Measurements close to users

Anna-Kaisa Pietilainen, Renata Teixeira
MUSE Team
Inria Paris-Rocquencourt

Network performance disruptions are frustrating

For users

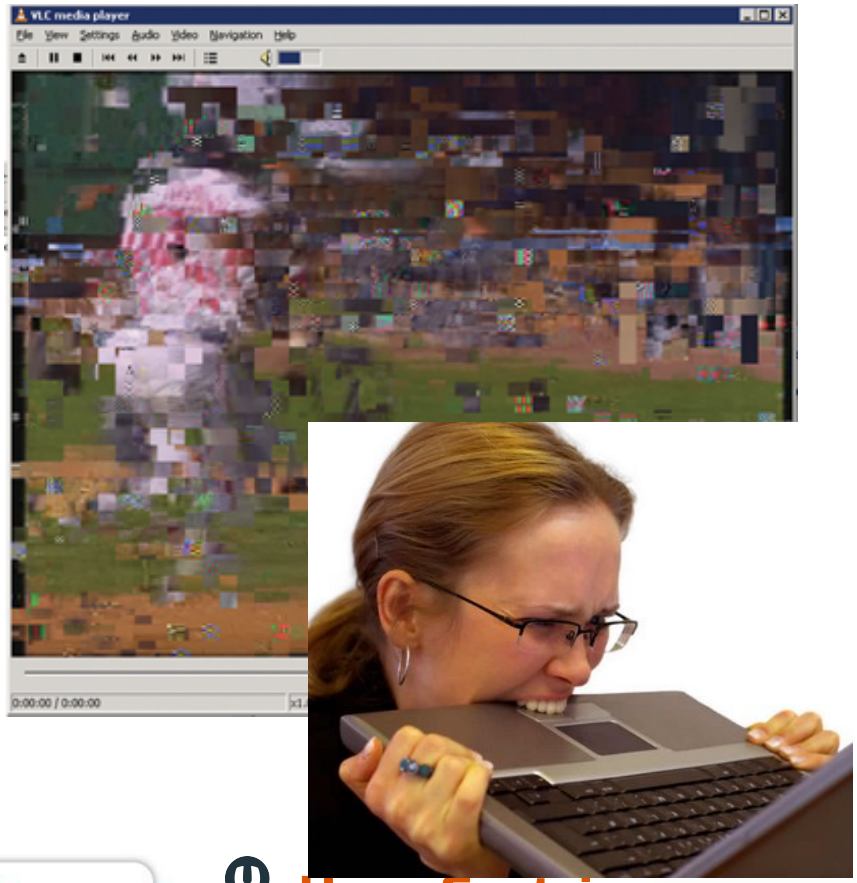

For ISPs

Home networks can cause performance disruptions

- Cross-traffic competes for bandwidth
- Large buffers and heavy uploads increase delays
- Poor WiFi increases jitter and reduce bandwidth
  - Poor placement of access point
  - Interference from other access points
  - Contention from other devices
  - Non-Wifi interference (e.g., microwaves, baby monitors)

Goal

Assist users to diagnose performance problems in the home network

- Automatic detection: Is there a problem?
  - Focus on performance disruptions that affect users
- Problem identification: where is the problem?
  - More detailed diagnosis when problem is local

Outline

- User experience of network performance
  - Measuring network performance close to users
  - Correlating with user experience
- Home network performance: Home vs. Access
  - Measurement vantage point: end-host vs. gateway
- Fathom: browser-based measurement platform

Approaches to measure performance close to users

- Active probing
  - Based on issuing probes, analyzing response
- Passive analysis of user's traffic
  - Tap incoming and outgoing traffic: tcpdump, pcap
  - Monitor status of TCP connections

RTT from active probes: ping

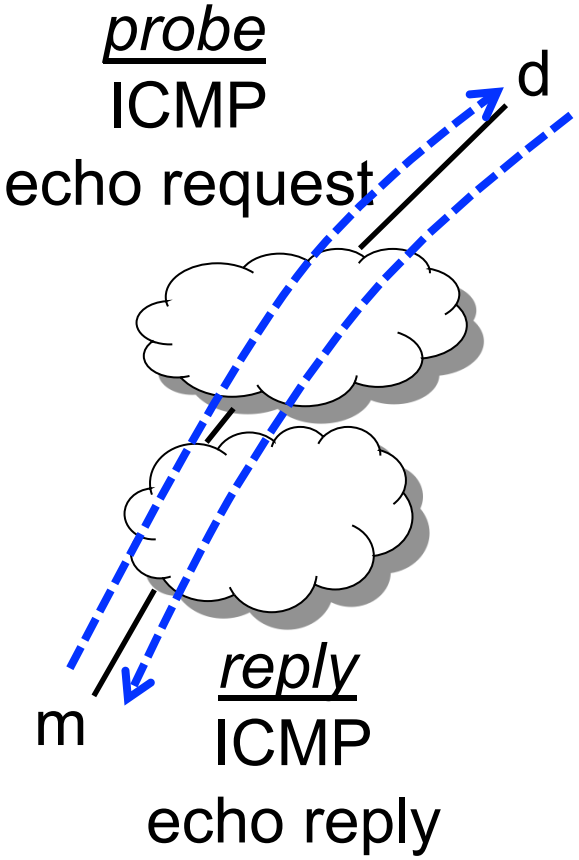

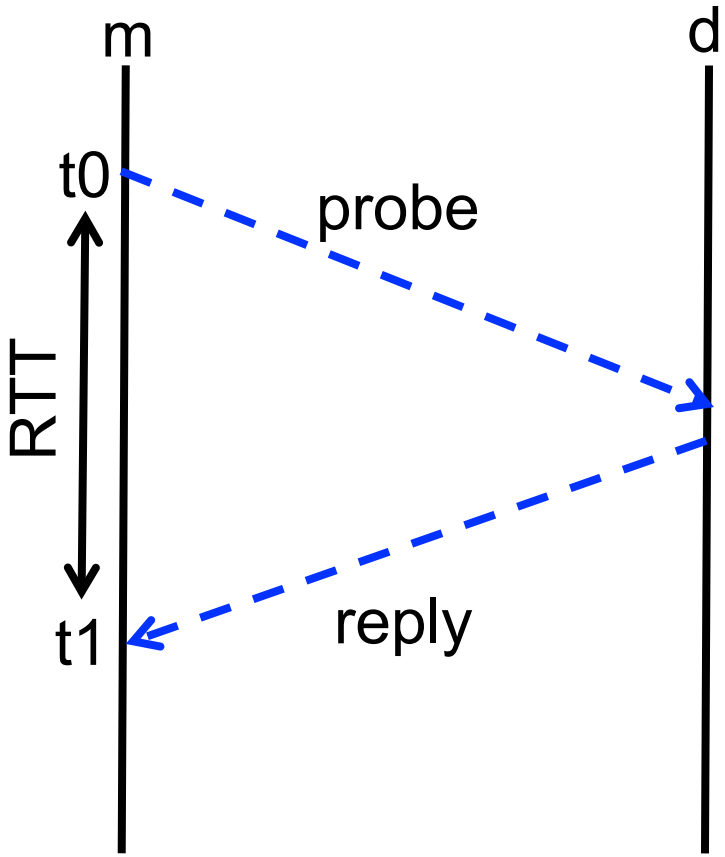

RTT from passive measurements: tcptrace

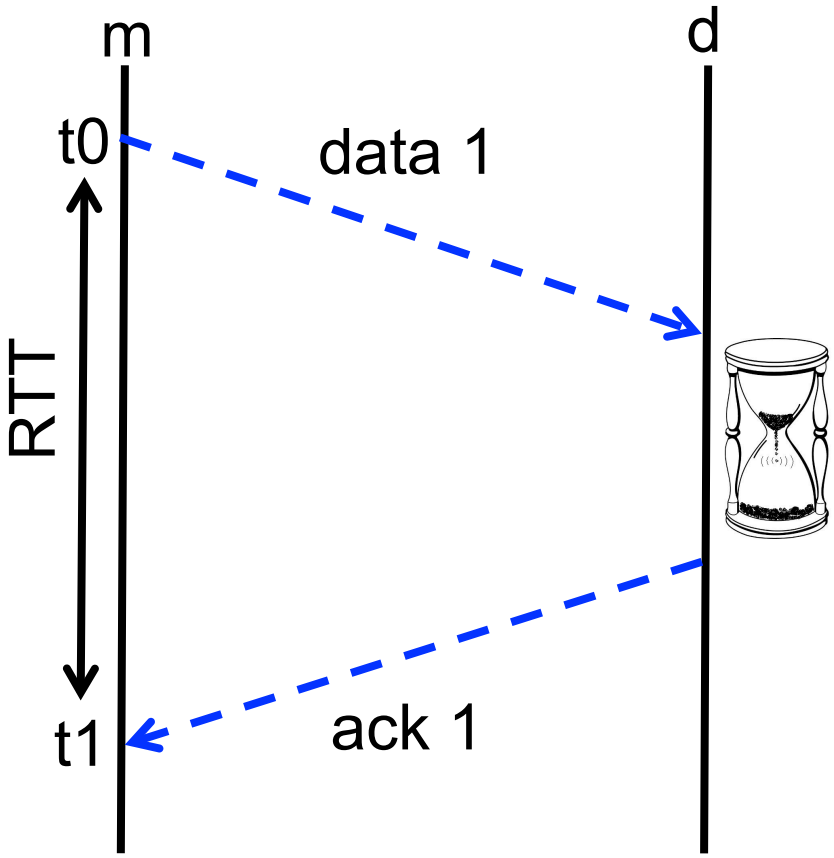

Other end-to-end performance metrics

Metric	Active	Passive
Loss	ping/iperf	TCP retransmissions
Throughput	iperf	TCP/UDP data rates
Delay variation/jitter	iperf	Difference between RTTs
Available bandwidth	pathload, spruce	
Capacity	ShaperProbe iperf UDP	

- More metrics
  - IETF IP Performance Metrics Working Group
- More tools
  - http://www.caida.org/tools/taxonomy/performance.xml
  - http://www.measurementlab.net/

Summary: passive vs. active

Passive

+ No need to inject traffic
+ Measures performance experienced by users
+ Measures destinations that don't respond to probes

– Privacy concerns
– Collection overhead
– Only measures paths with traffic

Active

+ No need to tap user's traffic
+ Measure performance of paths even without traffic
+ Often used for diagnosis

– Not direct measure of user experience
– Probing overhead
  – Cover a large number of paths
  – Continuous measurements

Outline

- User experience of network performance
  - Measuring network performance close to users
  - Correlating with user experience
- Home network performance: Home vs. Access
  - Measurement vantage point: end-host vs. gateway
- Fathom: browser-based measurement platform

Challenges in measuring user perception

- User perception varies
  - Per user, per environment, per application
  - For a given user according to external factors
  - Controlled environment versus field
- Can't ask frequent user feedback
  - At most ~10 per day
  - Orders of magnitude more network measurements (every millisecond)

Approaches to obtain user feedback

- Offline: out-of-bad feedback
  - Interviews, diaries
  - Pro: detailed feedback
  - Con: infrequent feedback; hard to correlate with network metrics
- Online: Integrated in measurement tool
  - System triggered, user triggered
  - Pro: more frequent feedback; automation is easier
  - Con: feedback can be harder to interpret

Online user feedback

- Which questions to ask?
  - Easy to fill, not to annoy users
  - Enough information to interpret results
- When to ask the questions?
  - User triggered: depends on user
  - System triggered: Experience sampling mechanism
    - Cover diverse levels of network performance

Example: HostView

- A data collection tool for laptops (Mac OS / Linux)
- Mixed methodology
  - Network traces
  - Application process names
  - Machine metrics
  - User feedback
- Deployment (Nov 2010 – Feb 2011)
  - 40 users (14 countries)
  - Most users ran tool for one month

HostView: User feedback

- System Trigged feedback
  - Experience sampling methodology (ESM)
  - Triggered based on state of machine
  - 5 short questions about network performance
  - At most 3 times a day
- User Triggered feedback
  - “I’m annoyed” button ☹
  - Same questions as in ESM
  - Can trigger as often as user wants

HostView: Example question

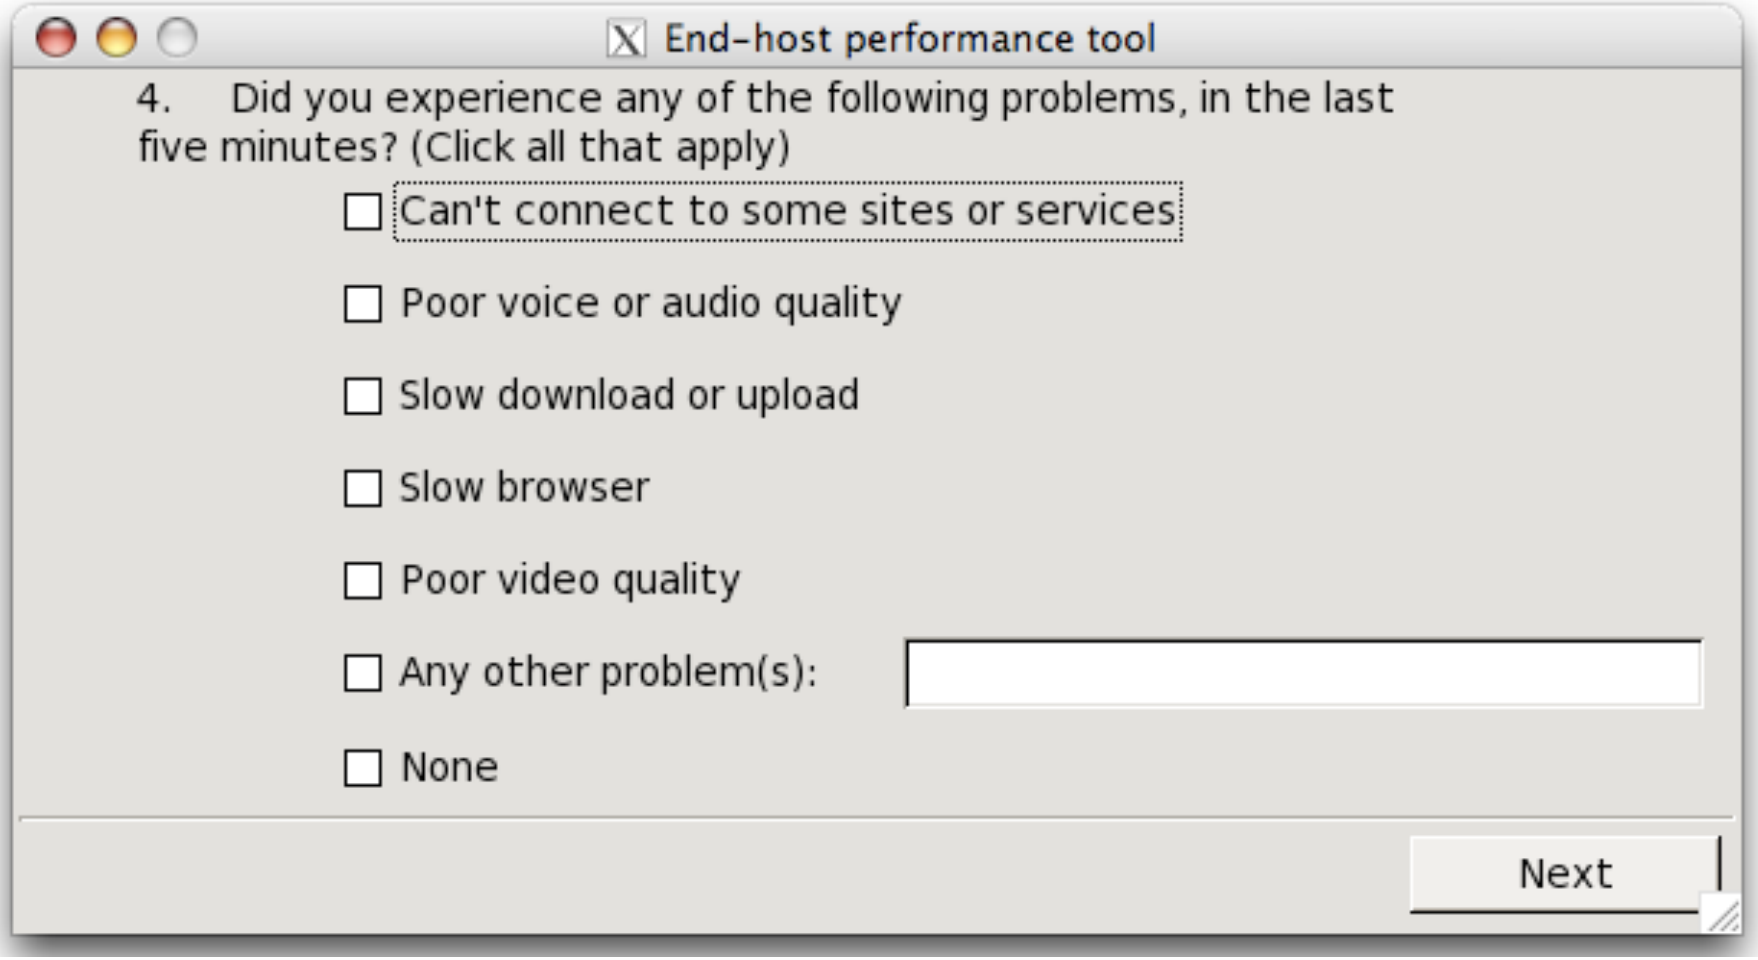

User vs. network reporting

- User perspective
  - Good/poor performance according to the user
- Network and system perspective
  - Good/poor performance according to network metrics
- Question: Do these co-occur?

Can’t connect to some sites or services

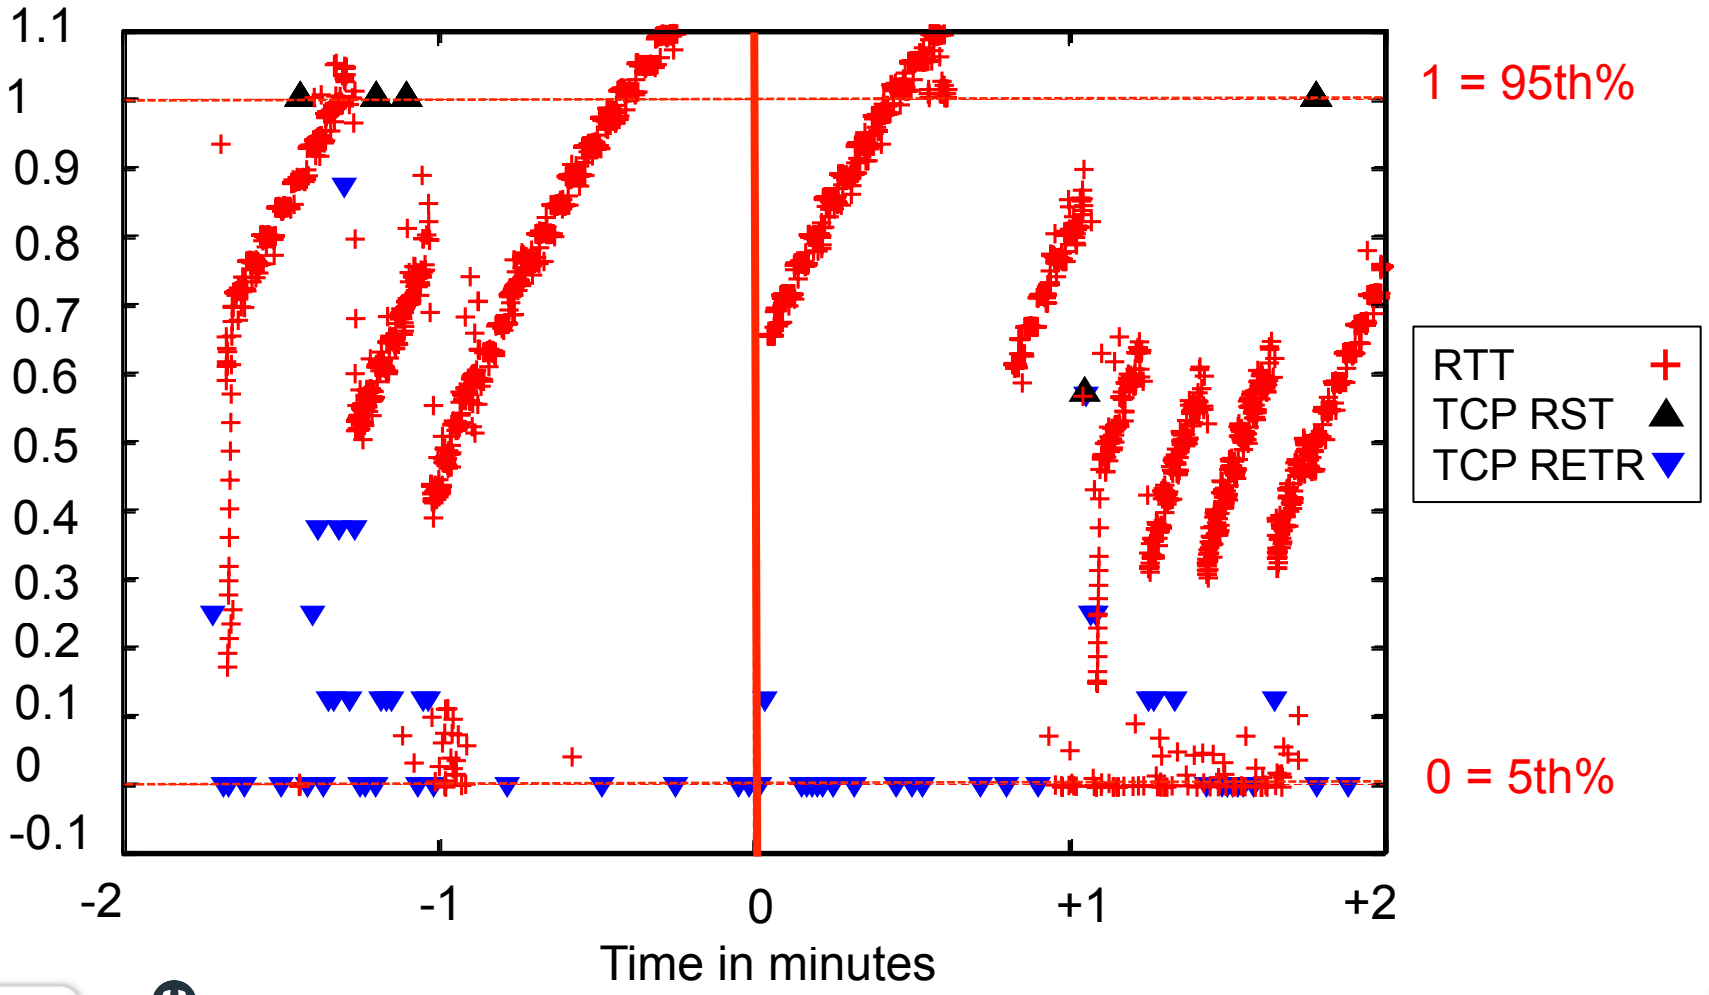

Everything is good!

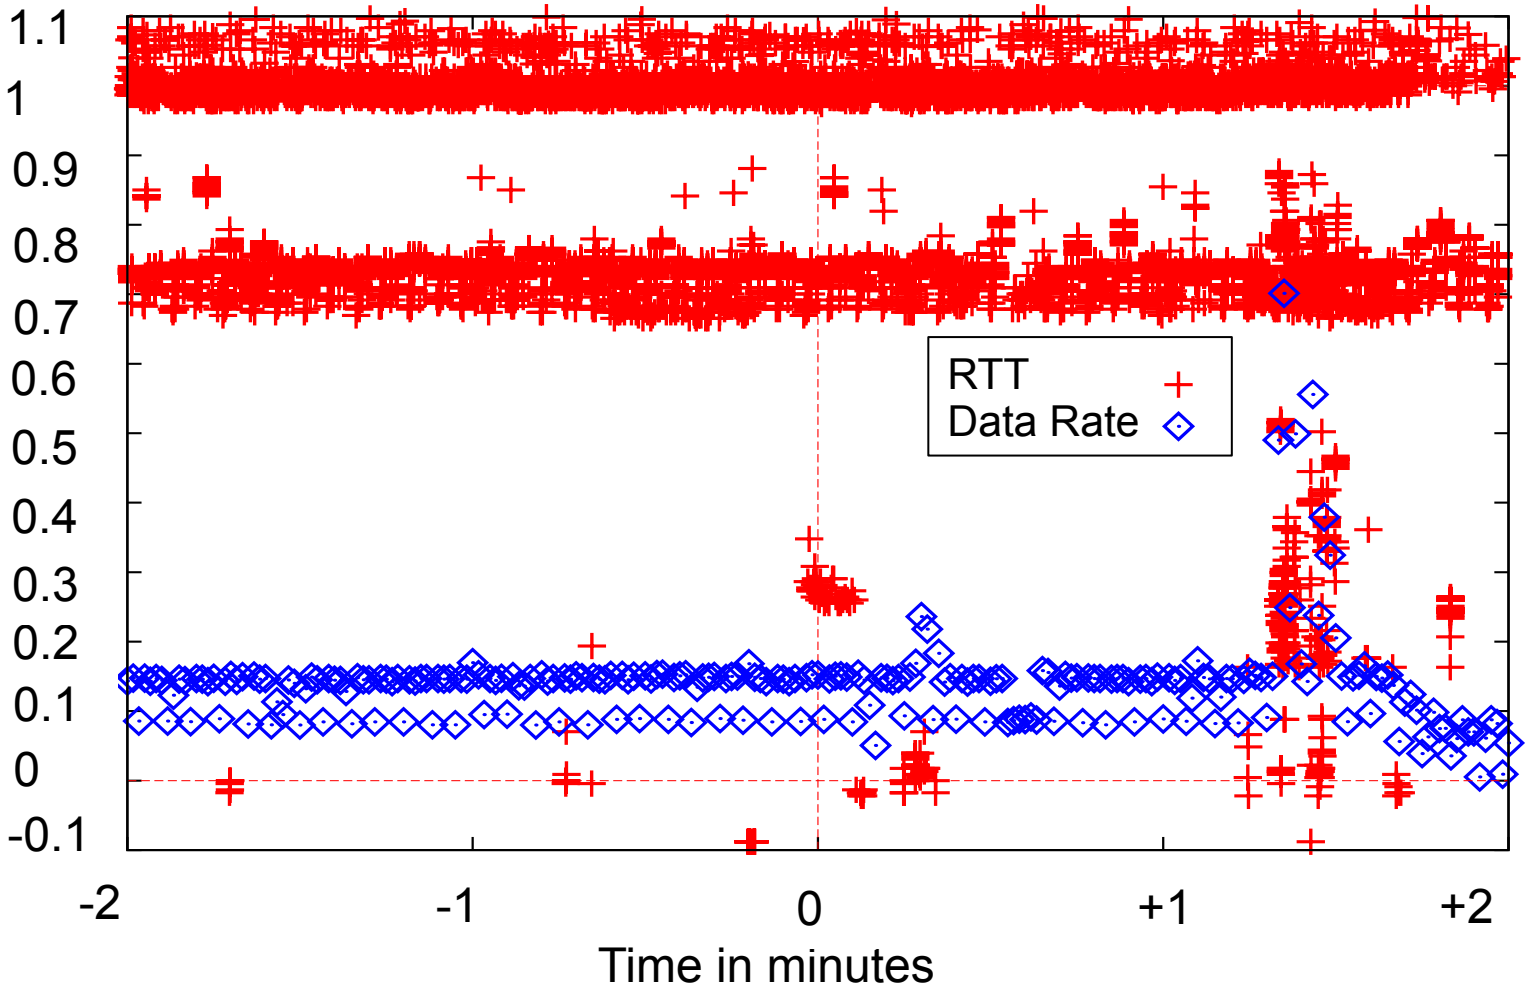

Summary: correlating user feedback with network performance

- Hard to get feedback from users
  - Many network performance samples without feedback
  - Users are diverse in how they report a problem
- Raw network metrics alone are not enough
  - Not all outliers affect the user perception

Outline

- User experience of network performance
  - Measuring network performance close to users
  - Correlating with user experience
- Home network performance: Home vs. Access
  - Measurement vantage point: end-host vs. gateway
- Fathom: browser-based measurement platform

What is the speed of my access link?

Home Network: AT&T DSL

6 Mbps Down, 512 Kbps Up

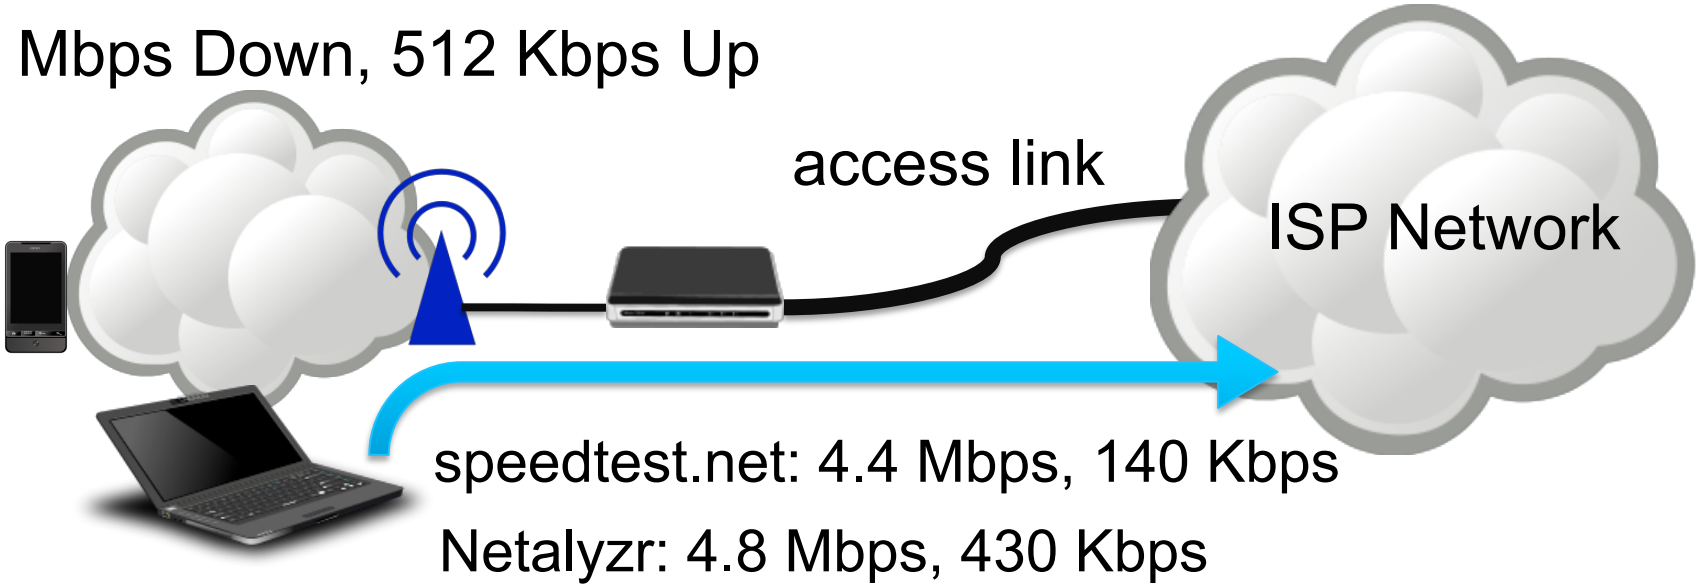

End host measurements are affected by *confounding factors*

Gateway better captures speed of access link

Home Network: AT&T DSL

6 Mbps Down, 512 Kbps Up

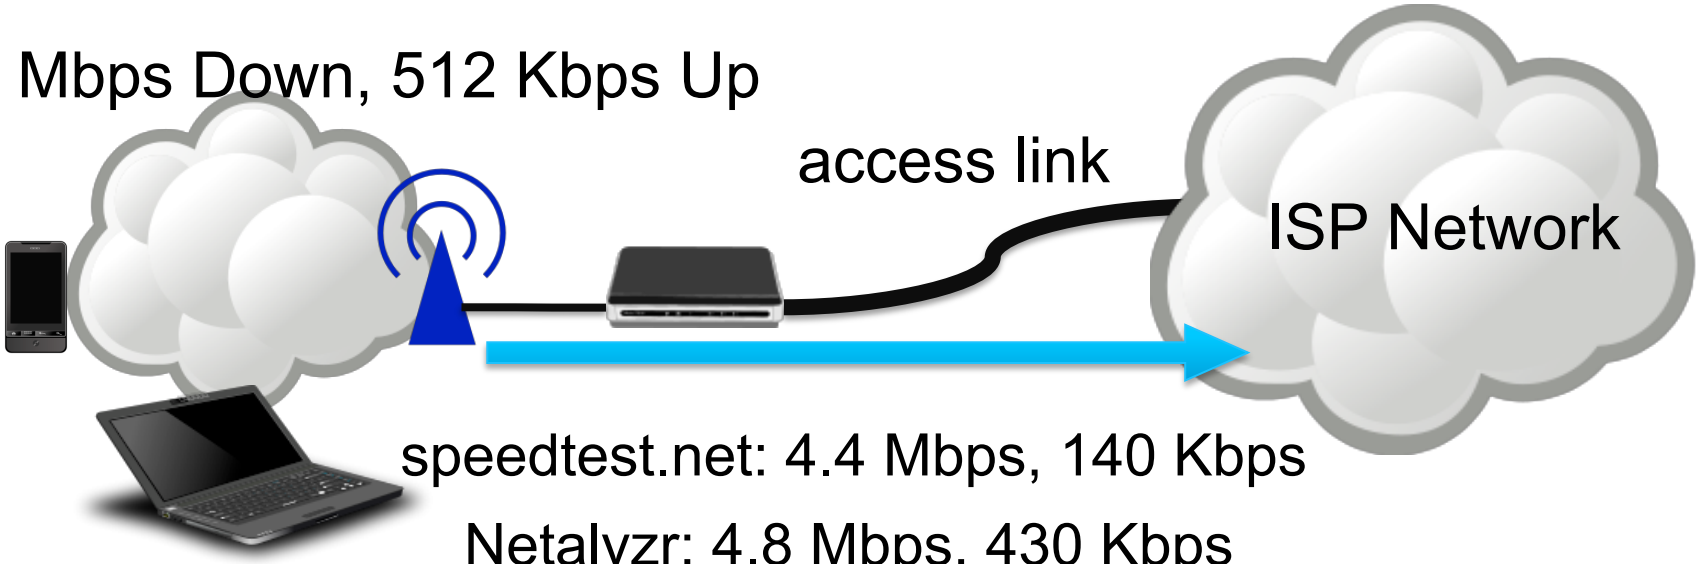

speedtest.net: 4.4 Mbps, 140 Kbps

Netalyzr: 4.8 Mbps, 430 Kbps

Gateway: 5.6 Mbps, 460 Kbps

Gateway can account for confounding factors

Gateway deployments

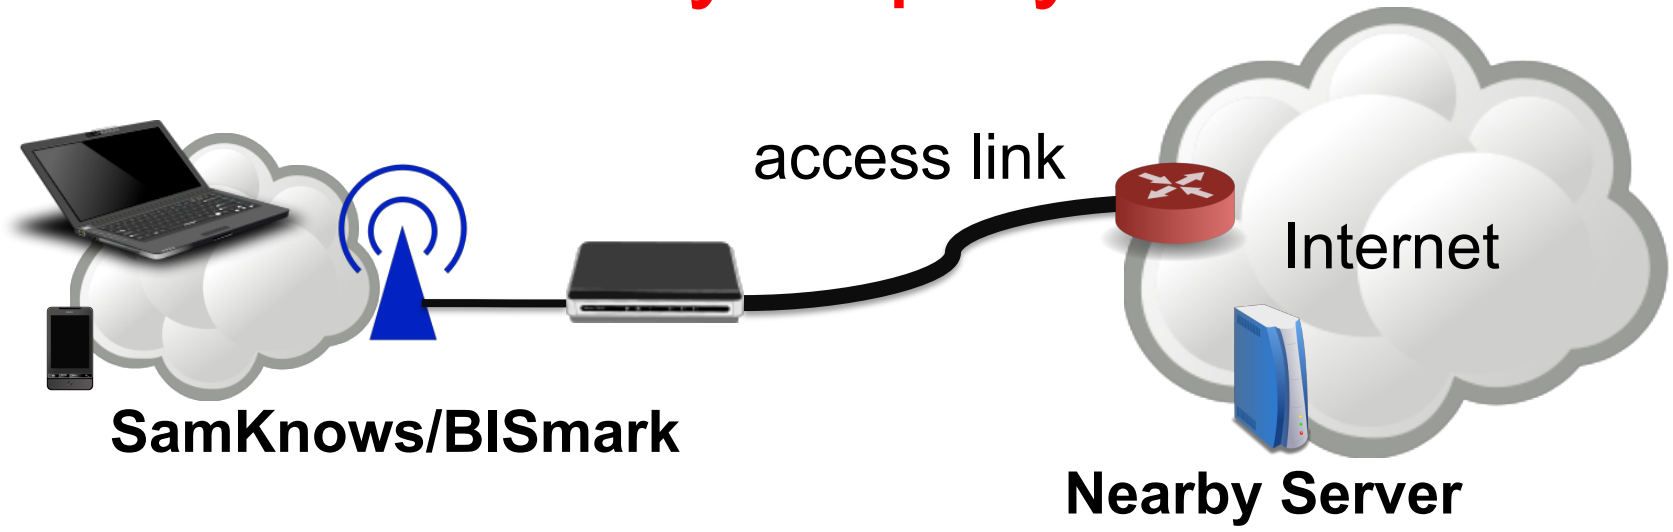

- SamKnows
  - Active measurements: throughput, delay, web performance, etc.
  - FCC deployment: ~10,000 gateways
- BISmark
  - OpenWRT router modified to perform active/passive measurements
  - Georgia Tech deployment: ~100 gateways

Interpreting throughput results

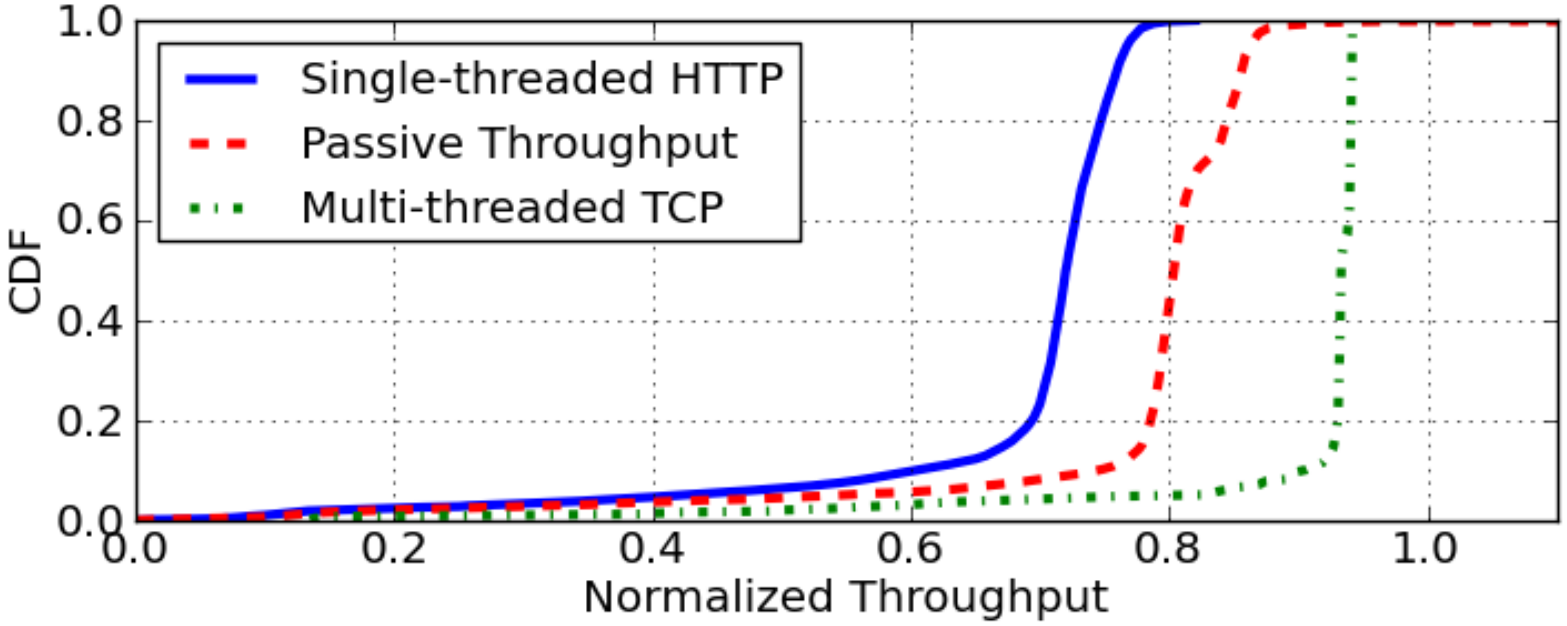

Different techniques measure different aspects of throughput

Summary: Gateway vs. end-devices

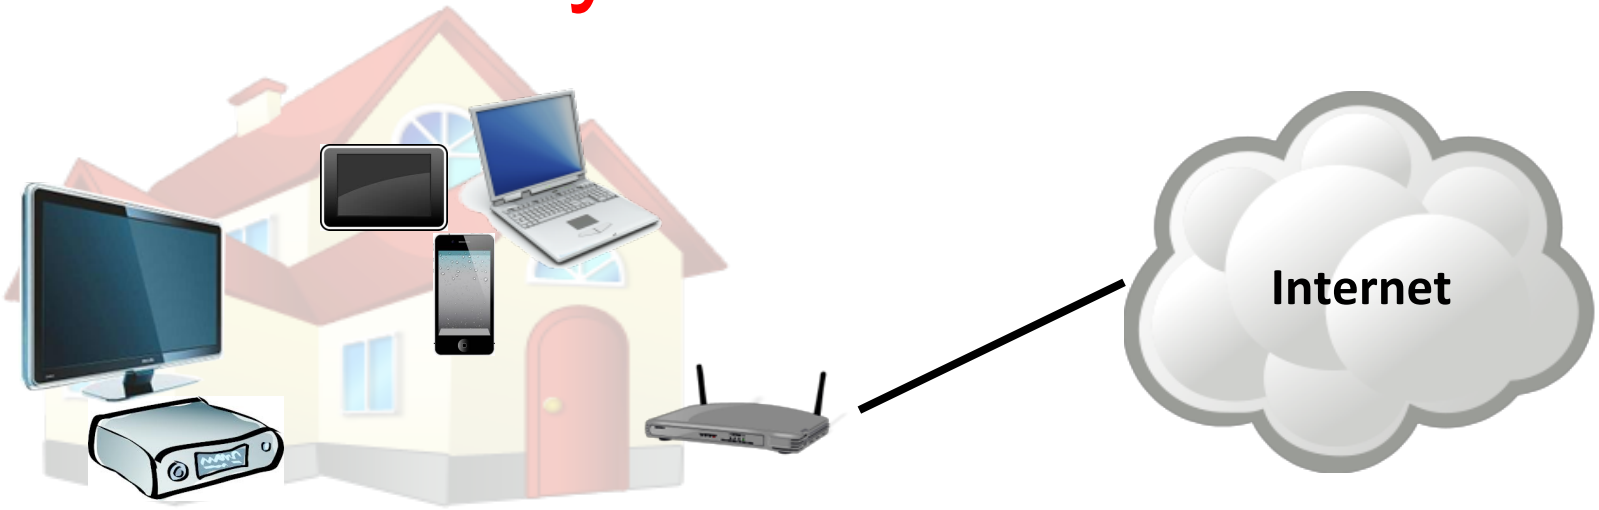

- Home gateway
  - Ideally placed between home devices and Internet
  - But, have limited resources and deployment is harder
- Instrument end-devices
  - Observe poor user experience
  - But, have limited view of home network and development is harder

Outline

- User experience of network performance
  - Measuring network performance close to users
  - Correlating with user experience
- Home network performance: Home vs. Access
  - Measurement vantage point: end-host vs. gateway
- Fathom: browser-based measurement platform

End-host measurements are challenging

- Measurement from end-hosts are vital
  - Researchers to understand Internet
  - Practitioners to diagnose user problems
- Hard to deploy measurements
  - Developers: Portability, safety
  - Users: need to install new software

A browser-based measurement platform

- Why browser?
  - Flexibility, deployability
  - Ubiquity of browser
- Fathom: Firefox extension
  - Measurement API in JavaScript
  - Web page performance
  - System performance
  - Active measurements