LARGE-SCALE INTERNET MEASUREMENTS FOR DIAGNOSTICS AND PUBLIC POLICY

Henning Schulzrinne

(+ Walter Johnston & James Miller)

FCC & Columbia University

Overview

- Quick overview
- What does MBA measure?
- Can we build measurement into the infrastructure
- Security and privacy issues

The role of network measurements

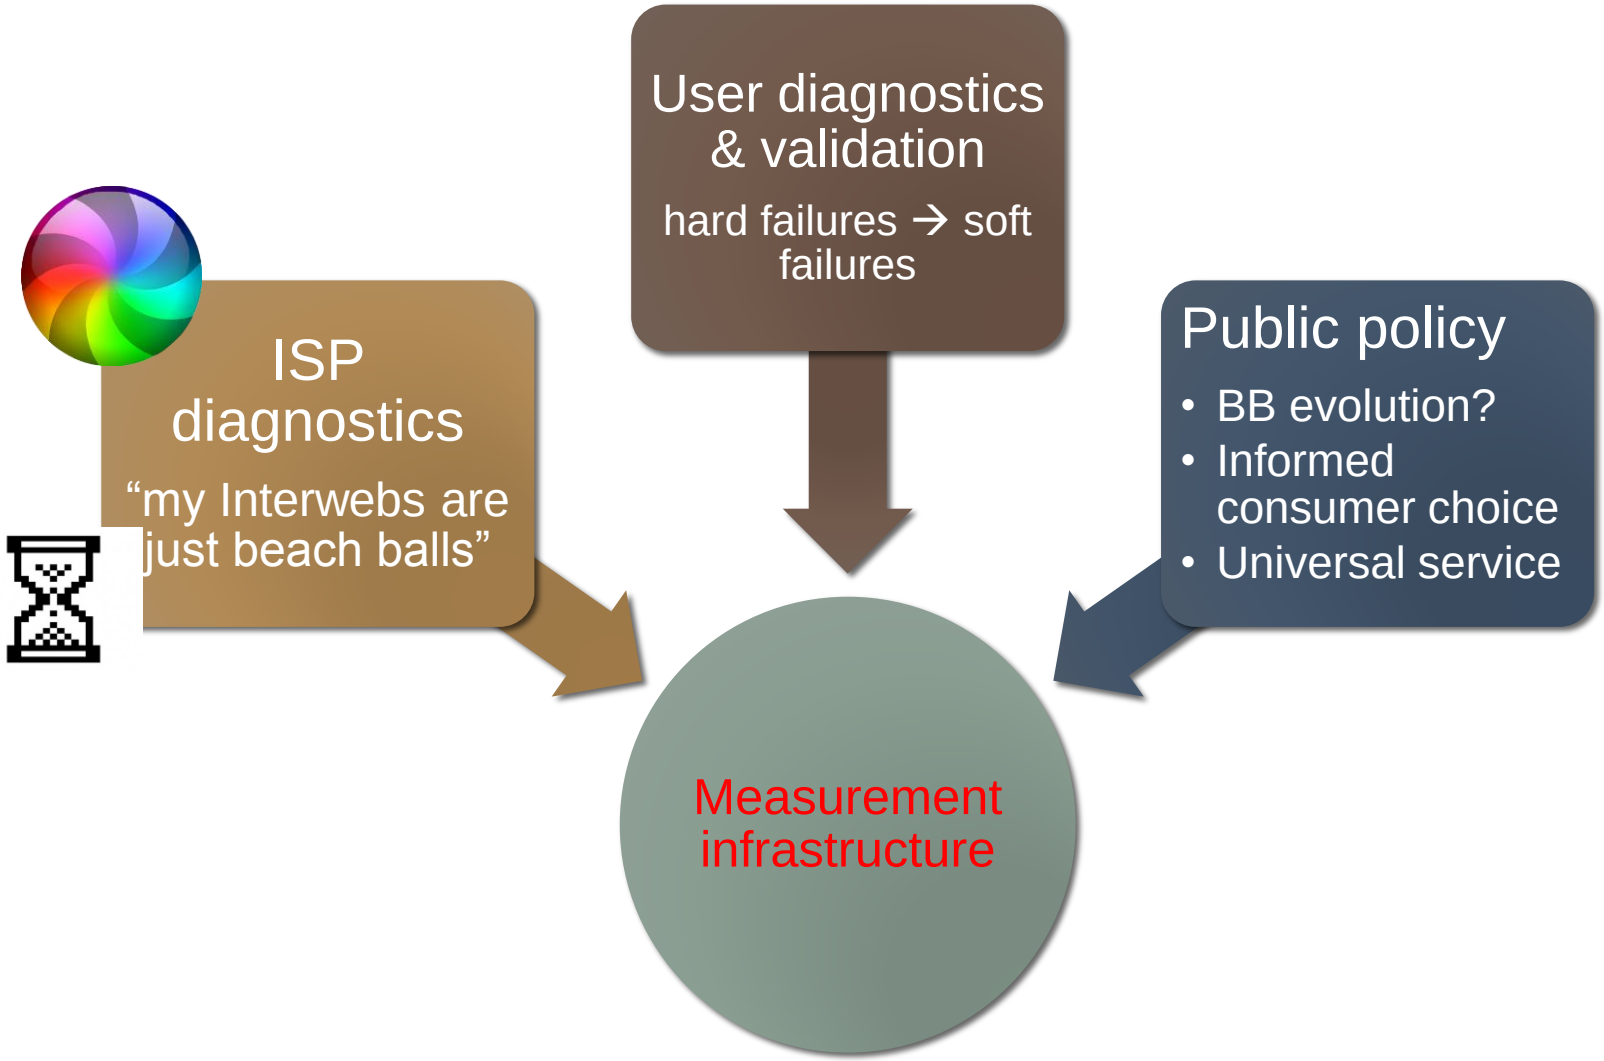

Measurement architecture

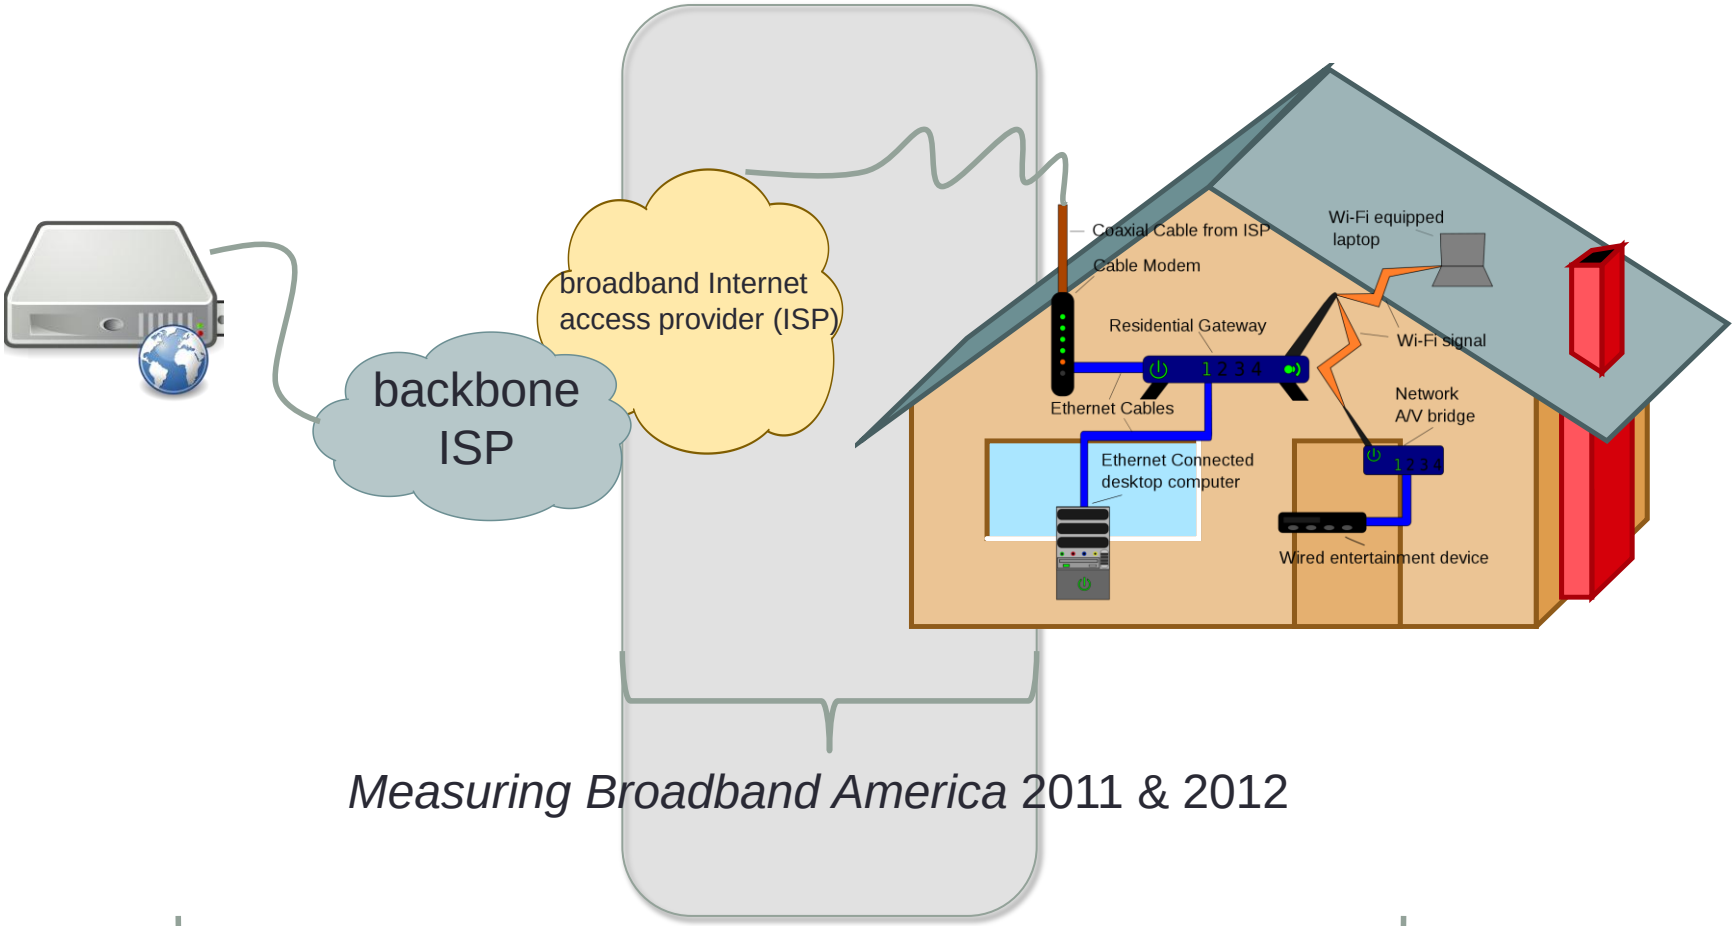

What was measured

Sustained Download	Burst Download
Sustained Upload	Burst Upload
Web Browsing Download	UDP Latency
UDP Packet Loss	Video Streaming Measure
VoIP Measure	DNS Resolution
DNS Failures	ICMP Latency
ICMP Packet Loss	Latency Under Load
Total Bytes Downloaded	Total Bytes Uploaded

L3/L4 measurements

Metric	Description
ICMP latency	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
ICMP packet loss	Percentage of packets lost in the ICMP latency test
UDP latency	Average round trip time of a series of randomly transmitted UDP packets distributed over a long timeframe
UDP packet loss	Fraction of UDP packets lost from UDP latency test
TCP download speed	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
TCP upload speed	Throughput in Mbps utilizing three concurrent TCP connections
Latency under load	Average round trip time for a series of regularly spaced UDP packets sent during downstream/upstream sustained tests

DNS measurements

Metric	Description
DNS resolution	Time taken for the ISP's recursive DNS resolver to return an A record for a popular website domain name
DNS failures	Percentage of DNS requests performed in the DNS resolution test that failed

Application measurements

Metric	Description
Web browsing	Total time to fetch a page and all of its resources from a popular website
Video streaming	Initial time to buffer, number of buffer under-runs and total time for buffer delays
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
Availability	Total time the connection was deemed unavailable for any purpose, which could include a network fault or unavailability of a measurement point
Consumption	A simple record of the total bytes downloaded and uploaded by the router

Application measurements

- Goal: rough estimate of consumer experience
  - relative ordering: “if metric is significantly better for ISP A than B, consumer experience should be similarly ordered”
- Non-goal: predict performance of specific browser or video viewer
- Try to capture latency and variability effects on QoE

Example: VoIP

- The Voice over IP (VoIP) test operated over UDP and, unlike the video streaming test, utilized bi-directional traffic, as is typical for voice calls.
- The Whitebox would handshake with the server, and each would initiate a UDP stream with the other. The test used a 64 kbps stream with the same characteristics and properties (i.e., packet sizes, delays, bitrate) as the G.711 codec. The test measured jitter, delay and loss. These metrics were measured by subdividing the stream into blocks, and measuring the time taken to receive each block (as well as the difference between consecutive times).
- Jitter was calculated using the PDV approach described in section 4.2 of RFC5481. The 99th percentile was recorded and used in all calculations when deriving the PDV.

Example: Web browsing

- The test recorded the averaged time taken to sequentially download the HTML and referenced resources for the home page of each of the target websites, the number of bytes transferred, and the calculated rate per second. The primary measure for this test was the total time taken to download the HTML front page for each web site and all associated images, JavaScript, and stylesheet resources. This test did not test against the centralized testing nodes; instead it tested against real websites, ensuring that the effects of content distribution networks and other performance enhancing factors could be taken into account.
- Each Whitebox tested against the following 9 websites:
  - http://www.cnn.com
  - http://www.youtube.com
  - http://www.msn.com
  - http://www.amazon.com
  - http://www.yahoo.com
  - http://www.ebay.com
  - http://www.wikipedia.org
  - http://www.facebook.com
  - http://www.google.com
- The results include the time taken for DNS resolution. The test used up to eight concurrent TCP connections to fetch resources from targets. The test pooled TCP connections and utilized persistent connections where the remote HTTP server supports them.
- The client advertised the user agent as Microsoft Internet Explorer 8. Each website was tested in sequence and the results summed and reported across all sites.

LMAP proposed architecture

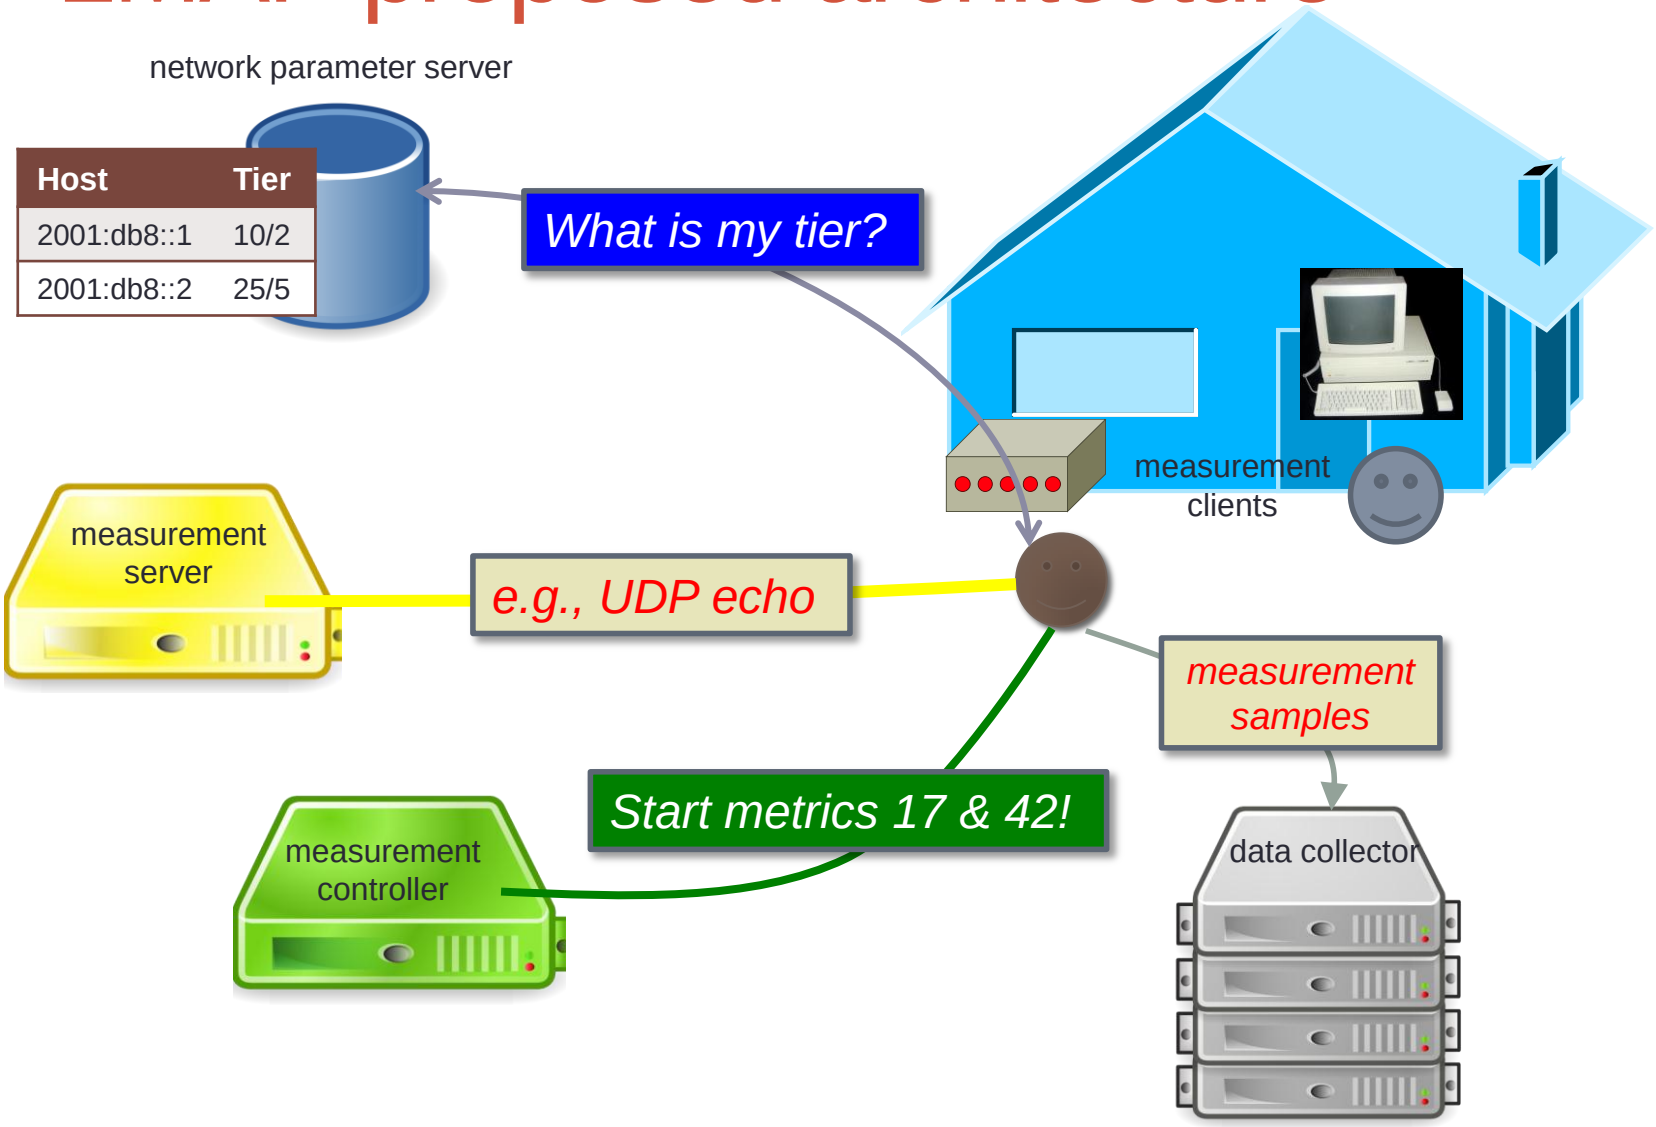

What's needed?

- Define higher-layer "QoE" metrics
  - active measurements for privacy
- Protocols for overall architecture
  - See also TR-69

Security and privacy issues

- Security
  - Prevent massive bot nets → authentication of controller
  - Logging and disclosure to consumers
  - Permission architecture – legal issue as much as technical:
    - explicit permission: “fcc.gov may request measurements”
      - measurement volunteers
      - informed consent standard
    - ISP contract: “You agree to let our service technician run tests”
    - explicit consumer request: “test my performance”
- Privacy concerns
  - Active measurements relatively safe
  - Do not measure when user traffic is active

Conclusion

- Measurements: from sometime experiment to built-in facility
- Good telecom policy needs good data
  - not just counting lines
  - PSTN transition to IP → there is no second network
- Re-use measurement for three purposes:
  - ISP diagnostics and planning
  - Consumer diagnostics
  - Public policy data gathering

BACKUP