

Mobile Network Performance from User Devices: A Longitudinal, Multidimensional Analysis

Ashkan Nikravesh, David R. Choffnes, Ethan Katz-Bassett
Z. Morley Mao, Matt Welsh



Northeastern

USC

Google

Problem

Mobile Network Performance:

- ✘ Poor visibility into user perceived performance.
- ✘ It is difficult to capture a view of network performance.

Problem

Mobile Network Performance:

- ✘ Poor visibility into user perceived performance.
- ✘ It is difficult to capture a view of network performance.

Why is that difficult?

- Performance depends on many factors

Problem

Mobile Network Performance:

- ✘ Poor visibility into user perceived performance.
- ✘ It is difficult to capture a view of network performance.

Why is that difficult?

- Performance depends on many factors

How to improve the visibility?

Problem

Mobile Network Performance:

- ✗ Poor visibility into user perceived performance.
- ✗ It is difficult to capture a view of network performance.

Why is that difficult?

- Performance depends on many factors

How to improve the visibility?

- Pervasive network monitoring is needed:
 - Continuous
 - Large-scale

Problem

Mobile Network Performance:

- ✗ Poor visibility into user perceived performance.
- ✗ It is difficult to capture a view of network performance.

Why is that difficult?

- Performance depends on many factors

How to improve the visibility?

- Pervasive network monitoring is needed:
 - Continuous
 - Large-scale
- Sampling performance of devices across:
 - Carriers
 - Access Technologies
 - Location
 - Time

Previous Works

Their Limitations:

- Passively collected from **cellular network infrastructure** (e.g. GGSN)
 - **One month** of data
 - Limited to a **single carrier**

Previous Works

Their Limitations:

- Passively collected from **cellular network infrastructure** (e.g. GGSN)
 - **One month** of data
 - Limited to a **single carrier**
- Collected from **mobile devices**, but **not continuously**

Previous Works

Their Limitations:

- Passively collected from **cellular network infrastructure** (e.g. GGSN)
 - **One month** of data
 - Limited to a **single carrier**
- Collected from **mobile devices**, but **not continuously**

Our work differs from previous related work:

- Longitudinal
- Continuous
- Gathered from mobile devices using controlled experiments.

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks
- ✓ Identify patterns, trends, anomalies, and evolution of cellular networks' performance.

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks
- ✓ Identify patterns, trends, anomalies, and evolution of cellular networks' performance.

We find:

- **Significant variance** in end-to-end performance for all carriers.

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks
- ✓ Identify patterns, trends, anomalies, and evolution of cellular networks' performance.

We find:

- **Significant variance** in end-to-end performance for all carriers.
- Part of the high variability is due to the **geographic** and **temporal** properties of network.

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks
- ✓ Identify patterns, trends, anomalies, and evolution of cellular networks' performance.

We find:

- **Significant variance** in end-to-end performance for all carriers.
- Part of the high variability is due to the **geographic** and **temporal** properties of network.
- **Routing** and **signal strength** are potential sources of performance variability.

Data Analysis

- ✓ Analyzing the data collected from:
 - ✓ 144 carriers
 - ✓ 17 months
 - ✓ 11 cellular networks
- ✓ Identify patterns, trends, anomalies, and evolution of cellular networks' performance.

We find:

- **Significant variance** in end-to-end performance for all carriers.
- Part of the high variability is due to the **geographic** and **temporal** properties of network.
- **Routing** and **signal strength** are potential sources of performance variability.
- Performance is **inherently unstable**.

Methodology

- User perceived performance:

Methodology

- User perceived performance:
 - HTTP GET Throughput
 - Ping RTT
 - DNS Lookup time

Methodology

- User perceived performance:
 - HTTP GET Throughput
 - Ping RTT
 - DNS Lookup time
- ✓ Traceroute
- ✓ Carrier + Cellular network technology
- ✓ Signal strength
- ✓ Location
- ✓ Timestamp

Methodology

- User perceived performance:
 - HTTP GET Throughput
 - Ping RTT
 - DNS Lookup time
- ✓ Traceroute
- ✓ Carrier + Cellular network technology
- ✓ Signal strength
- ✓ Location
- ✓ Timestamp

To **identify** and **isolate** the performance impact of each factor

Methodology

- User perceived performance:
 - HTTP GET Throughput
 - Ping RTT
 - DNS Lookup time
- ✓ Traceroute
- ✓ Carrier + Cellular network technology
- ✓ Signal strength
- ✓ Location
- ✓ Timestamp



To **identify** and **isolate** the performance impact of each factor

- Mainly collected by **Speedometer**:
 - 2011-10 to 2013-2 (17 months)
 - Internal android app developed by Google
 - Anonymized data

Dataset

- Mainly collected by **Speedometer**:
 - 2011-10 to 2013-2 (17 months)
 - Internal android app developed by Google
 - Anonymized data
- Mobiperf
 - 11 months
 - Only used for our signal strength analysis

Dataset

- Mainly collected by **Speedometer**:
 - 2011-10 to 2013-2 (17 months)
 - Internal android app developed by Google
 - Anonymized data
- Mobiperf
 - 11 months
 - Only used for our signal strength analysis
- Controlled experiments
Speedometer dataset: 4-5 measurements per minute

Dataset

- Mainly collected by **Speedometer**:
 - 2011-10 to 2013-2 (17 months)
 - Internal android app developed by Google
 - Anonymized data
- Mobiperf
 - 11 months
 - Only used for our signal strength analysis
- Controlled experiments
Speedometer dataset: 4-5 measurements per minute
- Code is open source and data is publicly available


Performance across Carriers

How observed performance matches with the expectations across access technologies?

Performance across Carriers

How observed performance matches with the expectations across access technologies?

- **Ping RTT Latency**




Performance across Carriers

How observed performance matches with the expectations across access technologies?

- **Ping RTT Latency**

- ① Latency varies significantly across carriers and access technologies

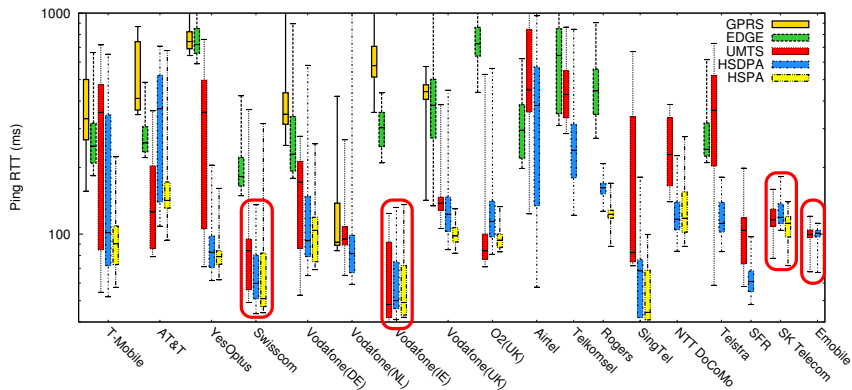


Performance across Carriers

How observed performance matches with the expectations across access technologies?


- **Ping RTT Latency**

- ① Latency varies significantly across carriers and access technologies
- ② Same performance for different access technologies



Performance across Carriers


- HTTP GET Throughput



Performance across Carriers

- **HTTP GET Throughput**


- ① Relatively smaller difference between the carriers



Performance across Carriers

- **HTTP GET Throughput**


- 1 Relatively smaller difference between the carriers
- 2 Download size (224KB) is not sufficiently large



Performance across Carriers

- **HTTP GET Throughput**


- 1 Relatively smaller difference between the carriers
- 2 Download size (224KB) is not sufficiently large



Performance across Carriers

• HTTP GET Throughput

- 1 Relatively smaller difference between the carriers
- 2 Download size (224KB) is not sufficiently large
- 3 Lower latency is generally correlated with higher throughput



Performance across different Locations


- Different topologies in different regions

Performance across different Locations

- Different topologies in different regions
- New York, Bay Area, and Seattle

Performance across different Locations

- Different topologies in different regions
- New York, Bay Area, and Seattle



Verizon LTE

Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability

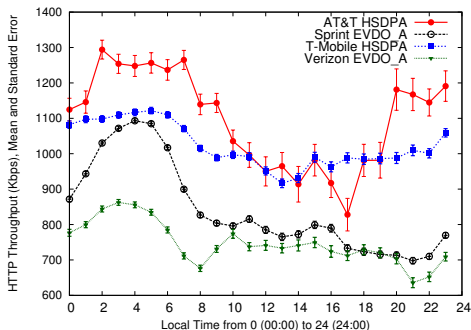
Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability
- **When** to measure the network?

Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability
- **When** to measure the network?


Time of the day



Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability
- **When** to measure the network?

Time of the day

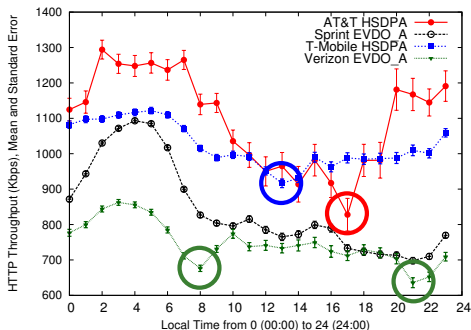


- 1 Throughput decreases during the busy hours of usage

Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability
- **When** to measure the network?

Time of the day




- 1 Throughput decreases during the busy hours of usage
- 2 Carriers experience minimum throughput at different times

Performance over Time

- **How much** performance depends on time?
 - time of the day
 - stability
- **When** to measure the network?

Time of the day



- 1 Throughput decreases during the busy hours of usage
- 2 Carriers experience minimum throughput at different times
- 3 Carriers experience different variation in performance during the busy hours

Performance over Time

Stability of Performance

- Users want a stable network!
- Measurements are expensive!

Performance over Time

Stability of Performance

- Users want a stable network!
- Measurements are expensive!

Two metrics:

- Auto-Correlation
- Weighted moving average

Performance over Time

Stability of Performance

- Users want a stable network!
- Measurements are expensive!

Two metrics:

- Auto-Correlation
- Weighted moving average



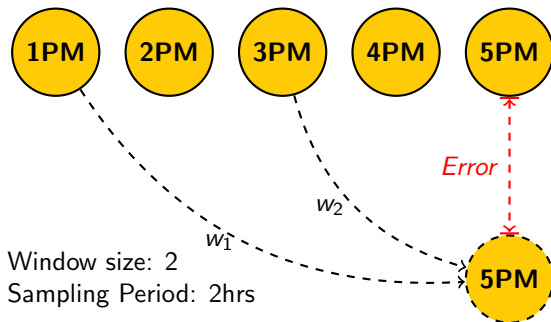
Performance over Time

Stability of Performance

- Users want a stable network!
- Measurements are expensive!


Two metrics:

- Auto-Correlation
- Weighted moving average




Performance over Time

Stability of Performance (Weighted Moving Average)



Performance over Time


Stability of Performance (Weighted Moving Average)



- 1 Prediction accuracy varies significantly by carriers

Performance over Time

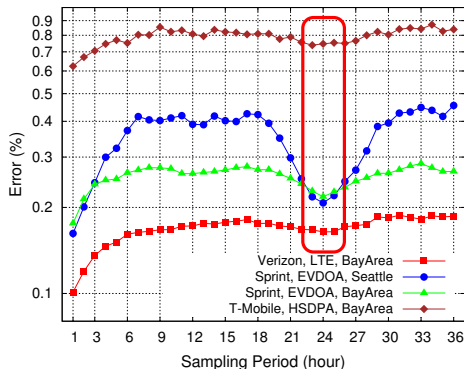
Stability of Performance (Weighted Moving Average)



- 1 Prediction accuracy varies significantly by carriers
- 2 Prediction error increases with longer sampling periods

Performance over Time

Stability of Performance (Weighted Moving Average)



- 1 Prediction accuracy varies significantly by carriers
- 2 Prediction error increases with longer sampling periods with the exception of 24hr sampling periods

Performance Degradation: Root Causes

Focus on the cases where persistent performance degradations were :

- observed in consecutive days
- affects both latency and throughput

Performance Degradation: Root Causes

Focus on the cases where persistent performance degradations were :

- observed in consecutive days
- affects both latency and throughput

Inefficient Paths

- Time evolution
- Impact on the performance


Performance Degradation: Root Causes

Focus on the cases where persistent performance degradations were :

- observed in consecutive days
- affects both latency and throughput

Inefficient Paths

- Time evolution
- Impact on the performance



(a) T-Mobile HSDPA (Seattle)


Performance Degradation: Root Causes

Focus on the cases where persistent performance degradations were :


- observed in consecutive days
- affects both latency and throughput

Inefficient Paths

- Time evolution
- Impact on the performance



(a) T-Mobile HSDPA (Seattle)



(b) Verizon LTE (Bay Area)

Performance Degradation: Root Causes



Signal Strength

How much it can affect performance?

Performance Degradation: Root Causes

Signal Strength

How much it can affect performance?



(a) Ping RTT and Packet Loss


(b) HTTP Throughput

AT&T HSDPA (Seattle), Arbitrary Strength Units (ASU)


Performance Degradation: Root Causes

Signal Strength

How much it can affect performance?



(a) Ping RTT and Packet Loss



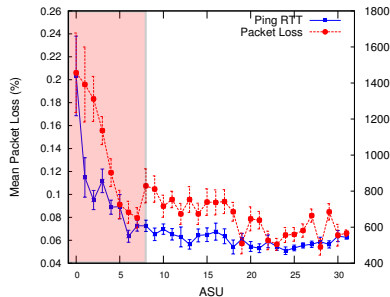
(b) HTTP Throughput

AT&T HSDPA (Seattle), Arbitrary Strength Units (ASU)


Performance Degradation: Root Causes

Signal Strength

How much it can affect performance?



(a) Ping RTT and Packet Loss



(b) HTTP Throughput

AT&T HSDPA (Seattle), Arbitrary Strength Units (ASU)

- ▶ Accounting for signal strength is important for interpreting measurement results.

Future Work

- We need for more monitoring and diagnosis.
- Data is difficult to get.

- We need for more monitoring and diagnosis.
- Data is difficult to get.

Mobilyzer

An Open Platform for Mobile Network Measurement

A comprehensive codebase for issuing measurements
for **researchers** and **developers**

Any Questions?

Thank You!