

Pitfalls for Measuring Q_o^* in a Middlebox-Infested World

David Choffnes
Northeastern University

Context

2

- This is a cautionary tale
- We started with some basic assumptions
 - ▣ A packet is forwarded to its destination if reachable
 - ▣ Receivers receive what senders send, or nothing
 - ▣ Middleboxes are present, but mostly harmless
- Then we started to try to understand **differentiation...**

Full details:

IMC 2015, and see <http://dd.meddle.mobi>

What Does This Have to Do with QoE?

3

- We wanted to measure **performance** from users/clients
 - ▣ Throughput
 - ▣ RTTs
 - ▣ Jitter
- ...and its impact on **real applications**
 - ▣ Ones that matter, not completely synthetic traffic
 - ▣ Like one that rhymes with *ShoeLube*
- Focus: What might go **wrong** with Qo* measurements

Lesson 1: What you send matters

4

- ❑ **DPI is everywhere** (at least in mobile)

- ❑ Triggers for DPI might include
 - ▣ IP addresses
 - ▣ ports
 - ▣ payload signatures
 - ▣ total number of connections
 - ▣ total bandwidth
 - ▣ time of day

- ❑ So what really matters?
 - ▣ ***We bought a DPI box to find out!***

What triggers detection?

5

Replayed YouTube through our DPI box in a fully controlled testbed

Original IPs	Original Ports	Other changes	Services detected as
✓	✓	-	You Tube
✗	✓	-	You Tube
✓	✗	-	You Tube
✗	✗	-	You Tube
✓	✓ / ✗	Begin connections with a packet containing '0'	HTTP / P2P
✓	✓ / ✗	Replaced (first packet only): youtube with amktube	HTTP / P2P
✓	✓ / ✗	Replaced (first packet only): youtube with amktube (HOST header only)	You Tube / You Tube
✓	✓ / ✗	Added 100 bytes to the beginning of first packet	HTTP / P2P
✓	✓ / ✗	Added 100 bytes to the end of first packet	You Tube / You Tube
✓	✓ / ✗	Replaced GET with AMK	HTTP / P2P
✓	✓ / ✗	Added "GET " to beginning of first packet	You Tube / You Tube

DPI + HTTPS?

6

- Just as an aside...
- We tested with SSL traffic from Netflix and others
 - ▣ Still detected correctly
 - ▣ Anyone know why?
 - ▣ SNI

Lesson 2: Beware of transcoders

7

- What you **send** is not necessarily what is **received**
 - ▣ Really difficult to do server-side QoE/QoS measurement
 - ▣ Can even affect client-reported statistics


- Examples:
 - ▣ Black Wireless (**compression proxy**)
 - ▣ Boost/Sprint
 - **Image transcoder**
 - But only if files are smaller than 700KB
 - ...and Boost is more aggressive at compression...
 - ...but results in worse download times

Lesson 3: Be careful how you compare

9

□ Too hot

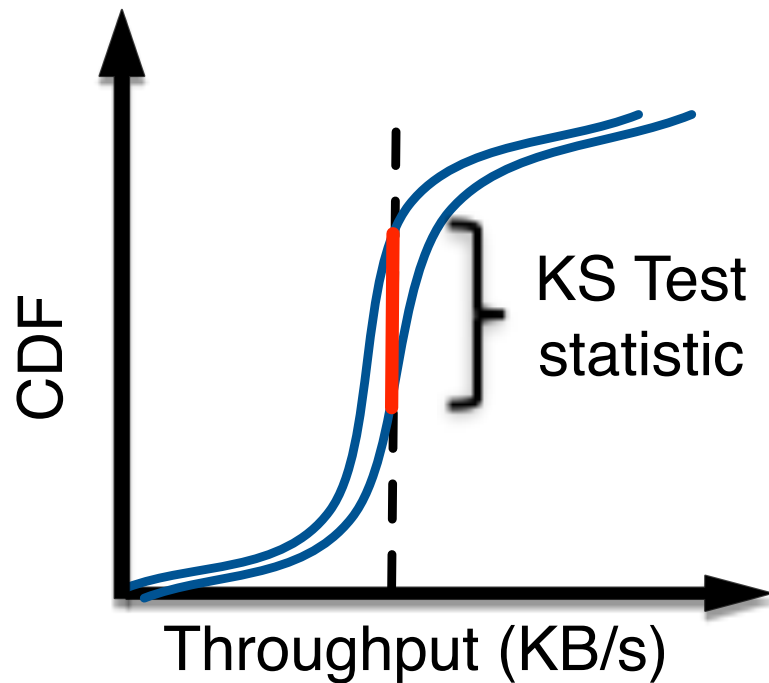
- ▣ Glasnost approach: send traffic as fast as you can



Lesson 3: Be careful how you compare

10

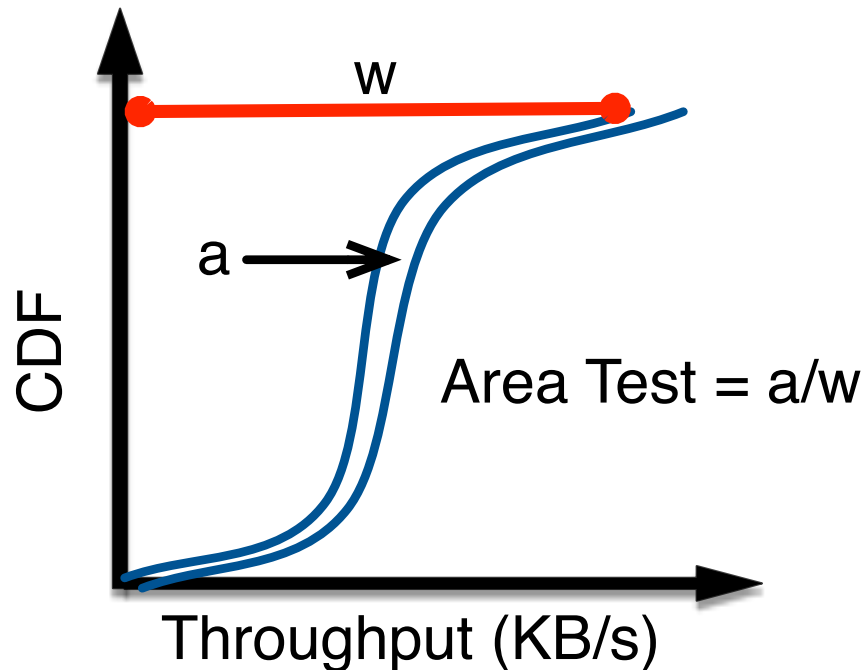
- Too noisy
 - ▣ KS-Test sensitive to small differences



Lesson 3: Be careful how you compare

11

- Area test (weighted KS test)
 - ▣ Detects difference only if large relative to distribution width
 - Good for detecting differentiation
 - ▣ **Open question:** What tests make sense for different apps?



A bunch of other lessons

12

- T-Mobile was **redirecting traffic to different IPs**
 - ▣ ...and then we found they **stopped** some time this year
- Black Wireless didn't have a **compression proxy**
 - ▣ ...until last week
- **Translucent proxies** mess with measurement
 - ▣ Same HTTP headers, different order
 - ▣ Unless it's Verizon, in which case they add a **tracking ID**
 - ▣ ...and sometimes **connection persistence** is changed

Open research challenges

13

- ❑ QoE measurements: What do we send and how?
 - ▣ Synthetic traffic **may not represent service** for real apps
 - ▣ But **state space** of possible classifiers/policies **enormous**
- ❑ What are the right **metrics**?
 - ▣ How do we compare performance across apps, networks?
 - ▣ What application **needs** vs. what it **receives**
- ❑ How to account for **interference** in the middle?
 - ▣ Requires techniques to identify when it's there
 - ▣ How do we **quantify impact** of modified content?

Open research challenges

14

- How do you do this for **mobile**?
 - ▣ Needs **large-scale** measurement, probably **crowdsourced**
 - ▣ Differentiation Detector, Proxy Detector, Mobilyzer, Insight, ...