NetMicroscope: Passive Measurements of Residential Internet Performance

Renata Teixeira
with
Francesco Bronzino, Sara Ayoubi, Israel Salinas (Inria)
Paul Schmitt, Guilherme Martins,Joon Kim, Nick Feamster (Princeton)

Who cares about residential Internet performance?

- Home users

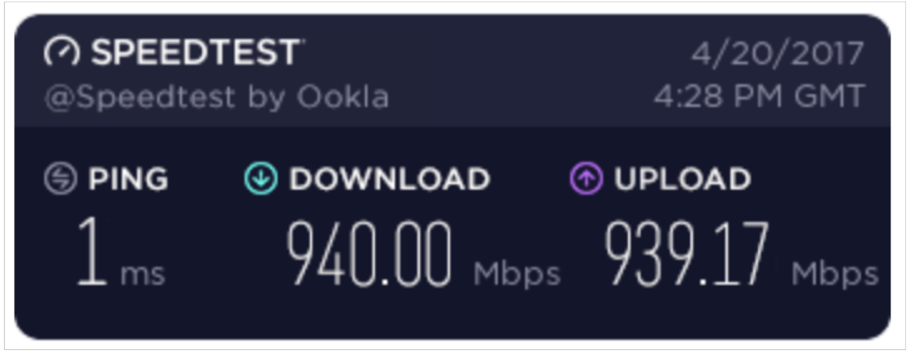

- ISPs, content providers

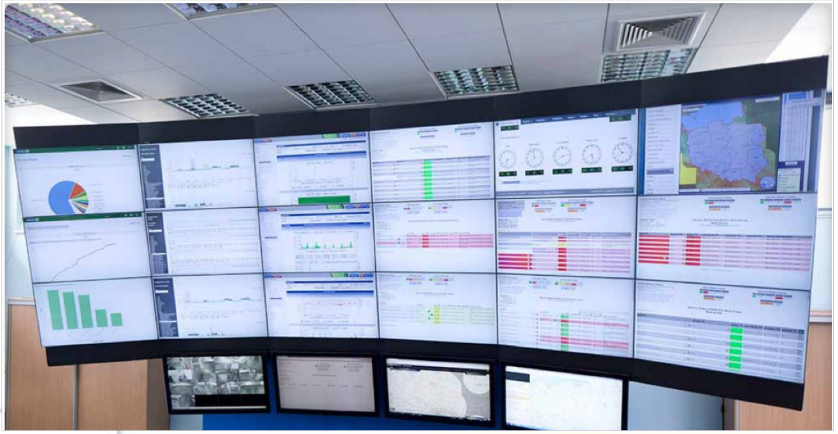

- Regulators, policymakers


Eighth
Measuring Broadband America
Fixed Broadband Report

A Report on Consumer Fixed Broadband Performance in the United States

Federal Communications Commission

Current approach: Active measurements

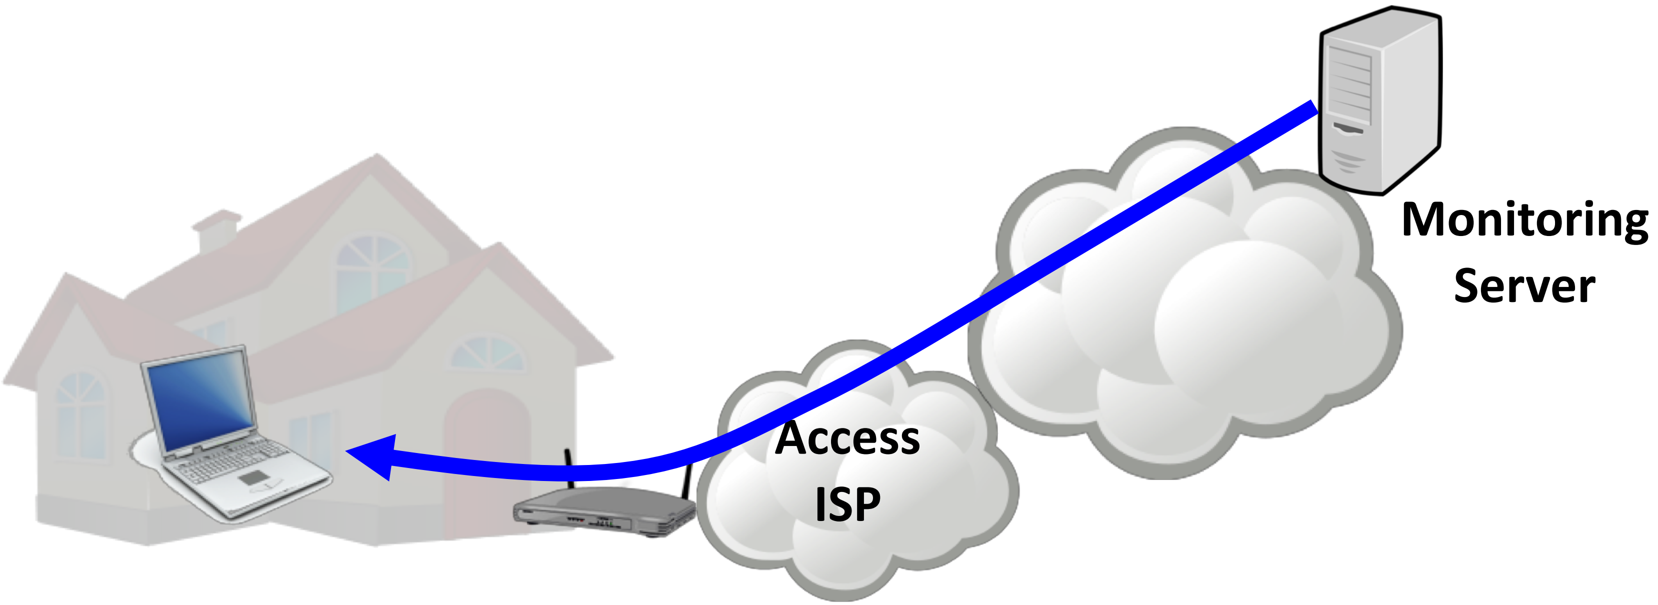

Active measurements are reaching their limits

- Access link may not be the bottlenecks
- “Filling up” path is disruptive
- Measured paths != application paths
- Per-application active measurements != user experience

NetMicroscope

- Measure traffic, infer application performance
  - Passive measurements to infer application quality
  - Targeted active probes to pinpoint bottlenecks

Challenges

- Infer application quality from network traffic
  - Applications have different communication patterns
  - Application traffic is often encrypted
- Passive measurements at increasing line rates
- Distinguish performance per network segments

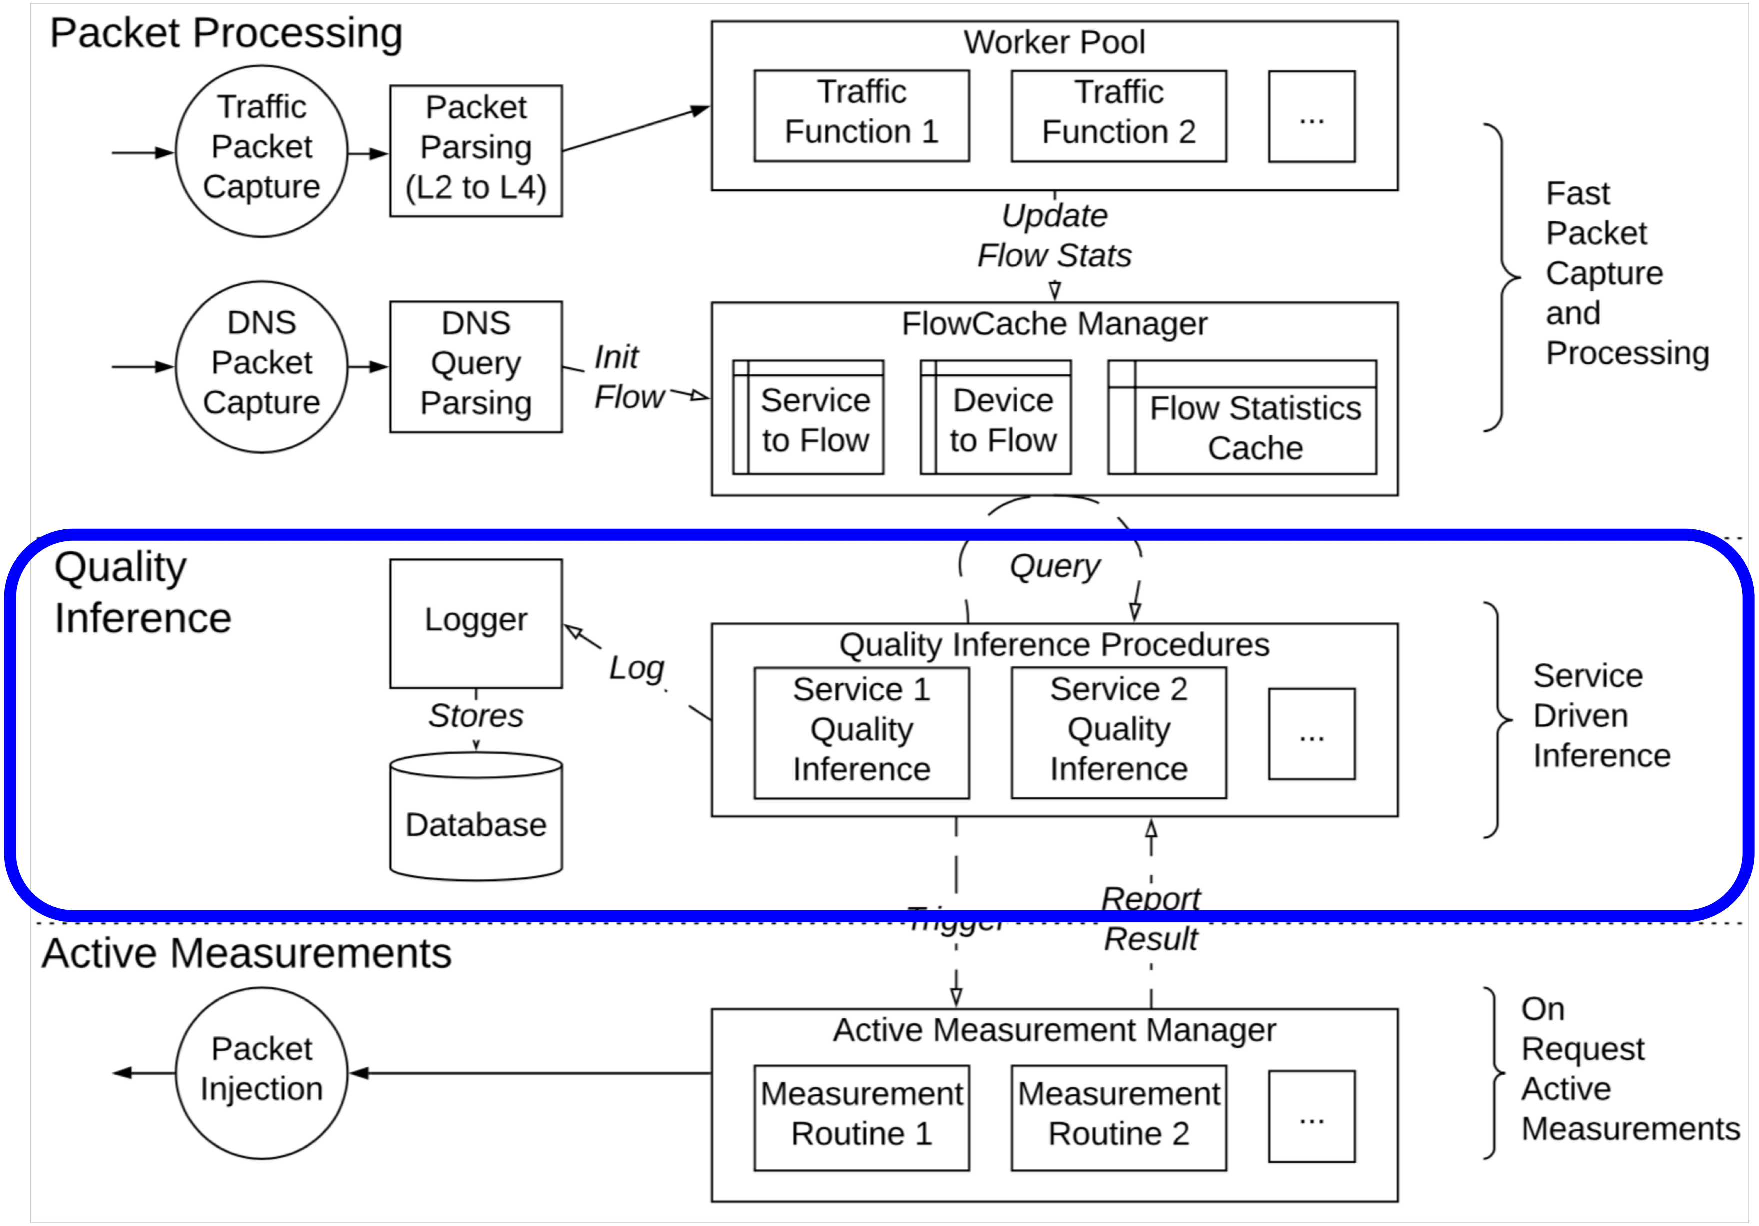
Packet Processing
Traffic Packet Capture
Packet Parsing (L2 to L4)
Worker Pool
Traffic Function 1
Traffic Function 2
...
Update Flow Stats
DNS Packet Capture
DNS Query Parsing
Init Flow
FlowCache Manager
Service to Flow
Device to Flow
Flow Statistics Cache
Fast Packet Capture and Processing
Quality Inference
Logger
Log
Stores
Database
Query
Quality Inference Procedures
Service 1 Quality Inference
Service 2 Quality Inference
...
Service Driven Inference
Trigger
Report Result
Active Measurements
Packet Injection
Active Measurement Manager
Measurement Routine 1
Measurement Routine 2
...
On Request Active Measurements

Use case: IP video

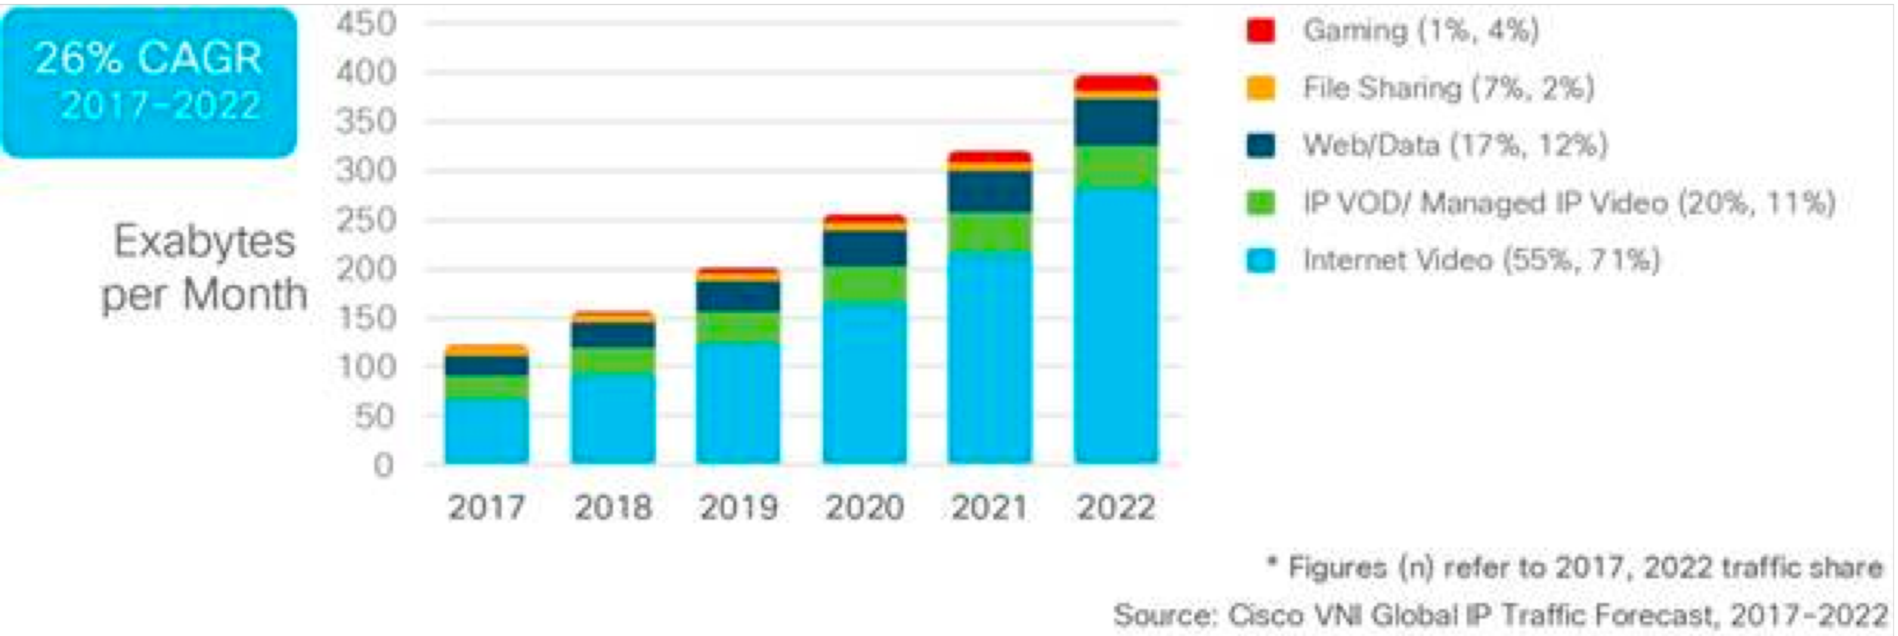

Video delivery is complex

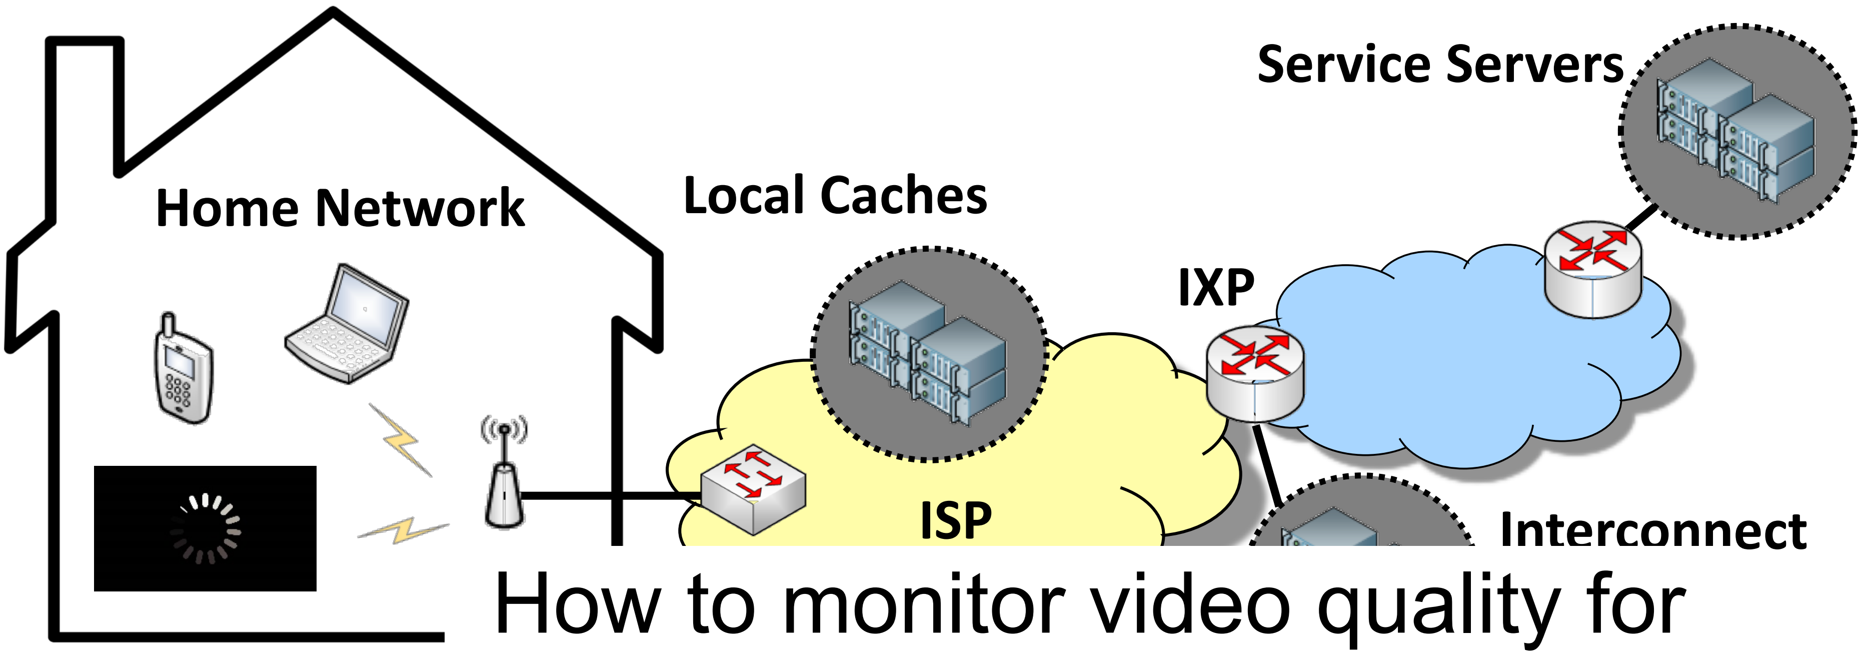

How to monitor video quality for encrypted video traffic?

Challenges of video quality inference

- Identify video streams within network traffic
- Online monitoring at increasing line rates
- Large diversity of video streaming services

Our approach

- Identification of video streams
  - DNS request/response
- Inference of video quality
  - Rely on statistical learning
  - Can we rely only on lightweight features?
  - Do models generalize across video services?
- Deployed in home networks
  - Between modem and WiFi router
  - Implemented for low-cost devices
    - Raspberry Pi, Odroid

Statistical learning to infer video quality

- Inference goal: Video quality metrics
  - Startup delay
  - Video resolution
  - Resolution changes
  - Rebuffering
- Training data with ground truth from browser
  - Services: Netflix, Youtube, Twitch, Amazon Prime
  - Controlled and in-home experiments
    - Over 11K video sessions

Input: Encrypted video traffic

Network layer	Transport layer	Application layer
throughput up/down	#flags up/down	seg. sizes (all previous, last-10, cumulative)
throughput down diff	rcv window size up/down	seg. request interarrivals
pkt count up/down	idle time up/down	seg. completions interarrivals
byte count up/down	goodput up/down	#pending requests
pkt interarrivals up/down	bytes per pkt up/down	#downloaded seg.
#parallel flows	round trip time	#requested seg.
	bytes in flight up/down	
	#retransmissions up/down	
	#out of order pks up/down	

Inria

Modeling approach

- Startup delay
  - Random forest regressor
- Video resolution
  - Random forest multi-class classifier
    - Classes: 240p, 360p, 480p, 720p, and 1080p
- Resolution changes
  - Random forest binary classifier
- Rebuffering
  - Random forest binary classifier

CAN WE RELY ONLY ON LIGHTWEIGHT FEATURES?

Feature importance: Video resolution

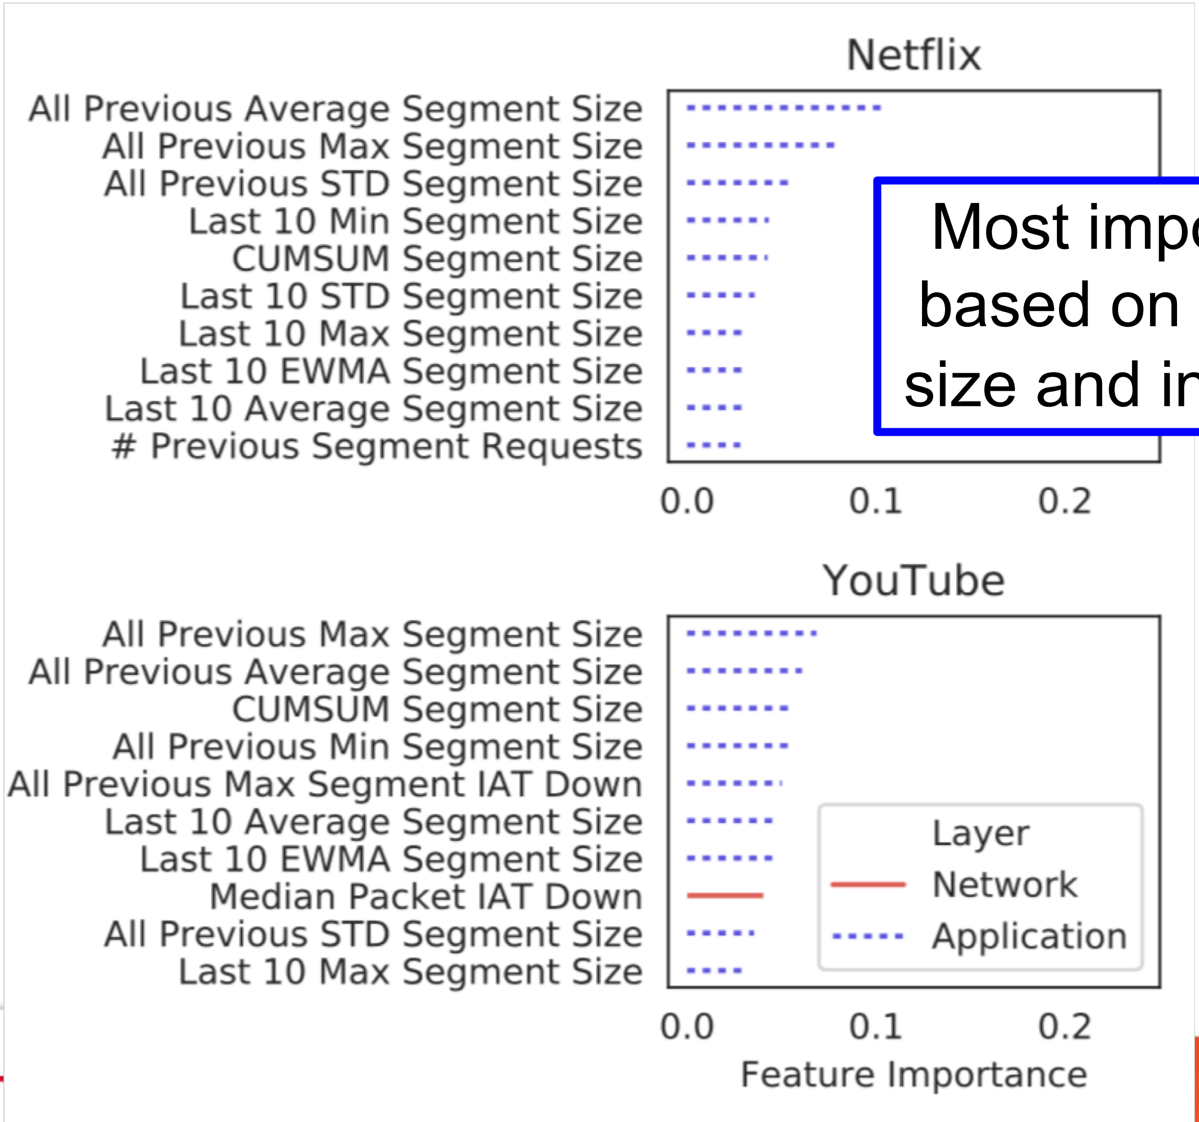

Most important features based on video segment size and interarrival times

DO MODELS GENERALIZE ACROSS VIDEO SERVICES?

General vs. specific models for video resolution

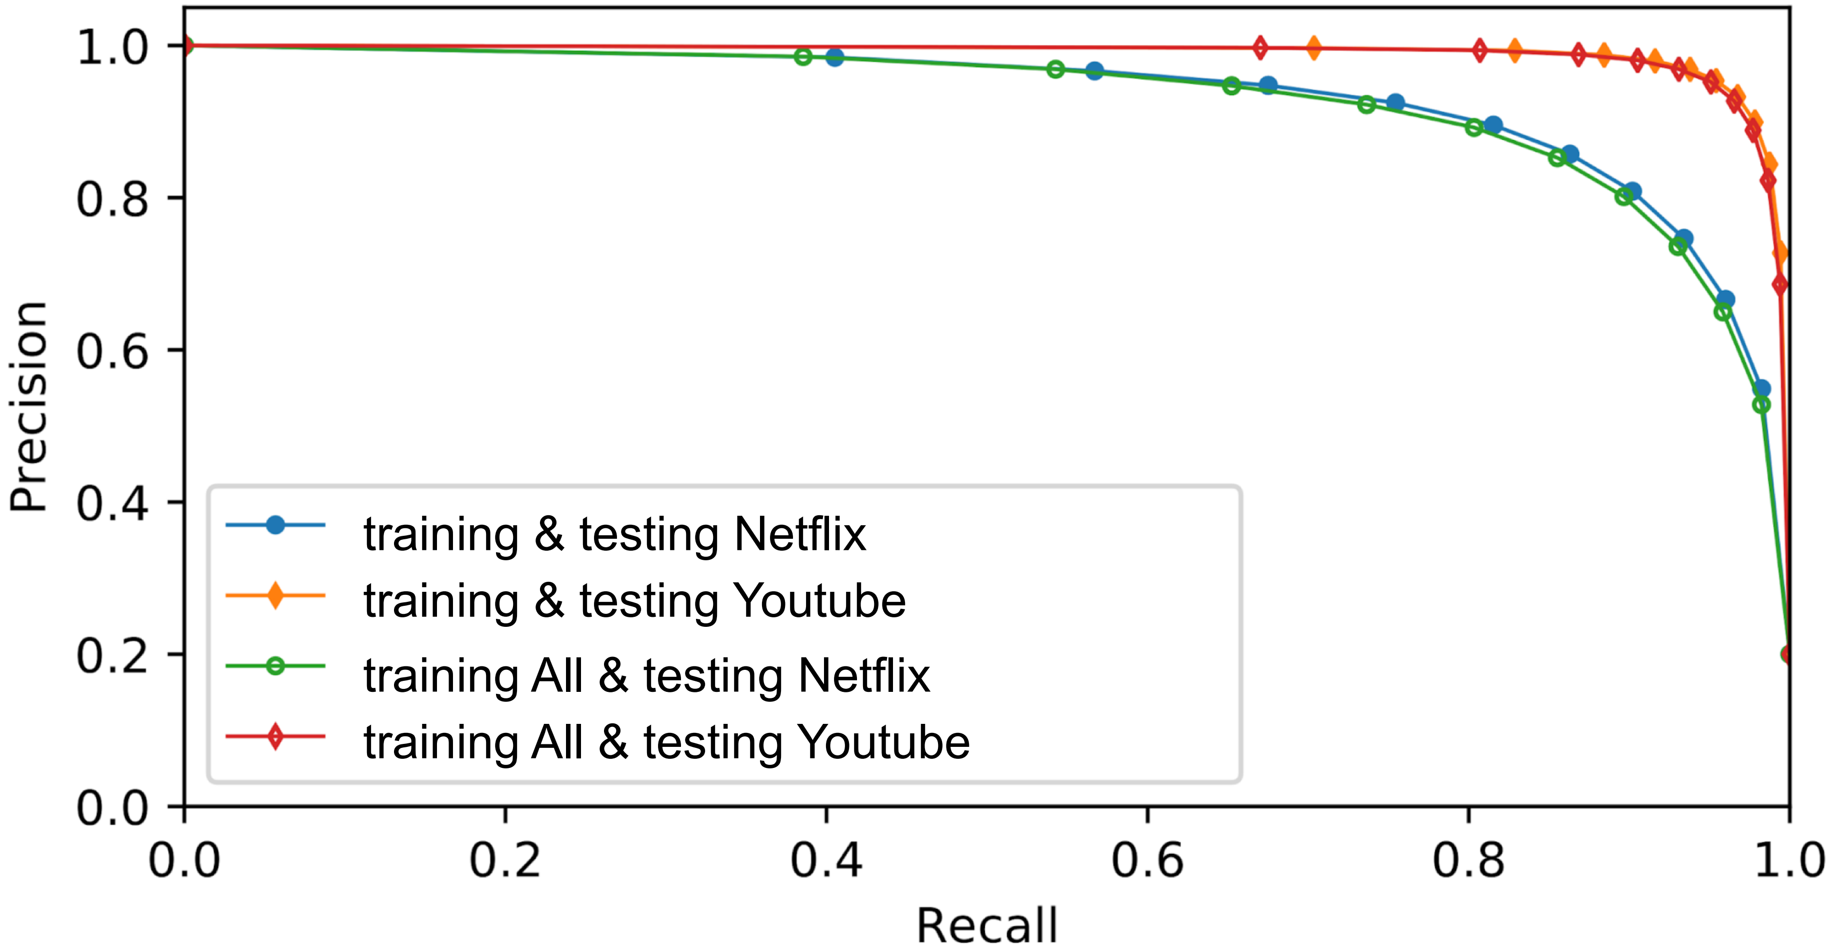

Deployment

- Instrumented homes
  - ~10 in Paris
  - ~50 in the US

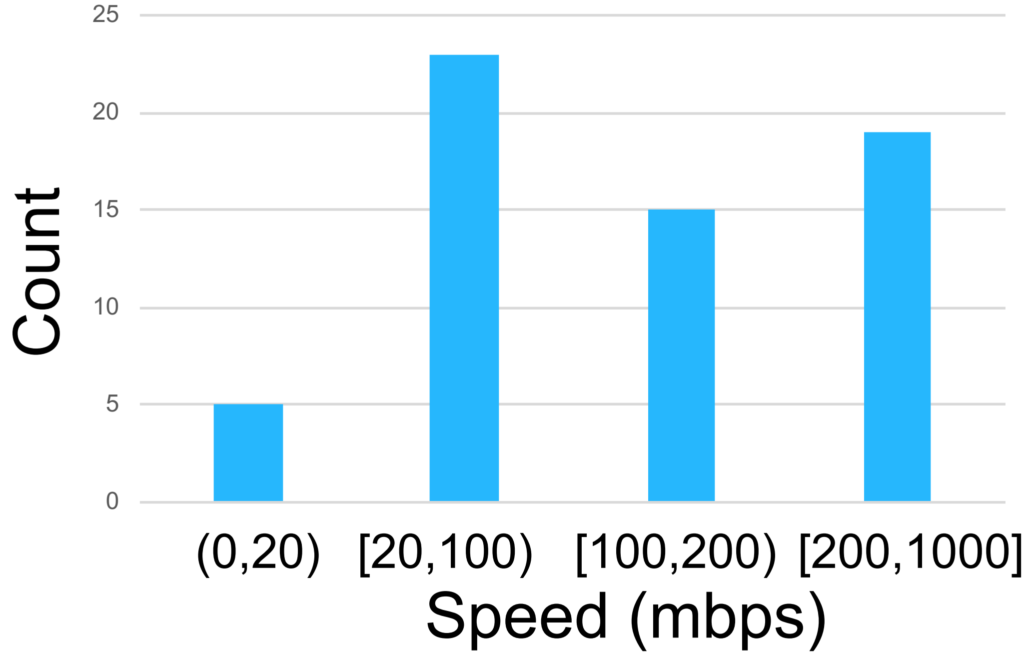

Preliminary lessons

- Identification of video sessions
  - Auto-play merges sessions
  - DNS method fails for some devices
- Inference of video quality
  - Harder to model rebuffering and resolution switches
  - Resolution model needs adjustment for more diverse set of devices

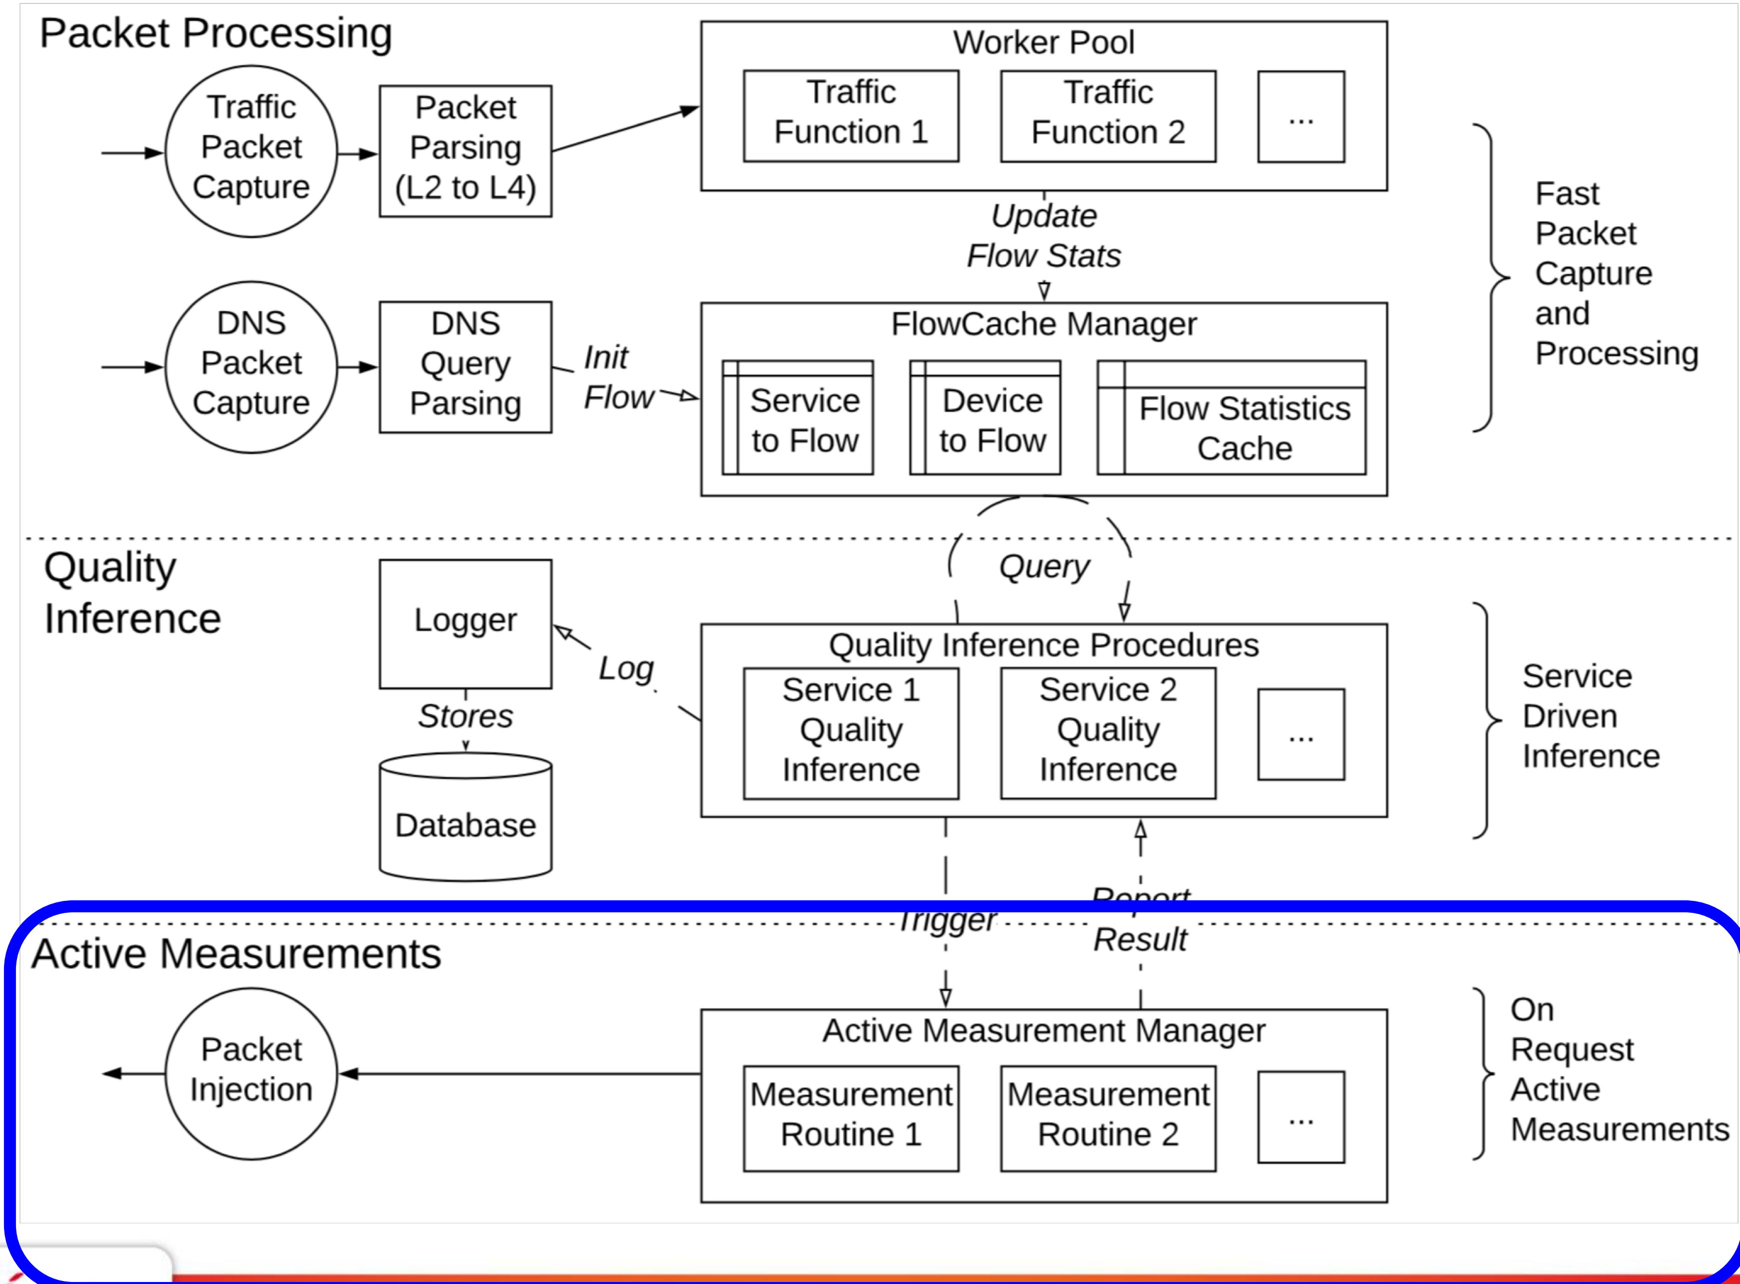
Packet Processing
Traffic Packet Capture
Packet Parsing (L2 to L4)
Worker Pool
Traffic Function 1
Traffic Function 2
...
Update Flow Stats
DNS Packet Capture
DNS Query Parsing
Init Flow
FlowCache Manager
Service to Flow
Device to Flow
Flow Statistics Cache
Fast Packet Capture and Processing
Quality Inference
Logger
Log
Stores
Database
Query
Quality Inference Procedures
Service 1 Quality Inference
Service 2 Quality Inference
...
Service Driven Inference
Trigger
Report Result
Active Measurements
Packet Injection
Active Measurement Manager
Measurement Routine 1
Measurement Routine 2
...
On Request Active Measurements

Tracing paths of application flows

- Problem
  - Traceroute may not capture application paths

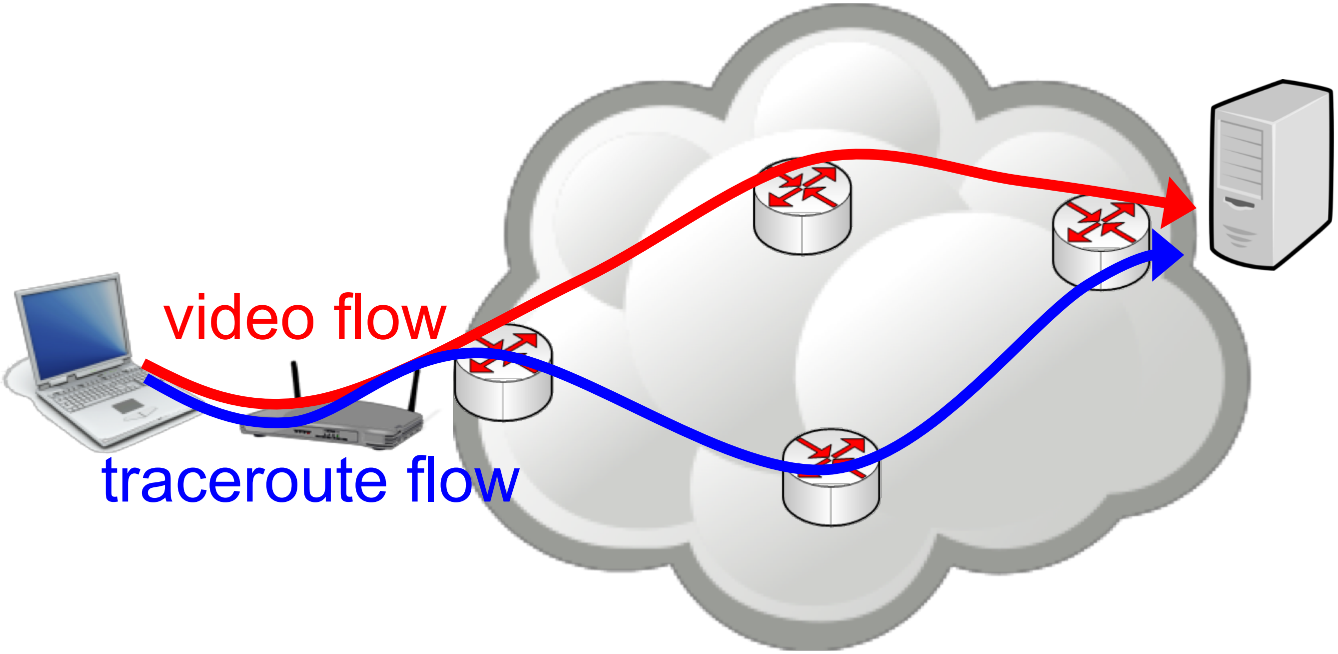

Service traceroute

- Basics
  - Listen to application traffic
  - Embeds traceroute probes within application flow
- New features
  - Signature DB to identify flows of given applications
  - Support for UDP
  - Support to trace multiple concurrent flows

“Service Traceroute: Tracing Paths of Application Flows”.
I. Morandi et al., to appear in PAM 19

Looking ahead

- How does speed relate to application quality?
- How to generalize quality inference to other applications?
- How to preserve privacy?
- How to regulate Internet access using application quality inference?