

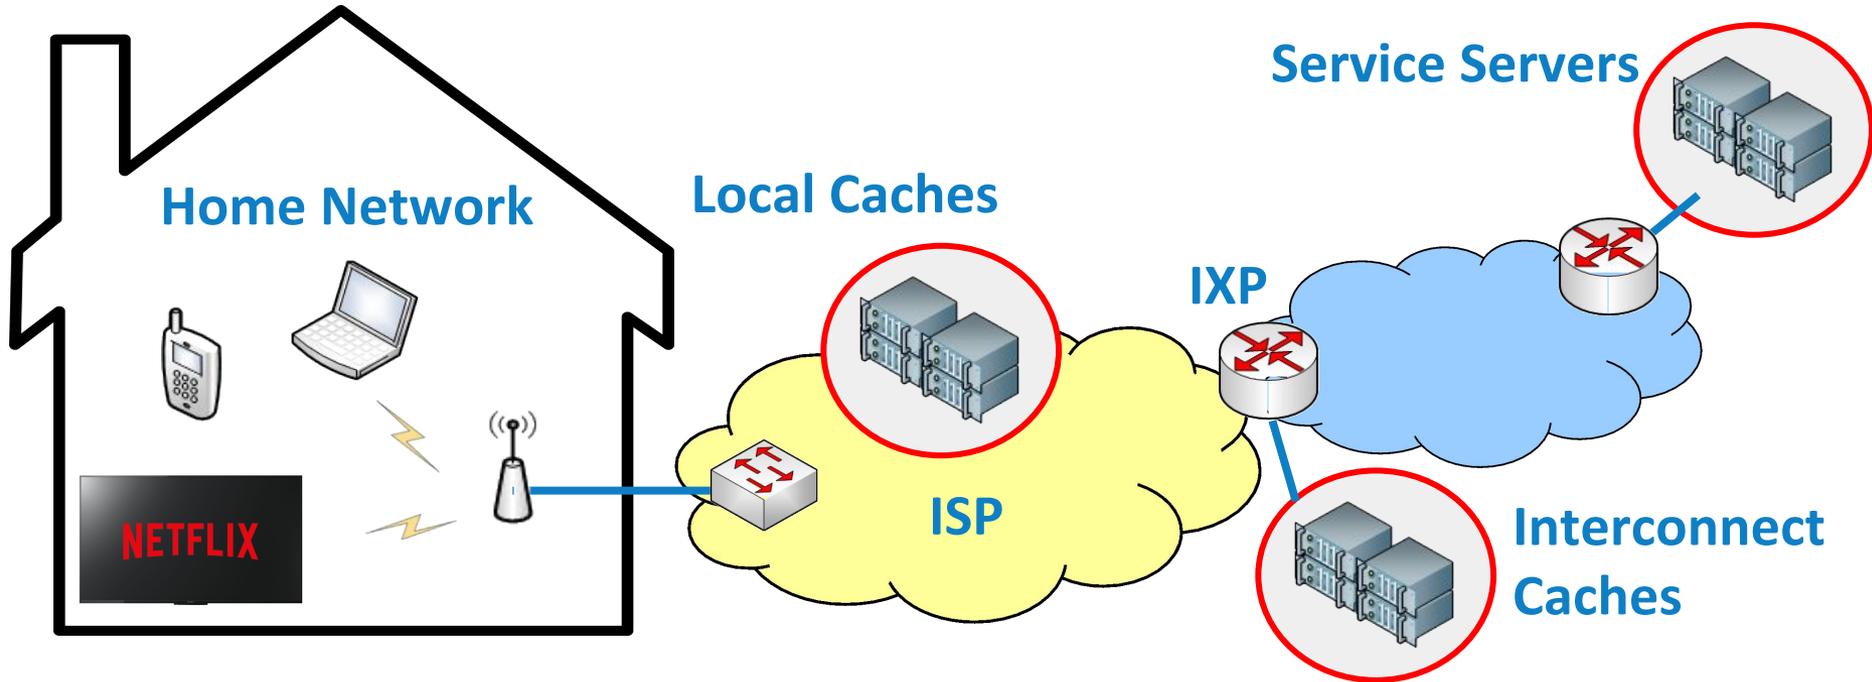
Correlating Network Congestion with Video QoE Degradation - a Last-Mile Perspective

Francesco Bronzino, **Paul Schmitt**, Renata Teixeira, Nick Feamster, Srikanth Sundaresan

AIMS
13 March 2018

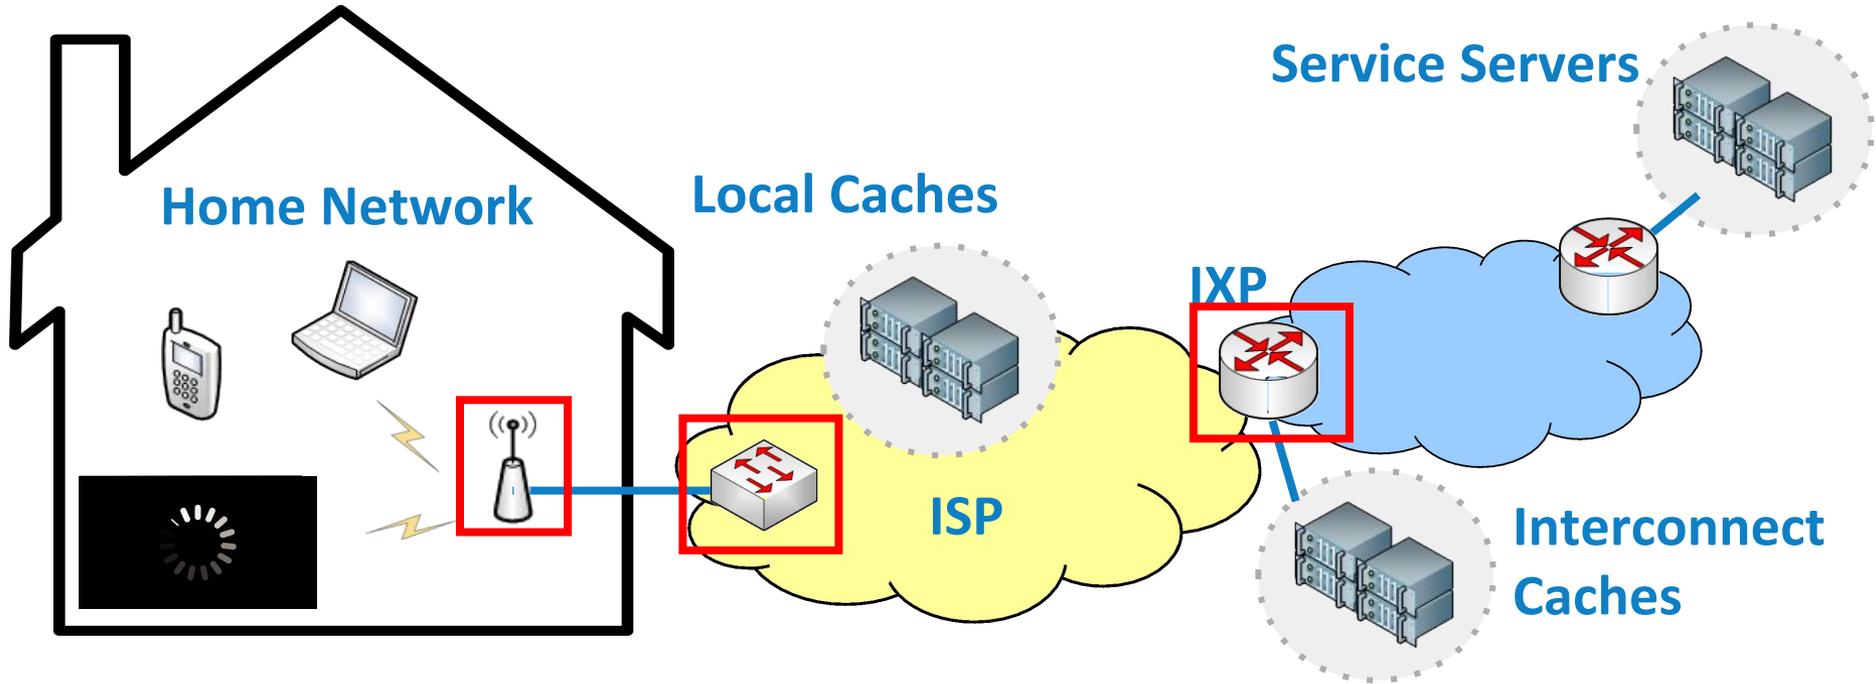


Last-Mile Measurement: Video



- Why is it an interesting use case?
 - Distributed service with diverse ecosystem of servers and clients

Last-Mile Measurement: Video



- Why is it an interesting use case?
 - Challenging to pinpoint and correlate the root causes of impairments from a single location.

Policy Implications

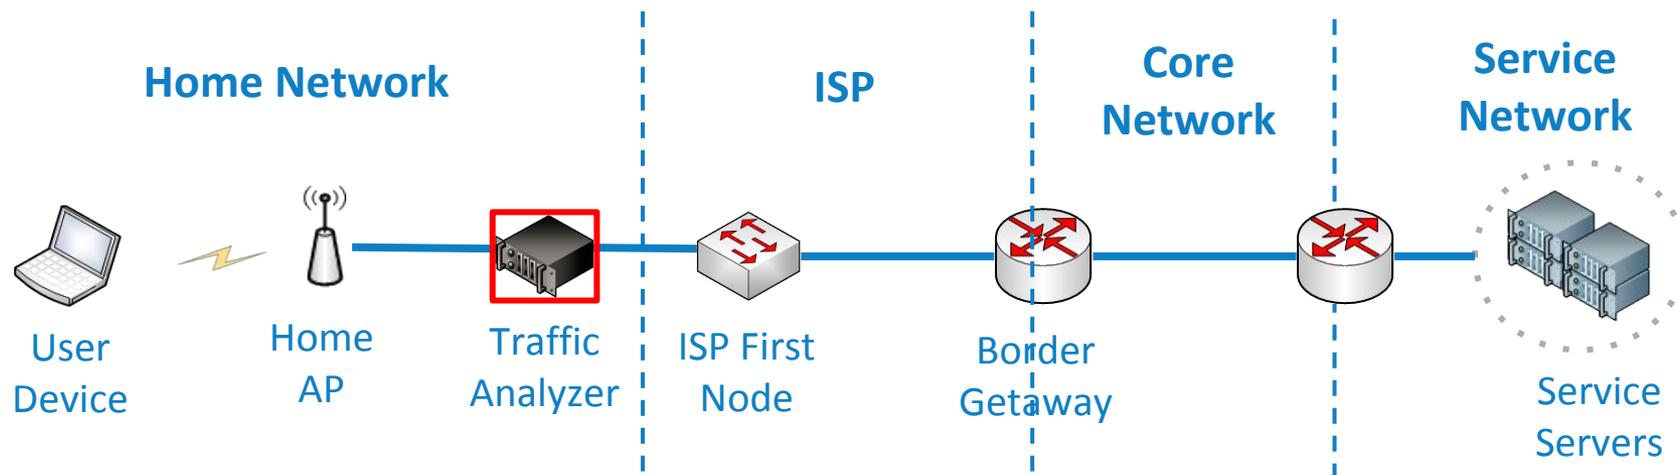
“After all, consumers have little understanding of what packet loss means; what they do want to know is whether their Internet access service will support real-time applications, which is the consumer-facing impact of these performance metrics.”

FCC Restoring Internet
Freedom Order (para. 226)

Real-Time Inference of Quality

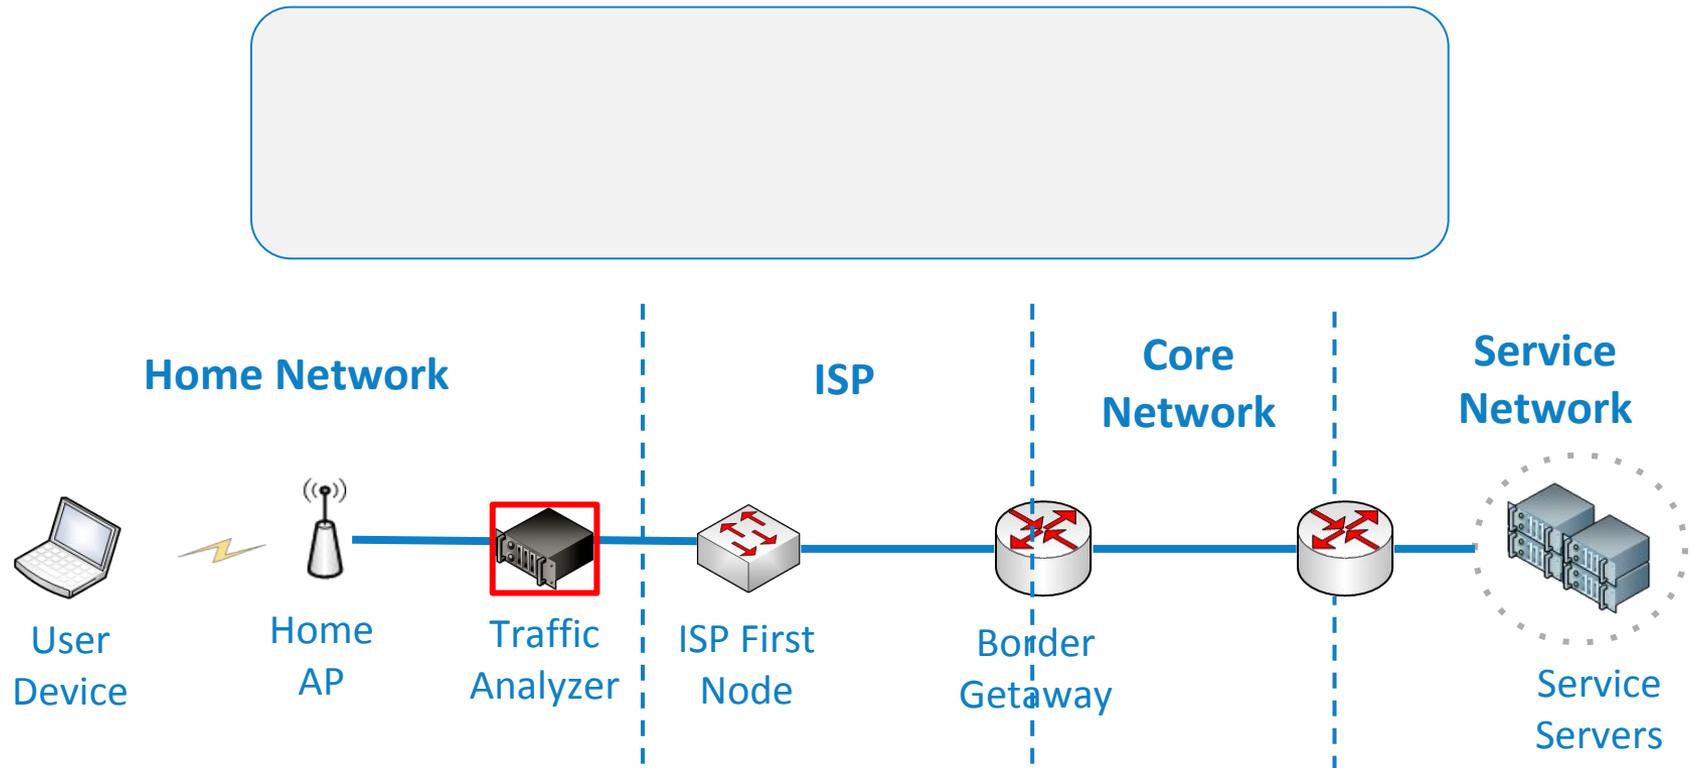
- Goal: detect impairments real-time and correlate them with their root causes
- Solution: build an all-in-one system working from the home vantage point
 - View of home network
- Challenges:
 - Traffic is encrypted
 - Pinpointing the root cause of impairments

Solution: Traffic Analyzer

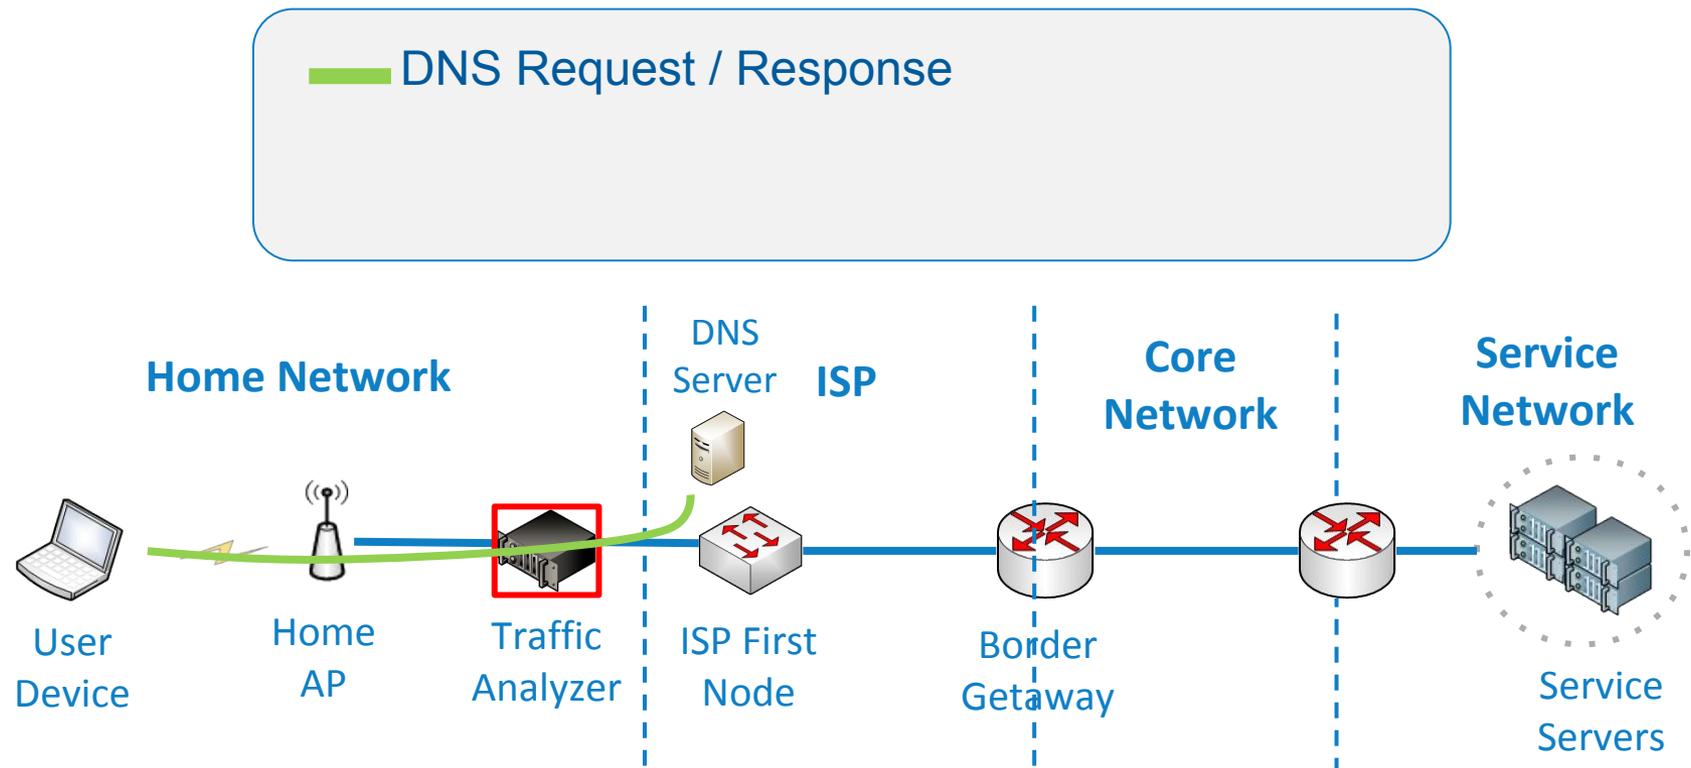


- Online traffic categorization & monitoring
 - Currently video, ads
- Good performance on cheap multipurpose hardware
- Requirement:
 - Wireless AP in bridge mode

Traffic Analysis: Core Techniques



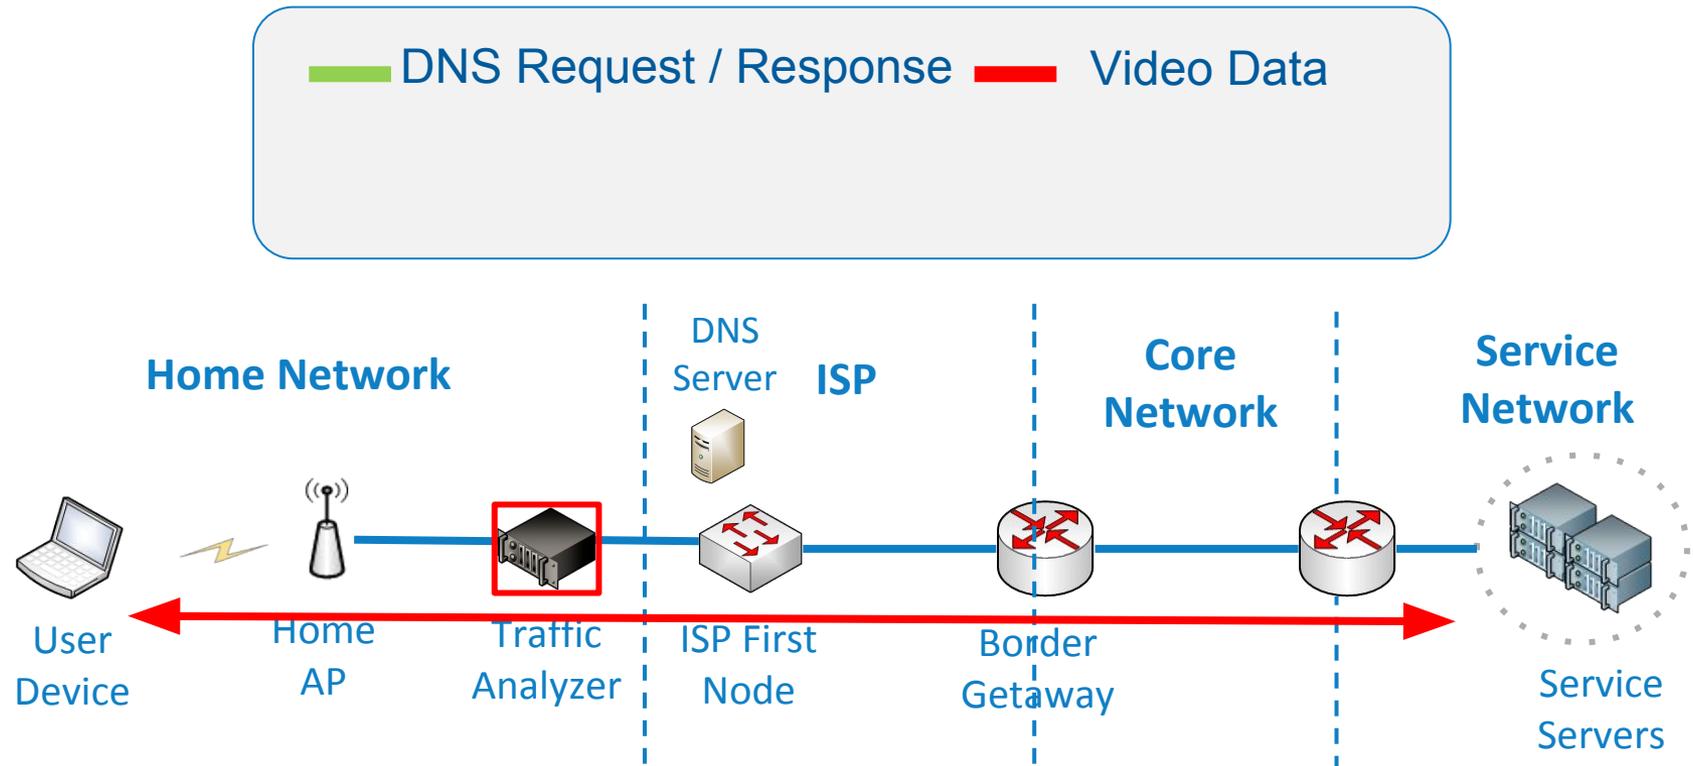
Traffic Analysis: Core Techniques



1 Use DNS data to categorize flows

- DNS response/request provide the mapping between domains and IP addresses
 - E.g. <nflxvideo.net, 198.38.120.155>

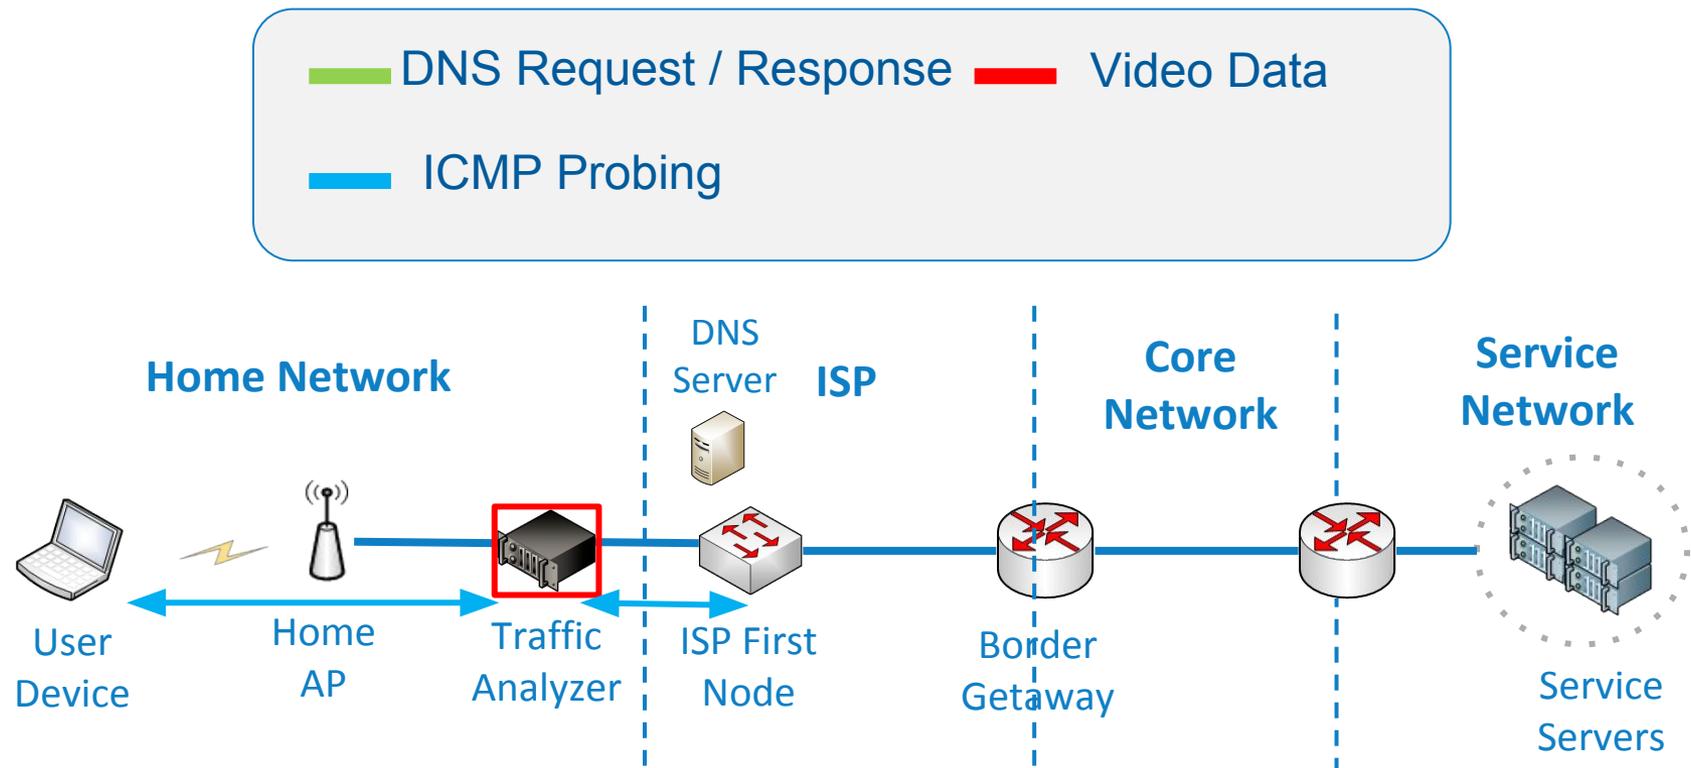
Traffic Analysis: Core Techniques



② Track flow characteristics

- Infer QoS from encrypted video flow traffic

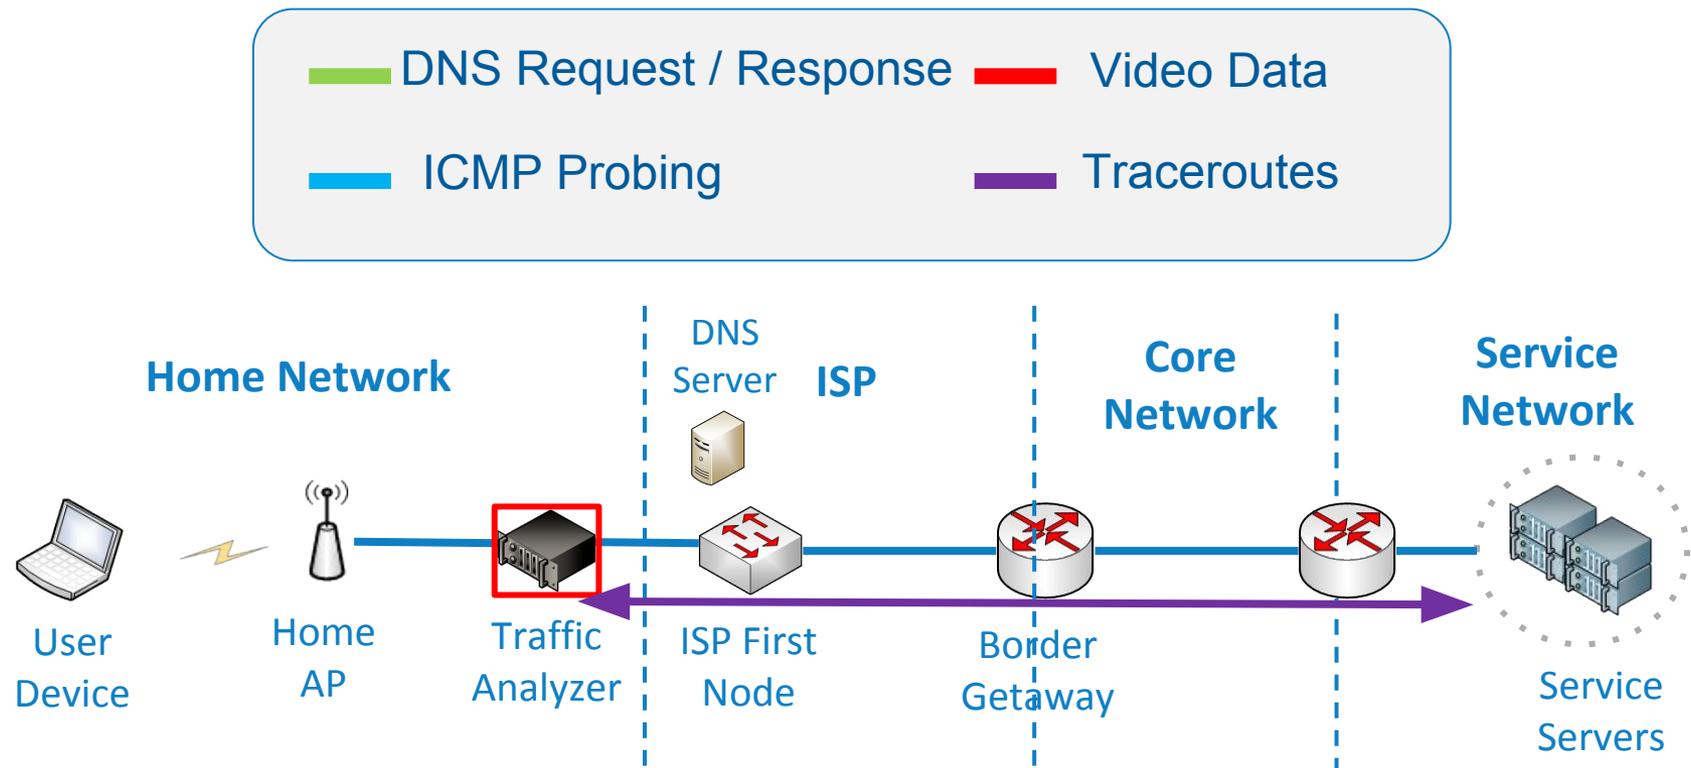
Traffic Analysis: Core Techniques



3 Active probing (ICMP Pings)

- Provides indicator of congestion in home/access networks and of wifi impairments

Traffic Analysis: Core Techniques

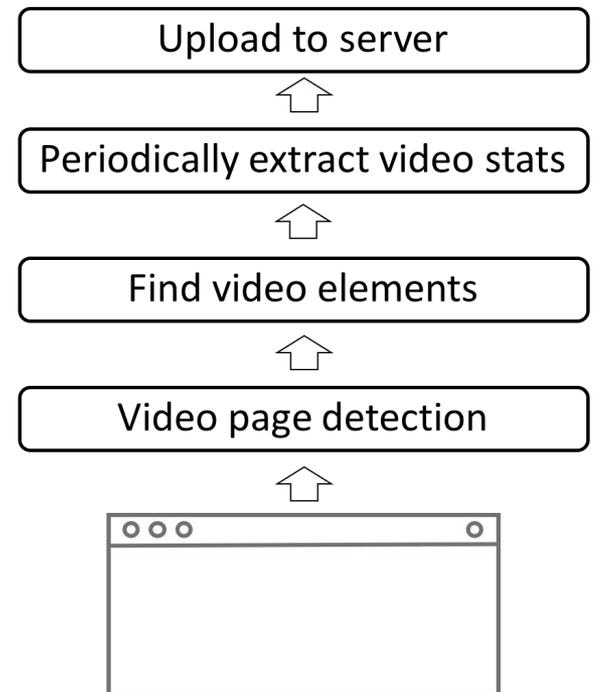


4 TCP-based traceroute to active services

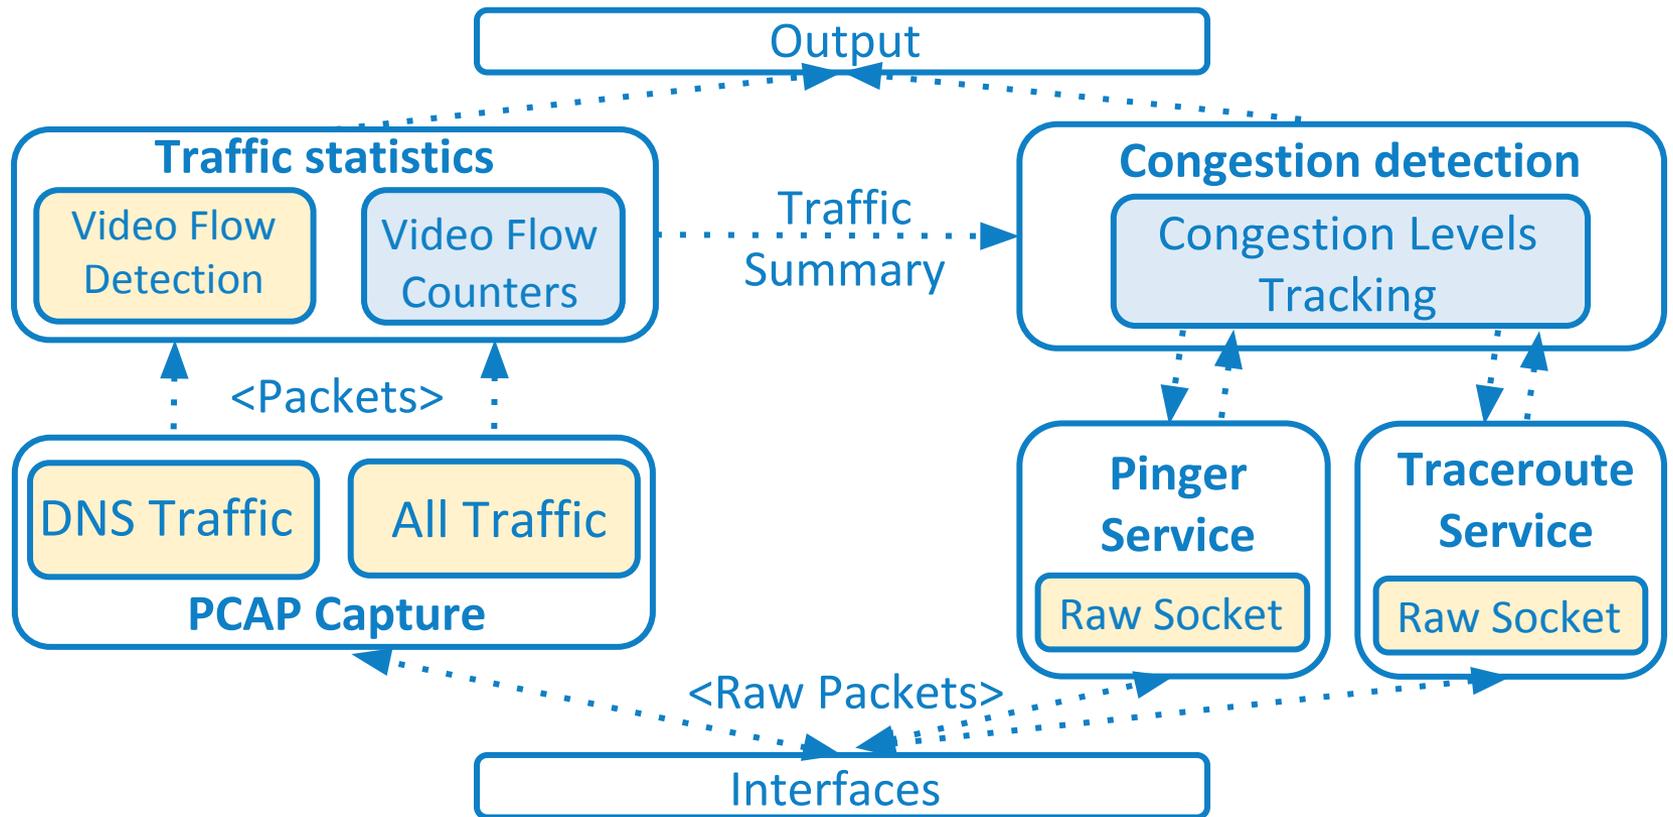
- Provides indicator of congestion at interconnect and change of paths

System Deployment

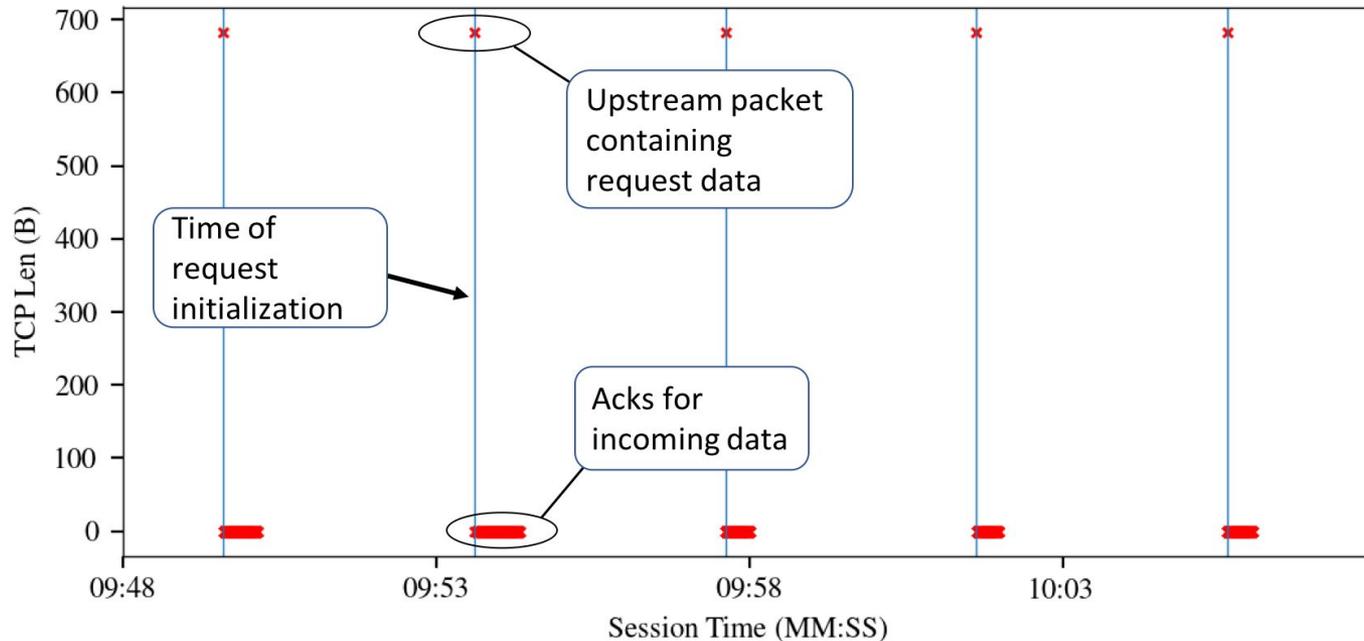
- Currently deployed in heterogeneous collection of homes
 - ~50 in the US
 - ~10 in Paris
- Ground truth collected via a browser extension
 - Information extracted from the HTML video tag
 - Information extracted from Netflix debug output



System Overview

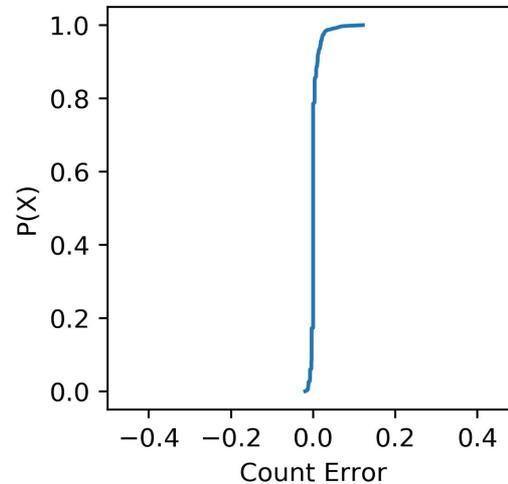
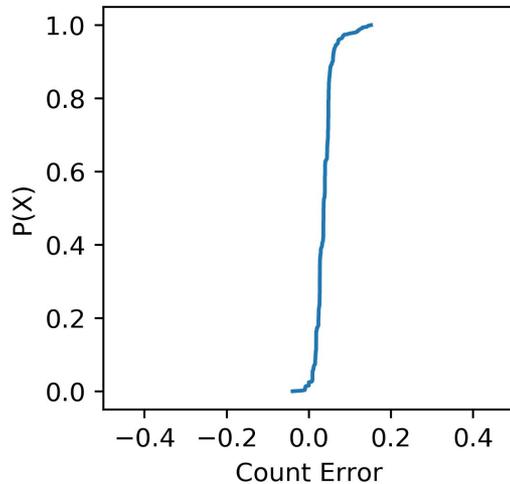
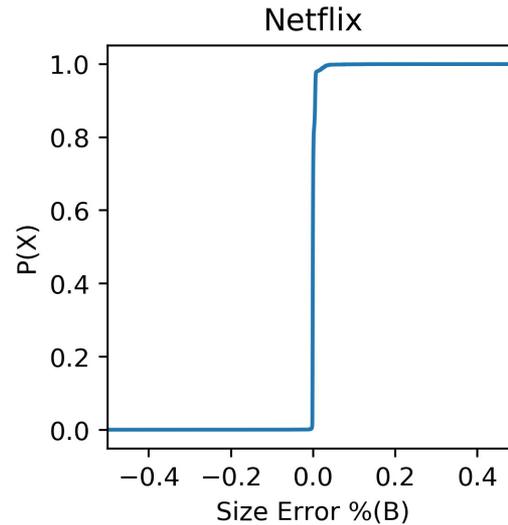
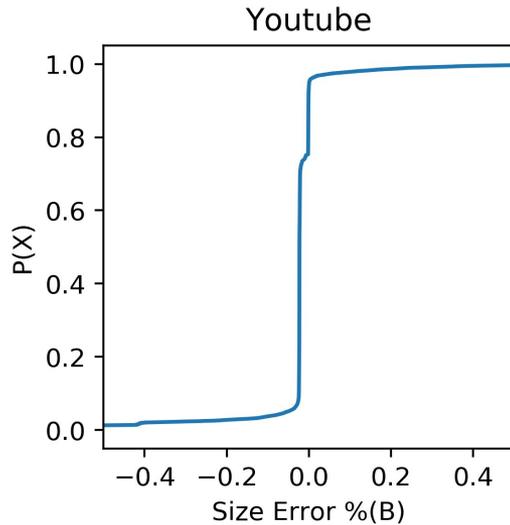


Video Segment Detection



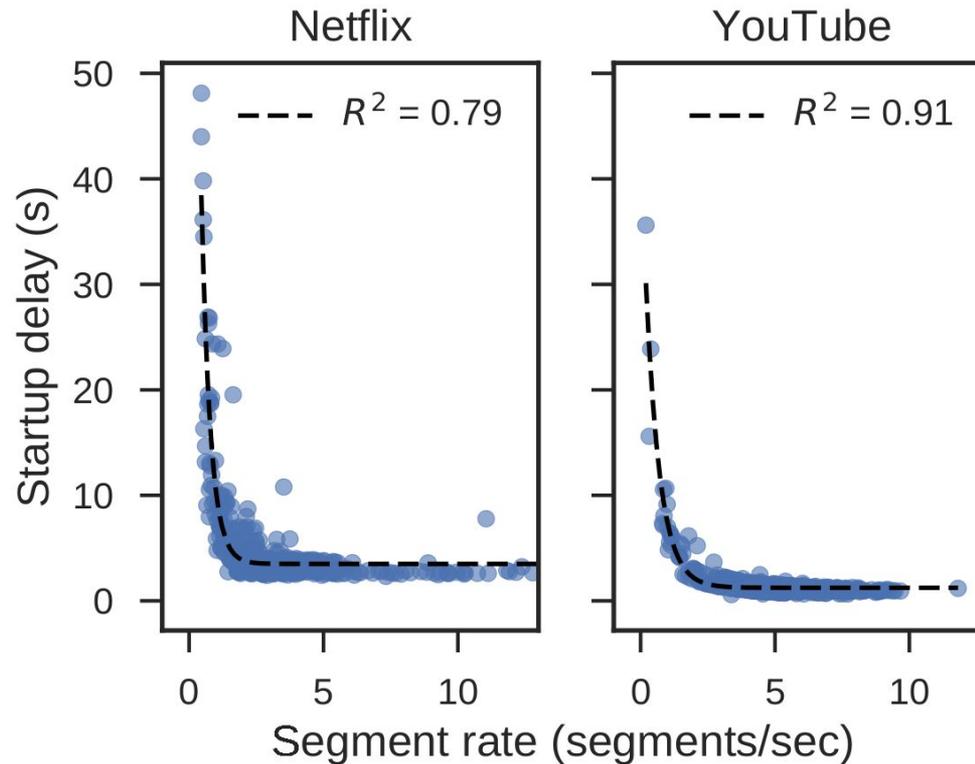
- HTTP traffic for video requests is sequential
- Upstream traffic reveals requests boundaries

Video Segment Detection



- More than 96% Netflix requests are correctly recognized within margin of error (0.21%)
- Challenges:
 - TLS
 - QUIC

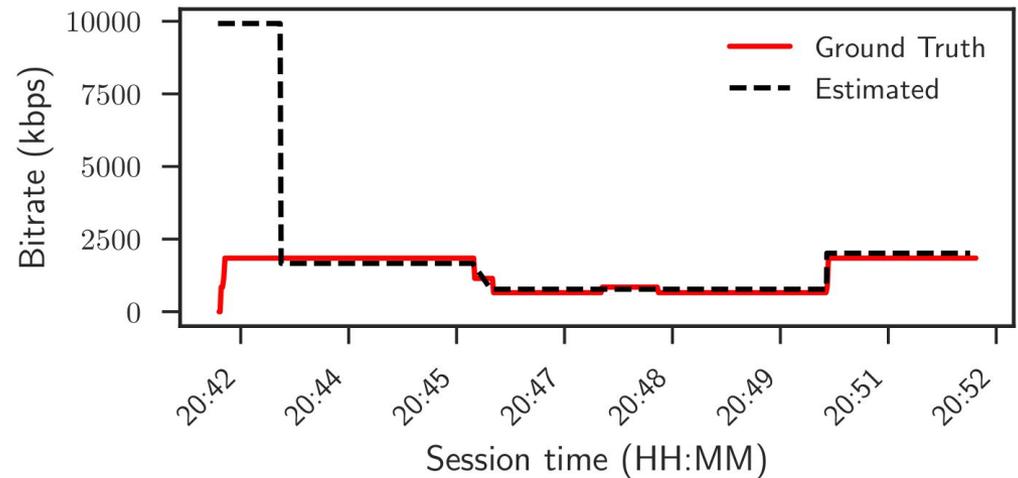
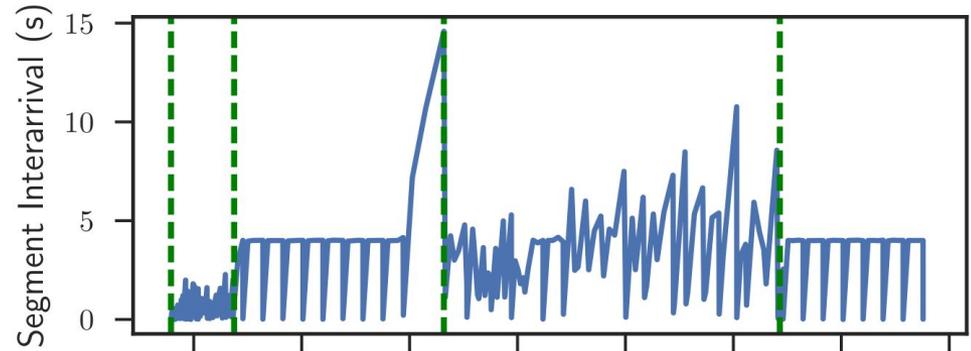
Video Startup Delay



- Exponential fit
- Median errors of 0.32, 0.81, and 1.91 seconds for startup delay ranges 1–2, 2–4, and 4–8 seconds, respectively

Video Bitrate Estimation

- Inter-arrival rate of segments
- Changes in inter-arrival rate allows for detection of quality switch events



Conclusions

- Presented a lightweight system running at the home gateway that allows for categorization and analysis of encrypted traffic
- Future
 - Congestion pinpointing
 - Refinement of buffering / switch detection
 - Proactive network optimizations
 - Modular extensibility - general purpose measurement platform

Questions?

pschmitt@cs.princeton.edu

