

Crowdsourcing ISP Characterization to The Network Edge

Zachary Bischof, John Otto, Mario Sánchez, John Rula,
David Choffnes*, Fabián Bustamante

Northwestern U., *U. of Washington



AquaLab

ISP Characterization

- *What is it?*
 - Understand how various factors affect performance
- *Who needs it?*
 - Subscribers shopping for alternative ISPs
 - Companies providing reliable Internet services
 - Governments surveying the availability of high-speed Internet services to their citizens

ISPs 'still mislead' on broadband

July 26, 2011



Sam Knows™



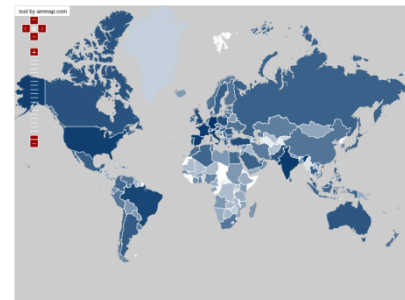
Broadband Speeds Are Largely as Advertised

August 2, 2011

ISP characterization – how should it be done?

- **At scale**

- To capture the diversity of providers and services



- **Continuously**

- To capture dynamic changes in management policies, and unscheduled events, ...






- **From end users**

- To guarantee accuracy, reduce bias




Existing approaches to characterization

- Web-based, user-initiated tests against dedicated or cloud servers
 - E.g. Netalyzr, Speedtest, ...
- End-host monitoring from dedicated servers
 - E.g. Dischinger et al. (IMC07) , Croce et al. (PAM09)
- Installing special monitoring devices at PoPs or home networks
 - E.g. BISMark, SamKnows, Keynote
- *An unavoidable tradeoff among vantage points, coverage and continuous monitoring?*



Our Approach

- *Crowdsourcing ISP Characterization to the Network Edge (C2E)*
 - Leverage the views of popular, network-intensive applications from the end-user
 - (e.g. VoIP, P2P, IPTV, gaming, ...)
 - Reduce number of active measurements
 - Based on experience of end users
 - Continuous monitoring achieved by combining views of multiple subscribers
 - Application usage can grow with the network edge

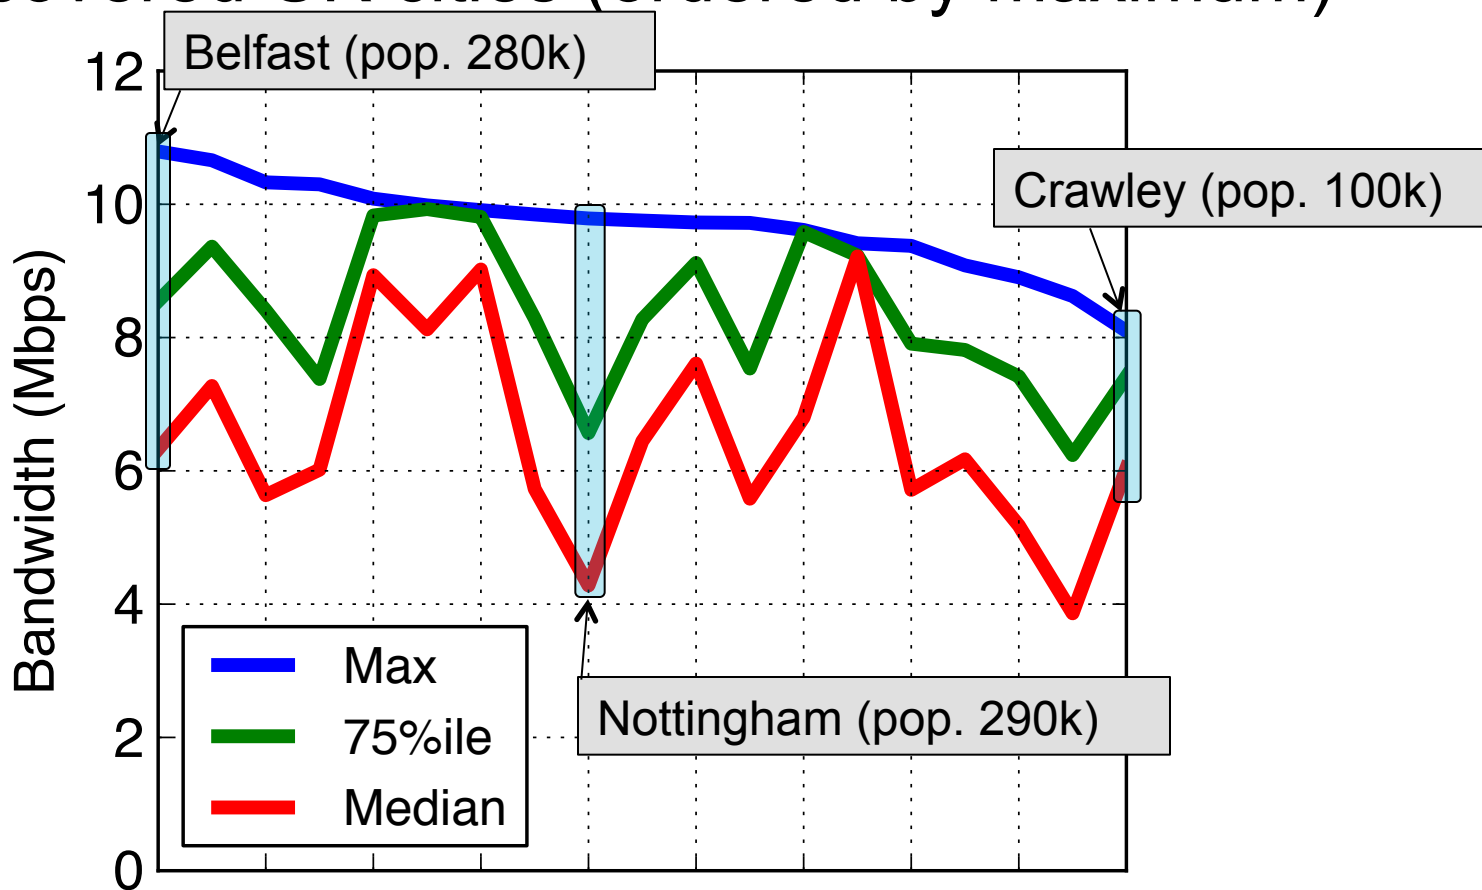


BitTorrent as a Hosting Application

- *BitTorrent*
 - Relatively long session times
 - High bandwidth usage
- *Ono* – client extension for Vuze
 - Aims to improve performance by suggesting “closer” peers
 - Users voluntarily contribute performance statistics
 - In total, about 1.3+ million users world-wide
 - Datasets from November 2009 and 2010

At Scale

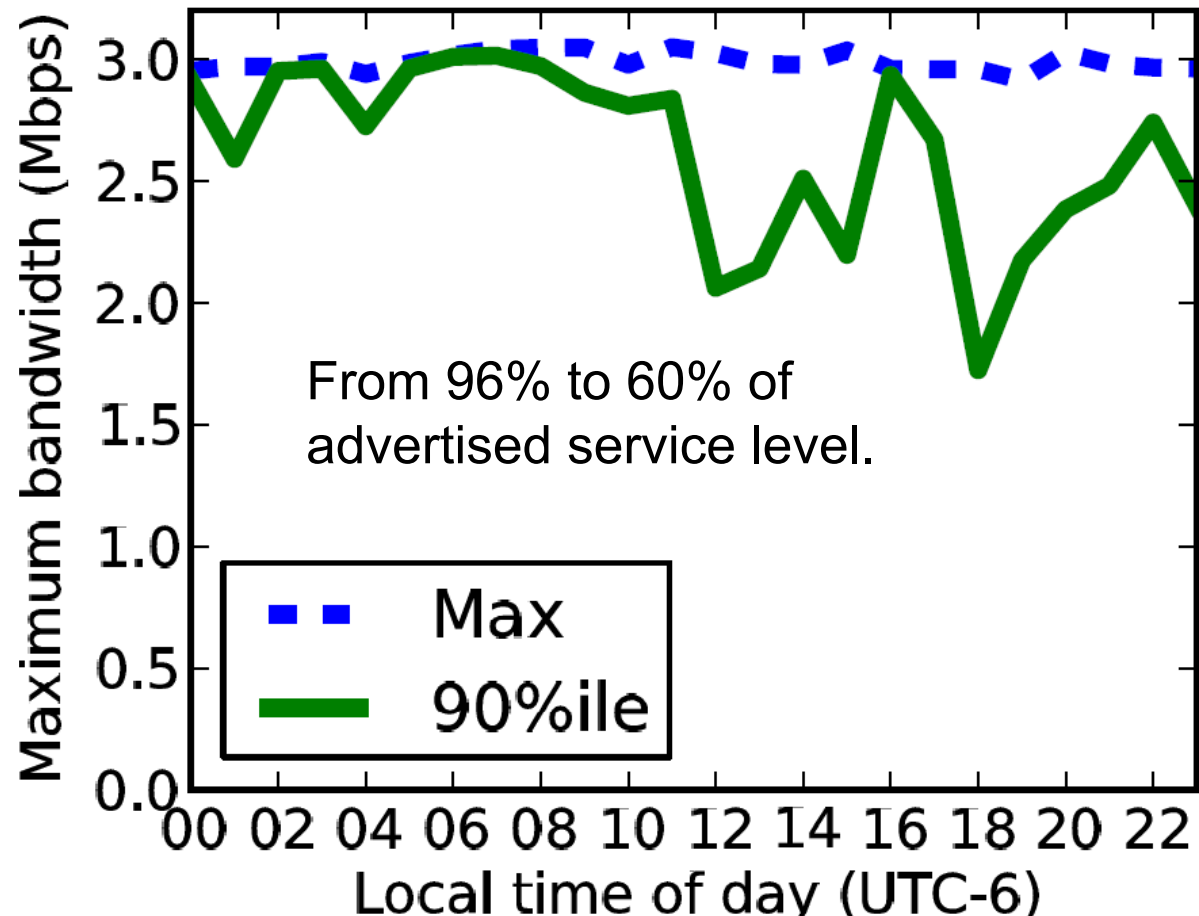
- Variations within a service level among 19 Virgin Media covered UK cities (ordered by maximum)



19 Cities with Virgin Media Peers

Continuous Monitoring

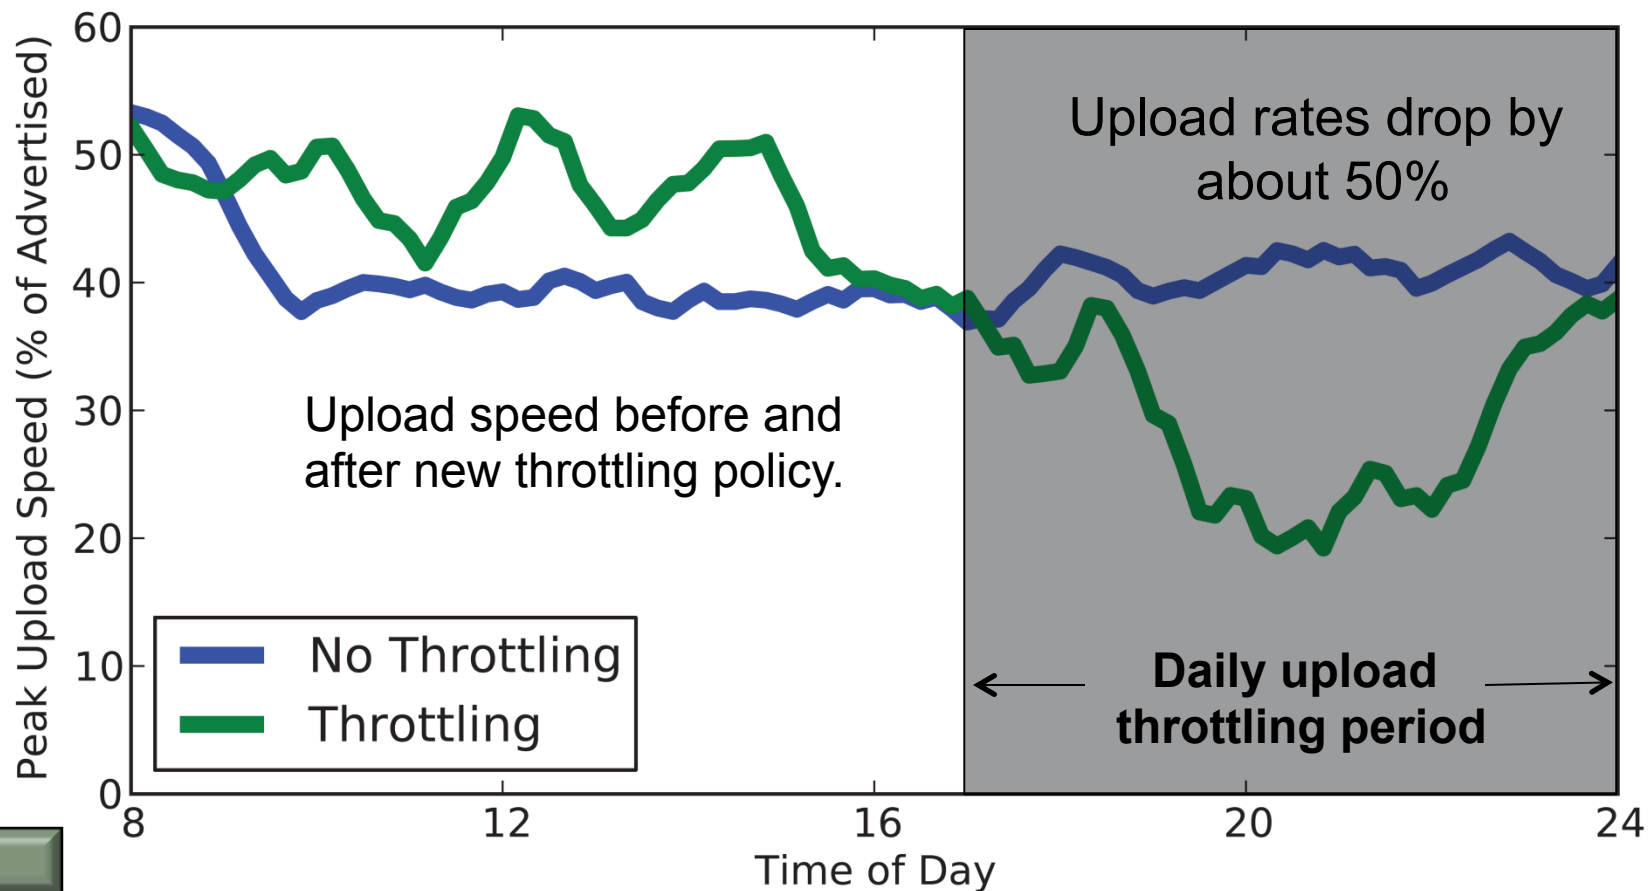
- Variations on Rogers performance during the day (aggregated over Nov. 2009)



Scale
Continuous
End-user

Continuous Monitoring

- Virgin Media new throttling policy in effect




Scale
Continuous
End-user

From the End-User

- Capturing the end-user's view


Normalized number of DNS timeouts



Scale
Continuous
End-user

Beyond single ISP characterization

- Percentage of cities containing at least one ISP providing each category of service



Nov. 2009; mapped with GeoLite

Current Status

- Looking to leverage other network-intensive applications
- Dasu – C2E implementation
 - Inform users of ISP performance
 - Eliminate confounding factors
 - Cross traffic
 - Wi-Fi or Ethernet?
 - Complementary to SamKnows/BISMark



The screenshot shows the Dasu v3 application window. The interface includes a menu bar (File, View, Community, Tools, Window, Help), a search bar, and a toolbar with icons for Download, Play, Device Launch, Up, Down, Start, Stop, Delete, and View. The left sidebar contains sections for VUZE (Getting Started, Vuze Plus, Vuze HD Network, Games, Notifications), FILES (Library, New, Downloading...), DEVICE PLAYBACK (beta, Turn On, In Progress, iTunes), SUBSCRIPTIONS (Subscriptions Overview), DVD BURN (Get Started), and PLUGINS & EXTRAS (Dasu, Network Status).

The main content area features the DASU logo and the AquaLab logo with the tagline "Planetary-scale distributed computing". Below the logo is a navigation bar with tabs: Summary, Interference, BT, Home Network, Internet, Settings, and About. The "Internet" tab is selected.

The "Internet" tab displays results for tests related to the user's Internet connection. It includes a description: "This tab displays results for tests related to your Internet connection, including the network latency to your configured DNS server(s), alternative public DNS services, and several popular websites."

DNS Server Latency

DNS Configuration Diagnostics

Status	Description
OK	Server #1 has lower latency than the best tested public DNS servers.
OK	Your primary server has the lowest latency.

Locally Configured DNS Servers

Order	IP Address	Responding IP	Domain	Latency (ms)
1	192.168.1.1	69.252.66.201	null	13

Latency Comparison to Public DNS Services

The bar chart shows a single bar for 'Configured #1' with a latency of 13 ms. The y-axis is labeled 'Latency (ms)' and the x-axis is labeled 'Server'.

The bottom status bar includes: Send Feedback, 2,392,216 users, 473.7 kB/s, [12K]* 224 B/s, and a help icon.