nsdi '21

Fast and Light Bandwidth Testing for Internet Users

Xinlei Yang, Xianlong Wang, Zhenhua Li, Yunhao Liu,
Feng Qian, Liangyi Gong, Rui Miao, Tianyin Xu

Outlines

1. Background
2. Motivation
3. State-of-the-Art
4. Novel Design
5. Evaluation
6. System Demo
7. Conclusion

1. Background

Bandwidth testing services (BTSes) are widely used

- Core component of many network applications
- Cited by government reports & trade press
- Handy measurement tools for Internet users

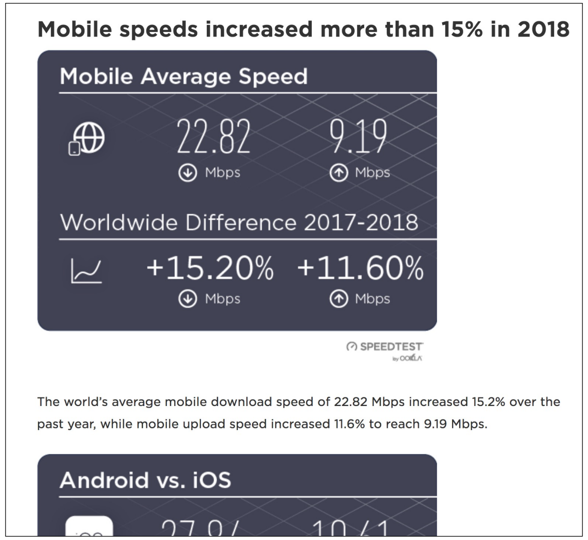

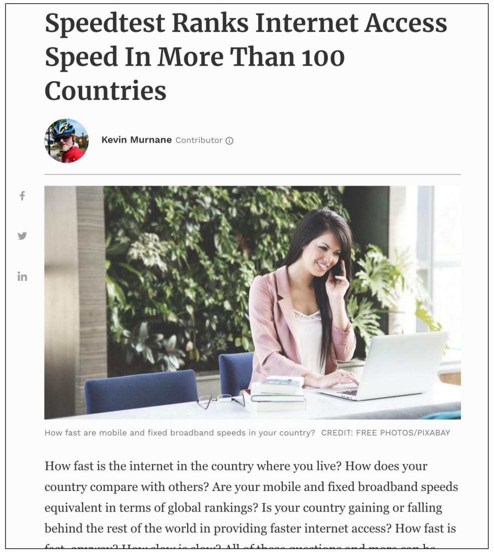

Speedtest Ranks Internet Access Speed In More Than 100 Countries

Kevin Murnane Contributor

How fast are mobile and fixed broadband speeds in your country? CREDIT: FREE PHOTOS/PIXABAY

How fast is the internet in the country where you live? How does your country compare with others? Are your mobile and fixed broadband speeds equivalent in terms of global rankings? Is your country gaining or falling behind the rest of the world in providing faster internet access? How fast is

SPEEDTEST

34,646,541,701

Tests taken with Speedtest to date

1. Background

BTSes are becoming increasingly important

- Virtual Network Operators (VNO) catching on
- Wireless access becoming ubiquitous
- Bandwidth-hungry apps (*e.g.*, UHD videos, VR/AR) emerging

Google Fi

Straight Talk wireless

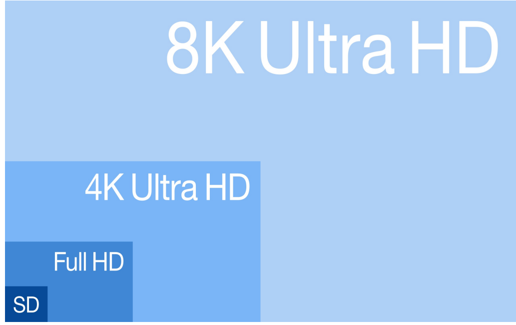

2. Motivation

Today's BTSes are not satisfactory

- Long test duration
- Excessive data usage
- Low accuracy for most BTSes

Example ***mmWave 5G, 1.15-Gbps downlink bandwidth***

BTSes	Duration (s)	Data Usage	Accuracy
Speedtest.att.com	19.1	1.37 GB	0.42
Sourceforge.net	20.8	2.75 GB	0.81
Fast.com	13.5	1.20 GB	0.68
SpeedTest.net	15.7	1.94 GB	0.87

2. Motivation

Today's BTSes are not satisfactory

- Long test duration
- Excessive data usage
- Low accuracy for most BTSes

Example ***mmWave 5G, 1.15-Gbps downlink bandwidth***

BTSes	Duration (s)	Data Usage	Accuracy

Can bandwidth testing be
fast, light, and ***accurate*** simultaneously?

3. State-of-the-Art

Popular Bandwidth Testing Websites	Alexa + Google	**18** popular bandwidth testing websites
Commercial Bandwidth Testing Apps	WiFi Master — WiFiMaster	A popular Android/iOS app with **800 million users**
Important Bandwidth Testing Interfaces	android — Android 11	**5G-oriented** bandwidth testing Android SDK APIs

3. State-of-the-Art

Research methodology

Small-scale study

1. Network traffic tracing
2. System reverse engineering

Device	Location	Network	Test Results
PC-1	U.S.	Residential broadband	88–96 Mbps
PC-2	Germany	Residential broadband	91–97 Mbps
PC-3	China	Residential broadband	90–97 Mbps
Samsung GS9	U.S.	LTE (60Mhz/1.9Ghz)	60–100 Mbps
Xiaomi XM8	China	LTE (40Mhz/1.8Ghz)	58–89 Mbps
Samsung GS10	U.S.	5G (400Mhz/28Ghz)	0.9–1.2 Gbps
Huawei HV30	China	5G (160Mhz/2.6Ghz)	0.4–0.7 Gbps

Large-scale benchmarking

Azure
DigitalOcean
Tencent Cloud
aws
Alibaba Cloud
VULTR

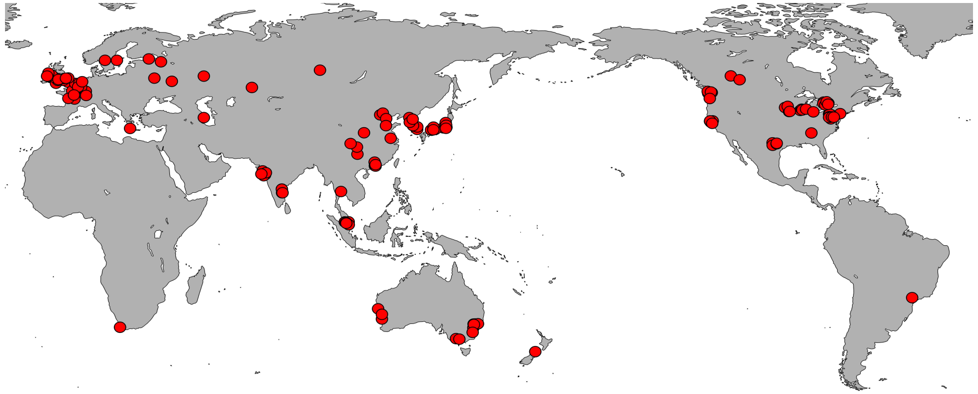

3. State-of-the-Art

Summarizing

BTS	# Servers	Bandwidth Test Logic	Duration	Accuracy (Testbed / 5G)	Data Usage (Testbed / 5G)
TBB	12	average throughput in all connections	10 s	0.59 / 0.31	42 MB / 481 MB
SpeedOf	116	average throughput in the last connection	8–230 s	0.76 / 0.22	61 MB / 256 MB
BWP	18	average throughput in the fastest connection	13 s	0.81 / 0.35	74 MB / 524 MB
SFtest	19	average throughput in all connections	20 s	0.89 / 0.81	194 MB / 2,013 MB
ATTtest	75	average throughput in all connections	15–30 s	0.86 / 0.53	122 MB / 663 MB
Xfinity	28	average all throughput samples	12 s	0.82 / 0.67	107 MB / 835 MB
FAST	~1,000	average stable throughput samples	8–30 s	0.80 / 0.72	45 MB / 903 MB
SpeedTest	~12,000	average refined throughput samples	15 s	0.96 / 0.92	150 MB / 1,972 MB
Android API-A	0	directly calculate using system configs	< 10 ms	NA / 0.09	0 / 0

TBB: *thinkbroadband.com,* ***SpeedOf:*** *speedof.me,* ***BWP:*** *bandwidthplace.com,* ***SFtest:*** *sourceforge.net,* ***ATTtest:*** *Speedtest.att.com,* ***Xfinity:*** *speedtest.xfinity.com,* ***FAST:*** *fast.com,* ***SpeedTest:*** *speedtest.net,* ***Android API-A:*** *getLinkDownstreamBandwidthKbps()*

3. State-of-the-Art

Reflection of bandwidth testing

Ideal case

Real case

User's access link bandwidth

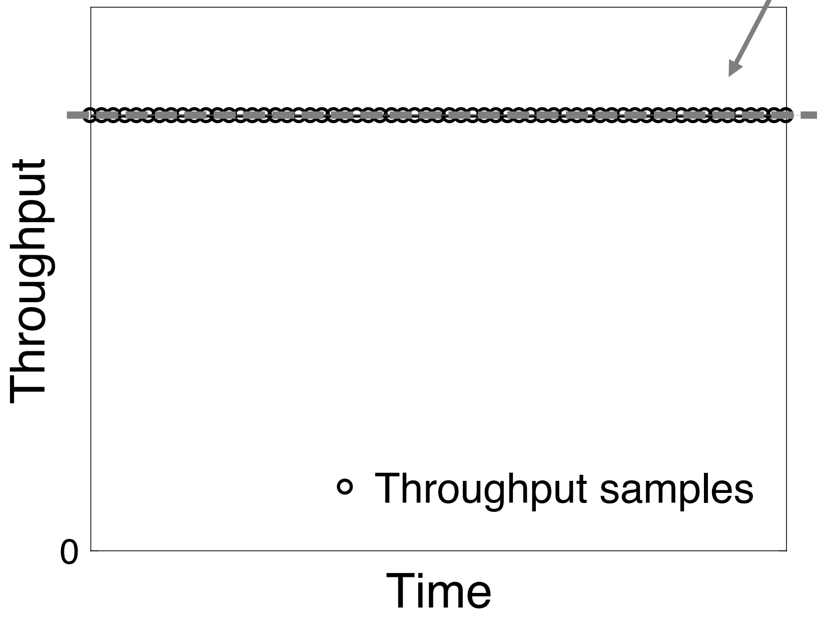

Noises

(congestion control, link sharing, *etc.*)

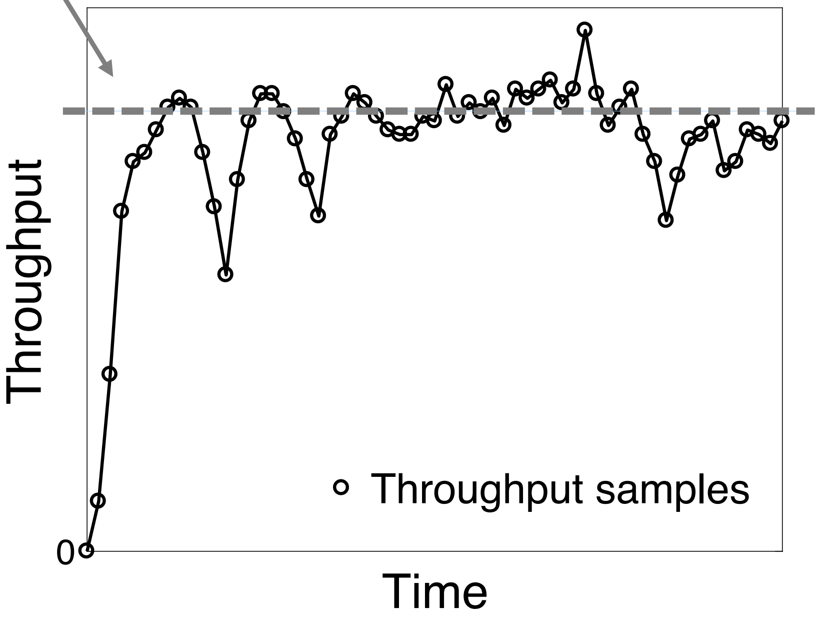

3. State-of-the-Art

Combating noises

Space Dimension

Speedtest.net

Our strength is in our hosted servers

The accuracy and high-quality performance of Speedtest is made possible through the 11,000+ servers around the world that host our Speedtest server daemon. This robust network of servers enables us to ensure that our users get local readings wherever they are on the planet.

*Using large-scale test server deployments (**spatial redundancies**) to ensure high-quality network connections, **largely reducing noises.***

Time Dimension

Fast.com

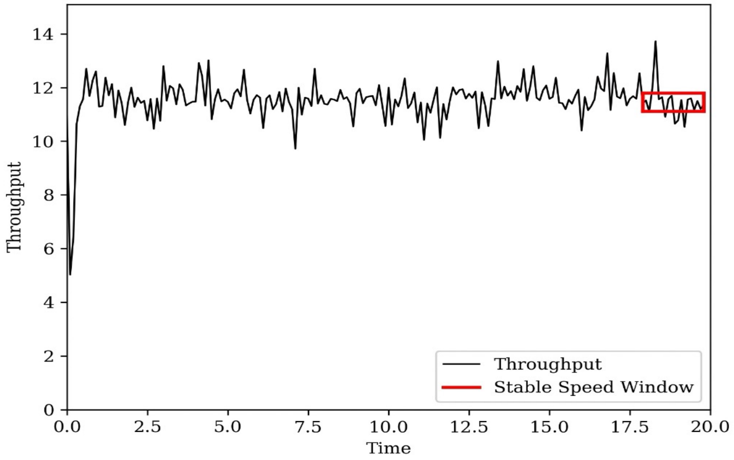

Test duration: often 20 - 30 s

*Using long test duration (**temporal redundancies**) to **wait for the coming of sufficient desired samples**.*

4. Novel Design

Most of today's BTSes use excessive **temporal and spatial redundancies** for **combating** noises

↓

Large-scale network deployments, long test duration, and excessive data usage

↓

Can we **accommodate and exploit** the noises rather than exhaustively suppress the impact of them?

4. Novel Design

- Re-consider BTS through rejection sampling

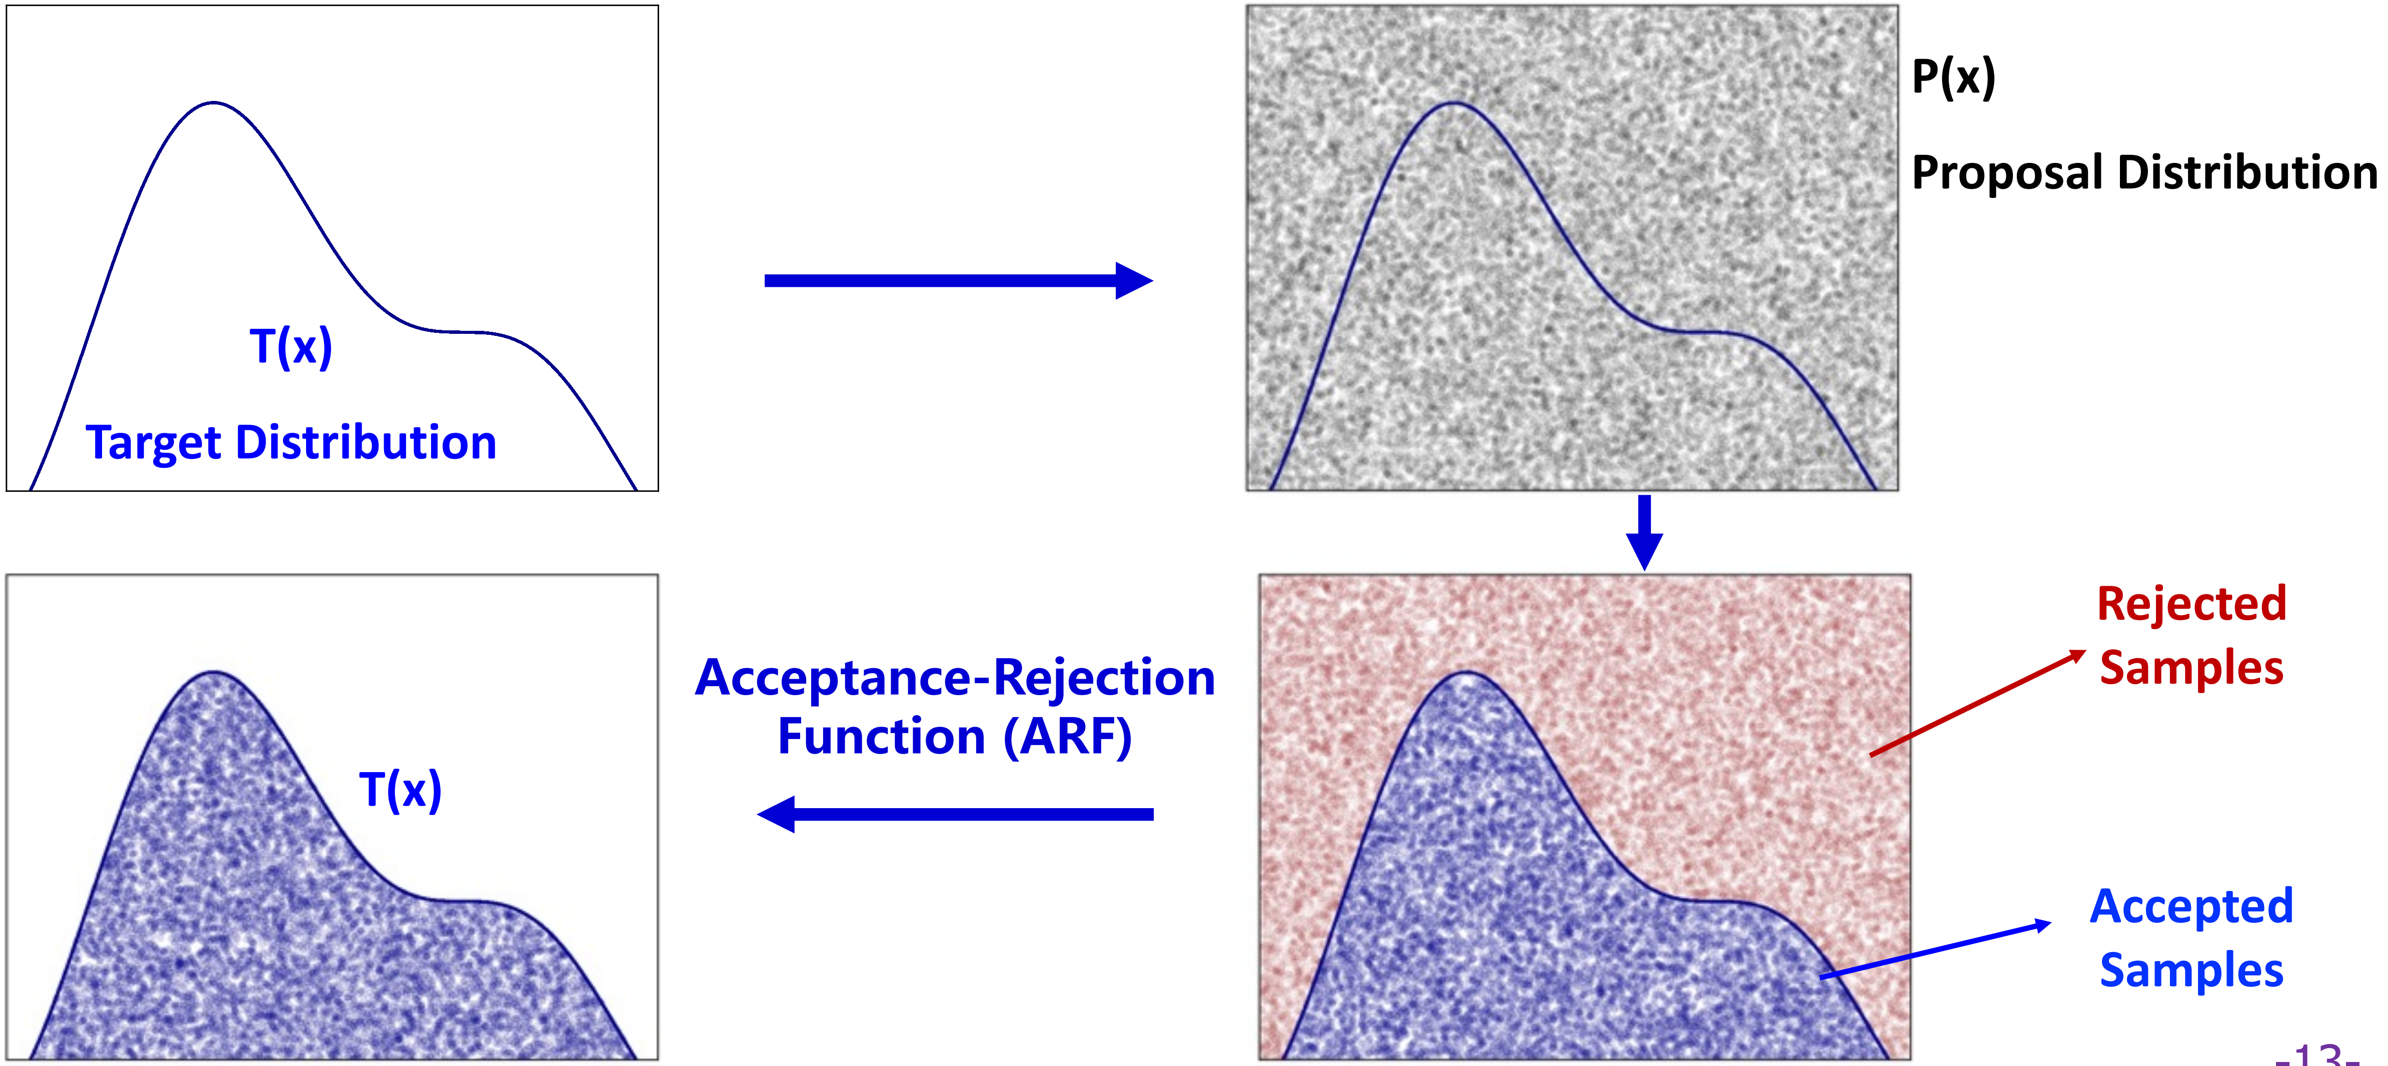

4. Novel Design

Modeling the bandwidth testing process

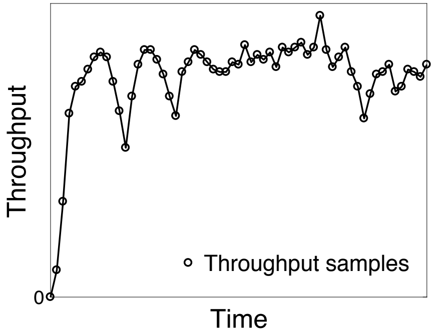

Samples following the P(x) distribution

Rejection Sampling

ARF

T(x)

Accept

Accepted samples

Reject

Rejected samples

Bandwidth Estimation

4. Novel Design

Modeling the bandwidth testing process

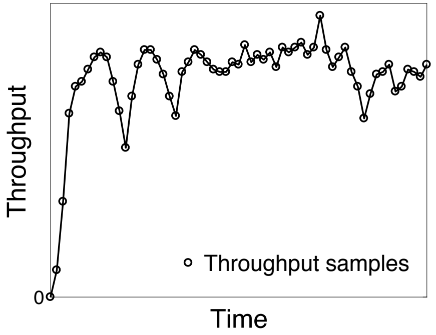

Samples following the P(x) distribution

Fuzzy Rejection Sampling

ARF

T(x) ?

Accept

Accepted samples

Bandwidth Estimation

Reject

Rejected samples

4. Crucial Interval Sampling (CIS)

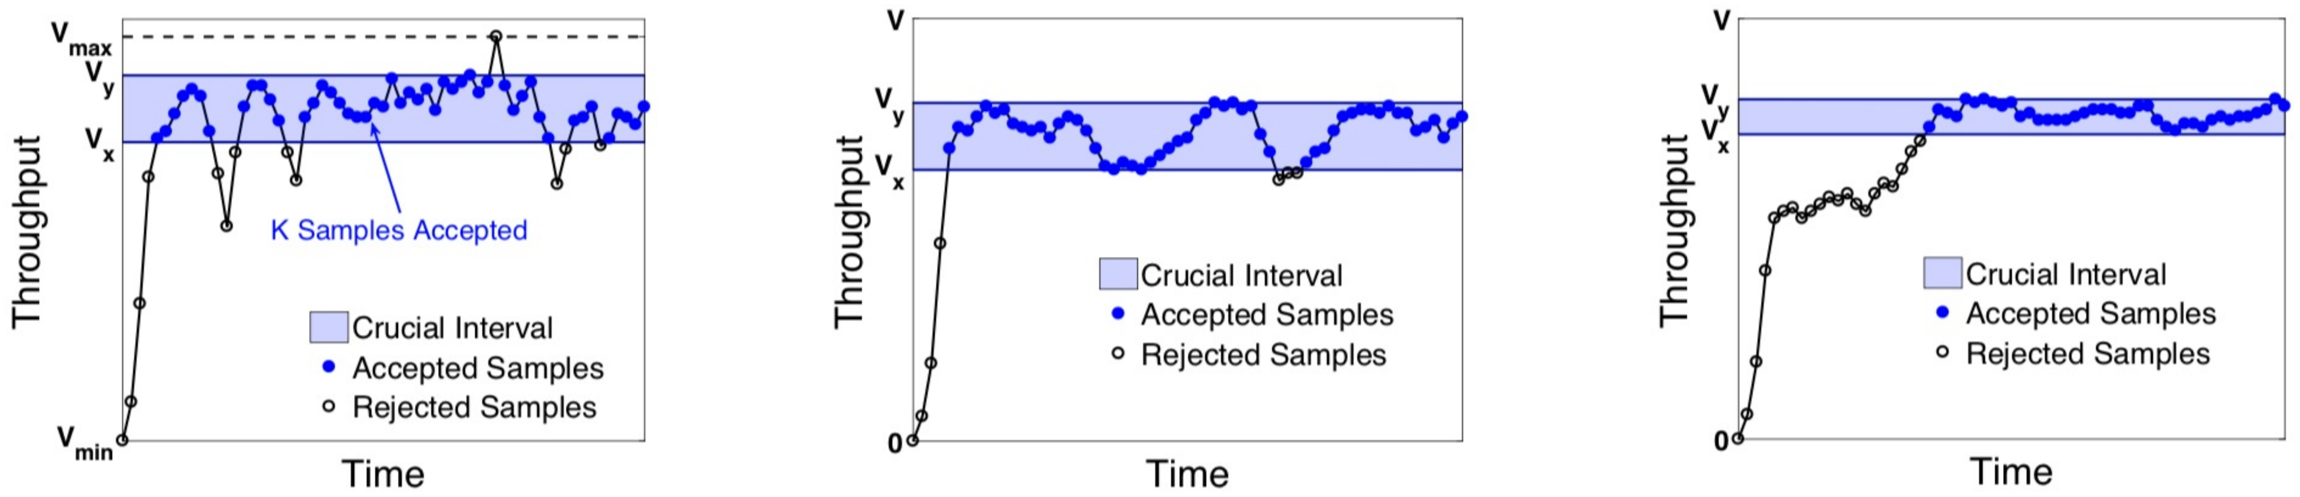

Key Findings:

- Noise samples are scattered **across a wide throughput interval**
- Desired samples tend to concentrate **within a narrow interval**

Crucial Interval: ***(ARF)***

$$F(V_x, V_y) = Density \times Size = \frac{V_{max} - V_{min}}{N} \cdot \frac{K^2(V_x, V_y)}{V_y - V_x}, \quad V_y - V_x \geq \frac{V_{max} - V_{min}}{N - 1}$$

*A throughput interval (1) whose **density is as high as possible**; and (2) which contains **as many samples as possible***

4. Crucial Interval Sampling (CIS)

Crucial intervals converge quickly

Accepted Samples

intensify the crucial interval by making it denser

Rejected Samples

help better "contrast" the crucial interval

Both accepted and rejected samples are exploited to make bandwidth tests fast and light.

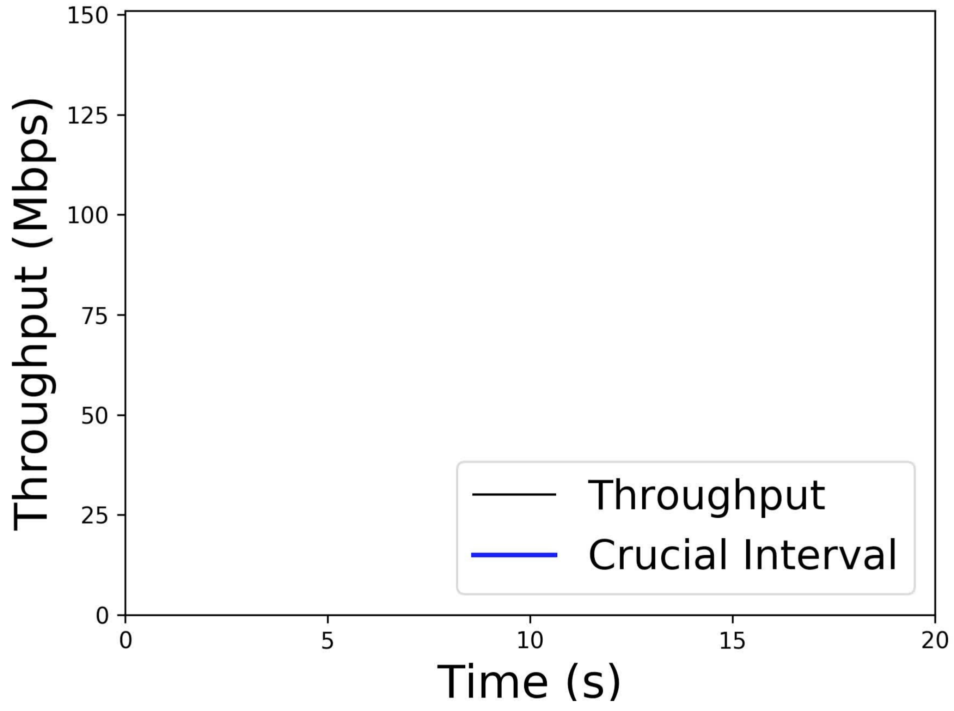

Video available at https://youtu.be/lgZOy59im7M

4. Crucial Interval Sampling (CIS)

Convex hull acceleration

Brute-force mechanism	Walking through all the throughput samples to find the crucial interval.	O(N^2)
Convex hull acceleration	Dynamically maintaining a convex hull for quickly finding the crucial interval.	O(NlogN)

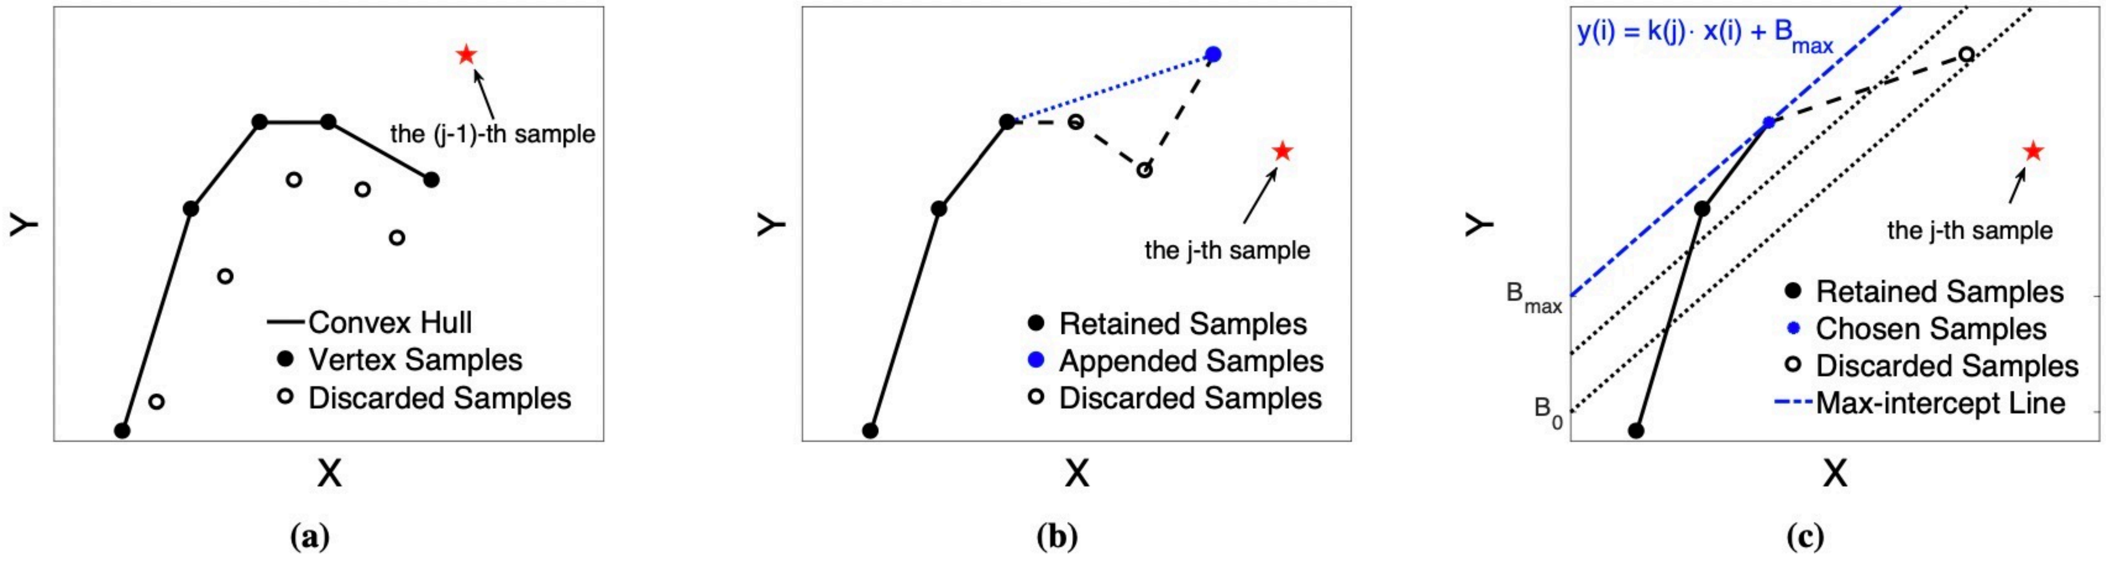

4. Elastic Bandwidth Probing (EBP)

Crucial interval not effective

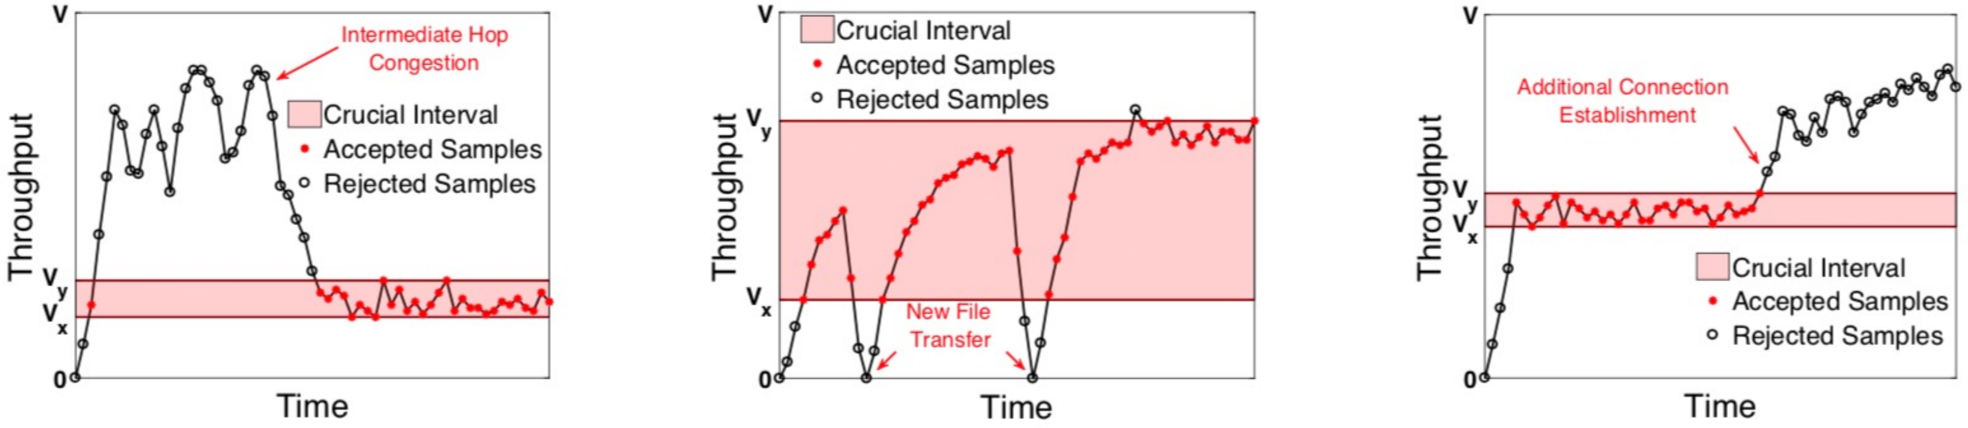

- **BBR:** emerging congestion control mechanism with a built-in bandwidth probing scheme
- Leveraging and improving BBR to realize elastic bandwidth probing
- Making crucial interval always effective

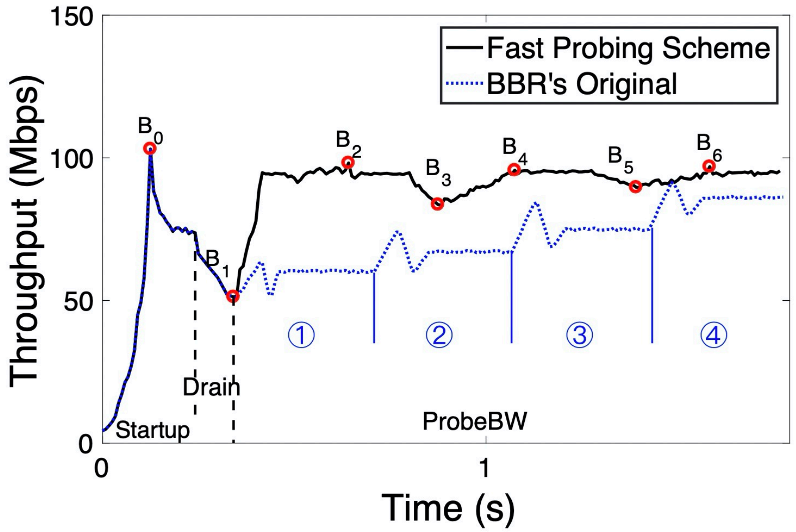

4. Novel Design

Architecture of FastBTS

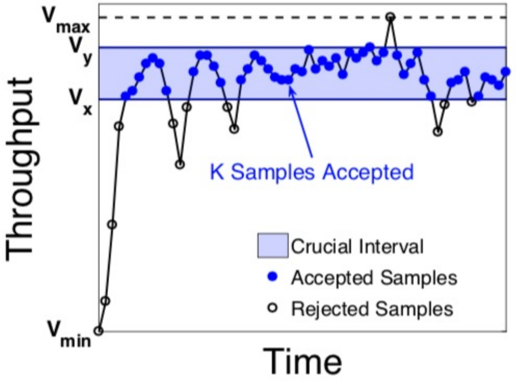

Crucial Interval Sampling

Fuzzy Rejection Sampling

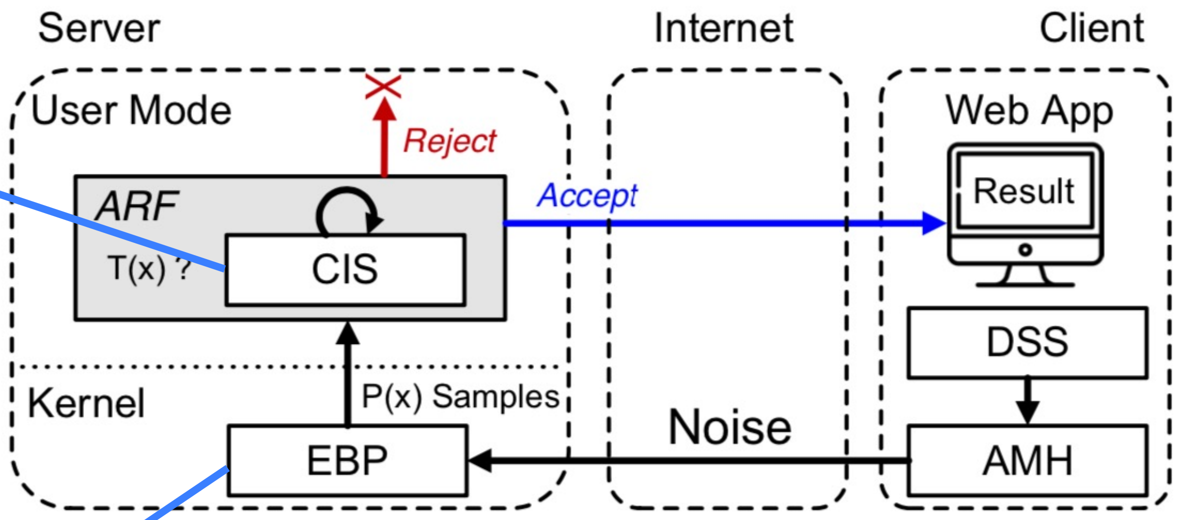

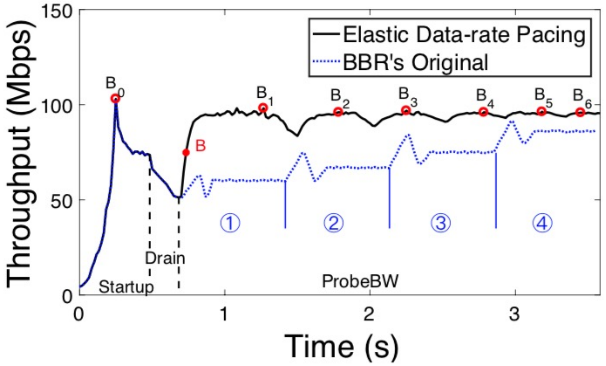

Elastic Bandwidth Probing

4. Novel Design

Architecture of FastBTS

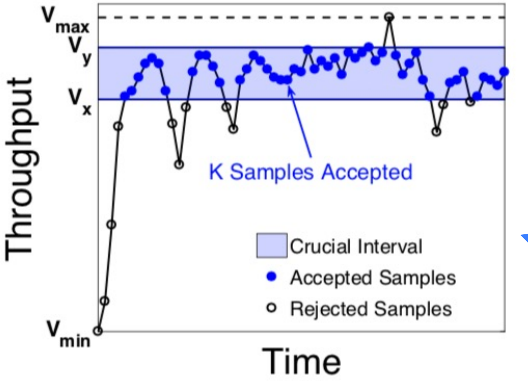

Crucial Interval Sampling

Fuzzy Rejection Sampling

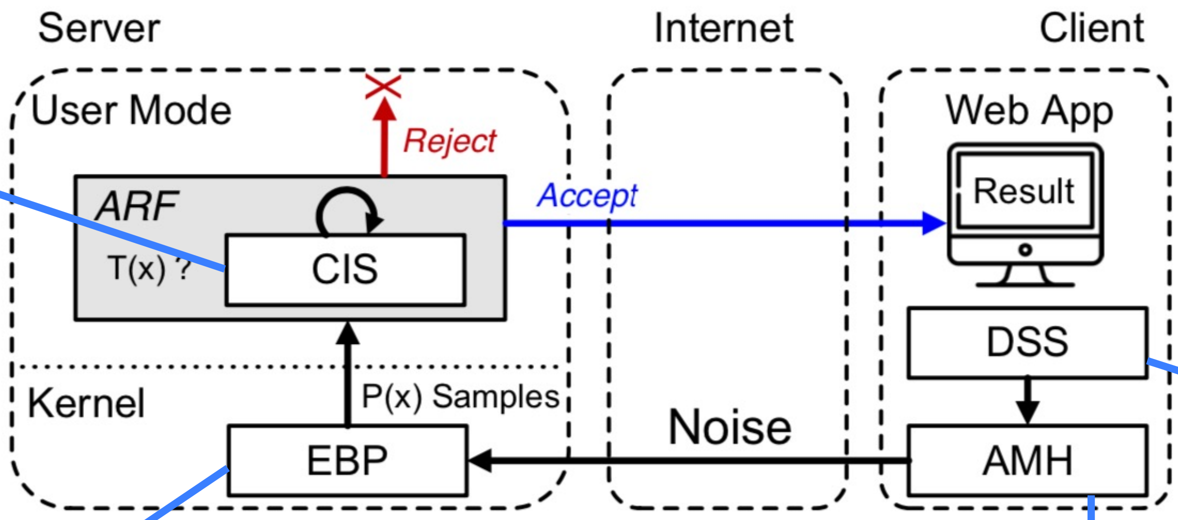

Cannot saturate user's access link bandwidth

Data-driven Server Selection

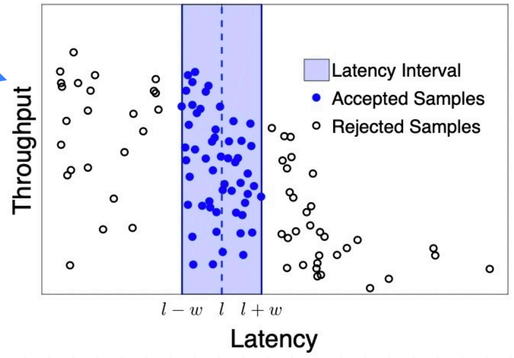

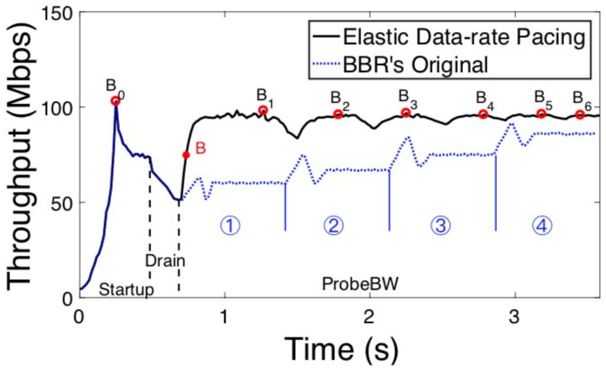

Elastic Bandwidth Probing

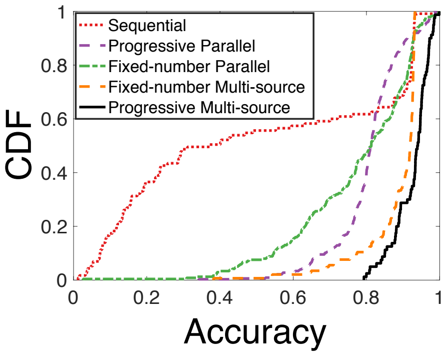

Adaptive Multi-homing

4. Data-driven Sever Selection (DSS)

Ping-based server selection

Low latency ≠ high throughput

Historical performance-based sever selection

Select servers with highest bandwidth estimations

Ping test

Historical Data

Crucial latency interval

max avg

Sorting servers

Select candidate servers

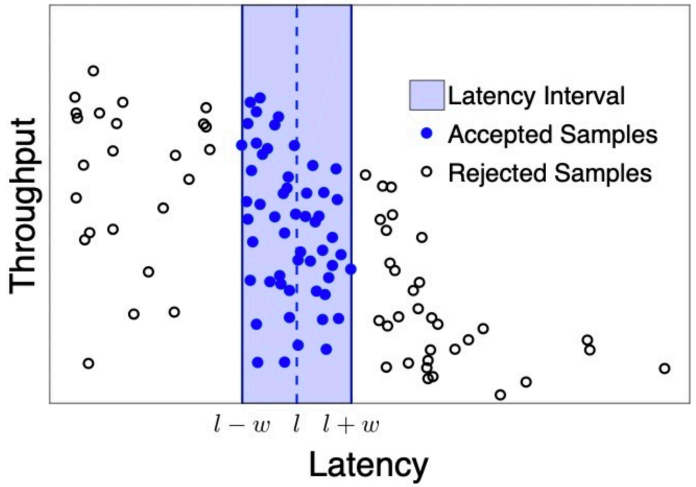

4. Adaptive Multi-Homing (AMH)

Adding concurrency level with fixed threshold

Under-estimating user's bandwidth (e.g., 5G)

Adaptive Multi-Homing

When shall we stop adding extra test servers?

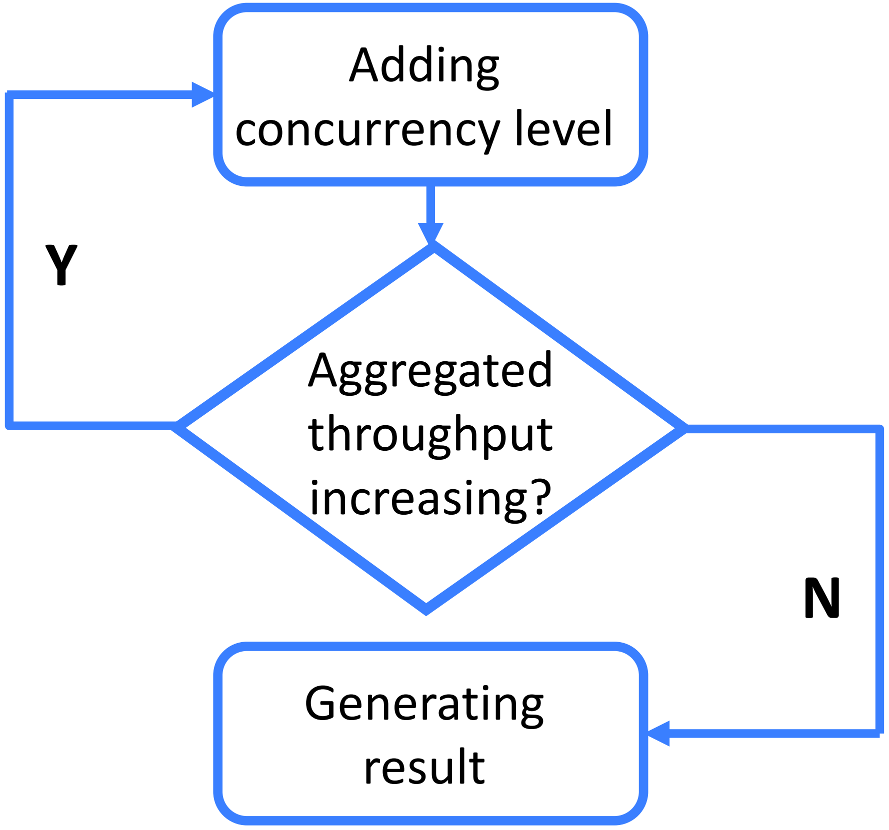

5. Evaluation

Testbed networks

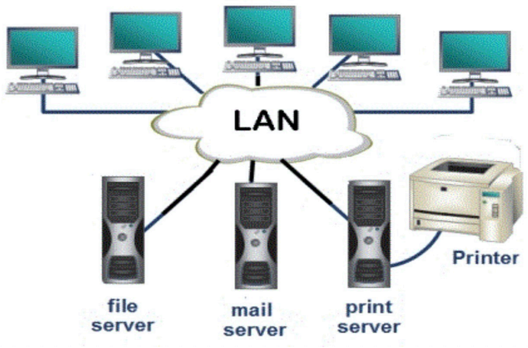

LAN

Residential broadband

Datacenter network

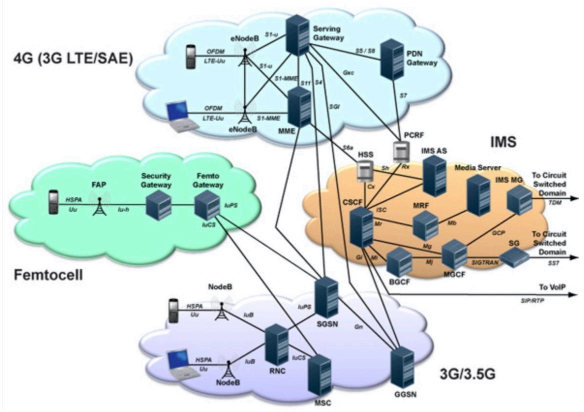

LTE network

mmWave & Sub-6Ghz 5G network

HSR cellular network

5. Evaluation

Major results

- **FastBTS vs. others on testbed networks: 5%–72%** higher average accuracy, **2.3–8.5×** shorter test duration, **3.7–14.2×** less data usage.
- **FastBTS vs. SpeedTest.net in real world:** FastBTS (with only 30 servers) achieves comparable accuracy compared with the production system of SpeedTest.net with ~12,000 test servers, incurring **5.6×** shorter test duration and **10.7×** less data usage on average.

6. System Demo

☐ Case 1 : PC + Wi-Fi (~100 Mbps)

SpeedTest.net

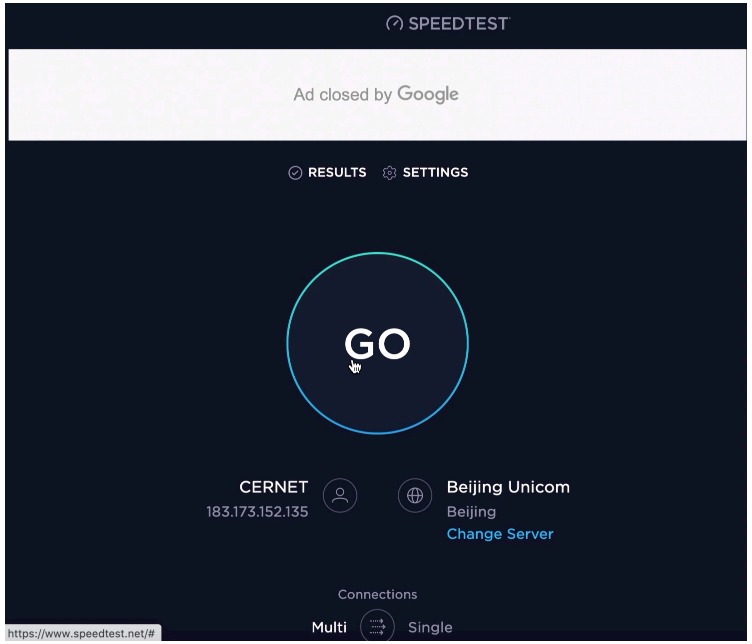

FastBTS.thucloud.com

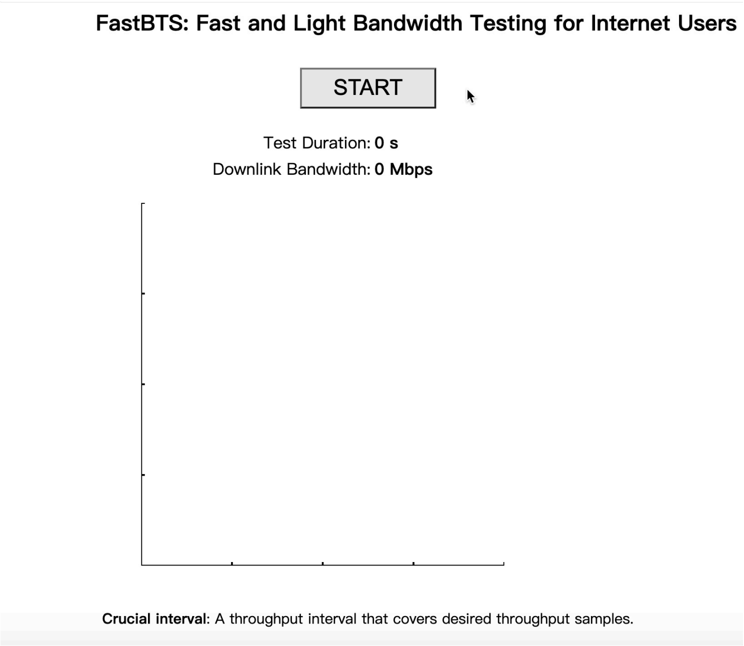

Duration: 15.0 seconds

Result: 95.18 Mbps

Data usage: 176 MB

Duration: 3.1 seconds

Result: 99.25 Mbps

Data usage: 37 MB

Videos available at: https://youtu.be/QbHO27RvzbU

6. System Demo

Case 2 : smartphone + Sub-6Ghz 5G (~500 Mbps)

SpeedTest.net

Duration: 15.0 seconds

Result: 484 Mbps

Data usage: 936 MB

FastBTS.thucloud.com

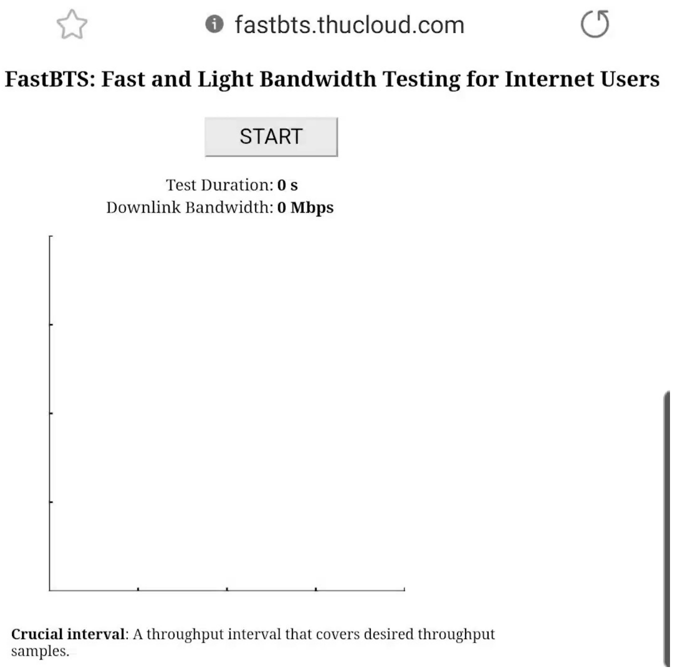

Duration: 4.1 seconds

Result: 543.07 Mbps

Data usage: 168 MB

Videos available at: https://youtu.be/VGN32d3dIAU

7. Conclusion

- We reveal how today's commercial bandwidth testing services actually work as well as their pros and cons based on in-depth investigations and large-scale benchmarking tests.
- We present FastBTS, a novel bandwidth testing solution that **accommodates and exploits network noises** to make bandwidth tests fast and light. With only 30 test servers, FastBTS achieves comparable accuracy compared with SpeedTest.net with ~ 12,000 servers, while incurring 5.6× shorter test duration and 10.7× less data usage on average.
- We have released all the source code at https://FastBTS.github.io and an online demo system at http://FastBTS.thucloud.com.

Thanks!

Q & A