SIGCOMM 2021

A Nationwide Study on Cellular Reliability: Measurement, Analysis, and Enhancements

Yang Li, Hao Lin, Zhenhua Li, Yunhao Liu

Feng Qian, Liangyi Gong, Xianlong Xin, Tianyin Xu

清華大學 Tsinghua University

xiaomi

University of Minnesota

University of Illinois Urbana-Champaign

Outline

1. Background
2. Methodology
3. Key Findings
4. Enhancements
5. Summary

Outline

1. Background
2. Methodology
3. Key Findings
4. Enhancements
5. Summary

1.1 Cellular Network Today

☐ Cellular Network Empowers Modern Mobile Ecosystem

☐ 5G Networks Drive Our Grand Vision of AI and IoT

- 10 Gbps bandwidth: 100× faster than 4G
- 1 millisecond latency: 50× faster than 4G
- 1 million devices/km² connection density : 100× more than 4G

Performance ↑ Availability ↑

1.2 What Have We Missed?

How Reliable Is Today's Cellular Data Network?

Performance Is Meaningless when
Cellular Data Connections Fail to Work

1.3 Current Knowledge of Cellular Reliability

For A Device, Cellular Data Connections Can Fail Mostly in 3 Ways

Failure Types	Signal	Connection	Data
Data_Setup_Error	√	×	×
Out_of_Service	√	√	×
Data_Stall	√	√	√, but the connection abnormally stalls

Unfortunately, Cellular Reliability Is Rarely Studied

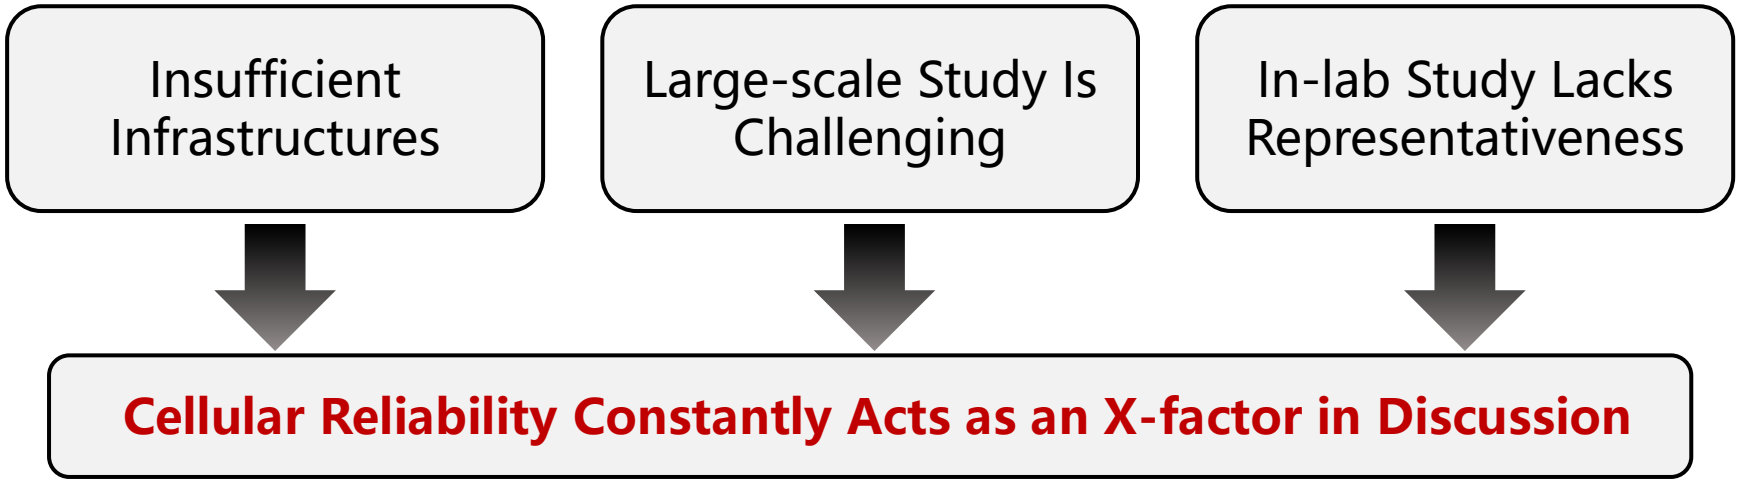

Outline

1. Background
2. Methodology
3. Key Findings
4. Enhancements
5. Summary

2.1 Collaborative Study

Collaboration & Objectives

- In collaboration with Xiaomi, a phone vendor with **~250M users**
- To conduct a large-scale study on cellular reliability problems
- **Our goal is to understand**
  - Prevalence
  - Severity
  - Root Causes
- **Xiaomi's business interests**
  - A Major Source of Complaints
  - Elusive Problems for Its Engineers

Well Aligned

2.2 Understanding Vanilla Android

Vanilla Android's Cellular Management Facilities

- Life cycle of a cellular connection is **modeled by a state machine**
- As state changes, failure-related problems are monitored

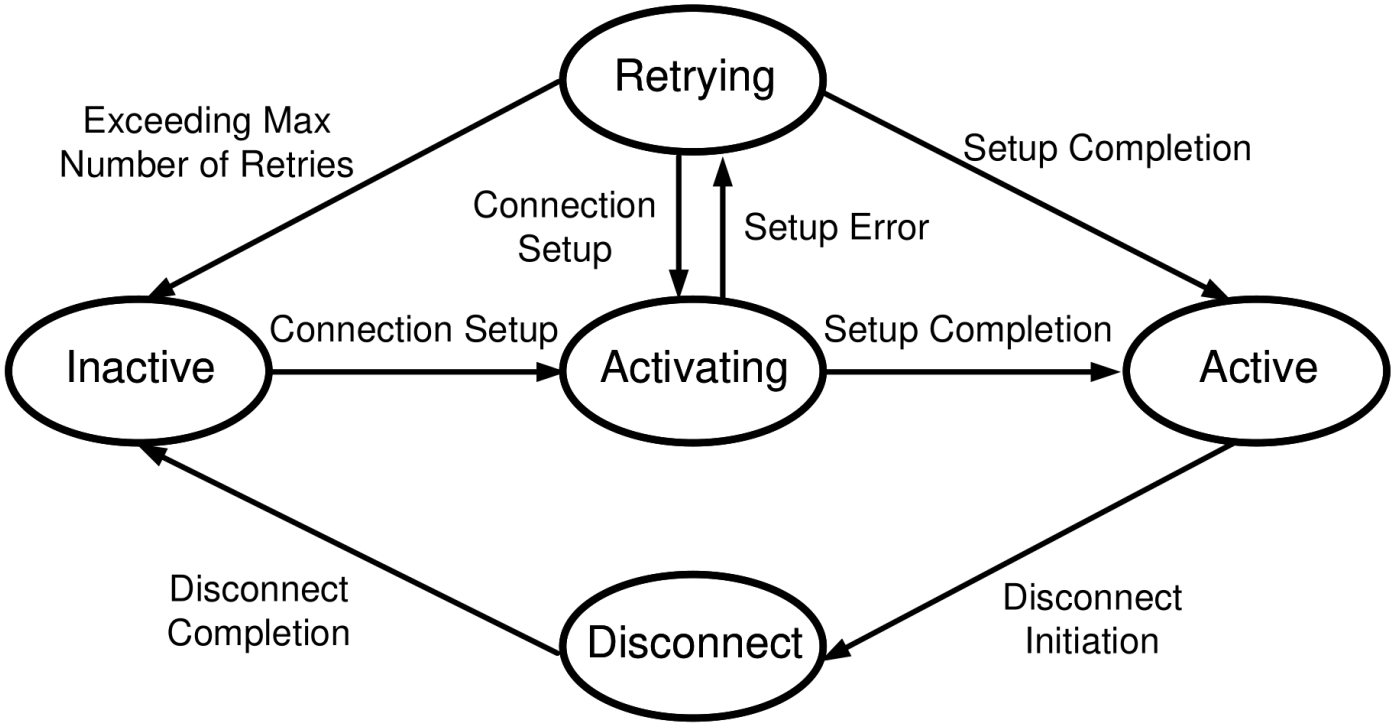

2.3 Limitations of Vanilla Android

- For concerned problems, Android provides
  - notification interfaces for system services
  - simple event logging, typically only for occurrences
- Major limitations & challenges

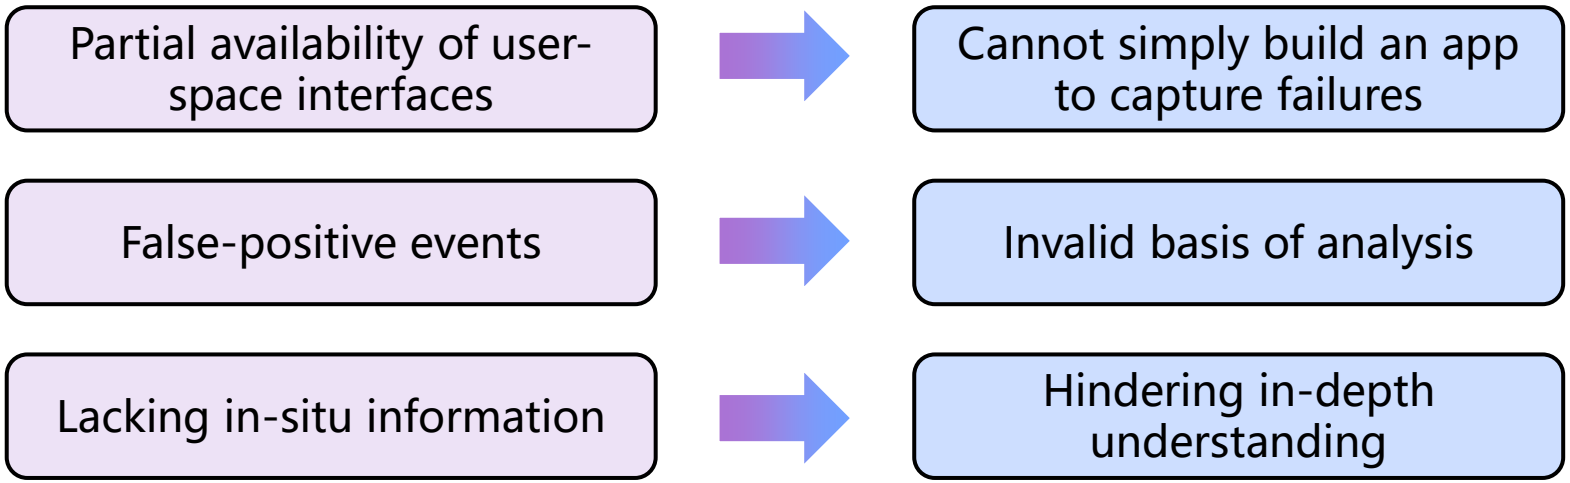

2.4 Continuous Monitoring Infrastructure

Android-MOD: Customized Android for Cellular Failure Capturing

- System service instrumentation
- Fine-grained system-level tracing
- Failure recovery monitoring

Service Instrumentation & Fine-Grained Tracing

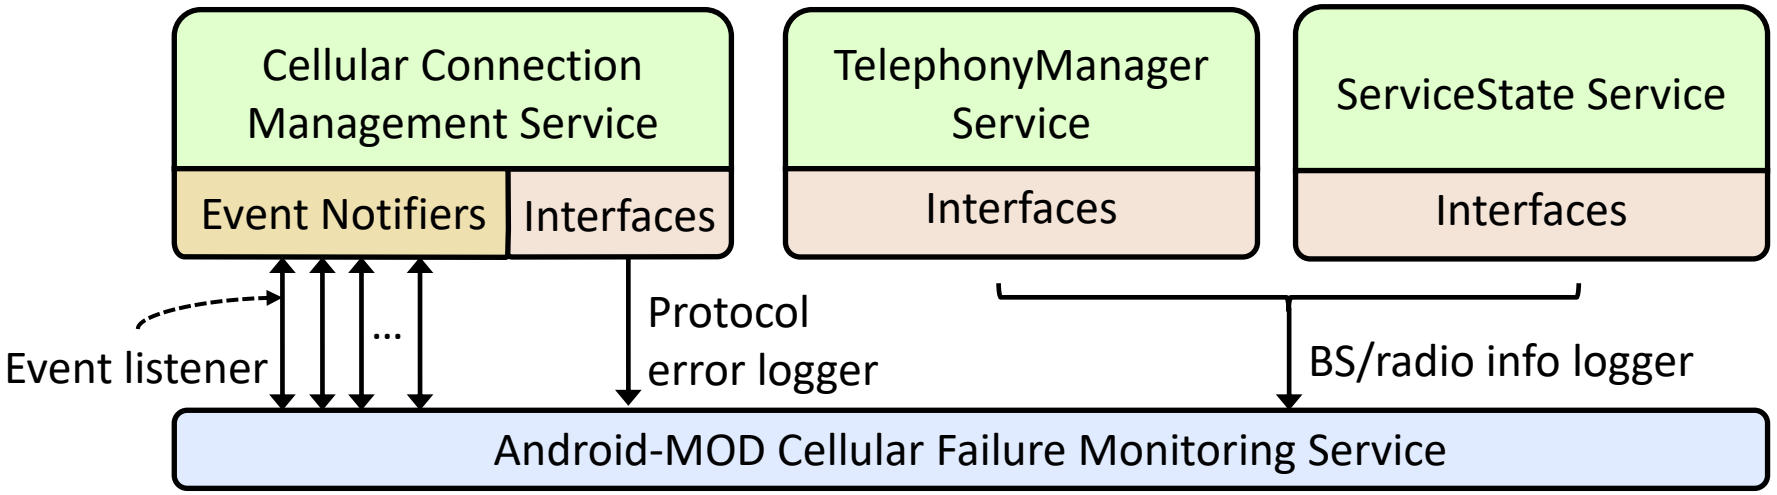

2.5 Failure Recovery Monitoring

- Data_Stall: 10 outbound TCP segments but no inbound in **1 min**

Fixed 1-min detection time	→	Up to 1 min of error
Lacking network stack information	→	Contain false positives

- Our approach: active and lightweight network-state probing

  - At most 5 seconds of detection time
  - Cross validation for ruling out false positives

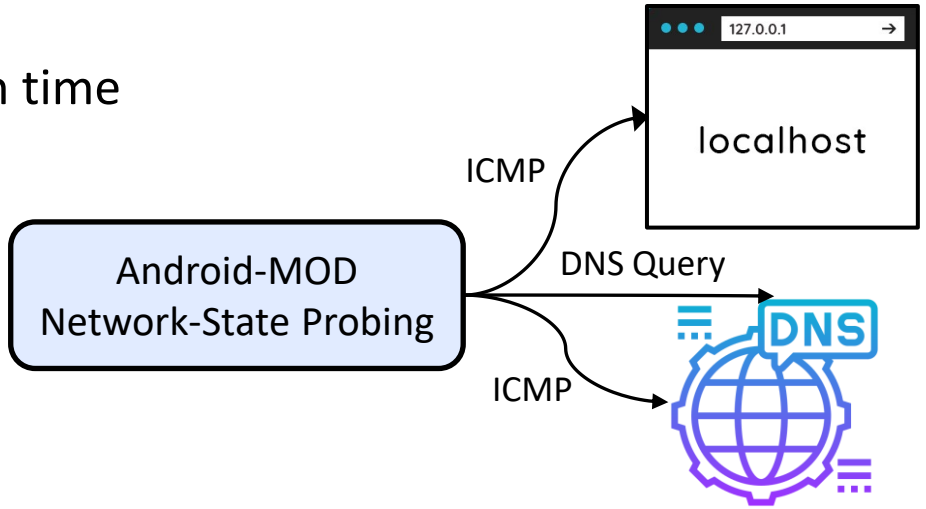

2.6 Large-Scale Deployment

Crowdsourcing Measurement

- Invited all the 250M Xiaomi's users to participate, **70M opted in**
- OS upgraded to Android-MOD
- Duration: **Jan.- Aug. 2020 (8 months)**
- Involved **34 device models**
- Captured **2 billion** cellular failures, involving **16 million** user devices, **3** mobile ISPs and **5 million** base stations

Model	CPU	Memory	Storage	5G	Android
1	1.8 GHz	2 GB	16 GB	–	10.0
2	1.95 GHz	2 GB	16 GB	–	9.0
3	2 GHz	2 GB	16 GB	–	9.0
4	2 GHz	3 GB	32 GB	–	9.0
5	2 GHz	3 GB	32 GB	–	9.0
6	2 GHz	3 GB	32 GB	–	10.0
7	2 GHz	3 GB	32 GB	–	10.0
8	2 GHz	3 GB	32 GB	–	9.0
9	2 GHz	3 GB	32 GB	–	10.0
10	2.2 GHz	4 GB	32 GB	–	9.0
11	1.8 GHz	4 GB	64 GB	–	10.0
12	2 GHz	4 GB	64 GB	–	10.0
13	2.05 GHz	6 GB	64 GB	–	10.0
14	2.2 GHz	6 GB	64 GB	–	9.0
15	2.2 GHz	4 GB	128 GB	–	10.0
16	2.2 GHz	4 GB	128 GB	–	10.0
17	2.2 GHz	6 GB	64 GB	–	10.0
18	2.2 GHz	6 GB	64 GB	–	10.0
19	2.2 GHz	6 GB	64 GB	–	10.0
20	2.2 GHz	6 GB	64 GB	–	10.0
21	2.2 GHz	6 GB	64 GB	–	10.0
22	2.2 GHz	6 GB	128 GB	–	9.0
23	2.4 GHz	6 GB	64 GB	YES	10.0
24	2.4 GHz	6 GB	128 GB	YES	10.0
25	2.45 GHz	6 GB	64 GB	–	9.0
26	2.45 GHz	6 GB	64 GB	–	9.0
27	2.8 GHz	6 GB	64 GB	–	10.0
28	2.8 GHz	6 GB	64 GB	–	10.0
29	2.8 GHz	6 GB	64 GB	–	10.0
30	2.8 GHz	6 GB	128 GB	–	10.0
31	2.84 GHz	6 GB	64 GB	–	10.0
32	2.84 GHz	6 GB	64 GB	–	10.0
33	2.84 GHz	8 GB	128 GB	YES	10.0
34	2.84 GHz	8 GB	256 GB	YES	10.0

Outline

1. Background
2. Methodology
3. Key Findings
4. Enhancements
5. Summary

3.1 Prevalence & Duration

- As many as **33 cellular failures** occur to a phone on average
- 77% report no failures; a device can experience up to 198K failures
- Average failure duration is **3.1 minutes**

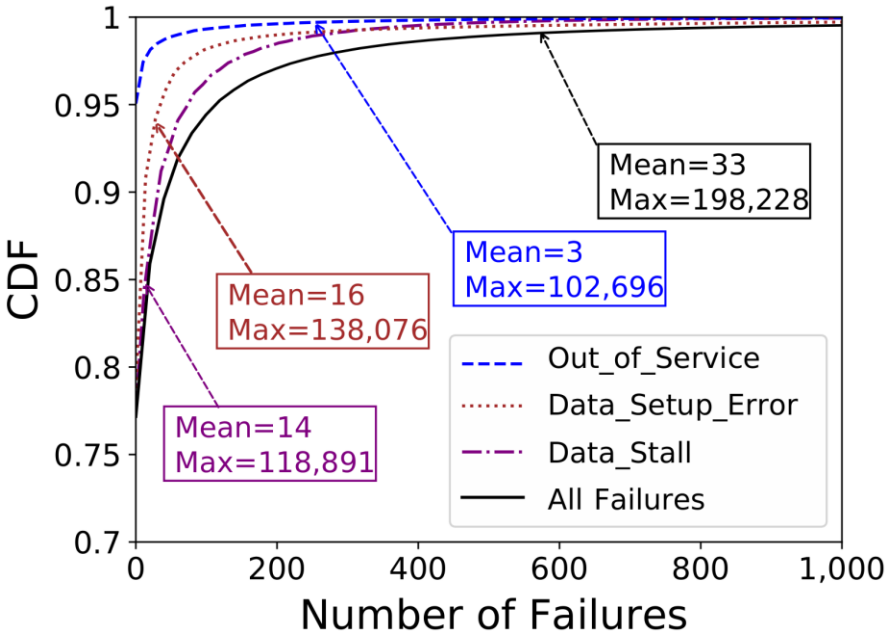

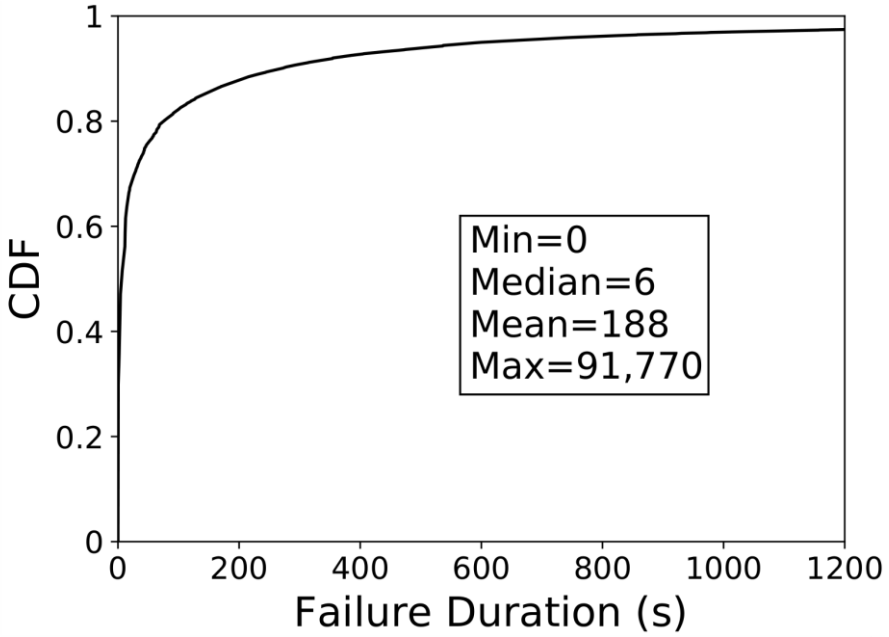

3.2 Hardware Configuration

- Failures occur prevalently on **all the 34** studied phone models
- Prevalence varies from 0.15% to 45% and averages at 23%
- Prevalence and frequency tend to **increase with better hardware!**

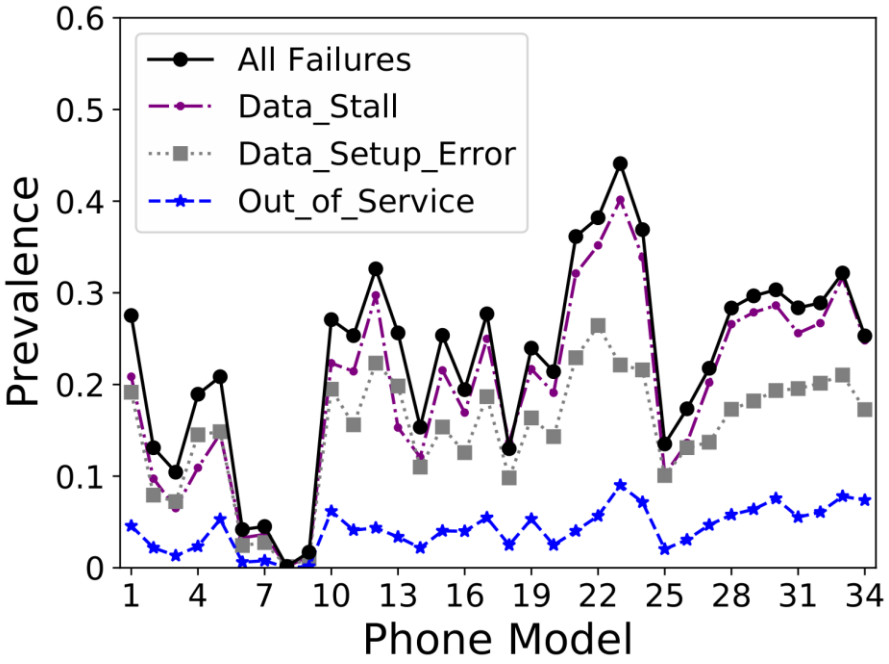

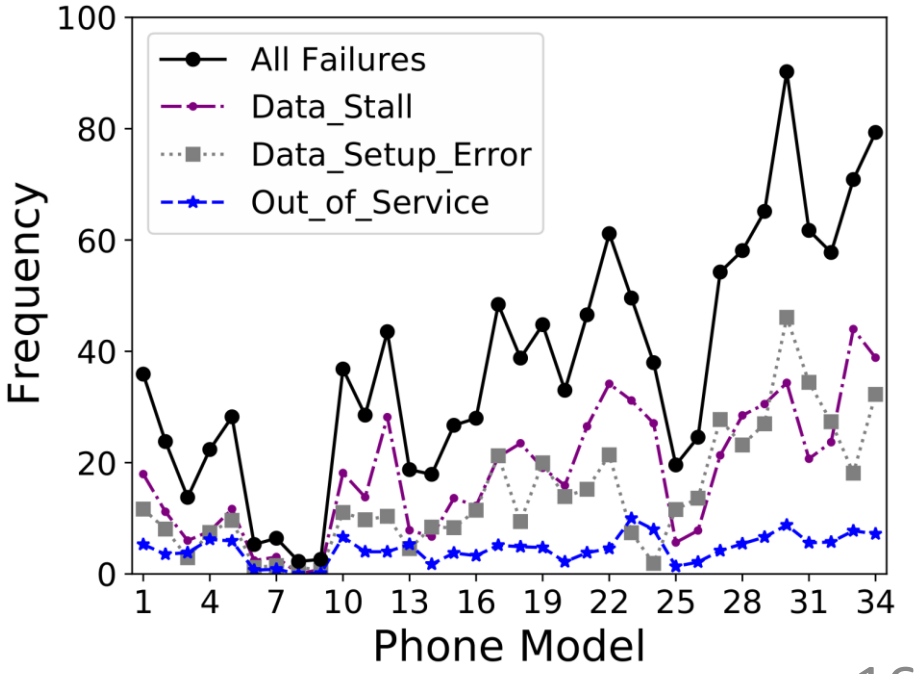

3.3 5G Capability

- Prevalence and frequency are higher on 5G devices
- 5G modules inflict high workload on the network stack of Android
- Today's production state of 5G modules is still immature

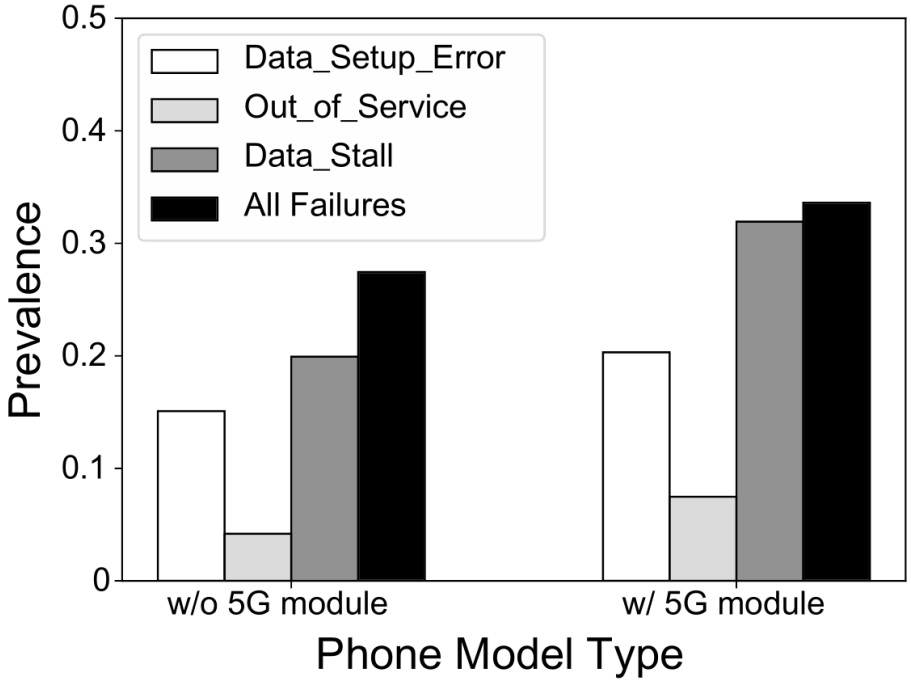

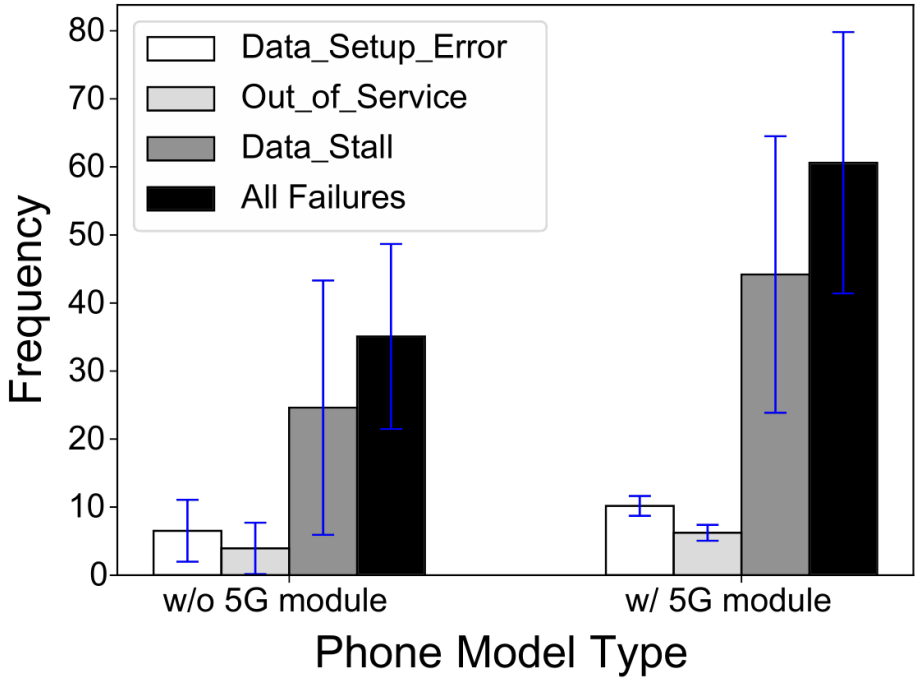

3.4 Android Version

- The 34 models run Android 9 and Android 10
- Despite quite a few improvements, Android 10 suffers more failures
- Android 10's **blindly prioritizing 5G** impairs connection stability
- **The situation of Android 11 is similar to that of Android 10**

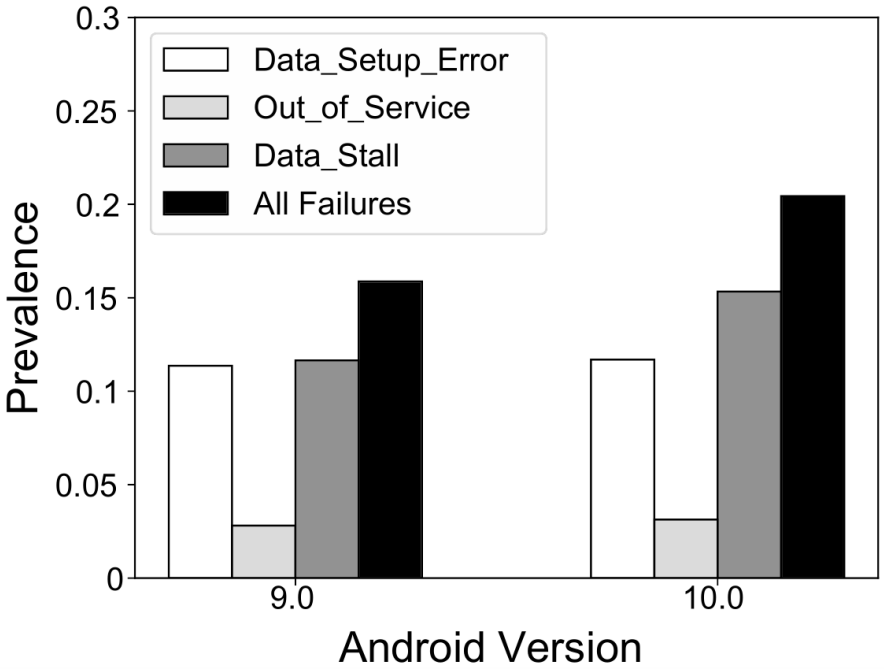

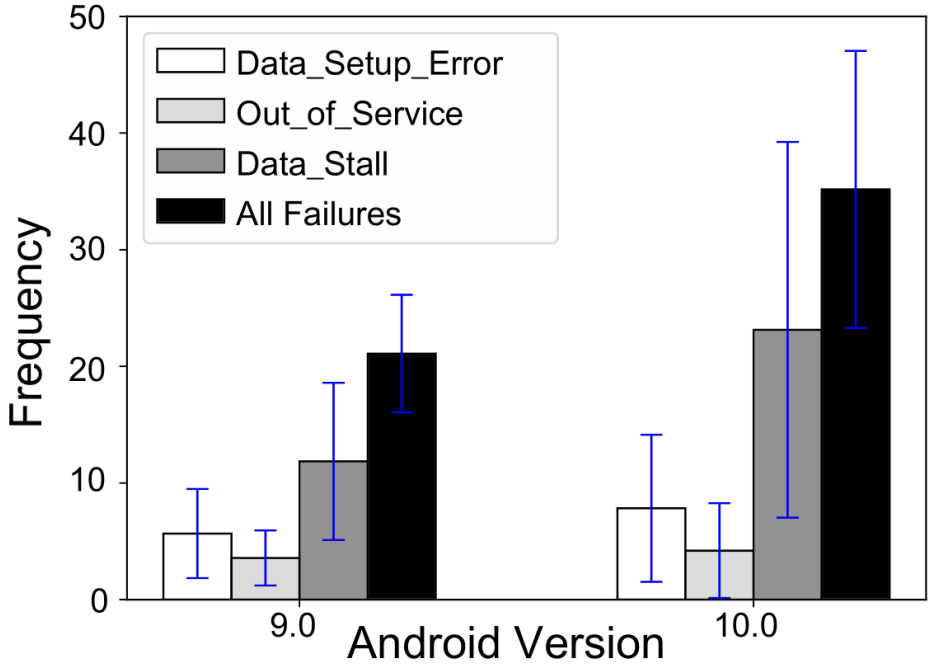

3.5 Data_Stall Recovery

- Data_Stall: 10 outbound TCP segments but no inbound in 1 min
- Three-Stage Progressive Recovery

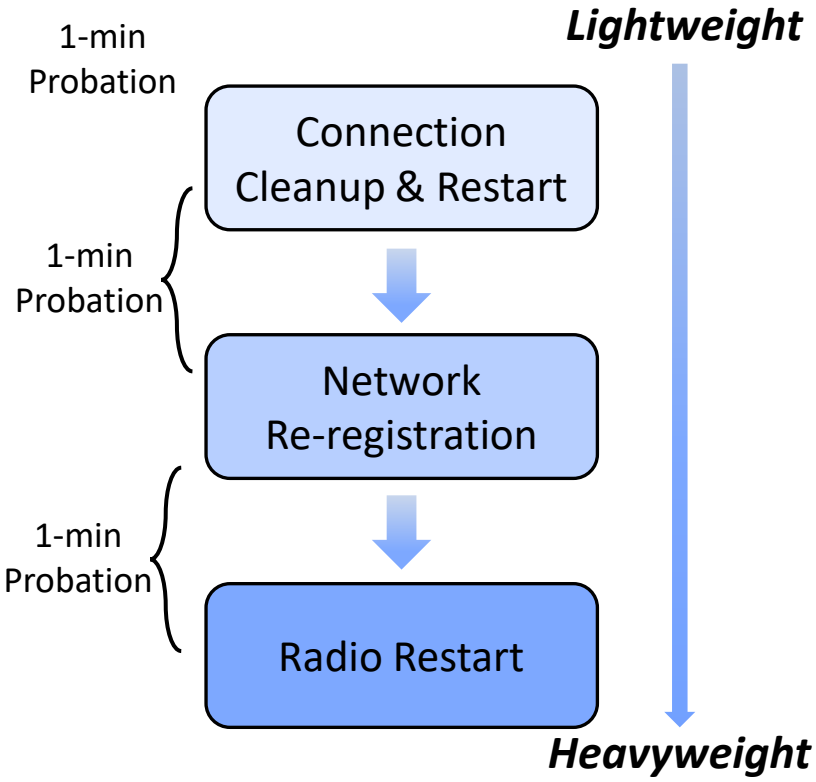

- Effective but **inefficient**
  - 60% Data_Stall failures are automatically fixed in just **10 seconds**
  - Victim users would manually reset the data connection within **~30 seconds**

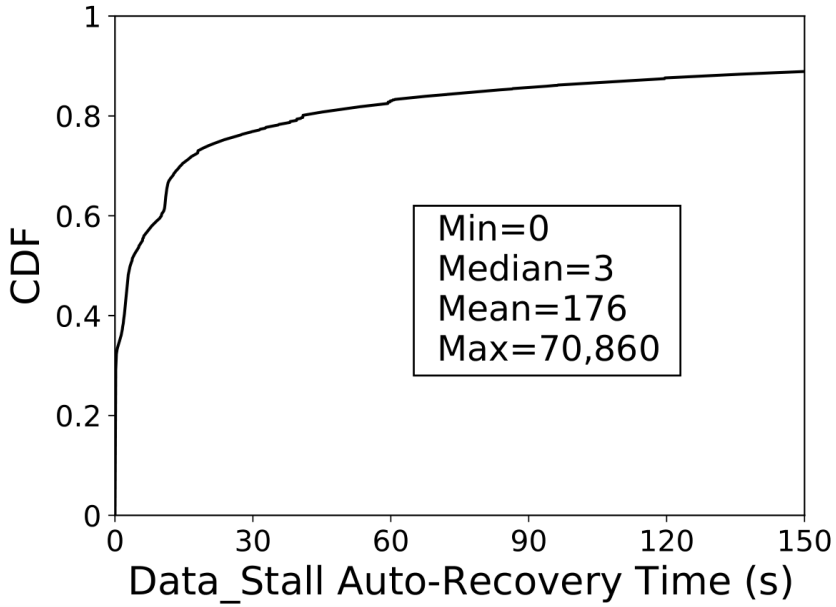

3.6 ISP & Base Station

- Failure distribution is quite skewed on BSes
- BSes with the most failures are mainly located in urban areas
- Failures are more prevalent on the users of ISP-B (China Telecom)

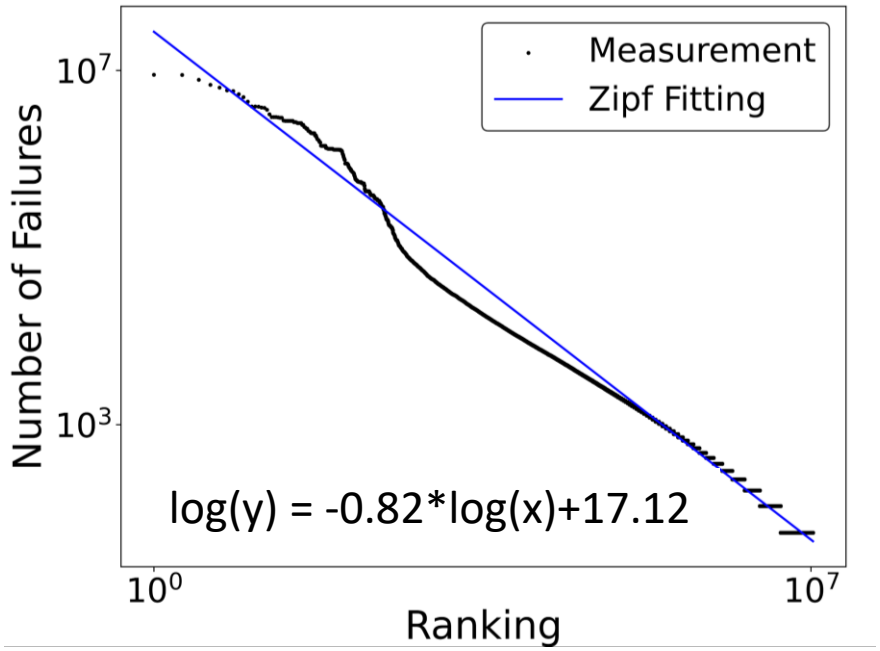

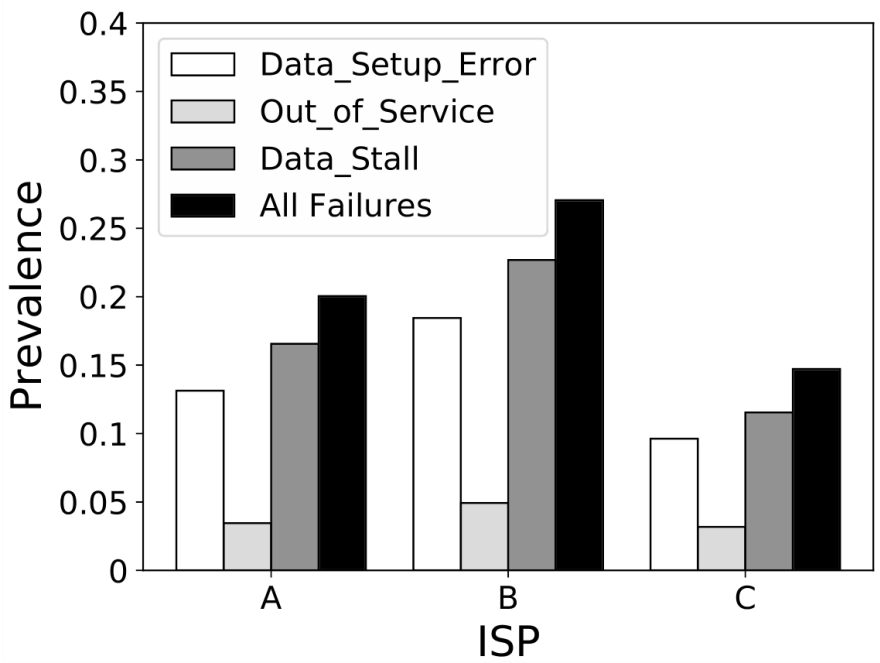

3.7 RAT & RSS

- Radio Access Technology (RAT): 3G BSes manifest fewer failures
- Received Signal Strength (RSS): **excellent signal ≠ reliable connection!**
- Excellent-RSS failures: densely-deployed BSes **around public transport hubs**, which increase **LTE mobility management overhead**

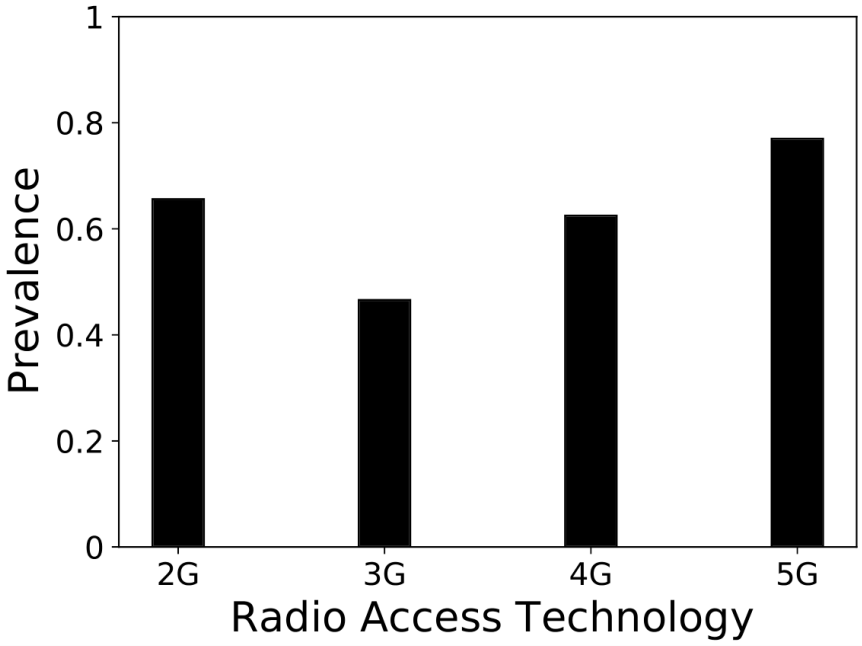

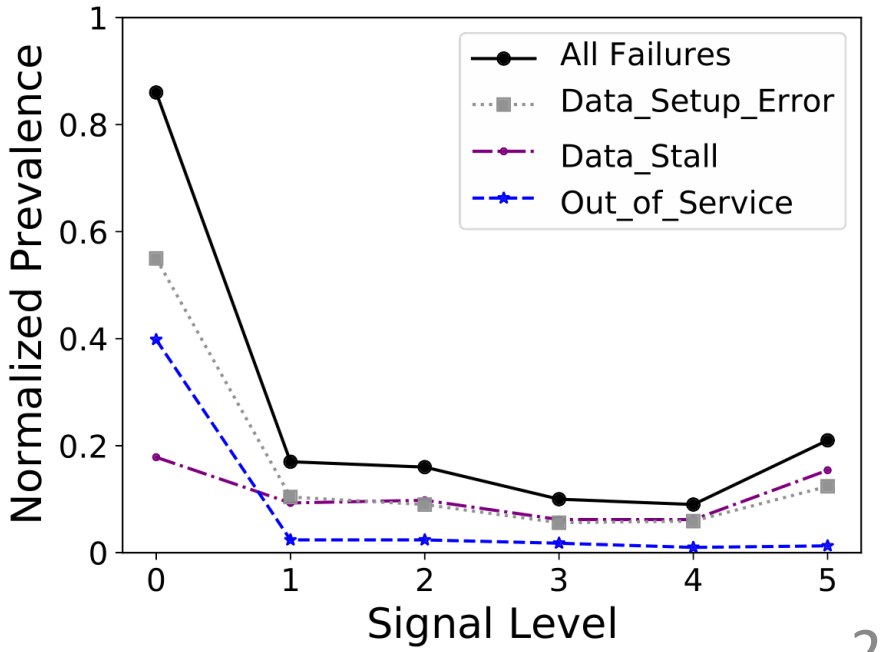

Outline

1. Background
2. Methodology
3. Key Findings
4. Enhancements
5. Summary

4.1 Guidelines in Principle

General Guidelines

Failures are severer on 5G devices	Vendors should carefully validate 5G modules
Android 10 is subject to more failures	Sufficient testing should be conducted for new OS features
3G BSes encounter fewer failures	Better utilization of the relatively "idle" infrastructures
More failures under excellent (level-5) RSS	Coordinated BS deployment, e.g., cross-ISP infrastructure sharing

4.2 Real-World Practice (1)

☐ Stability-Compatible RAT Transition

- Android 10's **blindly prioritizing 5G** impairs connection stability
- Taking the likelihood of cellular failures into account
- Better reliability without sacrificing data rate
- 4G/5G dual connectivity: smoother transition

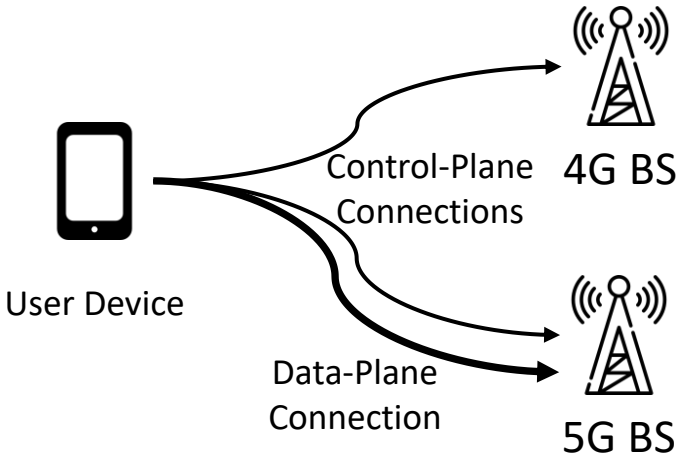

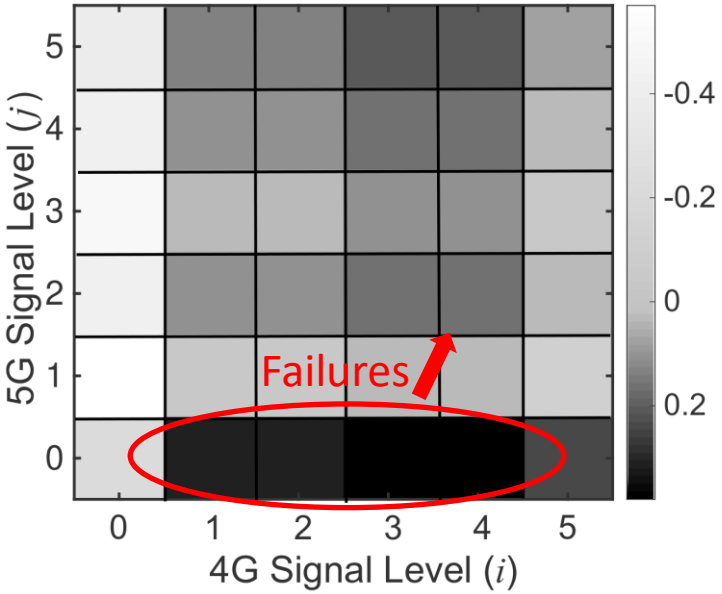

Increase of normalized prevalence of cellular failures.

Deeper color represents larger increase

4.3 Real-World Practice (2)

- Data_Stall: 10 outbound TCP segments but no inbound in **1 min**
- Android's Data_Stall Recovery: **effective but inefficient**

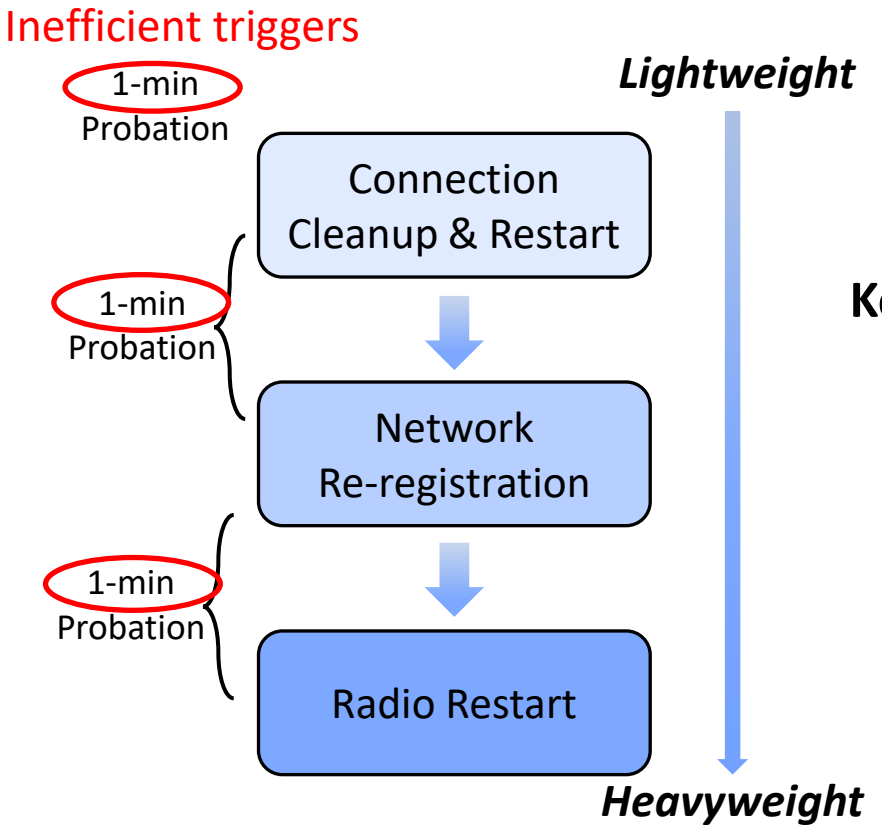

Android's ***three-stage progressive*** mechanism for Data_Stall Recovery

How to Find Proper Triggers?

Key insight: the conceptual three-stage recovery is essentially **a state transition process**

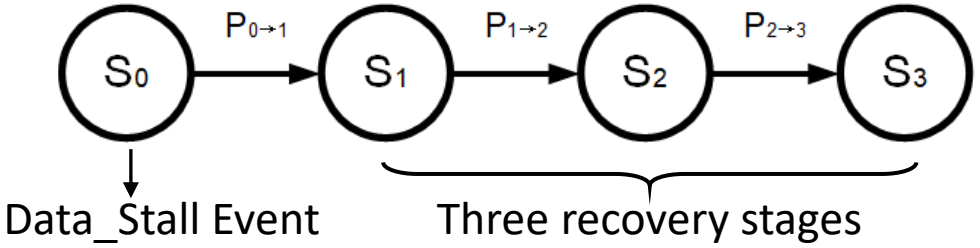

We can formalize the **overall recovery time** with a Markov model

4.3 Real-World Practice (2)

☐ TIMP-based Flexible Data Stall Recovery

Traditional Markov process

Transition Probability Is Constant

Real-world scenario

State Transition Depends on Elapsed Time

- ***Time-inhomogeneous Markov process* (TIMP)** models transition in a **time-sensitive** manner

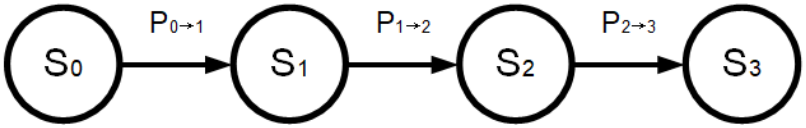

Traditional Markov Model

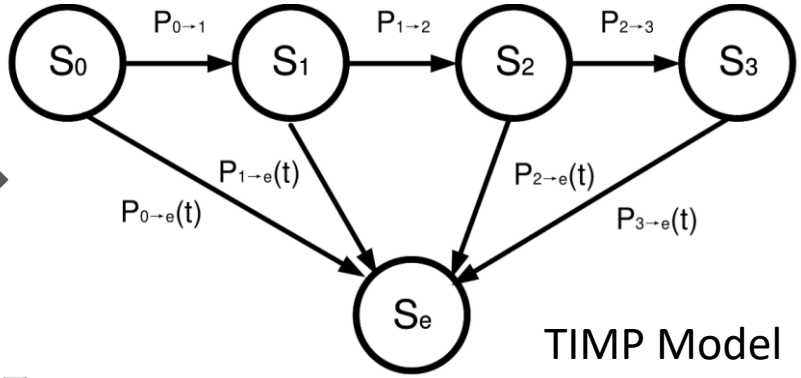

TIMP Model

- Optimizing overall recovery time modeled by TIMP with our data to **acquire more appropriate triggers**

$$min\ T_{recovery},$$
$$s.t.\ T_0 = T_{recovery},$$
$$T_i = \int_{\sigma Pro_{i-1}}^{\sigma Pro_i} \mathbb{P}_{i\to e}(t)dt + \mathbb{P}_{i\to i+1} \cdot T_{i+1} + O_i,$$
$$i = 0, 1, 2, 3.$$

4.4 Real-world Deployment and Evaluation

Patching the Two Enhancements in Android-MOD

- 40% of the 70M opt-in users upgraded, evaluated for two months
- Failures occur **40.3% less frequently** on the 5G phones
- **38% reduction** on the Data_Stall duration on average

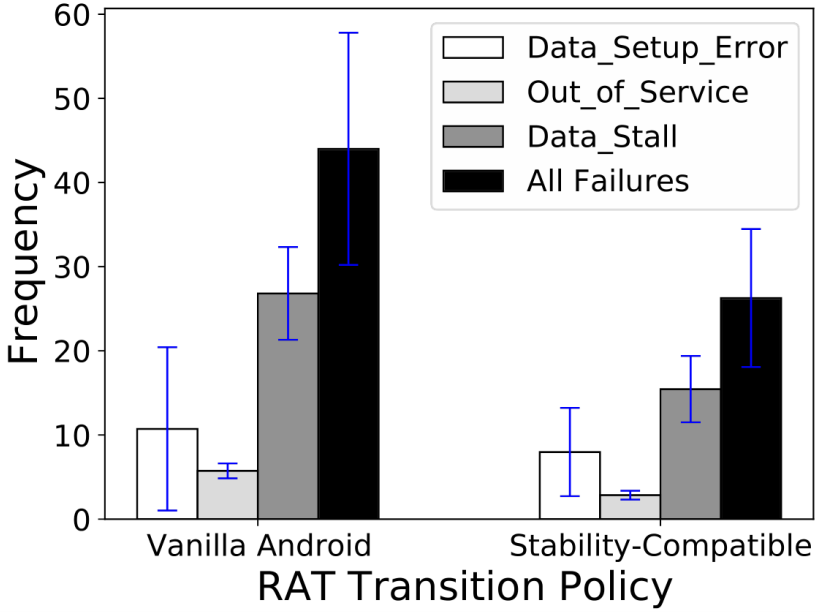

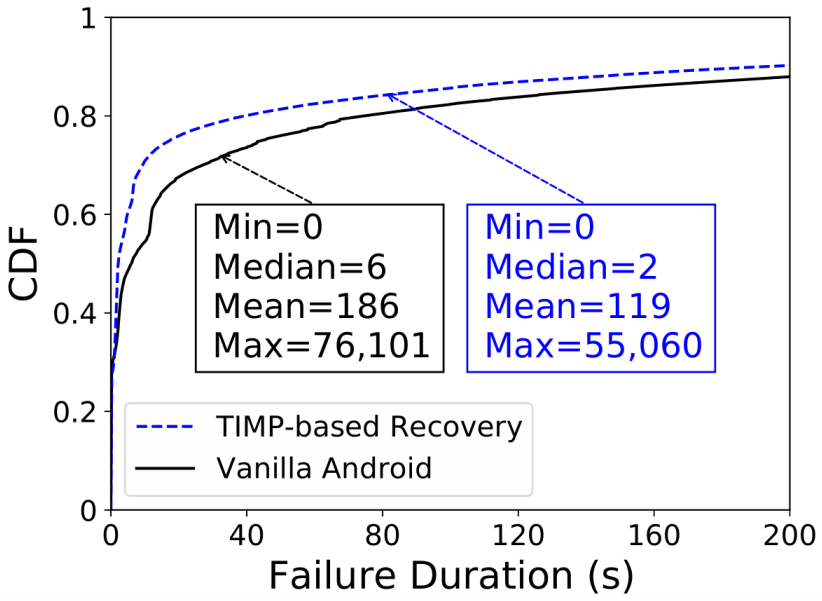

5 Conclusion

■ We conduct the first large-scale measurement of cellular reliability in the wild with more than 70M phones. We present our continuous monitoring infrastructure for capturing cellular failures on end devices.

■ We identify critical factors affecting cellular reliability. In particular, we pinpoint that software reliability defects are among the main root causes of cellular data connection failures.

■ We provide actionable insights for improving cellular reliability at scale. Most importantly, we have built on our insights to develop enhancements that yield remarkable real-world impact.

■ Source code released at https://CellularReliability.github.io