Linux**
  - *patches up to 2.6.32 available*
  - Raspberry Pi support almost there
- **Looking for partners**
  - to help push development along
  - to jointly address specific needs

Resources

- ***SyncLab website***
  - *http://www.synclab.org/radclock/*
  - *Papers*
  - *Patches*
- ***ACMQueue high level article***
  - *http://queue.acm.org/detail.cfm?id=1773943*
- ***Google TechTalk***
  - https://www.youtube.com/watch?v=o3nXgeh7v_U