
Active Measurement of Home Broadband Internet using SamKnows

Fraida Fund, NYU and U Thessaly
Sam Crawford, SamKnows and LEONE
Scott Kirkpatrick, HUJI
Martin Potts, LEONE and Martel

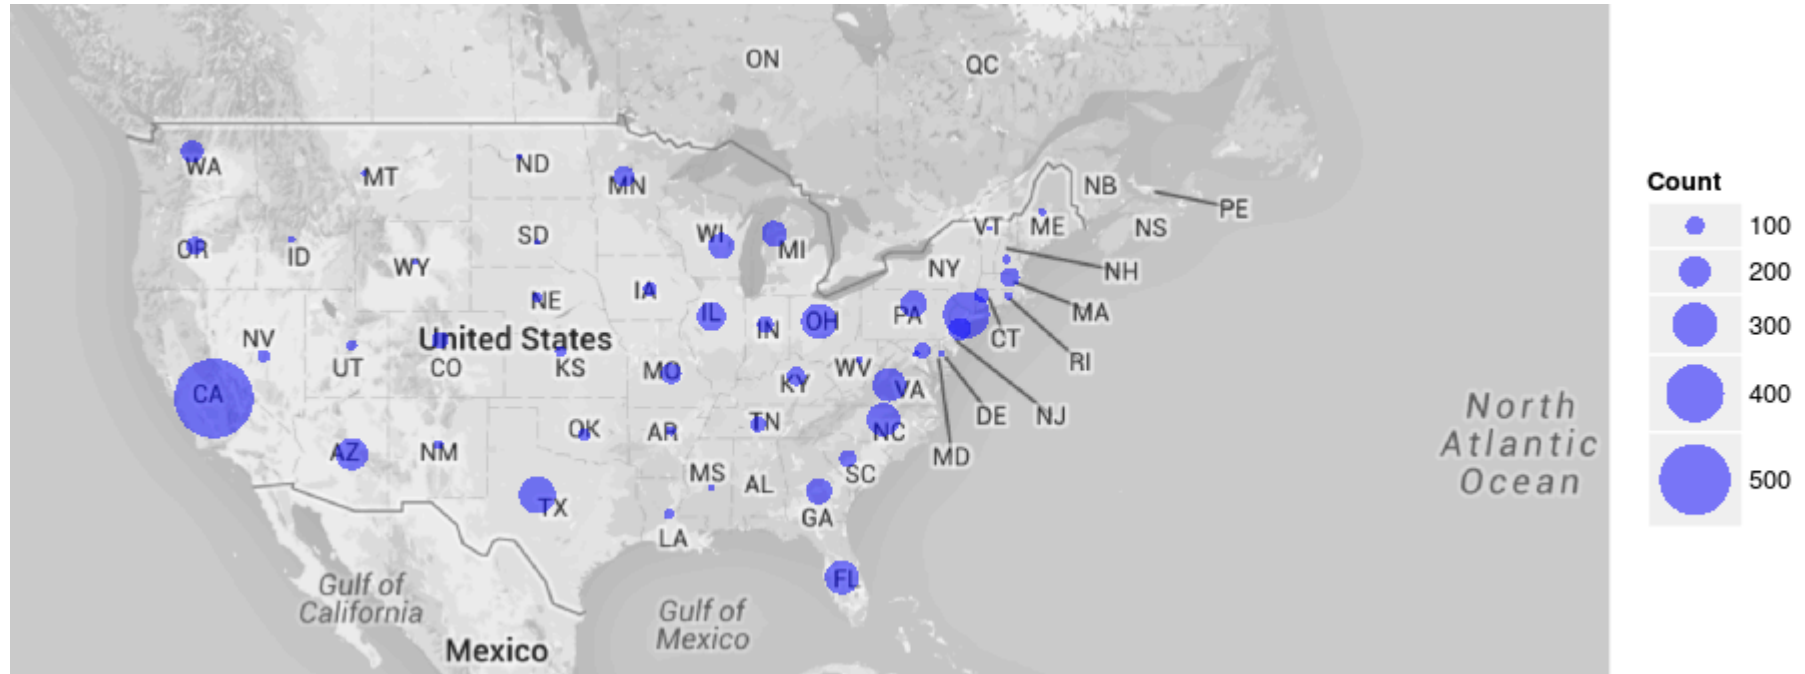
Measurements from >10,000 homes

US and EU

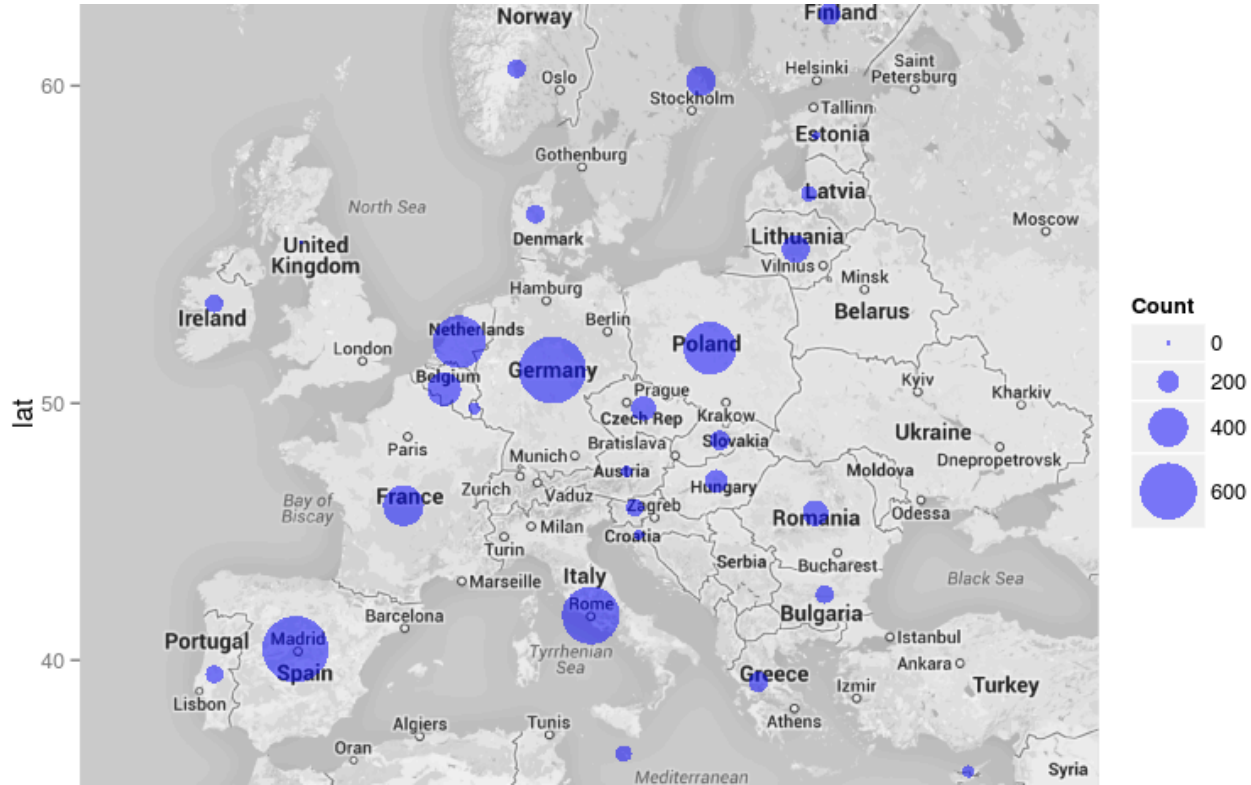
- Network metrics: download, upload, latency
- User experience metrics: DNS resolution, web page load
- **US only**: usage

(at different intervals, depending on how disruptive the test is)

United States



EU




General distribution of services


What portion of this sample has high-speed Internet?

Can we see stratification when looking at the entire population?

US



EU

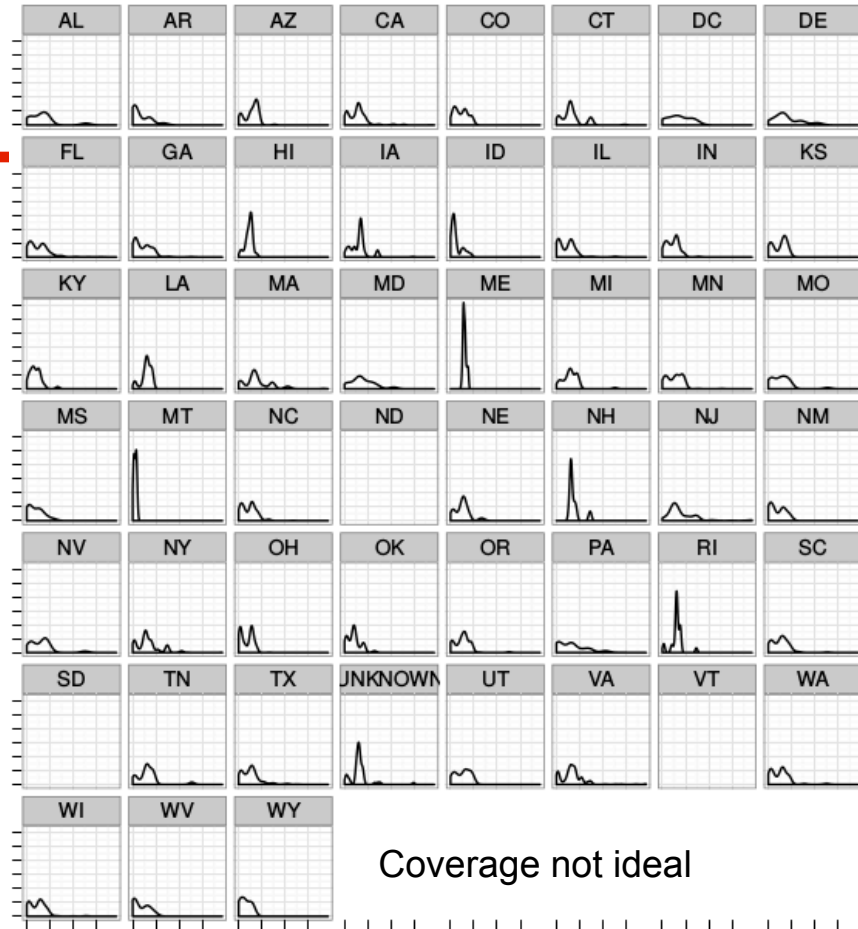



Geographic distribution of services

How is Internet service quality distributed geographically?

Is service in a given region stratified by technology?

In the following charts, compare Sweden with Italy



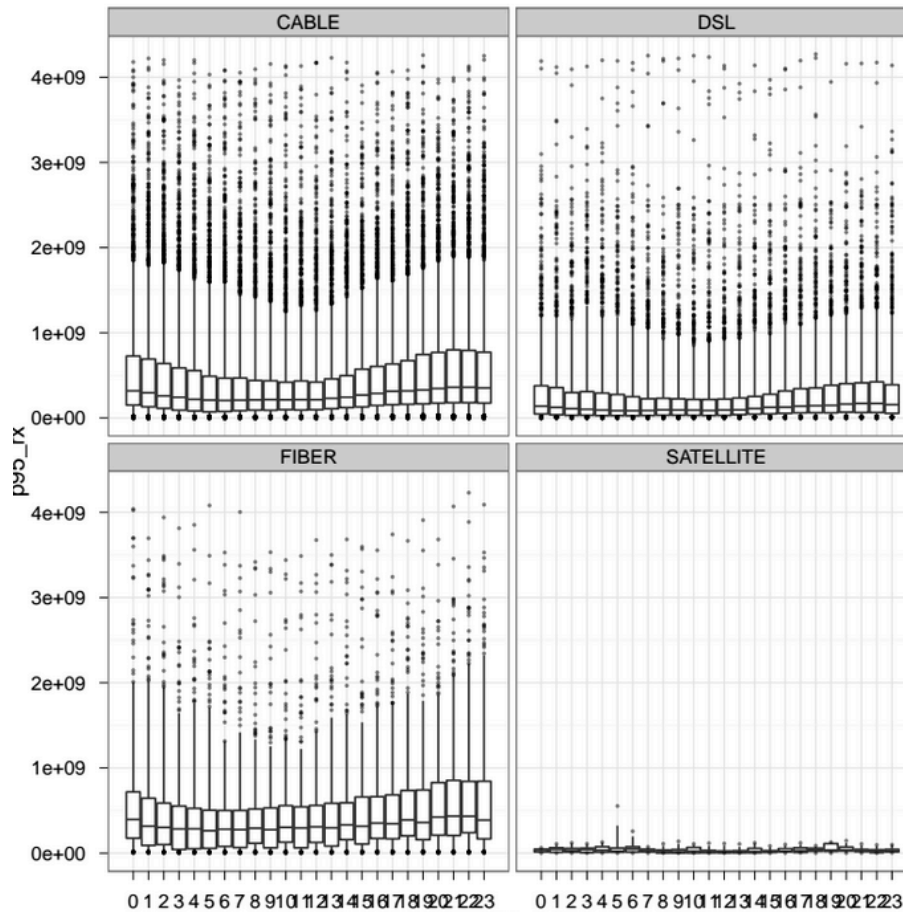
Coverage not ideal

Temporal distribution of services

How does Internet service usage vary in time?


How does Internet service quality vary in time?

Actual usage. Note 95% point shows the variation best.



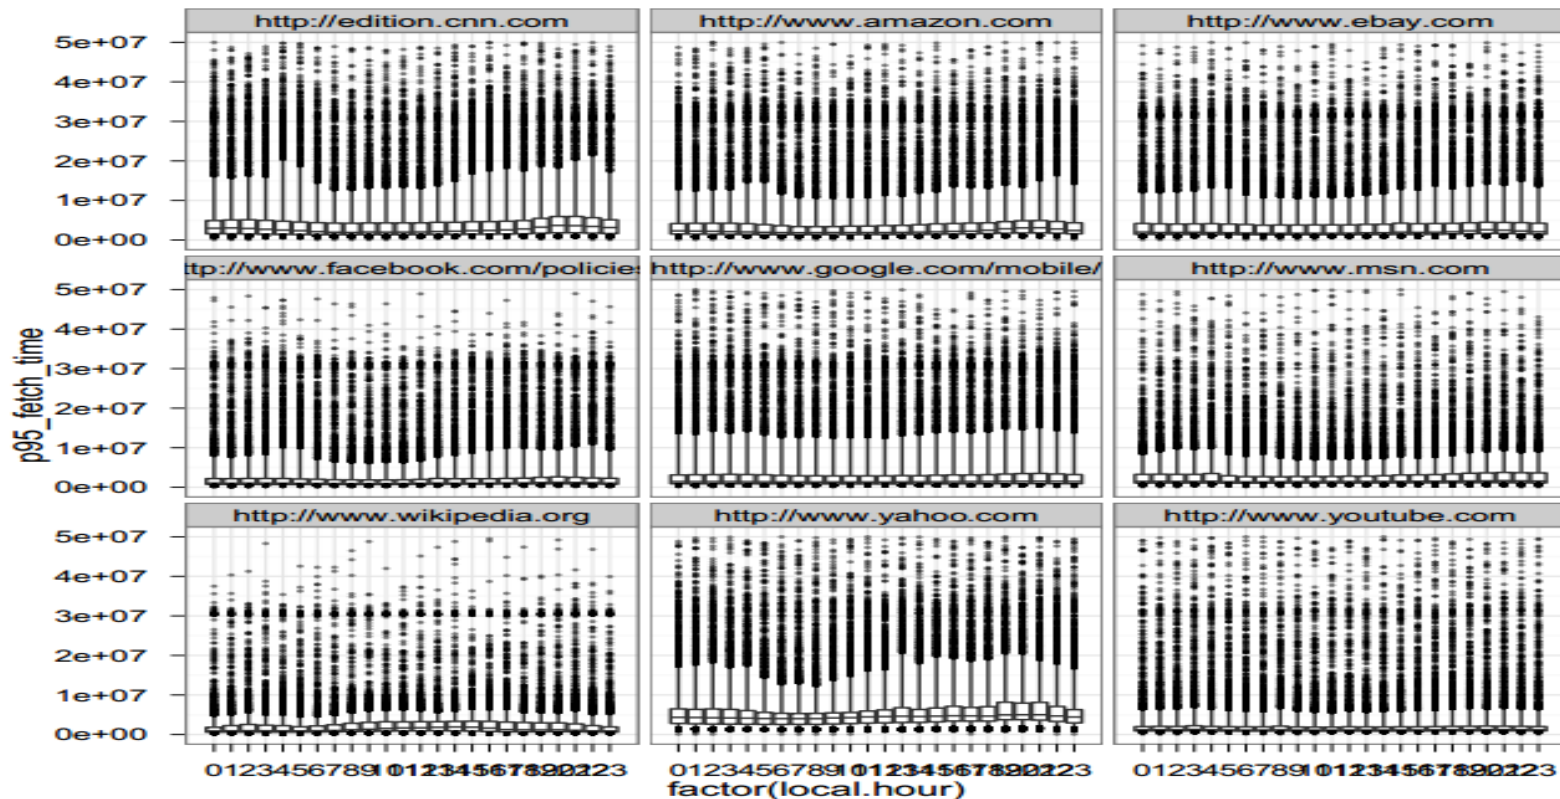
Usage
(US only)

This is a near-source test –
download from an MLab server




DL speed
(US only)

But performance of distant applications does vary with time – congestion?



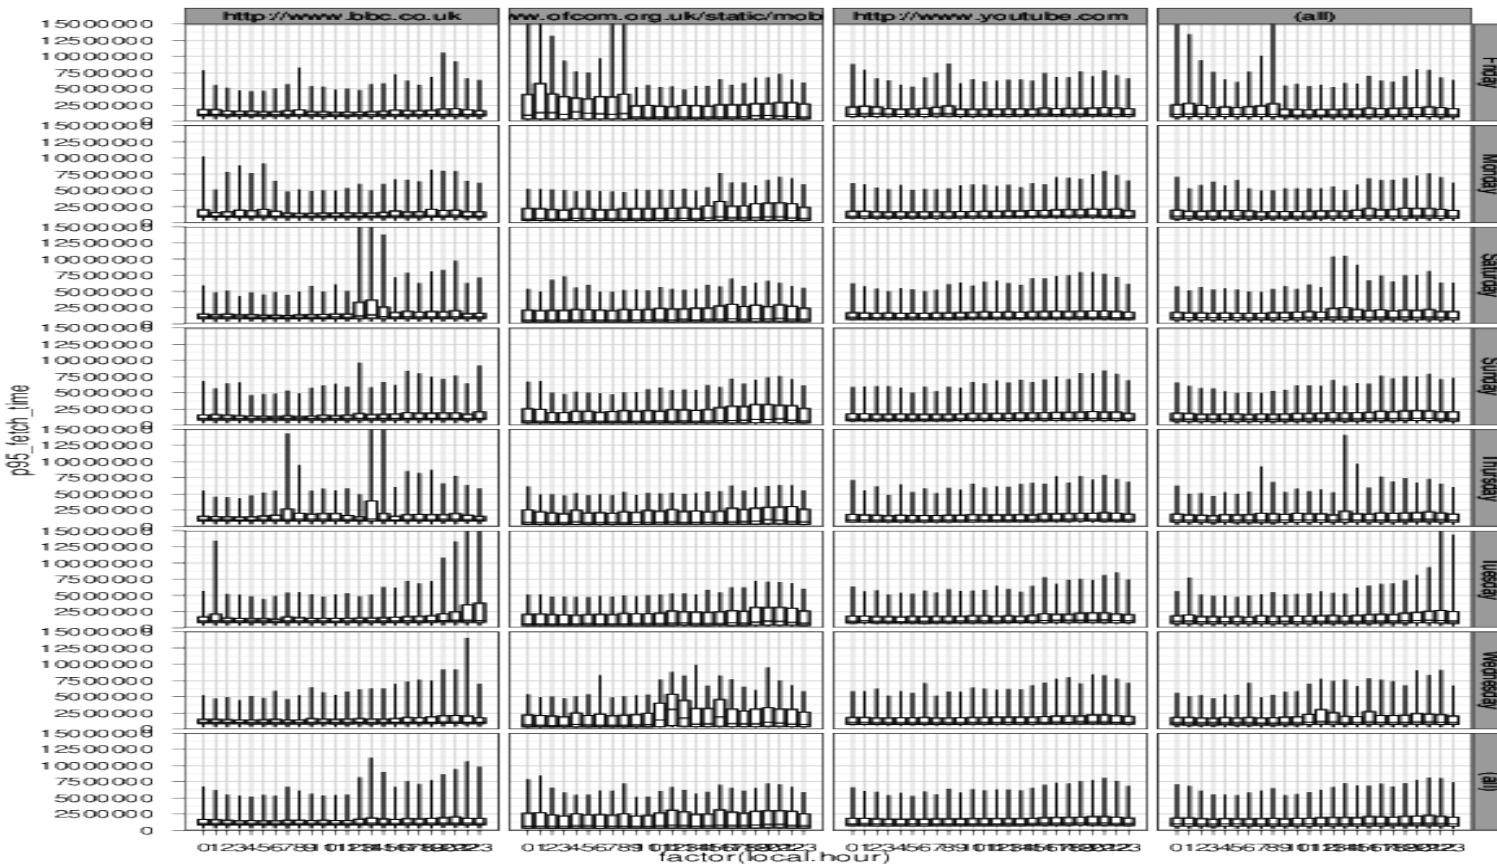
US

Round trip time to popular URLs vs. time of day shows significant variability.



EU

UK data also shows consistent local speeds, but varying web fetch delays



UK

Conclusions and Lessons Learned

- Observing from many points to a few important targets has advantages.
 - But 10,000 is not a large number, when spread over 50 states or 23+ countries.
 - Persistent anomalies observed by plotting against hour or day of week
 - Broad scale active measurement needs followup details to get to root causes
 - More frequent use of traceroute should give insights
-