

Glasnost: Enabling End Users to Detect Traffic Differentiation

Krishna P. Gummadi

Networked Systems Research Group
Max Planck Institute for Software Systems

High-level goal: Network Transparency

- Today, access ISPs are very opaque
- ISPs deploy middle-boxes to monitor & manage customer traffic
 - blockers, rate-limiters, firewalls, censors
- But, ISPs do not disclose their management policies
- This lack of transparency prevents:
 - end users from making an informed choice
 - application designers from adapting apps to ISP policies
 - regulators from monitoring ISPs and holding them accountable
- Can we use end-host based measurements to infer ISP policies?

Quick aside: Transparency vs. Neutrality

- Transparency: What management policies is an ISP employing?
- Neutrality: Is a particular management policy employed by an ISP acceptable?
 - is an ISP policy neutral? (according to some definition of neutrality)
- Neutrality has strong proponents and opponents
 - some argue that neutrality would hurt new network innovation
- But, it is hard to argue that transparency would be harmful
 - other than that some ISPs could benefit from a closed network

The Glasnost project


- Goal: Enable end users to verify whether their access ISPs are employing application-specific traffic differentiation

Application-specific traffic differentiation

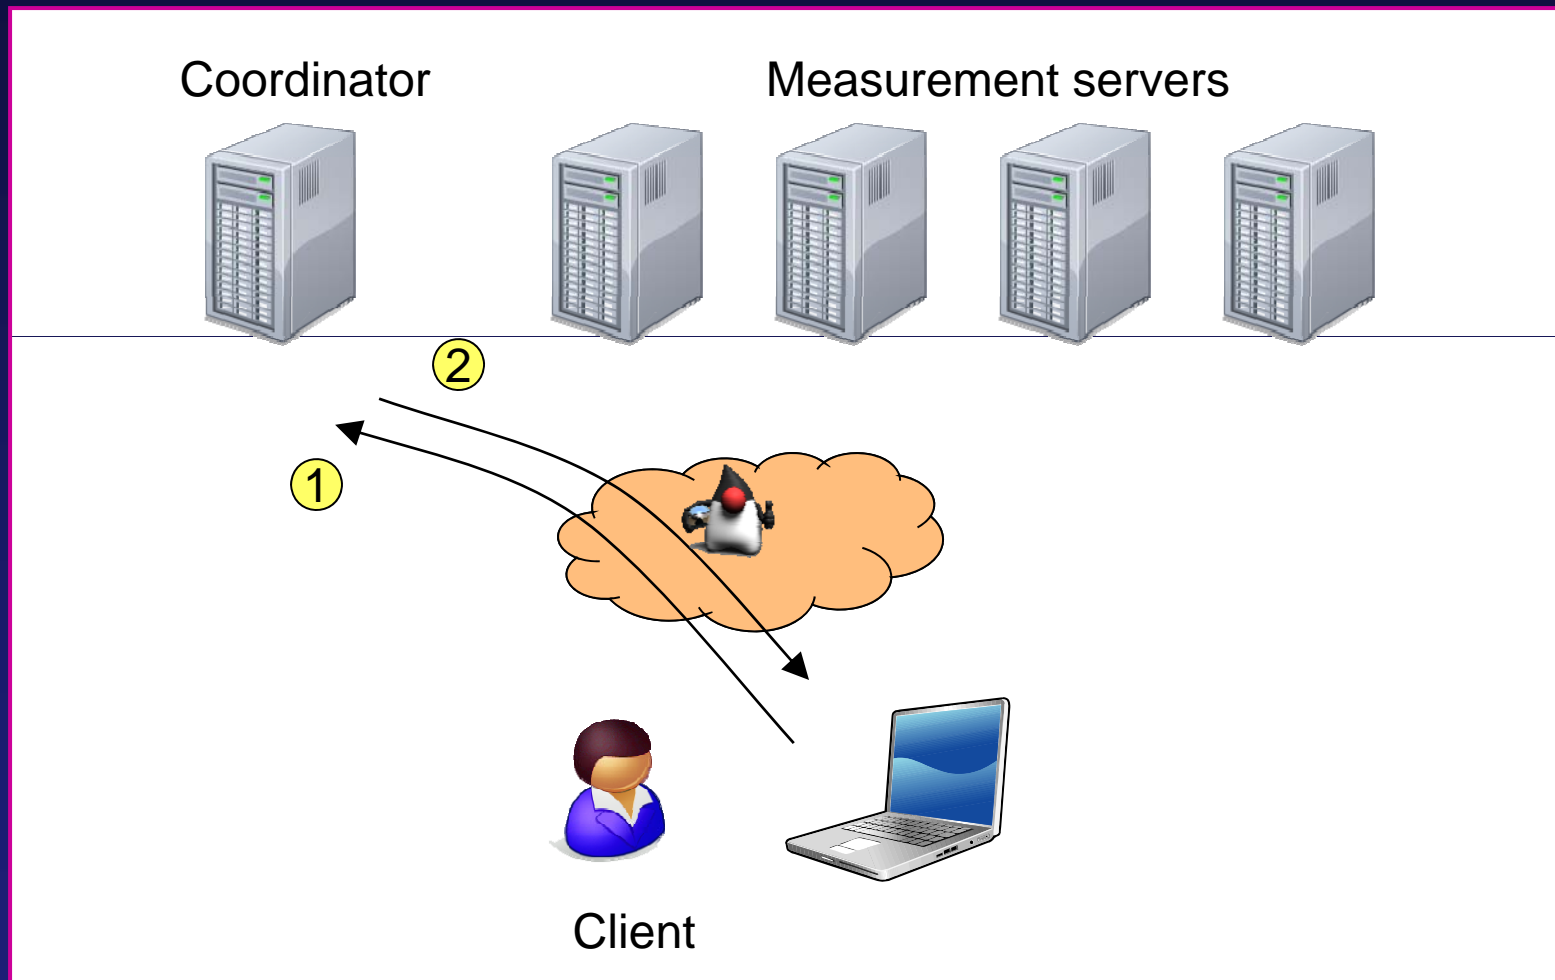
- Are packets from flow A being treated **differently** than packets from flow B?
- Primarily because they belong to different apps
 - because the flows carry different packet payloads
- Examples:
 - interrupting flows carrying BitTorrent protocol messages
 - rate-limiting flows to email ports vis-à-vis http ports

Key idea: Verifying traffic differentiation


- Run controlled active measurements between end hosts of a path to detect traffic differentiation by ISPs along the path
- Run back-to-back flows with same network-level characteristics, but with different packet payloads
 - i.e. replay traces changing only their payload
- Compare the behavior of flows



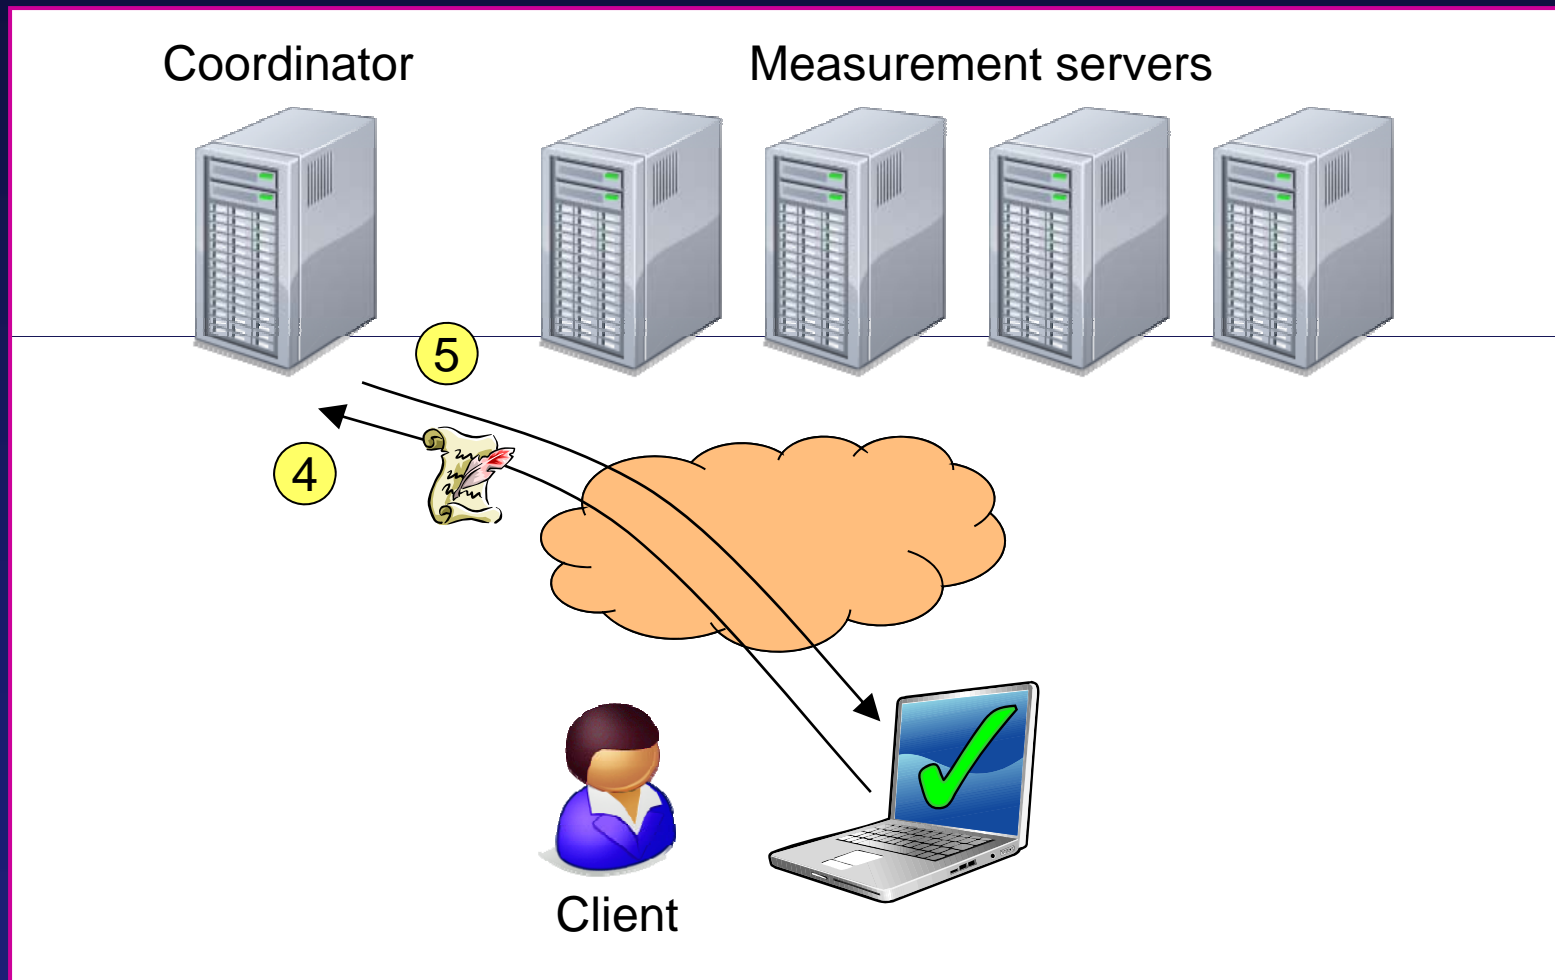
The Glasnost architecture



The Glasnost architecture



The Glasnost architecture



Enabling lay users to check their links



Max
Planck
Institute
for
Software Systems

Glasnost: Test if your ISP is manipulating BitTorrent traffic

■ [Home](#) ■ [BitTorrent Test](#) ■ [Results](#) ■ [Contribute](#) ■ [Updates](#) ■ [Contact](#) ■

- Our test runs BitTorrent and TCP downloads as well as uploads on a well-known BitTorrent port and a non-BitTorrent port.
- There are two test configurations available. The simple one takes approximately 4 minutes and each individual transfer is 10 seconds long. The full testset runs each transfer for 20 seconds and collect more data points and allows us to infer the status of your link with higher confidence.


- Full test (~ 7 minutes)
- Simple test (~ 4 minutes)

» [Start testing](#) «

Rest of the talk

- Detecting BitTorrent traffic differentiation with Glasnost
 - BitTorrent blocking
 - BitTorrent ratelimiting
- Measurement-lab

Detecting BitTorrent blocking




- A year ago EFF reported that Comcast was blocking BitTorrent
- This report set off a debate whether this is an acceptable policy
- EFF published their testing tool
 - requires networking expertise to run it
- Other easy-to-use tools like Vuze gathered insufficient evidence

BTTest: Easy-to-use BT blocking detector

MOZILLA FIREFOX

File Edit View History Bookmarks Tools Help

http://broadband.mpi-sws.mpg.de/transparency/bttest.php

 Max Planck Institute for Software Systems

Glasnost: Test if your ISP is manipulating BitTorrent traffic

■ Home ■ BitTorrent Test ■ Results ■ Contribute ■ Updates ■ Contact ■


- Our test runs BitTorrent and TCP downloads as well as uploads on a well-known BitTorrent port and a non-BitTorrent port.
- There are two test configurations available. The simple one takes approximately 4 minutes and each individual transfer is 10 seconds long. The full testset runs each transfer for 20 seconds and collect more data points and allows us to infer the status of your link with higher confidence.

Full test (~ 7 minutes)
 Simple test (~ 4 minutes)

» Start testing «

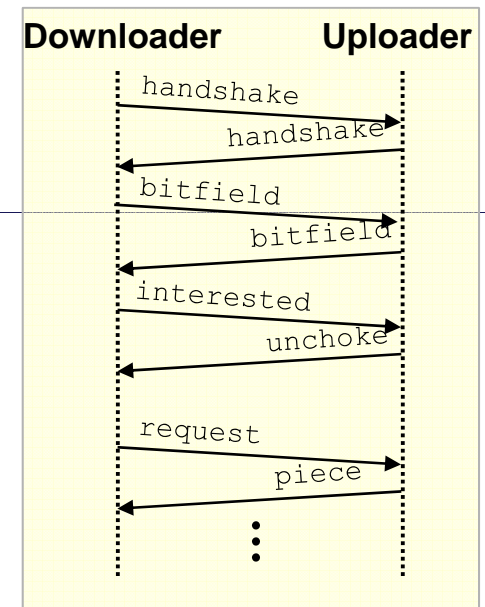
TCP RST packets. It seems like your ISP hinders you from uploading BitTorrent traffic to our test server.

How BTTest works



- BTTest runs flows between a user and a server
 - emulates BitTorrent transfers to trigger blocking
 - non-BitTorrent flows as controlled flows
- Transfers with different configurations
 - using BitTorrent and non-BitTorrent ports
 - testing uploads and downloads

BitTorrent packet exchange



How BTTest identifies BitTorrent blocking

- BTTest collects information from both end hosts
 - tcpdump on server side
 - information from applet on whether connection was reset
- BTTest reports blocking with forged resets when
 - user-side socket was torn down by a RST packet
 - server dump shows that the server did not send the RST

Results from BitTorrent blocking

- How prevalent is blocking?
- How are BitTorrent flows identified?
- At what times are BitTorrent flows blocked?
- Does greater transparency impact ISP behavior?

How prevalent is blocking?

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

- 98,530 hosts measured from 157 countries and 3,024 ISPs
- 4,575 hosts from 110 ISPs observed BitTorrent blocking
- Widespread blocking in the USA and Singapore
- Almost all blocking (99.5%) happens for uploads

How are BitTorrent flows identified?


- Based on TCP port
 - 15.8% of hosts experience blocking for all flows on the BitTorrent port
- Based on Protocol messages
 - For 98.2% of hosts blocking is based on BitTorrent messages
- Most ISPs seem to perform deep packet inspection
 - confirmed this with more controlled experiments for Comcast

At what times are BitTorrent flows blocked?

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

- ISPs claim to manage user traffic during times of peak utilization
- However, % of blocked tests stays rather high over time
- Similar results for other ISPs

At what times are BitTorrent flows blocked?



- ISPs claim to manage user traffic during times of peak utilization
- However, % of blocked tests stays rather high over time
- Similar results for other ISPs

Does transparency affect ISP policies?

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Future challenges for Glasnost

- How to design tests for the myriad traffic differentiation policies ISPs deploy?
- Which differentiation tests should an end user run?
- How to deal with noise in individual user measurements?

Measurement Lab
measurementlab.net

What is M-Lab?

- M-Lab is:
 - An **open, distributed server platform** on which researchers can deploy **active, client-server network measurement tools** that measure aspects of broadband Internet connections.

M-Lab's goal:

- To **advance research** and **empower the public with useful information** about their broadband connections. By enhancing Internet transparency, we aim to help **sustain a healthy, innovative Internet**.

Founded by: PlanetLab, New America Foundation, Google Inc., and a group of researchers (including me)

M-Lab and PlanetLab

- M-Lab = a "private" PlanetLab
- M-Lab's servers are separate and distinct from PlanetLab
 - **Narrower scope:** active, client-server measurements of broadband connections
 - **Servers:** all enterprise grade with 8 cores and 1 Gbps connectivity; three servers per site
 - **Allocation:** # tools on a server must be $< 1.5 * \#$ of cores
- M-Lab builds on PlanetLab
 - Depends on PlanetLab Consortium's OS/VM system (each tool gets a "slice") and OA&M

What does "open" mean?

- All collected **data to be made publicly available** either immediately or after an optional 1 year embargo.
- All researchers required to **publish client and server software source code** to allow for 3rd party review
- All researchers' tools will be operated and licensed in such a way as to **allow third-parties to develop client-side software** for measurements.
- **A collaborative effort**: welcomes support from all researchers, institutions, companies that want to make this succeed
M-Lab will not be used to collect and store data from other, passive monitoring of users' Internet activity.

M-Lab: present and future

- Currently "Proof of concept"
 - Limited number of tools for speed, diagnostic, and testing for BitTorrent throttling
 - NDT, NPAD, Glasnost
 - DiffProbe, NANO (coming soon)
 - Currently 6 servers in 2 location. Google will be rolling out 36 servers in 12 locations over first half of 2009
- Future:
 - Involve all researchers who want to participate
 - Host as large a variety of tools as possible
 - Expand server sites globally
 - Open data & open tools; data repository

Get Involved

- Want to deploy a tool?
 1. Look at instructions on our site:
<http://measurementlab.org/getinvolved>
 2. Read our discussion document
 3. Email the M-Lab steering committee
<http://measurementlab.org/contact>

M-Lab - policies

- M-Lab is still in development. There are a number of organizational and procedural issues that have yet to be formalized, including:
 - Processes for **adding to tools set**, and making changes to the set to produce optimal data
 - Processes for determining **which apps run where**
 - Processes for formally **admitting new M-Lab members**, whether providing server infrastructure, tools, data-storage/analysis or funding
- Going forward M-Lab wants to work with researchers to formalize these processes as necessary.

Got servers?

- M-Lab depends on contributions of servers and connectivity from other industry and institutional partners.
- Basic requirements:
 - Provide at least 3 enterprise-grade servers and/or 1 Gbps connectivity and rackspace
 - A /26 of IPv4 address space is required to support a site with three servers. IPv6 connectivity is a plus.
 - Administration: initial setup and physical maintenance of the host.

Questions, concerns, contacts

- M-Lab site:
<http://measurementlab.net>
- Contact me:
Presenter Name: gummadi@mpi-sws.mpg.de
- Contact the M-Lab Steering Committee:
<http://measurementlab.net/contact>

Thank you!

**For more information, please visit:
<http://broadband.mpi-sws.org/transparency/>**