DipZoom: A Marketplace for Internet Measurements

Michael Rabinovich, Sipat Tiukose, Zhihua Wen Limin Wang
EECS Department
Case Western Reserve University

Internet Measurements

A-priori measurement platforms (e.g. IDMaps)

- Great for large-scale characterizations

On-demand measurements (Keynote, Scriptroute/PlanetLab, traceroute servers)

- Hard to deploy sufficient platform to serve unpredictable needs
- Vulnerable to being “gamed”
- Limited choice of measurements

Example 1: Choosing a CDN

- CDN measurements
  - Chinese readership
  - With/without persistent connections
  - Repeat/new users
- Keynote limitations
  - Only one location in China
  - Not all measurements offered
  - Losing CDN cries foul

Example 2: Latency Measurements

- Measuring latency between a laptop connected through ***VZACCESS BroadbandAccess*** and a Linux PC on Case network
- RTT measured by average of 2000 pings: 280 msec

King Measurements

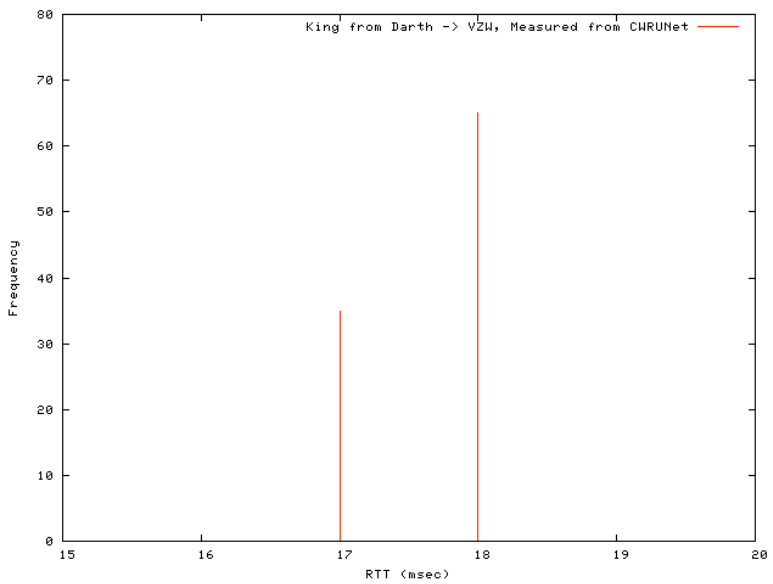

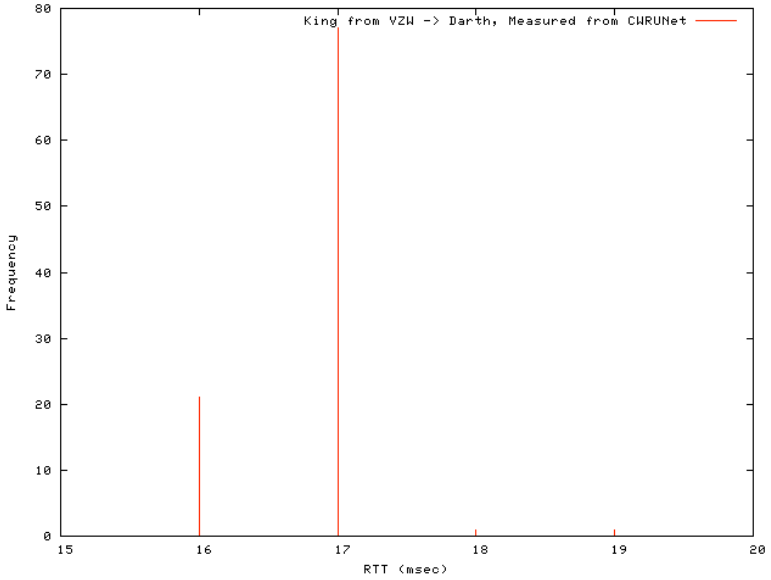

Needs:

- Focused on-demand measurements
- Infrastructure that scale with Internet
  - –Measuring host (MH) location
  - –MH type (platform, connectivity)
  - –Measurement type and regime

Our Approach:

- Coax Internet users to become measurement providers
- Deploy a *matchmaking service* instead of measurement infrastructure
- Use market approach with real money as the means to control the system

DipZoom: Deep Internet Performance Zoom

- Anyone can offer measurements
- Anyone can request measurements
- Anyone can offer measuring software
- Participants are free to set their prices, compete for requests, bit and solicit bids, etc.
- Facilitates open echosystem, “ebay for Internet measurements”

Some Questions

- Will anyone want to become a provider?
  - Seti@home, upromise.com, gomez.com suggest “yes”.
- Will anyone pay for the measurements?
  - Keynote and Gomez suggest “yes”

Related Work

- Gomez.com
  - Closed system
- DIMES, Traceroute@home
  - No incentives
  - Users participate in a particular measurement experiment
  - Users can't ask for a measurement

System Overview

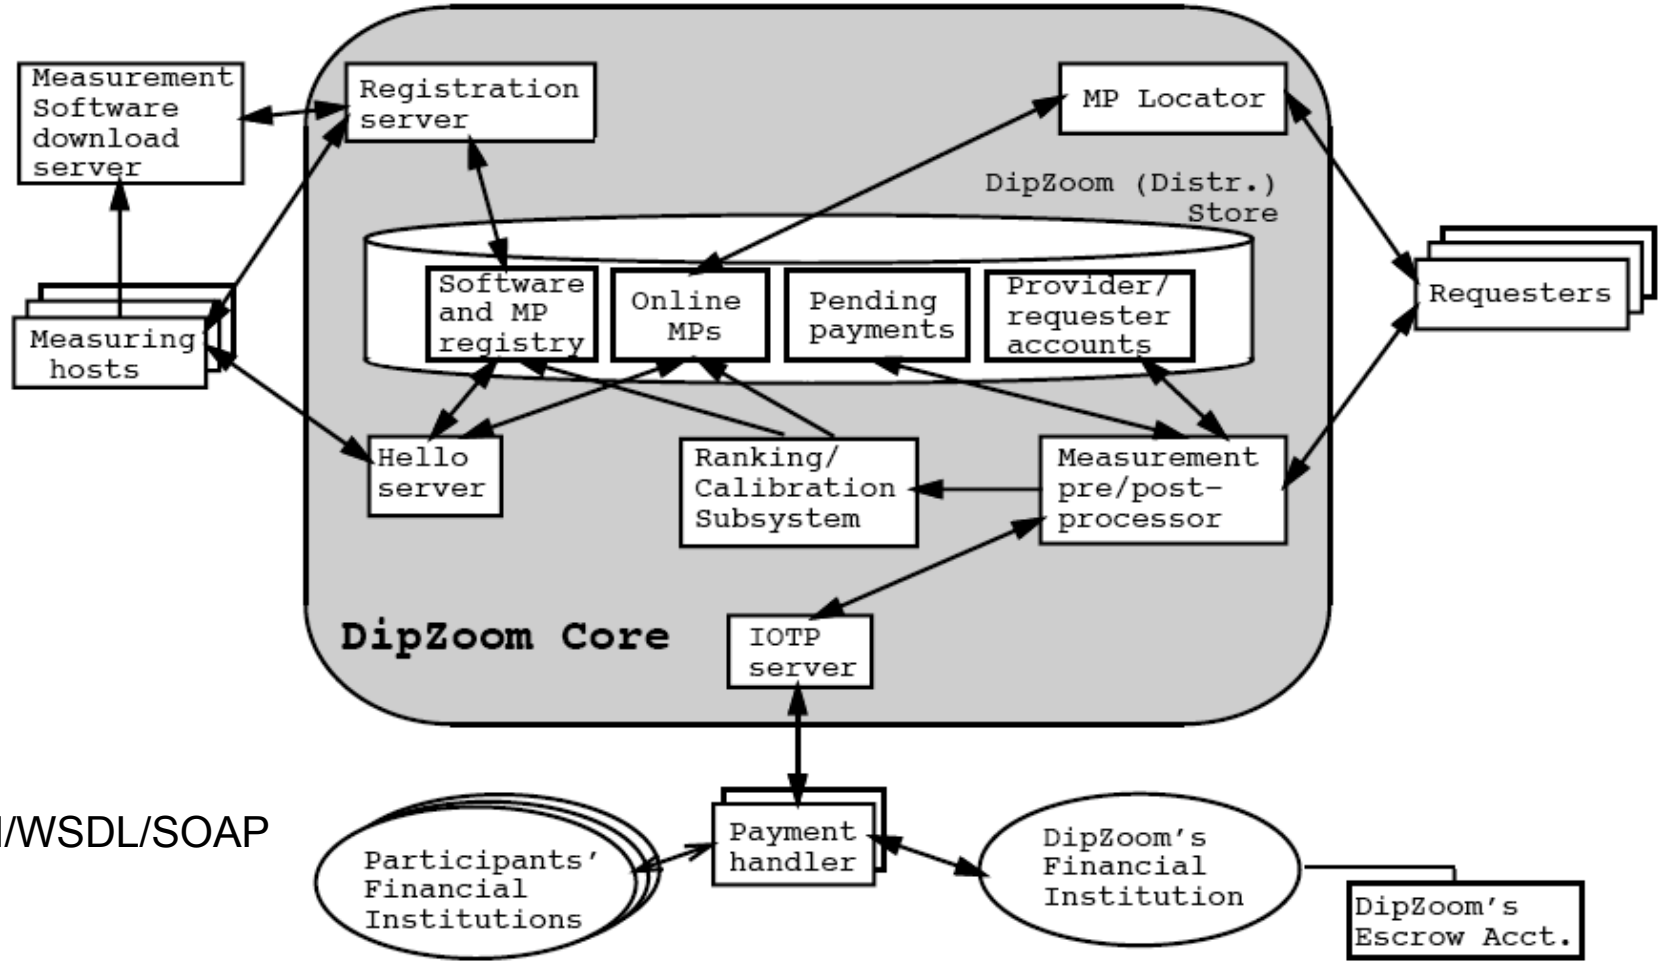

UDDI/WSDL/SOAP

SSL

IOTP

Issues

- Security
  - Induced DoS attacks against measurment target
  - Open ports on measuring hosts
  - High-level DoS attack against measuring host
  - Measurement side-effects
- Payment trust
  - Trusted core helps
  - Replay-based cheating
- Measurement trust
  - Fake MH registrations
  - MH impersonation
  - Fake measurements
- Traversing firewalls and NATs

Core Needs

- Integrity of measuring software
- Globally unique ID of measuring host (MHID)
- Duplicate detection + request/response matching
- Measurement rate limiting

- Building blocks of a solution:
  - –Unique embedded secret
  - –GUID or MAC address or hostID
  - –Nonces
  - –Ranking and calibration

DipZoom Request Credential

- DipZoom core returns an encrypted credential *<MHID, nonce>* with response to requester's query
- Requester includes the credential with request
- Measuring host
  - Decrypts nonce
  - Modifies it using a well known operation (nonce + 1)
  - Return encrypted nonce with response
  - Caches nonces for early duplicate detection
- Nonce/modified nonce addresses request replay and response replay, and third-party response replay attacks
- MHID addresses the random nonce attack

Ranking and Calibration

- Security measures raise the bar but do not guarantee protection from malicious MH.
- If can't protect - detect and blacklist!
  - Deploy calibrating measurement targets
  - Purchase measurements from suspect MHs
  - Compare responses with passive measurements by calibrating hosts
- Can calibrating hosts be gamed?
  - Keynote advertises its measuring hosts location
  - Calibrating hosts are secret
  - The risk of blacklisting deters data mining

Status

- Pre-alpha is hereby released!
  - http://www.eecs.case.edu/~sxt85/dipzoom/index.html
  - No payments yet
  - Includes NAT/firewall traversal
  - Either measuring software or a client and MH bundle
  - Just ping and wget for now

Summary

- Growing Internet diversity (devices, links, applications) entails growing needs for focused measurements
- Proprietary platforms are insufficient
- DipZoom: a facilitator instead of infrastructure
  - Open system (pricing, measurements, participants)
  - Market approach to system control
  - Based on P2P principles
- Many challenges ahead

Temporary Web site for more information:
http://www.eecs.case.edu/~sxt85/dipzoom/index.html