Leveraging Context-Triggered Measurements to Characterize LTE Handover Performance

Shichang Xu ✉, Ashkan Nikravesh, and Z. Morley Mao

University of Michigan{xsc,ashnik,zmao}@umich.edu

Abstract. In cellular networks, handover plays a vital role in supporting mobility and connectivity. Traditionally, handovers in a cellular network focus on maintaining continuous connectivity for legacy voice calls. However, there is a poor understanding of how today's handover strategies impact the network performance, especially for applications that require reliable Internet connectivity. In this work, using a newly designed context-triggered measurement framework, we carry out the first comprehensive measurement study in LTE networks on how handover decisions implemented by carriers impact network layer performance. We find that the interruption in connectivity during handover is minimal, but in 43% of cases the end-to-end throughput degrades after the handover. The cause is that the deployed handover policy uses statically configured signal strength threshold as the key factor to decide handover and focuses on improving signal strength which by itself is an imperfect metric for performance. We propose that handover decision strategies trigger handover based on predicted performance considering factors such as cell load along with application preference.

1 Introduction

Mobile devices rely on cellular networks to get network access to support data services. Since the coverage of each cell[1] is limited, handover between cells is essential for ensuring continuous connectivity and mobility. In addition, when the device is in the coverage of multiple cells, a proper policy should handover the mobile device to a cell that provides good performance.

There has been little work to understand how the deployed cellular network handover policies affects network layer performance in the wild. Specifically, questions such as what is the interruption in the network during handover and whether network performance consistently improves after handover are not well understood. In this paper, we perform the first large-scale study of handovers in LTE network using crowd-sourced measurements of over 200 users across three major carriers for the purpose of evaluating the performance implications of existing handover algorithms and policies. Performing measurements to capture transient handover events efficiently is challenging. To address such challenges, we develop a novel context-triggered measurement framework that dynamically initiates performance measurements of interest only when handover is likely to occur to reduce the measurement overhead.

[1] Each cellular base station has one or more set of antennas and it communicates with the mobile devices in one or more sectors called cells each of which has a unique ID [7].

Based on our measurement results, we identified fundamental limitations in the current design and deployment of handover algorithms: the use of static configurations on signal strength difference with neighboring cells and a lack of awareness of network performance. As a consequence, in 43% of cases the throughput degrades after the handover. By analyzing physical layer information in LTE network, we found that the cause of the performance degradation is that target cells have higher load and allocate less physical resources.

Our findings help motivate the need for handover algorithms based on network performance considering both signal strength and cell load. The measurement also shows the opportunity to improve the handover decision (§4): (a) currently handovers do not occur only when devices experience poor signal strength, indicating that the time of handover could be potentially changed without risk of link failures ; (b) the dense deployment of cells provides more than one candidate target cell the device could be potentially handed over to in many cases.

We summarize the main contribution of our work below.

– We designed a context-triggered measurement framework to support lightweight and accurate handover measurements. Using this setup, we collected 5 months' data from 200 users across three major cellular carriers in the U.S. to investigate performance impact of handover in LTE network.
– Using cross-layer analysis to incorporate radio link layer visibility with our data collector, we found that the current deployed handover policy relies on statically configured thresholds on signal strength. It focuses on improving signal strength and leads to potential performance degradation after the handover.
– We found that the interruption caused by intra-LTE handover is usually minimal. However, the median performance improvement after the handover is close to 0 in metrics including latency, throughput and jitter. The current handover algorithms do not appear to optimize performance.
– We identified that the performance degradation after handover is caused by higher load in the target cell and less allocated physical resources to the devices. We proposed cells predict performance after handover based on signal strength and cell load information and make handover decisions based on performance.

2 Background and Related Work

In this section, we first provide some background on handovers (§2.1). The related terminologies are summarized in Table 1. Then we summarize related works (§2.2).

2.1 LTE handover and data transmission procedure

Handovers within LTE networks are initiated by the cells and they can maintain ongoing network connections. We first define some basic terminology. *Serving cell* is the cell a user equipment (UE) is currently connected to. *Target cell* becomes the serving cell after the handover. We also refer to the original serving cell after the handover completes as the *source cell*.

Terminology	Definition
Reference Signal Received Power (RSRP)	The average power received from the reference signals. It is a metric of the downlink signal strength.
Physical Resource Block (PRB)	The basic unit of allocation of resources to the UE.
Event A3	The signal strength of neighbor cell becomes better than the serving cell by a *relative* threshold value.

Table 1. Related terminologies in LTE network

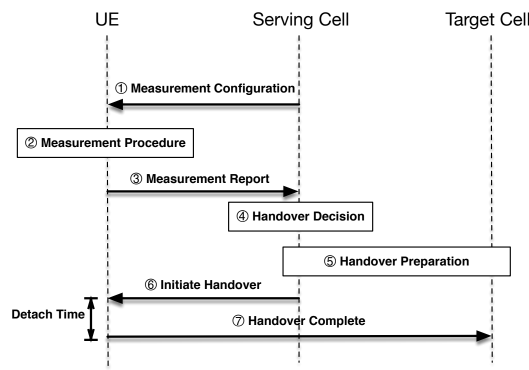

Fig. 1. The general handover procedure in a LTE network

As Figure 1 shows, to help decide when to trigger handover, the mobile device measures radio signal strength of both the serving cell and neighbor cells periodically. A commonly used metric of signal strength is *Reference Signal Received Power* (RSRP), *i.e.,* the average power received from the reference signals of the cell. The serving cell sends the *measurement configuration* to the device to specify when measurement results should be reported back. Depending on the measurement configurations, the *measurement reports* can be either event-triggered or periodical. Event-triggered reports are sent only when the link quality satisfies certain conditions. From previous work, a common trigger for intra-LTE handover is *event A3* [8, 11], where the signal strength of neighbor cell becomes better than the serving cell by a relative threshold value. Based on factors including measurement reports and load information, the serving cell makes decisions on handover [13]. The actual handover decision algorithm depends on the implementation at eNodeB or LTE base-station. After the source cell determines to perform a handover for a UE, it conducts a negotiation with the target cell to ensure enough resource at the target cell. Then it sends a message to the UE to initiate the handover. The UE disconnects with the source cell and connects to the target cell. After it successfully connects to the target cell, it notifies the target cell of the completion of the handover. These signaling messages between the cell and the UE are exchanged using radio resource control (RRC) protocol.

To understand how the network layer performance is determined by the lower layer in LTE network, we also briefly describe the data transmission procedure in the physical layer. Wireless communication requires radio spectrum resources. In LTE, cells dynamically allocate the physical radio resources in the unit of *physical resource blocks* (PRBs) to UEs and transmit data to the UEs using the allocated PRBs. The allocation strategy is not standardized in the specification and depends on vendor-specific implementation. However, typically cells use proportional scheduling algorithm [7], which optimizes cell efficiency while maintaining fairness across all UEs in the long term. When the *cell load* increases, i.e. more UEs connect to the cell and the total traffic volume increases, the allocated PRBs for each UE reduce. We denote the ratio between the allocated PRBs to a UE and the total PRBs of a cell as *PRB ratio*. The number of bytes transmitted by each PRB is determined by the signal strength, i.e., with strong signal strength and good channel quality, the cell could use coding schemes with high efficiency and thus transmits more data on each PRB. In summary, the performance in LTE network is affected by both the cell load which determines the PRB allocation and the signal strength which determines the transmission efficiency.

2.2 Related work

The problem of handover in cellular networks has attracted significant attention in both academia and industry. However, there is little work on understanding the performance impact of handover decisions in operational LTE network.

Previous work measured intra-LTE handovers using simulation [8, 9] and testbeds [13, 22] to understand the performance of applications during handover. Our work differs in that we measured the handover performance in the wild. Recent work [14, 16, 19, 23] study persistent handover loops caused by misconfigurations. We also identify such misconfigurations for a few cells in the wild but find they are not dominant. Our focus is to characterize the interruption caused by handover regardless of handover loops and compare the network performance before and after handover to understand the performance implications of deployed handover policies. Some other work [4, 12, 20, 21] studied handovers between different technologies, e.g. 3G and 4G. Our work studies handover between different cells in LTE network.

3 Methodology

To understand the impact of handover on performance in the wild, we crowd-source our measurement using a context-triggered measurement framework.

3.1 On-device measurement support

We combine passive monitoring with active measurements to study handover with minimal measurement overhead.

Passive Monitoring We keep track of device context including network type, signal strength and location. Through a novel use of the built-in diagnostic interface from Qualcomm communication chips, we also collect (1) lower layer RRC layer information including measurement configurations and handover messages and (2) physical layer information including PRB allocations. Our lower-layer message collection builds upon SnoopSnitch [2] which is an open-source Android app aimed to detect attacks such as fake base station using data from the Diagnostic Interface. The collector requires root privilege and reads the raw radio messages from the character device `/dev/diag` when DIAG_CHAR option is activated in Android kernel. The collector also collects fine-grained signal strength information from the diagnostic interface every 40ms, while the signal strength information from Android API updates only every two to three seconds. We are one of the first to crowd-source LTE radio-link layer messages.

Active Measurement To understand how handover impacts network performance, we use the Mobilyzer measurement library [18], a principled mobile network measurement platform, to measure network performance. We issue ping, TCP throughput, and UDP burst measurements to capture network characteristics using metrics including latency, throughput, jitter, packet loss.

Compared with passive monitoring, active measurements consume data resources and can cause significant impact on battery life. As data and battery resources are scarce on mobile devices, we need to capture performance during handover events efficiently. Towards this goal, we develop a context-triggered framework to trigger measurements only when a handover is predicted to occur in the near future.

Context-triggered Measurement Framework In general, deciding when to issue measurements is a challenging task. If we simply perform measurements periodically, the interval is difficult to configure. A small interval leads to large amount of unnecessary measurements that fail to capture interesting phenomena, wasting valuable data and battery resources on the device; while a large interval can miss the phenomena we are interested in.

One approach to solving this problem is to trigger measurements [3] based on context that specifies the conditions of interest. We estimate the likelihood of observing relevant events based on the device context and trigger measurements only when the probability of capturing desired events is high. This helps reduce unnecessary measurements while capturing more events of interest.

We design a context-triggered measurement framework atop Mobilyzer [18]. We send the devices measurement tasks with specific context requirements. The devices keep monitoring related context and trigger measurements once the context conditions are met. Note that different contexts contain different information with different cost. Even querying the same context with different granularity requirement has different cost implications. These considerations motivate our design of supporting a multi-level triggering procedure. At the first level, we monitor a context with the lowest cost. If the context indicates that the possibility of desired event occurrence is high, we monitor another context with higher cost or the same context with higher accuracy. This can be done with multiple layers until we reach high confidence that the event will occur. There is certainly a trade-off from using many levels to reduce measurement overhead but at the cost of introducing delay in capturing the event of interest which could lead to fewer events captured. We argue that the number of levels need to be adjusted depending on the type of events and the overhead of triggering at each level.

We apply this framework to understand the performance impact of handovers. Based on the passively collected lower layer messages, we find that LTE handovers usually happen when a cell with a stronger signal strength than the connected cell is discovered for a mobile user (§4). We implement the framework as follows. We first use sensors to detect user movement, as the power consumption of movement detection sensors such as accelerometers is only 5mW for an active device. Once we detect that the user is moving, we start to read fine-grained signal strength data from the diagnostic interface, which consumes around 200mW. If the neighbor cell signal strength is stronger than the serving cell, a ping, throughput or udp burst measurement is triggered, consuming more than 1500mW to activate the radio [10].

To evaluate the effectiveness of the framework, we run simulation on all traces collected from PhoneLab deployment [5] as explained later. Figure 2 compares the overhead and accuracy of measurements with and without context triggered framework. $T = x$ means the measurement are triggered when the signal strength of neighbor cell is stronger than the serving cell by the threshold of x dBm. $P = y$ denotes periodic measurement every y minutes. We calculate the overhead as the average number of measurements for identified handovers, and accuracy as the percentage of measurements that capture handovers. As shown, our framework can reduce the overhead to 1% while increasing the accuracy by 10 times. In the actual deployment, labeled as "Real" in Figure 2, we use $T = 0$. Compared to the simulation results, the real deployment has a slightly lower overhead and accuracy, because we imposed constraints on the resource usage of active measurements to reduce impact on user experiences.

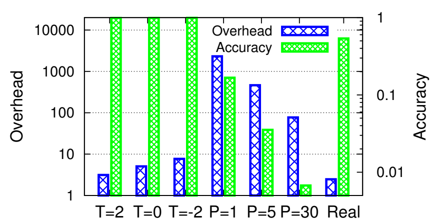

Fig. 2. Context-triggered measurement improves efficiency and accuracy of measurements for intra-LTE handovers (PhoneLab). ($T = x$: triggered measurement using x dBm threshold. $P = y$: periodic measurement every y min).

3.2 Crowd-sourced measurement

PhoneLab Deployment PhoneLab [5] is a smartphone testbed located at the University at Buffalo with more than 200 participants. Each participant receives a Nexus 5 device running Android Lollipop with unlimited Sprint data plan. Developers can deploy experiments on the devices by modifying the Android system.

To understand handover policies and performance impacts, we add a system service called *HandoverTrackerService* in Android system. This service monitors context information of the device and triggers active measurements. The lower layer information and measurement results are uploaded to servers. At the beginning of deployment, we keep collecting lower layer messages and perform active measurement periodically to avoid bias in the collected data. After analyzing the data and understanding when handover is triggered, we update the deployment and leverage the context-triggered measurement framework to reduce the measurement overhead.

To guarantee minimal influence of active measurements on user experience, heavy-weight measurements such as throughput are performed only when the screen is off and users are not interacting with the device. To control the power consumption of issued active measurements, we build an energy model for all measurements, and stop all measurements when the power consumption reaches 10% of total battery resources after the device is unplugged from the power source. We also enforce a limit on the daily data usage generated by the active measurements.

We deploy the measurement system on PhoneLab testbed and collect a dataset PHONELAB for 5 months from January 2016 to May 2016[2]. In total we observe 8403 cells and 283,556 intra-LTE handover events. For active measurements, we collect 49,594 throughput measurements, 159,210 ping measurements and 50,409 UDP burst measurements.

Local deployment We also deploy our measurement setup to 20 local users with unlimited AT&T data plans. We install an app called *HandoverTrackerApp* on their devices. The app collects the same information as *HandoverTrackerService* in PhoneLab deployment. We also collect data from a local device with T-Mobile service.

Both the crowd-sourced measurements and local deployment were IRB approved. The descriptions of the experiment and collected data are presented to the participants and they have the option to opt-out the experiment data collection.

MobileInsight dataset MobileInsight dataset [1] is a publicly available dataset containing lower layer cellular messages[3] collected from more than 8 US/Chinese network carriers spanning 3 years from year 2015 to 2018 using the tool MobileInsight [15]. The types of lower layer information collected by MobileInsight is similar to our data collection deployed on PhoneLab.

4 Handover Policy Inference

In LTE networks, the cells make decisions on when to initiate handovers. The handover decision process is not standardized in the 3GPP specification and is left to be defined using carriers' network configurations. To infer handover trigger policies in practice, we implement an RRC stack emulator that keeps track of the current device information, such as RRC connected state, connected cell ID, measurement configurations, and processes handover related messages. We feed RRC messages from each device to the emulator and output information including recent measurement reports and corresponding measurement configurations when processing handover initialization commands.

We first characterize the deployment of the cells and analyze how many cells the device usually observes from the signal strength measurements of neighboring cells in

[2] We confirm the inferred handover policy from PHONELAB are still current with the newer MobileInsight dataset as described later in §4.

[3] The MobileInsight dataset does not have active measurements on network performance.

Carriers	Sprint	T-Mobile	AT&T
Handovers triggered from Events	99.84%	100%	99.40%
Handovers triggered from Event A3	98.39%	89.51%	94.41%
Event A3 triggering handovers	91.87%	88.73%	87.58%
Handover count	283,556	286	330
Handover-involved cell count	6,304	33	45

Table 2. Overall statistics of handovers.

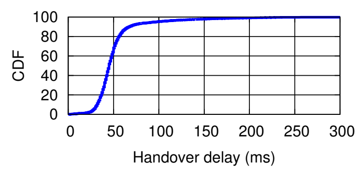

Fig. 3. Measured delay from the last measurement report to the time when handover starts (Sprint).

the PhoneLab dataset. We find that in 77.4% of cases the device observes at least one neighboring cell. In 41.9% of cases, the devices observe multiple neighboring cells. These neighboring cells can be of good signal strength. Among all the measurements, in 18.5% of cases, there is at least 1 neighbor cell with RSRP no worse than 5dBm lower than the serving cell. This indicates that carriers deploy cells densely to ensure connectivity and the potential chances of performing handovers between cells are high.

We find a strong correlation between measurement reports sent from the UE to the cell (shown in Figure 1) and handover events observed on the UE. As illustrated in Figure 3, 95.4% of handovers in Sprint occur within 100 ms after the measurement report is sent. If we consider a measurement report helps *trigger* the handover when a handover occurs within 500ms after a measurement report is sent, for all carriers studied, more than 99.4% of handovers are triggered by measurement reports, as shown in Table 2. Such close timing proximity implies potential causality.

We find the majority type of measurement reports that triggers handovers is event A3. Event A3 indicates that $\text{signal}_{\text{neighbor}} - \text{signal}_{\text{serving}} > threshold$. 98.39% of all handovers in Sprint network are triggered by event A3 reports. On the other hand, event A3 measurement reports have a high success rate of triggering handover. In Sprint network, for 91.9% of event A3 reports, the cell initiates a handover within 500ms. For 98.1% of event A3 reports, the cell initiates a handover within 2s. One reason why some reports fail to trigger handovers is that the device releases the RRC connection or the data collector stops collecting data before the handover occurs.

We find for each pair of cells, the threshold value of event A3 that triggers handover is statically configured and does not change over a long period. In the Sprint network, for the pairs of cells that have more than 100 handovers, 10.4% always set the threshold to 2 dBm and 2.9% always set it to 4 dBm. The other cells used the threshold of 2 dBm at the first 2 months of our data collection period and then changed to 4 dBm. The threshold of event A3 can affect how often handover happens. A lower threshold can be met more easily, thus encouraging more frequent handovers.

To understand whether handovers occur mostly for devices experiencing poor signal strength, we plot the distribution of the serving cell's RSRP values right before

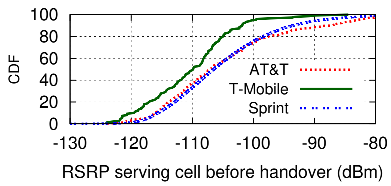

Fig. 4. RSRP of serving cell before the handover. Despite strong signal strength, handovers still occur.

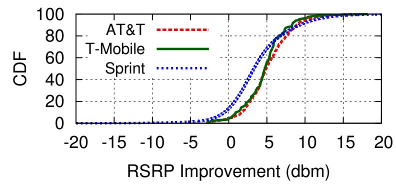

Fig. 5. In 90% of cases, signal strength improves after intra-LTE handover.

handovers in Figure 4. As shown, there is no direct relationship between current serving cell signal strength and handover occurrence. Handover occurs even when serving cell signal strength is already strong. For AT&T and Sprint, more than 20% of handovers happen when serving RSRP is stronger than -100 dBm. This is due to the fact that most of the handover events are triggered by event A3 using the relative signal strength threshold.

From the definition of the measurement report events, we know that handovers triggered by event A3 are likely to improve the signal strength of the UE. This is confirmed by Figure 5, showing the signal strength improvement after the handover.

We also validate our observation using the MobileInsight dataset. Among the 4873 observed handover events in the dataset, 86.5% are triggered by event A3 measurement reports. For 99.1% of the cell pairs, the A3 threshold value triggering handover is fixed. This confirms similar handover policies are used across time across different carriers. We next study the performance implications of such handover policies.

5 Performance Impact of Handover

We characterize the disruption during handover and the performance change after handover.

5.1 Performance disruption during handover

Due to the underlying physical radio layer transmission mechanism, during intra-LTE handovers, the device has to disconnect from the currently connected cell before connecting to the target cell, thus introducing a period during which the device is detached from the network preventing any data exchange. This unavoidably generates an interruption to ongoing traffic during handover. In intra-LTE handovers, the detach time is defined as the interval from the time when the device receives handover initialization message from the source cell and the time when the device successfully connects to the target cell.

To maintain good user experience during handover, the detach time needs to be kept low. In our observation, the detach time of successful handovers is within 35ms, which is quite minimal. However, handovers can fail due to various reasons such as insufficient radio resources in the target cell. When a handover fails, detach time can increase

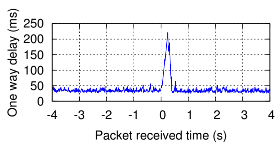

Fig. 6. An example of data pause caused by handover. 0 is the time when the handover occurs (Sprint).

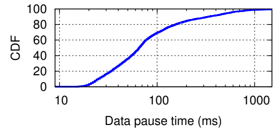

Fig. 7. Data pause in UDP burst measurement during handover (Sprint)

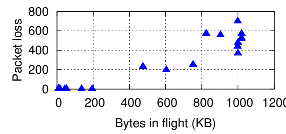

Fig. 8. TCP packet loss during handover due to limited buffer size in source cell for lossless handovers (AT&T)

dramatically. If the UE fails to connect to the target cell, the UE aborts the handover process and initiates connection re-establishment procedure with the source cell instead, which increase the detach time up to 775 ms from our observations. Moreover, in some cases, the re-establishment request is rejected by the source cell, and the UE is forced to release the connection and establish a new connection. This can further increase the detach time to 2.7 sec. The handover failure rate observed in Sprint is 0.18%.

Low detach time does not necessarily mean low impact on application-layer traffic. After the UE connects to the target cell, it may not resume data transmission from the new cell immediately. We use UDP burst measurement to characterize the *data pause time* during handover. Figure 6 shows an example of UDP measurement results. After the handover, the first few packets are delayed for about 200 ms. As Figure 7 shows, the median data pause time is 66 ms, which increases traffic jitter and may degrade real-time applications such as VoIP.

We examine packet losses during handover. During the handover, some data may be buffered in the source cell if the device is receiving data. Depending on how such data is handled, intra-LTE handovers can be categorized as *seamless*, which discards all data in the PDCP retransmission queue in the source cell, or *lossless* which forwards such data to the target cell. Recent work [17] shows that seamless handover is better in terms of goodput while lossless handover is better in terms of latency. We find that all cells in Sprint network use lossless handover, as no packet loss is found after the handover.

In AT&T network, we found three cells drop packets when there are handovers between them. In order to understand the underlying cause of this phenomena, we carry a Nexus 5 device that keeps downloading data from a local server while moving in the coverage area of these three cells. Server throughput is throttled at different values using the *tc* tool. All tcpdump traces from both the server and the client are captured.

Figure 8 shows bytes in flight right before the handover and the corresponding number of lost packets during the handover. We find that the number of lost packets has a strong correlation with the number of bytes in flight. We infer that there is a buffer in the cell that buffers packets during the data transmission between the server and the device. When a handover happens between the source cell and the target cell, the source cell tries to forward packets in the buffer to the target cell. However, during handover, the device cannot receive packets from the source cell in time, thus the number of accumulated packets at the cell may exceed the buffer size. In that case, the source cell has to drop packets during the handover.

We infer the cause of this unusual behavior of these three cells during handover is the poor configuration of their buffer size. From the experiment results, the buffer size

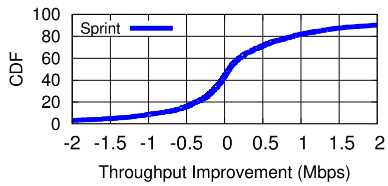

Fig. 9. Throughput improvement after handover (Sprint).

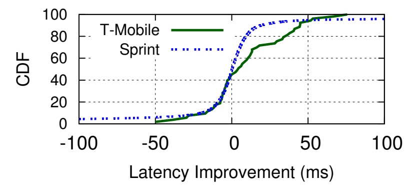

Fig. 10. Latency improvement after handover (Sprint).

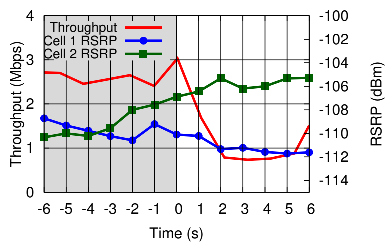

Fig. 11. An example of performance degradation after handover even through signal strength gets improved (Sprint). At time 0, a handover occurs from Cell1 to Cell2.

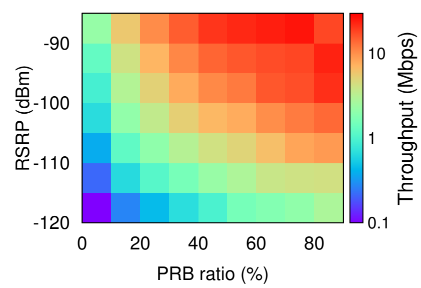

Fig. 12. The performance depends on both signal strength and allocated PRB (Sprint).

of these cells is between 250 KB to 400 KB. For normal TCP connections, the small buffer size does not cause packet loss due to the flow control in TCP. However, during handover, the small buffer size can easily lead packets loss, further degrading the performance of handover. In the worst case of our experiment, the duration of retransmission for the lost packets is 2.27 sec, which can greatly degrade user experiences.

5.2 Performance change after handover

One desirable goal of handover is to improve performance after switching to a new cell. We analyze the data to compare performance before and after the handover.

We filter out the throughput measurements that include handovers and calculate the average throughput value in the 5 seconds before the handover occurs and the average value in the 5 seconds after the handover. As shown in Figure 9, we find that the throughput does not improve consistently after the handover. In 43% of cases, the throughput decreases after the handover. Similar to throughput, neither latency (Figure 10) nor jitter improves consistently based on the ping and UDP burst measurements. The median improvement is close to 0 for all these metrics. The current handover algorithms do not appear to optimize performance.

As mentioned in § 4, the currently deployed handover decisions of all the carriers use signal strength as one of the key metrics and focus on improving signal strength after the handover. However, signal strength is an imperfect metric for performance, as performance also depends on other factors such as allocated PRBs determined by the cell load [6]. We calculate the throughput values each second using the crowd-sourced

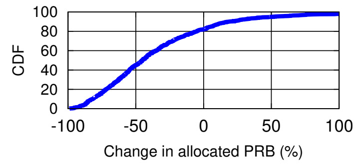

Fig. 13. The performance degradation is associated with higher cell load and less allocated PRBs (Sprint).

data and associate them with the RSRP and obtained PRB ratio. As shown in Figure 12, the achieved throughput is determined by the PRB ratio as well.

Figure 11 shows an example where signal strength increases after the handover, but the performance degrades. To confirm the root cause of such performance degradation, we look into the change in allocated PRBs after handover for the cases where throughput decreases. As shown in Figure 13, in most of such cases the allocated PRBs of the UE decreases after the handover, indicating that the target cell has a higher load.

Instead of making handover decisions simply based on signal strength, we propose that each cell maintains a 2-dimensional performance map from signal strength value ranges and load value ranges to performance ranges including throughput, loss rate, and delay. The performance values are updated by passively monitoring ongoing traffic at the cell. The cells exchange the performance information of the UE and UEs perceived signal strength of the potential target cell with its neighbors and trigger handover only if there is significantly enough performance improvements. The type of performance metric to be considered depends on user traffic demand inferred from its traffic types.

6 Conclusion

Handovers are essential for maintaining connectivity as users move with their devices. With the introduction of small cells in the incoming 5G network, handovers will become more frequent and critical. In this paper, we conduct the first comprehensive empirical study to investigate the decision strategies of intra-LTE handover in the wild and analyze their impact on performance. Our study examines currently deployed decision policies by carriers and sheds light on opportunities for improving the handover decision process with respect to application performance.

Our analysis shows that the policies enforced by carriers are not optimized in terms of performance. The key factor to decide handover is signal strength, and the handover trigger thresholds are found to be statically configured. We discover that the performance can degrade after the handover. We propose that the handover decision should depend on performance information predicted using both signal strength and cell load information.

7 Acknowledgements

This work is partially supported by NSF under the grants CCF-1628991 and CNS-1629763.

Bibliography

[1] MobileInsight Data Sharing. http://www.mobileinsight.net/data.html

[2] SnoopSnitch. https://opensource.srlabs.de/projects/snoopsnitch

[3] Allman, M., Paxson, V.: A Reactive Measurement Framework. In: Proc. Passive and Active Measurement Conference (2008)

[4] Balachandran, A., Aggarwal, V., Halepovic, E., Pang, J., Seshan, S., Venkataraman, S., Yan, H.: Modeling web quality-of-experience on cellular networks. In: Proceedings of the 20th annual international conference on Mobile computing and networking. pp. 213–224. ACM (2014)

[5] of Buffalo, U.: PhoneLab: A Smartphone Platform Testbed. https://www.phone-lab.org/

[6] Chakraborty, A., Navda, V., Padmanabhan, V.N., Ramjee, R.: Loadsense: Passively estimating cellular load. In: Communication Systems and Networks (COMSNETS), 2014 Sixth International Conference on. pp. 1–3. IEEE (2014)

[7] Cox, C.: An introduction to LTE: LTE, LTE-advanced, SAE and 4G mobile communications. John Wiley & Sons (2012)

[8] Dimou, K., Wang, M., Yang, Y., Kazmi, M., Larmo, A., Pettersson, J., Muller, W., Timner, Y.: Handover within 3GPP LTE: design principles and performance. In: Vehicular Technology Conference Fall (VTC 2009-Fall), 2009 IEEE 70th. pp. 1–5. IEEE (2009)

[9] Herman, B., Petrov, D., Puttonen, J., Kurjenniemi, J.: A3-Based Measurements and Handover Model for NS-3 LTE. In: MOBILITY 2013, The Third International Conference on Mobile Services, Resources, and Users. pp. 20–23 (2013)

[10] Huang, J., Qian, F., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O.: A close examination of performance and power characteristics of 4G LTE networks. In: Proceedings of the 10th international conference on Mobile systems, applications, and services. pp. 225–238. ACM (2012)

[11] Jansen, T., Balan, I., Turk, J., Moerman, I., Kurner, T.: Handover parameter optimization in lte self-organizing networks. In: Vehicular Technology Conference Fall (VTC 2010-Fall), 2010 IEEE 72nd. pp. 1–5. IEEE (2010)

[12] Javed, U., Han, D., Caceres, R., Pang, J., Seshan, S., Varshavsky, A.: Predicting Handoffs in 3G Networks. In: Proceedings of the 3rd ACM SOSP Workshop on Networking, Systems, and Applications on Mobile Handhelds. p. 8. ACM (2011)

[13] Kantola, R.: Performance of Handover in Long Term Evolution. Ph.D. thesis, Aalto University (2011)

[14] Li, Y., Deng, H., Li, J., Peng, C., Lu, S.: Instability in distributed mobility management: Revisiting configuration management in 3g/4g mobile networks. In: Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science. pp. 261–272. ACM (2016)

[15] Li, Y., Peng, C., Yuan, Z., Li, J., Deng, H., Wang, T.: Mobileinsight: extracting and analyzing cellular network information on smartphones. In: MobiCom. pp. 202–215 (2016)

[16] Li, Y., Xu, J., Peng, C., Lu, S.: A first look at unstable mobility management in cellular networks. In: Proceedings of the 17th International Workshop on Mobile Computing Systems and Applications. pp. 15–20. ACM (2016)

[17] Nguyen, B., Banerjee, A., Gopalakrishnan, V., Kasera, S., Lee, S., Shaikh, A., der Merwe, J.V.: Towards Understanding TCP Performance on LTE/EPC Mobile Networks. In: Proc. of AllThingsCellular Workshop (2014)

[18] Nikravesh, A., Yao, H., Xu, S., Choffnes, D., Mao, Z.M.: Mobilyzer: An Open Platform for Controllable Mobile Network Measurements. In: Proceedings of MobiSys (2015)

[19] Peng, C., Li, Y.: Demystify undesired handoff in cellular networks. In: Computer Communication and Networks (ICCCN), 2016 25th International Conference on. pp. 1–9. IEEE (2016)

[20] Shafiq, M.Z., Erman, J., Ji, L., Liu, A.X., Pang, J., Wang, J.: Understanding the impact of network dynamics on mobile video user engagement. In: The 2014 ACM international conference on Measurement and modeling of computer systems. pp. 367–379. ACM (2014)

[21] Tu, G.H., Peng, C., Wang, H., Li, C.Y., Lu, S.: How voice calls affect data in operational LTE networks. In: Proceedings of the 19th annual international conference on Mobile computing & networking. pp. 87–98. ACM (2013)

[22] Zhang, L., Okamawari, T., Fujii, T.: Performance Evaluation of TCP and UDP during LTE Handover. In: Wireless Communications and Networking Conference (WCNC), 2012 IEEE. pp. 1993–1997. IEEE (2012)

[23] Zhao, X., Ma, H., Jin, Y., Yao, J.: Measuring instability of mobility management in cellular networks. IEEE Network (2018)