

Accounting for Roaming Users on Mobile Data Access: Issues and Root Causes

Guan-Hua Tu, Chunyi Peng, Chi-Yu Li, Xingyu Ma, Hongyi Wang,
Tao Wang⁺, Songwu Lu

UCLA Computer Science, Los Angeles, CA 90095, USA;
CECA, School of EECS, Peking University, Beijing 100871, China⁺
{ghtu, chunyip, lichiyu, maxy12, hywang, slu}@cs.ucla.edu, wangtao@pku.edu.cn⁺

ABSTRACT

In this paper, we study how mobility affects mobile data accounting, which records the usage volume for each roaming user. We find out that, current 2G/3G/4G systems have well-tested mobility support solutions and generally work well. However, under certain biased, less common yet possible scenarios, accounting gap between the operator's log and the user's observation indeed exists. The gap can be as large as 69.6% in our road tests. We further discover that the root causes are diversified. In addition to the no-signal case reported in the prior work [23], they also include handoffs, as well as insufficient coverage of hybrid 2G/3G/4G systems. Inter-system handoffs (that migrate user devices between radio access technologies of 2G, 3G, and 4G) may incur non-negligible accounting discrepancy.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Network*; C.4 [Performance of Systems]: *Design Studies*

General Terms

Measurement, Experimentation, Performance

Keywords

Cellular Networks, Mobile Data Services, Data Accounting

1. INTRODUCTION

Mobile data accounting records the data usage volume for each roaming user over wireless cellular networks. It has become an increasingly important problem in recent years. On one hand, most cellular carriers use metered accounting, which charges mobile users based on the data volume they consume. It thus departs from the typical, flat-rate based charging over the wired Internet and requires accurate accounting. On the other hand, user mobility is prevalent in reality and mobile data traffic increases with the popularity of smartphones [22, 27]. The global mobile data traffic

grew 2.3 times in 2011, and is expected to double over the next four years [3]. To meet the rapidly increasing demand, mobile operators have been accelerating the deployment of higher-speed 3G/4G (third-generation/fourth-generation) cellular infrastructure from the legacy 2G/3G networks [2].

Although the cellular network accounting system has been generally working well, recent studies [23] have shown that, accounting gap, which measures the difference between data usage accounted by operators and data volume recorded at mobile devices, does exist under certain extreme conditions in reality. Overcharging may occur when a mobile user enters a no-signal or weak-signal zone, in which the device does not receive data packets but the operator charges the user for the undelivered volume. [23] models the accounting gap for no-signal scenarios with the packet lost/undelivered volume being $s \times t$, where s is the data transmission speed and t is the duration the user stays in the no-signal zone, and shows that the accounting gap is up to 450 MB. However, prior work studies indoor, static scenarios only, and the impact of user mobility is not explored. In this work, we examine how *mobility* affects data accounting in operational 2G/3G/4G systems. We focus on *outdoor, roaming* scenarios. Note that, it is quite common that cellular users access data services “on the go,” e.g., searching Google Map for route directions while driving on the road.

At a first glance, we think that mobility seems a non-issue for data accounting. Anyway, seamless mobility support has been an appealing highlight of cellular networks. In fact, it is the only large-scale, wide-area mobility management solution deployed in practice. After two-decade technology evolution, it has been operating pretty well. The network-initiated handoff solutions in 2G/3G/4G systems remain intact despite the evolution trend toward all IP-based design of the overall architecture. Data losses due to handoff may occur but the volume is expected to be generally negligible. However, our experiments contradict our initial belief and reveal some interesting, yet not necessarily common cases.

We address three aspects in mobility-oriented accounting: (1) Can non-negligible accounting gap be ever observed in the mobility scenario? (2) What are the root-cause factors beyond the known no or weak signals [23]? (3) Which factor plays the dominant role if new factors are indeed uncovered? To this end, we run experiments and perform analysis in two large cities over three major US carriers. Before proceeding with experiments and results, we rush to clarify what this work is not about. Our study does not seek to characterize the common cases. We identify what *possible* cases, even rare/worst cases, may go wrong in accounting for mobility. Though we tested 13 routes and hundreds of runs in two major areas, they may not capture the average-case performance in the statistical sense. Our results are biased to certain degree. Moreover, though we discover quite large gaps in certain settings, these

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiSys'13, June 25–28, 2013, Taipei, Taiwan

Copyright 2013 ACM 978-1-4503-1672-9/13/06 ...\$15.00.

somewhat biased results should not be interpreted as the failure of operational cellular systems. On the contrary, we show that current systems are generally successful despite few observed glitches.

Our study yields several findings. First, we discover that accounting gap indeed exists for the mobility scenario. The gap can be as large as 69.6% in our road tests. The volume discrepancy is route dependent and operator specific. Second, we further find that, the root causes are also diversified. Accounting gap is observed even with strong signal (measured in RSSI) settings. The key factor is the associated handoff, which may play a dominant role in certain scenarios. Third, various types of handoffs are triggered in operational 2G/3G/4G networks with distinctive quality of packet delivery. Operators have been using all deployed systems simultaneously whenever they can, and thus various handoffs are triggered. This practice is partly for offloading traffic from the high-end 4G/3G systems, partly due to partial deployment of high-end technology. As expected, intra-system handoff (across cells with the same radio access technology) works well and incurs little or no loss in 3G/4G systems. However, inter-system handoff (across radio technologies, e.g., between 3G and 2G, 4G and 3G) is problematic for data accounting. It leads to visible overcharging in many test routes; we observe the accounting gap ratios greater than 5% in 9 test routes. In one discovered setting, a popular handoff implementation, which uses buffering to improve wireless link performance, may negatively increase the observed accounting gap (in the range of tens to hundreds KBs). Fourth, low mobility speed may incur more occurrence of inter-system handoff (e.g., three times of that observed at high mobility speed), which leads to a larger accounting gap than the high-speed case. Fifth, we uncover a slightly new case for insufficient coverage. Certain regions are covered by legacy 2G/3G networks but not by high-end 3G/4G systems. The observed gap thus differs from that in the no-coverage case. Finally, we see the average accounting gap ratio between 0.0-40.1% when using five real applications: Web browsing, Email, FTP, Youtube and PPS Streaming, and 0.0-3.6% for a few mobile-phone users in their daily commute.

The rest of the paper is organized as follows. Section 2 introduces the data accounting process and mobility support in cellular networks. Section 3 describes the study methodology. Sections 4 and 5 summarize the results and root causes. Sections 6 and 7 discuss the root causes of accounting gap observed in handoff and insufficient coverage, respectively. Section 8 describes other factors contributing to the accounting gap. Section 9 provides possible solutions to the accounting gap. Section 10 compares with the related work, Section 11 concludes this paper and Appendix documents the trace processing details.

2. BACKGROUND

We now introduce the cellular network architecture, and its accounting and mobility management.

Cellular network architecture: Figure 1 illustrates a conceptual architecture for packet-switched cellular networks [19], which are widely used in the 3G and 4G systems. It consists of core network (CN), radio access network (RAN), and mobile devices. The major components in RAN are base stations (BSes)¹, which provide wireless access to the mobile devices and relay packets between CN and mobile devices. CN is composed of three network elements: (1) gateway (GW), which forwards packets between external networks and RAN, acting like routers in the Internet, as well as records per-device data usage and transfers accounting records to the charging

¹In practice, there shall be BS controllers. BSes in this work also provide functions that BS controllers offer.




Figure 1: Packet-Switched cellular network architecture and its accounting system.

| Acronym | Term | Generation | Predecessor | MaxRate |
|---------|---|------------|-------------|--------------|
| GSM | Global System for Mobile communications | 2G | - | 9.6Kbps |
| GPRS | General Packet Radio Service | 2.5G | GSM | 56-114 Kbps |
| EDGE | Enhanced Data Rates for GSM Evolution | 2.5G | GRPR | 384 Kbps |
| UMTS | Universal Mobile Telecommunications System | 3G | EDGE | 2 Mbps |
| HSPA | High Speed Packet Access | 3.5G | UMTS | 14.4-42 Mbps |
| HSPA+ | Evolved HSPA | 3.5G | HSPA | 84 Mbps |
| cdmaOne | Code Division Multiple Access One | 2G | - | 14.4 Kbps |
| EVDO | Evolution-Data Optimized or Evolution-Data Only | 3G | cdma2000 | 2.4 Mbps |
| eHRPD | Evolved High Rate Packet Data | 3.5G | EVDO | tens of Mbps |
| LTE | Long Term Evolution | 4G | HSPA/eHRPD | 150-300 Mbps |

Table 1: Major cellular network technologies.

server; (2) Databases (DBs), which store the mobile device’s registration, location and profile information (e.g., the used data plan); (3) Charging server, which charges mobile users based on their data usage.

In the past two decades, cellular networks have been evolving to provide higher speed, e.g., from 9.6 Kbps (2G GSM) to 2 Mbps (3G UMTS) and further to 300 Mbps (4G LTE). Different generations of cellular networks mainly vary in their RAN technology. Table 1 summarizes the major operational cellular technologies [13, 19, 20, 28]. In practice, an operator continues to upgrade its cellular network technologies; a hybrid network is usually deployed at any given time. In our measurement, all 2G/3G/4G technologies are observed in the same area, and all technologies except GSM/cdmaOne are observed in use.

Accounting: Accounting is a critical management feature to enable the cellular operators to realize their profit. Most cellular operators adopt a usage-based scheme to charge mobile users. As shown in Figure 1, the core gateways record the volume of data packets that traverse them in both uplink (i.e., from the mobile device to the Internet) and downlink (i.e., from the Internet to the mobile device) directions. The status of data packets turns from “unaccounted” to “accounted” after they pass those gateways.

Mobility support: The transmission range of one BS is limited (e.g., ranging from several hundred meters to several kilometers). To provide seamless mobile access, the cellular operator deploys a large number of BSes, each covering a small area. When a mobile device roams from the coverage of one BS to another, it performs a handoff (HO) to switch its associated, serving BS. It is one crucial technique for mobility support. The handoff procedure works as follows. It is first triggered by the mobile device or the serving BS when needed (e.g., when the perceived signal strength is too weak); then, the serving BS (say, BS1) finds another BS (say, BS2) that probably provides better performance (e.g., with stronger signal strength); it sends BS2 a handoff request to reserve radio

resources for the mobile device. After that, the mobile device disconnects from BS1 and connects to BS2. During this process, internal gateways also update mobile location and adjust the forwarding path accordingly. The ultimate goal for HO is to not disrupt ongoing services. Though its concept is relatively straightforward, the implementation in reality is quite complex. With various cellular technologies and BS types involved, there are multiple types of handovers. We will address their performance impact, design issues and accounting effect later in this work.

3. EXPERIMENTAL METHODOLOGY

We now describe our experimental methodology to study the impact of mobility on data accounting.

Experiment setting: We test all three major US operators, which serve about 243 million mobile subscribers and cover 75.3% of the US market [18]. We denote them as OP-I, OP-II and OP-III for privacy concerns. Hybrid 2G/3G/4G networks are used in these carrier networks. For instance, all three technologies of LTE (4G), UMTS (3G) and EDGE (2G) are observed in the same area.

We perform experiments in two largest metropolitan areas in the US, New York (NY) and Los Angeles (LA). The test area covers 16 towns and 4 major freeways in two 29×64 square kilometers and 48×58 square kilometers regions. We have tested with 13 routes, which cover four types of roads: (1) *local* in *rural* areas, (2) *local* in *urban* areas, (3) *freeway* in *rural* areas and (4) *freeway* in *urban* areas. The basic route information is shown in Table 2. The first five routes are located in NY while the others are in LA. The route distance ranges from 1.9 to 40.9 km, with the median value being 15 km. The short route (i.e., 1.9-km Route-10) is explored because of its interesting network deployment. Difference also exists between the routes in NY and those in LA. The NY routes are located in the rural area around a medium-sized town (north NYC). Six routes in LA are close to the downtown area, and the other two are near the mountain and coastal areas. We select routes mainly by their importance to mobile users, e.g., roads with heavy traffic or necessary pathways between a rural area and an urban area. For example, Routes 7 and 13 are major freeways connecting north/south and west/east LA areas, respectively. Route 12 is a major route between Malibu city and Westwood area.

Note that, we do not intend to use these 13 routes to represent all possible cases (i.e., the statistics may be biased). Constrained by the time spent on each experiment (we need to wait for the data volume charged by the operators before we continue another experiment), we use real traces to analyze a few cases and demonstrate what happens for accounting on the go. These routes sample diverse geographic regions and different network deployment, thus shedding light on how mobile accounting works in reality.

We have run driving tests during three months (from August 1 to October 31, 2012). While driving on a test route, we establish a data session from the mobile device to our deployed server (via UDP or TCP) or popular Internet services (e.g., Youtube and PPStream). We then collect the data volume recorded by operators and mobile devices, as well as log network status traces. In our tests, we use six Android phone models, including Samsung Galaxy S1/S2/Note/Stratosphere, and HTC Incredible S/Sensation that run on 2.3.5, 2.3.6, 4.0.3 or 4.0.4 OS versions. To ensure clean runtime environment (i.e., no more background services), we conduct factory reset first and disable “Background data” and “Auto-sync” features before each test [23, 24].

Collected results: For each test, we record data volume observed by different parties. In particular, we collect (1) V_{ue} , the data volume perceived by mobile devices, (2) V_{op} , the data volume ac-

| Name | City | Area | Type | Distance (km) |
|----------|------|-------|---------|---------------|
| Route-1 | NY | Rural | Freeway | 28.5 |
| Route-2 | NY | Rural | Freeway | 19.8 |
| Route-3 | NY | Rural | Local | 11.7 |
| Route-4 | NY | Urban | Local | 8.8 |
| Route-5 | NY | Rural | Local | 9.8 |
| Route-6 | LA | Rural | Freeway | 31.7 |
| Route-7 | LA | Urban | Freeway | 19.2 |
| Route-8 | LA | Urban | Local | 9.4 |
| Route-9 | LA | Urban | Freeway | 7.2 |
| Route-10 | LA | Urban | Local | 1.9 |
| Route-11 | LA | Rural | Freeway | 15.0 |
| Route-12 | LA | Rural | Local | 28.3 |
| Route-13 | LA | Urban | Freeway | 41.0 |
| Total | | | | 232.3 |

Table 2: Route information.

counted by the operator, and (3) V_{sr} , the one recorded by our deployed server if used. To ensure that the V_{ue} is accurately recorded, the mobile data usage is collected from two tools. One is Traffic Monitor [9], an Android application in Google Play to collect data usage for WiFi and cellular interfaces. It records data volume for each application with 0.01 KB accuracy. The other is our developed tool that uses the TrafficStates class [10] in Android SDK to retrieve the data volume of mobile devices on a per-application basis. Note that the data volume recorded by both tools contains the headers of both network layer (i.e., IP) and transport layer (i.e., TCP/UDP) in our experiments. We use both to record the mobile data volume and verify whether the volume is consistent or not.

The data volume V_{op} charged by the operator is obtained via two methods [23]. One is to dial a special number to retrieve the current monthly data usage and calculate the used data volume during experiments. It usually takes 5–30 minutes for the operators to update the usage record. We disable the network access (i.e., packet-switched service) of mobile devices until they update data records. To further mitigate the impact of the updating latency, we have multiple mobile user accounts (e.g., multiple sim cards) for each operator. Before the operator finishes updating data usage of mobile user account A , we use another mobile user account B to run experiments. The other is to log onto the operators’ Web sites and access itemized data usage records. All three operators support the DIAL-IN method, while OP-I and OP-III also support the second online method. All three support data usage with 1 KB accuracy. [14] specifies that the data usage recorded by operators covers both application data volume as well as network-layer and transport-layer headers.

In order to verify the accuracy of traffic monitor tools and how the operators account data usage (whether they consider network layer and transport layer headers or not) in current practices, we conduct an experiment to send/receive several UDP datagrams, each of which carries 1 byte UDP payload. If carriers do not account IP/UDP headers, the data usage recorded by operators should increase by 1 KB after 1024 UDP datagrams are sent/received. However, we observe that the recorded data usage already achieves 1 KB after tens of UDP datagrams are sent/received. The data usage collected by both TrafficMonitor and our tool also exceeds 1KB; the volume is consistent with those accounted by operators.

In addition to data volume, we also collect real-time cellular network status and packet delivery logs at the phone. In the network trace, we *periodically* log the following information: *timestamp*, *operator*, *network type*, *cell identifier*, *signal strength*, and *location* (i.e., GPS latitude and longitude). The record interval is 250 ms. Table 3 shows an example of mobile network traces using OP-I

| TIME(ms) | OP | TYPE | CID | RSSI | LAT | LON |
|----------|------|----------------------|-----------|------|-------------|---------------|
| 7590 | OP-I | EDGE | 37605 | -103 | 34.0513862 | -118.50484915 |
| 8777 | OP-I | UMTS | 58873657 | -109 | 34.05124171 | -118.50496962 |
| ... | ... | ... | ... | ... | ... | ... |
| 72194 | OP-I | UMTS | 56645543 | -113 | 34.04644347 | -118.50847563 |
| 73221 | OP-I | UMTS | 588735920 | -113 | 34.04644347 | -118.50847563 |
| ... | ... | ... | ... | ... | ... | ... |
| 157982 | OP-I | unknown ^a | n/a | -113 | 34.03924701 | -118.50924827 |

Table 3: An example of cellular network trace.

^aIt can be “search for *”. Both imply that the phone perceives no networks (no coverage) and we use “unknown” to denote both cases hereafter.

network. We use relative time² to record *timestamps*. The *network type* (TYPE) denotes the used radio access network, and our data set covers eight 4G/3G/2G technologies: LTE, HSPA+, HSPA, HSDPA, UMTS, EVDO, EDGE and GPRS, introduced in Section 2. The *cell identifier* (CID) is the associated cell ID. We use TYPE and CID to determine whether a handoff occurs; details are given in Section 4. The *signal strength* (RSSI) records the strength of perceived radio signals from the associated cell; it may vary greatly upon a handoff.

The packet delivery trace is logged in an *event-triggered* manner. When a new packet is sent/received, the mobile phone logs the following attributes: *timestamp*, *sequence number* of the packet received/sent, and the accumulative delivery information including the number of received/sent bytes or packets. To obtain the sequence number and timestamp of sent packets, we insert them in each packet that is sent from our deployed server. For popular Internet services, packet traces do not contain such information.

4. ACCOUNTING GAP FOR ROAMING USERS

In this section, we first offer an assessment of the mobility impact on accounting over all tested routes. We then use an example to illustrate where the charging gap occurs. We seek to answer two key questions:

- Does nonnegligible accounting discrepancy exist for roaming users?
- Is there any other factor contributing to such gap beyond the no-signal factor identified in prior work [23]?

4.1 All Tested Routes

We first run simple experiments to study whether accounting gap exists over test routes or not. We download UDP datagrams from our deployed server at a constant rate (say, 200 kbps). Note that, our test does not seek to capture the common-case scenario, but identify accounting issues in simple mobility settings. We drive at the full speed (e.g., 104 km/h (65 mph) on freeways or 56 km/h (35 mph) at local) in the absence of heavy vehicle traffic. The results depend on several factors, e.g., the adopted applications, source rate, driving speed and operator policy. We will address their impacts in later sections. We run experiments at least three times on each route. For those routes with small or even no gap (< 1MB), we observe that the gap results are stable. For the routes of interest, e.g., those with large accounting gap, we repeat 10-15 runs and still observe that the gap be consistently large. Figure 2 plots the median accounting gap ($V_{OP} - V_{ue}$), gap ratio (Gap/V_{OP}) and unit-distance gap ($Gap/Distance$) from top to down. Note that it only shows results under the given experiment setting (i.e., the mobile

²The time starts recording once the experiment begins.




Figure 2: Median accounting gaps, ratios and unit-gaps on all the routes in preliminary experiments. The dash lines denote 3 MB gap, 10% ratio and 500 KB gap per km.

device is constantly transferring data during the test runs) and does not plot results for all mobility scenarios and real applications. The other experimental settings and applications, e.g., Youtube, will be elaborated in Section 8.

We make three observations. First, our experiments show that *the accounting gap caused by user mobility indeed exists and may affect many people*. The gap is observed in both rural and urban areas, and on local roads and freeways with heavy traffic (e.g., annual average daily traffic reaches 374,000 vehicles). In some routes, the accounting gap ratio even reaches up to 69.6%. For instance, on Route 12, the data volume accounted by OP-II is 44.3 MB while that recorded by the mobile device is only 13.5 MB.

Second, *the accounting gap is route dependent*. For example, in OP-I, the accounting gap varies from 0.0 MB to 20.2 MB (49.0%). The unit-distance gap ranges from zero (or several KB) to 2.7 MB per km (OP-II, Route-10). Significant accounting gaps do not exist in most routes. For instance, only 6 out of 39 cases (i.e., route plus operator) have more than 500 KB gap per km or 10 cases have the gap ratios larger than 10%. However, large accounting gap does exist in certain routes, e.g., Route 12 for all the three operators, and Routes 9 and 10 for OP-II and OP-III.

Third, *the accounting gap is operator specific*. For example, on Route 3, the accounting gaps and ratios for OP-I, OP-II and OP-III are 0.06 MB (0.2%), 0.7 MB (2.4%) and 6.6 MB (22.7%), respectively. Nine routes have ratio differences larger than 5% among three operators. They are Routes 1, 2, 3, 5, 6, 9, 10, 12 and 13 in our test. In terms of the accumulative volume gap on all test routes, OP-I, OP-II and OP-III yield the gap as large as 20.6 MB (5.3%), 48.9 MB (12.7%), and 52.4 MB (13.4%), respectively. OP-II and OP-III perform worse than OP-I.

4.2 Case Study on An Example Route

To better understand what is going on, we first use an OP-I trace on Route-12 for a case study. The route takes about 30-minute drive. Our measurement shows that OP-I charges the mobile user of $V_{op} = 41.1$ MB while the phone only sends and receives $V_{ue} = 21.0$ MB. The accounting gap reaches 20.1 MB, about 49.0% of the accounting volume. For OP-II and OP-III, the measured discrepancy turns into 30.8 MB (69.6%) and 23.4 MB (51.9%), respectively.

We seek to find answers to three issues: (1) Why does accounting




Figure 3: An example of network status trace and packet delivery log on Route 12 using OP-I.

gap occur? (2) What factors contribute to the gap? (3) Are there other factors in addition to the no-signal case discovered before [23]? To this end, we plot the phone traces of packet reception, RSSI and detected events over time (minute) in Figure 3.

The top plot of Figure 3 shows the reception rate at the phone in one-second bin. We find that, the accounting gap is incurred by the failure in packet delivery, which is the norm rather than an exception in the mobility setting. In the case, the client reception rate is expected to be 200 kbps, the same as the source. However, the actual link rate fluctuates and even falls to zero occasionally. For example, packet delivery pauses more than six times (e.g., during [1, 3] and [19.4, 24.5] minutes). Unfortunately, our prior study shows that the accounting system is based on the local view at the core network [23]. Operators count the data packets traversing the core network gateway, no matter whether those packets have been successfully delivered to end users or not. Failure to deliver those packets that have been accounted by the core network results in accounting gap.

The next issue is on which factor leads to failure of packet delivery in the mobility scenario. We first observe that packet loss occurs when the RSSI is low. The middle plot of Figure 3 shows that RSSI fluctuates along the route. The minimal observed RSSI value is -113 dBm, which infers that the phone enters into a dead zone³. As expected, the time window of low (or no) packet reception matches well with the one with low RSSI, for example, during the intervals of [1.5, 3], [14.5, 16], [19.4, 23.2] and [28, 30] (min). This finding also confirms the no-signal/weak-signal case of [23].

However, we find that the accounting gap may also occur with relatively high RSSI setting, e.g., during the interval of [3, 4.5]. Our trace analysis shows that the decisive factor is handoff. Surprisingly, even though handoff generally works well in 2G/3G/4G networks, it may occasionally incur large amount of packet delivery loss. The bottom plot of Figure 3 marks all the handoff events learned from the network trace. It turns out that handoff can be classified into two categories: intra-system handoff and inter-system handoff. Upon an intra-system handoff, the mobile device still uses the same RAN and CN, but different base stations. In contrast, the mobile device switches to different RAN and CN after an inter-system handoff. In this example, we observe 4 inter-system handoffs and 31 intra-system handoffs. We also note that the events of weak-signal, no-signal coverage and handoff are not

³ [15] indicates that the lowest signal strength measured at phones is -113 dBm, which is too weak to enable data links.

| | Handoff (HO) | Non-Handoff (NH) |
|------|------------------|---|
| RSSI | various HO types | SC: RSSI > -105 dBm WC: RSSI $\in (-113, -105]$ dBm NC: RSSI $= -113$ dBm or network TYPE is "Unknown" |

Table 4: Event classification.

orthogonal. Handoff often occurs when the RSSI value is small (e.g., around -113 dBm). It is usually triggered in a weak-signal or no-signal zone. A phone performs handoff to another base station with stronger RSSI when it leaves the coverage of its original base station.

5. DIVERSIFIED ROOT CAUSES

In this section, we show that the root causes for overcharging are more diversified beyond the no-signal/weak-case factors identified in prior work [23]. We further identify which factor plays a dominant role.

5.1 Root-Cause Event Classification in Traces

To identify root cause events and their impact, we classify all events into two nonoverlapping categories: *Handoff* (HO) and *Non-Handoff* (NH), based on the occurrence of handoff. In the *HO* category, more sub-events are defined based on the handoff type (Section 6). The *NH* category is further divided into three sub-cases based on RSSI values: strong coverage (SC), weak coverage (WC) and no coverage (NC). Table 4 lists our event classification. We mainly use two RSSI thresholds, -113 dBm and -105 dBm. -113 dBm is the minimal RSSI observed on test phones, or when the network operation mode turns into "UNKNOWN"; regarding weak signal strength, there is no agreed definition in the literature. We define it based on the end-user perception. When RSSI is smaller than -105 dBm, the signal strength icon on the test phones is retreated to the weakest signal strength level, e.g., Level 1 of four levels.

We next seek to compute the accounting gap for each event. Based on the operator's record, it is easy to calculate the *total* accounting gap between mobile users and the operator. To further understand which factor or event is more crucial, we need to learn the accounting gap during each cause event. There are two methods to do that. The first is to re-do experiments for each event. To this end, we first identify when and where each event happens; for

| | SC | | WC | | NC | | HO | |
|--------|-------|-------|-----|-------|------|-------|------|-------|
| | Dur | Ratio | Dur | Ratio | Dur | Ratio | Dur | Ratio |
| OP-I | 190.4 | 90.7% | 5.3 | 2.5% | 0.9 | 0.4% | 13.4 | 6.4% |
| OP-II | 202.8 | 88.1% | 3.8 | 1.7% | 15.5 | 6.7% | 8.0 | 3.5% |
| OP-III | 201.0 | 86.9% | 3.0 | 1.3% | 1.1 | 0.1% | 27.2 | 11.8% |

Table 5: Time durations (minute) for four events.

| | SC | | WC | | NC | | HO | |
|--------|------|-------|-----|-------|------|-------|------|-------|
| | Gap | Ratio | Gap | Ratio | Gap | Ratio | Gap | Ratio |
| OP-I | 1.5 | 7.1% | 0.6 | 2.8% | 1.9 | 9.4% | 16.6 | 80.4% |
| OP-II | 13.0 | 26.5% | 0.7 | 1.5% | 25.4 | 51.9% | 9.8 | 20.0% |
| OP-III | 15.1 | 28.7% | 0.0 | 0.0% | 0.0 | 0.0% | 37.4 | 71.3% |

Table 6: Accounting gaps (MB) for four events.

example, the first inter-system handoff happens at the 3rd minute in the example of Section 4.2; then we re-run this experiment only on this sub-route. This method looks reasonable but not feasible in practice. First, it is hard and even impossible to guarantee to cover only a single event in any experiment. The start and end of the experiments cannot be accurately controlled; some events only last several seconds (e.g., the intra-system handoff around the 13th minute). Moreover, the action depends on the historical status. The phone performs handoff at certain time instant because it is associated with another BS before but the signal strength from that BS degrades later. By repeating experiments only around handoff, the phone may not even connect to the original BS. The second method is to decouple the accounting gap for each event from the complete route trace. This is our processing choice. There are two steps. We first extract all time windows of each event, and single out those for handoffs and non-handoffs. We then classify them according to RSSI values or handoff types, and finally calculate the accounting gap (i.e., packet loss) during each event window based on the packet reception log. The details are given in Appendix.

5.2 Findings

Table 5 shows the total time duration for four events on all thirteen test routes. Note that, the number of experimental runs differs on each route. Thus, for each single event, we calculate the average time duration in all experiment runs performed on one route. It shows that these three operational cellular networks have good coverage; SC occupies more than 86-90% of the test time and OP-I is slightly better than OP-III in this test case. WC and NC are rare for OP-I and OP-III, because they usually trigger handoffs when signal strength degrades and both operators have good BS deployment coverage. For OP-II, NC time is longer because of the dead zone where the phone cannot connect with any OP-II base station (more details will be discussed in Section 7). OP-III has longer handoff duration; this is due to its radio access technology. It will be further discussed in Section 6.

Table 6 summarizes the accounting gap for four events on all test routes. From this table, we can find out which event plays an important role to affect accounting gap for mobile users. We observe that major root-cause events for the accounting gap vary with operators. For OP-I and OP-III, HO contributes to the majority of

| | SC | WC | NC | HO |
|--------|-----|-----|-----|-----|
| OP-I | 0.0 | 0.1 | 2.1 | 1.2 |
| OP-II | 0.1 | 0.2 | 1.6 | 1.2 |
| OP-III | 0.1 | 0.0 | 0.0 | 1.4 |

Table 7: Unit-time gap (MB/min) for four events.

the accounting gap (80.4% and 71.3%). For OP-II, in addition to HO, NC contributes to the main portion (51.9%). It is deployment specific and OP-II fails to provide sufficient coverage for mobile users. It is not surprising to see that the SC event also contributes a relatively large portion of the accounting gap (about 26-28% for OP-II and OP-III). In most times, mobile users stay in the SC zone, even when driving. In terms of unit-time gap (Gap/duration), the gap is much smaller in SC, see Table 7. Moreover, we see that NC and HO have higher unit-time gap than the other two events. This is not difficult to understand. Mobile users experience their worst packet delivery when there are no signals, or handoffs are triggered when signals fluctuate. However, three more issues remain to be addressed: (1) How do HO and NC events affect packet loss and thus accounting misalignment? (2) Are there hidden mechanisms or insights to improve the current system? (3) How do other factors, such as mobility speed, traffic types/source rates, and network deployment, affect the result? We next elaborate on these aspects.

6. WE PAY FOR HANDOFF

In this section, we explore how handoff affects accounting gap and leads to overcharging for users. We identify different cases and their root causes. Handoff incurs accounting gap because data transmission suspends during it. The mobile device must disconnect with the serving BS to connect with another BS, because the device is usually unable to concurrently connect to both BSes. The data transmission suspension starts from the time when last packet is received before a handoff, and continues until a new packet is received after a handoff. Different types of handoffs result in distinctive suspension durations.

6.1 Impact of Handoff Types

Handoff can be broadly classified into two categories: *intra-system* handoff and *inter-system* handoff. During an intra-system handoff, the mobile device still uses the same RAN and CN. In contrast, after an inter-system handoff, the mobile device switches to different RAN and CN. Using this criterion, we have four network sets: $S_{2G} = \{EDGE, GPRS\}$, $S_{3G1} = \{HSPA, UMTS\}$, $S_{3G2} = \{EVDO\}$, and $S_{4G} = \{LTE\}$. Inter-system handoff implies that users move from one set to another. For intra-system handoff events, there are two cases: (1) users move from one network to another network within the same set, e.g., from EDGE to GPRS; (2) users stay at the same network but move to another cell, e.g., within GPRS RAN.

Handoff Occurrences: From the traces of the 13 test routes, we have discovered 22 inter-system handoffs (OP-I: 5, OP-II: 9, OP-III: 8) and 554 intra-system handoffs (OP-I: 46, OP-II: 50, OP-III: 458). Note that the number of handoff occurrence is the average of all DL-UDP-200kbps experimental runs performed on each test route. The number of handoff occurrence is influenced by other factors to be discussed in Section 8.

We make two observations. First, most handoff events are intra-system handoffs. The number of intra-system handoffs represents 90.2%, 84.7%, 98.3% of the total number of handoffs within OP-I, OP-II and OP-III, respectively. Therefore, inter-system handoff events are not widely observed, since operators deploy the same radio access networks in most of 13 test routes.

Second, the number of intra-system handoff events within OP-I and OP-II is much smaller than that in OP-III. We find that, our test phones are mainly using OP-I and OP-II's 3G HSPA networks on Routes 2, 3, 6-11 and 13. When 3G HSPA networks are used, the cell identifier is not going to be changed. However, this scenario is not observed on other radio access technologies including 4G LTE, 3G EVDO/UMTS and 2G EDGE/GPRS. It may be caused by




Figure 4: An example of composite handoff (1 inter-HO, 5 intra-HOs).

specific implementations of phone vendors or operator deployment policy. [12] states that HSPA operators such as Telstra in Australia are reporting mobile broadband downlink speed of 2.3 Mbps with the range up to 192 km from the cell site. Our longest test route is 41 km and the sending rate of UDP is 200 kbps. Thus, this observation may be caused by operator deployment policy. In Section 8, we conduct another experiment to study the root cause.

Composite Handoff: From our experimental results, we find that more than one handoff may be performed within the same data transmission suspension period. We define it as a composite handoff. In contrast, if only one handoff occurs, we denote it as a single handoff in this work.

We observe that the composite handoff occurs in three scenarios: (1) RSSI is close to (or equal to) -113 dBm; (2) high mobility speed; (3) upgrade or downgrade of RAT. For the first scenario, we find that when the phone is unable to associate with a cell that has strong signals, it keeps switching RATs within the same cell, e.g., alternating between HSPA and UMTS, or switching among a set of cells (e.g., two or three cells) using the same RAT. When users are stuck in such a “jumping” scenario, users do not receive any packets and incur large accounting gap until they move away from this area. We illustrate this scenario in Figure 4, which plots the RSSI and network type of a 112.9s composite handoff. It contains one inter-system handoff and five intra-system handoff events. The RSSI degrades to -113 dBm during [10, 104] seconds. The phone is initially stuck in the “jumping” scenario between HSPA and UMTS networks during [32,35] seconds, followed by in two HSPA cells (53 and 57) within [35, 78] seconds. This situation disappears when the phone moves away from this area and uses the GPRS networks at the 105th second. In the second scenario, a user moves to the next new cell and triggers another handoff before data suspension incurred by the previous handoff completes.

In the third scenario, we find that the inter-system handoff between 3G HSPA and 2G EDGE networks does not always directly move users to the target RAN. For example, users may traverse some intermediate RANs before reaching the final RAN. An inter-system handoff from 3G HSPA to 2G EDGE may go through HSPA, UMTS, GPRS and EDGE networks. The sequence of traversed RANs is highly dependent on operator deployment history. Most operators would upgrade their existing GPRS and UMTS base stations to EDGE and HSPA BSes, respectively, to offer higher rate. Users are thus able to access four RATs at the same place. The selection of RANs is determined by the signal strength of each RAN measured at the mobile device.




Figure 5: Accounting gap (MB) and time duration (s) with handoff types.

In this paper, we further define a composite handoff as a composite inter-system handoff if an inter-system handoff is observed among the associated handoff events; otherwise, we denote it as a composite intra-system handoff.

6.1.1 Accounting Gap and Duration of Single/Composite Handoff

In addition to the regular three experimental runs on each test route, we perform the DL-UDP-200kbps experiments on where inter-system handoffs occur with extra 10 runs. For intra-system handoffs, we only analyze the results collected from three runs, since the number of intra-system handoffs observed for each operator is more than 120. Figure 5 plots the accounting gap and the transmission suspension time for single/composite inter-system and intra-system handoffs, which are denoted as inter-S/inter-C and intra-S/intra-C, respectively. The bar, upper line, and lower line mark the median, maximum, and minimum values of accounting gap or transmission suspension duration in each case, respectively.

We make three observations. First, the accounting gap is always observed when inter-system handoffs occur no matter they are single or composite handoffs. The minimum gaps for single and composite inter-system handoffs within in OP-I, OP-II and OP-III are 1.3 MB, 0.1 MB, 0.2 MB, 0.6 MB, 0.5 MB and 0.3 MB, respectively. Second, the accounting gaps for most intra-system handoffs are almost zero, i.e., the median accounting gaps for single intra-system handoffs within OP-I, OP-II and OP-III are all 0.0 MB and those for composite intra-system handoffs within OP-I, OP-II and OP-III are 0.5 MB, 0.1 MB and 0.0 MB, respectively. Third, the time duration of most inter-system handoffs is longer than that of most intra-system handoffs. For example, the median time durations of single and composite inter-system handoffs within OP-I are 13.2s and 42.6s, respectively, but those of single and composite intra-system handoffs are merely 0.0s and 0.1s, respectively. Hence, we believe that the inter-system handoffs play an important role in terms of accounting gap.

6.1.2 Accounting Gap and Duration of Handoffs with Same or Different RATs

To understand why inter-system handoffs cause larger accounting gap and longer time duration than most intra-system handoffs,




Figure 6: Accounting gap varies with handoff types. Left: Inter-system handoff; Right: intra-system handoff.




Figure 7: Suspension duration varies with handoff types. Left: Inter-system handoff; Right: intra-system handoff.

| From \ To | 4G | | 3G | | | 2G | |
|-----------|------|--------|---------|---------|--------|-------------|---|
| | LTE | HSPA | UMTS | EVDO | EDGE | GPRS | |
| 4G | LTE | OP-III | × | × | OP-III | × | × |
| 3G | HSPA | × | OP-I,II | × | OP-II | OP-I, OP-II | |
| | UMTS | × | OP-I,II | OP-I | × | OP-II | × |
| | EVDO | OP-III | × | × | OP-III | × | × |
| 2G | EDGE | × | OP-I,II | OP-I,II | × | OP-I, OP-II | × |
| | GPRS | × | × | × | × | OP-I, OP-II | × |

Table 8: List of handoff types observed.

we further study handoff events with the same or different RATs. Table 8 lists the observed handoff types in terms of RATs⁴. Both OP-I and OP-II support S_{2G} and S_{3G1} networks, but GPRS is not observed in OP-II. Handoffs are observed within and between LTE and EVDO networks for OP-III, but OP-III does not deploy S_{2G} and S_{3G1} networks. Though both OP-I and OP-II claim to support LTE (S_{4G}), we do not observe it due to phone hardware constraint.

Figures 6 and 7 plot the accounting gap and the duration of data transmission suspension with respect to different handoff types. The bar, upper line, and lower line denote the median, maximum and minimum values, respectively. Note that we do not differentiate single or composite handoffs, since they have similar behaviors in terms of accounting gap and time duration of inter-system and intra-system handoffs. The composite handoffs only contribute 21.1%, 16.9% and 3.4% of all handoff events in OP-I, OP-II and OP-III, respectively. Our results show that they may significantly affect the maximum accounting gap, but not the median value, for each handoff type.

We make three observations. First, we find that most 3G/4G intra-system handoffs do not incur accounting gap and have almost zero data transmission suspension duration. However, 2G intra-system handoffs contradict this finding. Second, the accounting gap with an inter-system handoff is larger than that with an intra-

system handoff. For example, for OP-II, the accounting gap for inter-system handoff from 2G to 3G networks is about 1 MB while the intra-system handoff within 2G networks is about 0.22 MB. The intra-system handoff within 3G network even does not incur any volume gap. Third, longer data transmission suspension time usually leads to larger accounting gap. For instance, for OP-III, the data transmission durations for inter-system handoff from 3G to 4G networks and from 4G to 3G networks are about 24s and 66s, leading to 0.61 MB and 1.34 MB volume gap, respectively. However, there exists an exception and counter-intuitive finding in OP-I. The data transmission suspension for inter-system handoff from 2G to 3G networks is longer but the gap is smaller.

Based on the above findings, we would like to ask three questions: (1) Why do all intra-system handoffs within 2G networks have accounting gap while most intra-system handoffs within 3G and 4G network have no gap? (2) Why do all inter-system handoffs cause accounting gap? (3) Why may shorter data transmission suspension time create larger accounting gap?

6.2 Certain Intra-System Handoff Incurs Accounting Gap But Others Do Not

Table 9 shows the accounting gap and data transmission suspension for eight intra-system handoffs. We make two observations. First, for the intra-system handoffs within 4G LTE, 3G EVDO, 3G HSPA and 3G UMTS networks, no accounting gap is seen in OP-I, OP-II and OP-III. The data transmission suspension time is around 0.04 to 0.05 seconds, close to the packet inter-arrival time at the sending rate of 200 kbps with 1KB packet size, i.e., $1/(200/8) = 0.04$ s. The underlying reason is that current 3G/4G standards support soft handover/handoff [20], which enables to simultaneously connect to multiple base stations and send/receive data to/from them. In contrast, in hard handoff, the mobile device breaks the connection to the original base station before the connection to the new base station is established. The soft handoff consequently provides seamless access to the original and new base stations. How-

⁴In the current practice, carriers deploy several RATs simultaneously, which affect what types of handoffs are observed.

| | Type | OP-I | | OP-II | | OP-III | |
|----|------------|------|-------|-------|------|--------|------|
| | | Gap | Dur | Gap | Dur | Gap | Dur |
| 4G | LTE<>LTE | n/a | | n/a | | 0 | 0.04 |
| 3G | EVDO<>EVDO | n/a | | n/a | | 0 | 0.04 |
| | HSPA<>HSPA | 0 | 0.04 | 0 | 0.04 | n/a | |
| | HSPA<>UMTS | 0 | 0.05 | 0 | 0.05 | n/a | |
| | UMTS<>UMTS | 0 | 0.04 | 0 | 0.04 | n/a | |
| 2G | EDGE<>EDGE | 0.19 | 7.95 | 0.23 | 8.12 | n/a | |
| | EDGE<>GPRS | 0.5 | 18.71 | 0.28 | 8.16 | n/a | |
| | GPRS<>GPRS | n/a | | n/a | | n/a | |

Table 9: Median accounting gap (MB) and data suspension duration (second) for intra-system handoffs.

ever, soft handoff requires extra signal processing and radio resources. For example, the UMTS soft handoff [20] requires different codes for downlink transmissions, so that the mobile device is able to distinguish signals from different base stations. In summary, 3G/4G mobile devices are able to connect to multiple base stations simultaneously. The data suspension duration during intra-system handoff is thus significantly reduced.

However, in some exception cases, intra-system handoffs within 3G/4G networks still incur accounting gap. Since soft handoff cannot be applied to two base stations using different frequency bands [20], mobile users suffer from the issues similar to the inter-system handoff. In practice, to improve the spatial diversity of cellular networks, operators configure base stations to use different frequency bands. If the mobile device is moving to a new base station using a different frequency band, it has to disconnect with the current base station and then connect to the new base station. Accounting gap is thus observed in this scenario.

Moreover, intra-system handoffs within 2G networks lead to larger accounting gap and longer data transmission suspension than those within 3G/4G networks. 2G networks are using TDMA [29], where a mobile device only sends or receives packets at given time slots over a specific frequency band, and soft handoff is not supported. The 2G mobile device cannot receive packets from multiple base stations concurrently. Moreover, GPRS and EDGE adopt different Modulation and Coding Schemes [29] so that the device requires extra time to synchronize with new base stations. In summary, the intra-system handoff between EDGE and GPRS networks takes longer time than that within pure EDGE or GPRS networks.

6.3 Inter-System Handoff Always Incurs Accounting Gap

Table 10 shows the median accounting gap and data transmission suspension for three types of inter-system handoffs. We can see that inter-system handoffs always incur larger accounting gap than its intra-system counterparts. For example, the accounting gap of inter-system handoff between 3G HSPA and 2G EDGE within OP-II is 0.6 MB, whereas the accounting gap of 3G HSPA and 2G EDGE intra-system handoffs is 0.0 MB and 0.23 MB, respectively. The major cause is that the mobile device uses different radio access technologies, such as 3G EVDO and 4G LTE for OP-III, TDMA-based 2G and CDMA-based 3G for OP-I and OP-II. Most mobile devices cannot connect to two base stations using different RATs due to hardware constraints (e.g., only a single antenna to receive data on one frequency). They have to disconnect with the original BS during handoff. This leads to inevitable suspension time. Another reason is that the core network also varies. It takes time to update/modify states in core network components before establishing a new radio access bearer [17].

| | Type | OP-I | | OP-II | | OP-III | |
|--------|------------|------|------|-------|------|--------|------|
| | | Gap | Dur | Gap | Dur | Gap | Dur |
| 4G<>3G | LTE<>EVDO | n/a | | n/a | | 1.3 | 65.6 |
| 3G<>2G | HSPA<>EDGE | 0.5 | 17.5 | 0.6 | 10.8 | n/a | |
| | UMTS<>EDGE | 0.5 | 18.0 | 0.6 | 10.5 | n/a | |

Table 10: Median accounting gap (MB) and data suspension duration (second) for inter-system handoffs.

6.4 Shorter Suspension Time May Incur Larger Accounting Gap

Table 10 further shows a counter-intuitive result: shorter suspension time may incur larger accounting gap. According to the related work [23], the accounting gap would be given by $data_suspension_time \times application_sending_rate$ in principle. However, this equality does not hold for the mobility scenario. Its root cause is related to buffering and the operator’s policy on buffer management.

We first verify that buffering indeed exists in 2G/3G networks. We set the sending rate at our UDP server slightly higher than that can be accommodated by the receiving mobile device, and observe the reception behavior at the mobile. Figure 8 (Left) and (Middle) plots the sequence number of received packets and packet travel time (that measures the time spent from the server to the device during the packet travel process) using an OP-II 3G network, respectively. At the beginning, the device receives packets without any loss and the packet travel time gradually increases, possibly due to queueing delay at the buffer. It then receives packets intermittently (since the buffer becomes full) and the packet travel time is limited by the receiving speed and stabilizes around 10 seconds. If the buffer were nonexistent, packets would be received intermittently since the beginning. We thus infer that buffering does exist.

We next show that buffering does not help data accounting during inter-system handoffs. To analyze the impact of buffer management during an inter-system handoff on accounting gap, we repeat the experiment in the inter-system handoff area and configure the sending rate of our UDP server at 400kbps and 800kbps. Intuitively, the buffer cannot be observed if the receiving rate at the mobile device is higher than the sending rate of UDP server. Before the experiment, we use the Speedtest.net [8] to measure the maximum transmission rate at the mobile device. If the rate is higher than 800 kbps, we go to the place with weaker signal strength within the coverage of the same base station. Once an experiment starts, we wait for 5 seconds (to ensure full buffer) and drive through the handoff area. Figure 8(Right) plots packet reception before and after an inter-system handoff. We make two observations. First, the buffer is full when the inter-system handoff occurs. The mobile device receives all packets without loss before the 180th packet. Upon receiving the 180th packet, the mobile device starts receiving packets intermittently until the 885th packet is received. Based on the packet reception trace, we infer that the buffer is full when data transmission suspension starts. Second, the user does not receive the packets that are stored in the buffer at the time of the inter-system handoff occurrence. The data suspension starts from the time when the 885th packet is received to the instant when the 1400th packet is received. The mobile device does not receive any packets from 886 to 1399. If the packets in the buffer were not lost, the mobile device would intermittently receive packets from 886 to 1399. However, we do not observe such events after data transmission suspension completes. We believe that the user has lost all packets stored in the buffer when the inter-system handoff occurs.

The mobile user loses all packets in the buffer when the inter-system handoff occurs, even though the data transmission suspen-




Figure 8: Packet reception. Left: during buffer experiment; Middle: packet travel time during buffer experiment; Right: with an inter-system handoff.




Figure 9: Buffer size in 2G, 3G and 4G networks

sion is short. This is the reason why shorter data transmission suspension may incur larger accounting gap. It depends on how many packets are buffered within the RAN when an inter-system handoff occurs. Along this direction, the buffer size does affect the accounting gap. Larger buffer can potentially store more packets, thereby incurring larger accounting gap upon inter-system handoffs.

We further quantify how bad the accounting gap due to buffering can be in practice. Figure 9 shows that, the estimated buffer size varies from 18KB to 356KB in operational networks. Such a buffer size also offers a worst-case upper bound for the accounting gap. The buffer size for 4G is not observed because we do not find a spot where the rate in 4G is lower than 800kbps⁵.

Note that the accounting operations specified by the standard [14] do not take into account the buffer drops triggered by the inter-system handoff events. Therefore, operators should not be held accountable for the incurred accounting gap. The handoff of a mobile device is triggered when its signal strength is weak. During the handoff, it may be charged for the packets that it does not receive. Although carriers may be aware of such handoff-incurred accounting gap, they are unable to avoid billing users based on current accounting operations without new mechanisms, which will be discussed in Section 9.

7. INSUFFICIENT COVERAGE

Insufficient coverage is commonly observed [4, 6], and hybrid networks (e.g., 2G and 3G) in the same region are also common practice in reality. We now study their impact on data accounting for mobile users.

Insufficient Coverage: Prior work demonstrates that accounting gap exists in no-signal zones in static cases [23]. It also holds true in mobility scenarios. Due to insufficient coverage, mobile users may cross no-signal zones on a regular basis while driving. For example, on Route 12, we discover no coverage on an 8-km sub-

route in a residential area near mountains; RSSI is -113 dBm and no handoffs occur within OP-II. The no-coverage area contributes 71.1% of accounting gap on Route 12 using OP-II. Among all three operators, OP-II experiences severe issues due to insufficient coverage (shown in Section 5.2).

Hybrid Network: One more interesting case is insufficient coverage due to hybrid network deployment, where both high-speed (e.g., 3G/4G) and low-speed (e.g., 2G) cellular technologies coexist. Uncovered by high-speed networks, but covered by low-speed networks, mobile users might experience accounting gap due to the improper switching between these two networks. When users leave the coverage of the high-speed network, operators migrate them to the low-speed network through inter-system handoff. If the application source rate is higher than that can be accommodated by the low-speed technology, packets have to be dropped, thus incurring accounting gap. Our tests show that, the accounting gap strongly depends on how long the user stays in the low-speed network and the receiving rate at the mobile device. We present the results in OP-II, and the other two carriers are similar. Figure 10 shows the average accounting gap of different rates and durations. We make two observations. First, given the same application source, longer duration leads to larger accounting gap. For instance, the accounting gap for 1-min and 2-min is 2 MB and 4 MB, respectively, when the rate is 400kbps. Second, given the same duration, faster application source rate leads to larger accounting gap. For example, the accounting gap of 400kbps and 800kbps is 2 MB and 4.9 MB, respectively, when the duration is 1 minute.

We note that the root cause differs from the conventional case of insufficient coverage. In this case, the gap ratio grows faster than the source rate increase. For example, the accounting gap increases to 245% (i.e., $(800 - 123)/(400 - 123) = 2.45$) if the application source rate increases from 400kbps to 800kbps. The reason is that, the mobile device is able to receive packets in the low-speed network. Thus, we have to consider the receiving rate in this example. In this experiment, the average transmission rate in OP-II is about 123kbps. We find that the accounting gap is roughly given by $(Application\ Source\ Rate - Mobile\ Receiving\ Rate) \times duration$.

The accounting gap caused by hybrid networks is not as large as that with insufficient coverage. However, it is more often observed than the former case, and it may last for longer time. In current practice, even 4G LTE device still makes a voice call through the legacy 2G network.

8. FACTOR IMPACT

We assess how six factors affect mobile data accounting.

Real Applications: We now quantify the impact of applications on the accounting gap. We conduct experiments with five applications including Web browsing (Webkit [11]), Email (Gmail [5]),

⁵This is the maximum transmission rate supported by our ISP.




Figure 10: Accounting gap in OP-II hybrid network

| | Web Browsing | Email | FTP | Youtube | PPS |
|--------|--------------|-------|------|---------|-------|
| OP-I | 0.0% | 0.0% | 0.6% | 0.7% | 24.8% |
| OP-II | 0.0% | 0.0% | 0.6% | 1.6% | 40.1% |
| OP-III | 0.0% | 0.0% | 0.6% | 0.7% | 21.3% |

Table 11: Average accounting gap ratio ($Gap/V_{op}(\%)$) with real applications on Route 12.

FTP (AndFtp [1]), Youtube and PPS [7], on Route 12. The first four applications are running over TCP, whereas the last one, a popular peer-to-peer video streaming application [7], is using UDP. During driving, we keep on fetching the homepage of CNN.com in Webkit and refreshing the inbox in Email. For the FTP test, we download a 2.9GB file. We play a 1-hour 360p video in both Youtube and PPS. The applications are stopped once we arrive at the destination. Each application is performed three runs.

Table 11 shows the average accounting gap ratio (Gap/V_{op}) for real applications. We make three observations. First, there is no accounting gap observed on both Web browsing and Email applications for all three operators. The reason is that, the downlink data traffic is triggered by the phones' request messages. If the requests are not successfully delivered to the Web or Email server, web pages or emails will not be sent to the phones. Second, only small accounting gap exists for both FTP and Youtube applications. This is because they rely on the TCP flow and congestion control to adapt their sending rates. The reason that the gap still exists is that the ongoing session is not immediately stopped when a handoff occurs or a no-coverage area is encountered. Moreover, prior to the inter-system handoff event, as the signal strength gets worse, the TCP sender congestion window and transmission rate become smaller and slower, respectively. It thus leads to a smaller number of packets stored in the buffer than that in UDP. Mobile users do not experience severe accounting discrepancy like UDP. We also observe that the frozen session is then resumed by the TCP retransmission mechanism. Figure 11 shows the traces of one FTP test over time for OP-I. They include the TCP sequence numbers observed at the phone, RSSI, and network type. We see that a handoff does suspend the data transmission, which is during [210, 290] seconds. However, the accounting gap is very small, since the FTP server does not keep on sending packets during handoff. Third, the accounting gap ratios incurred by PPS are 24.8%, 40.1%, 21.3% for OP-I, OP-II and OP-III, respectively. We discover that the ratios are much smaller than those in our DL-UDP-200kbps experiments, though the PPS streaming rate is higher than 200 kbps. Despite a rate control mechanism, PPS responds much slower than those TCP-based applications, thus leading to larger accounting gap. The volume gap is closely related to the application/transport-layer control, so applications with an inert rate control may degrade. This is consistent with the results in the static scenarios [23].




Figure 11: TCP sequence numbers with handoff occurrences in time scale (s).

| User | OP-I | | OP-II | | OP-III |
|---------------|-------------|---------------------------------|---------------------------------------|--|--|
| | 1 | 2 | 3 | 4 | 5 |
| Apps | LINE, Gmail | Whatsapp, Gmail, WeatherChannel | FaceBook, Messenger, PPS, LINE, Gmail | Pandora, Radio, Gmail, Whatsapp, Stock | Facebook, Whatsapp, Skype, LINE, Gmail |
| Route Dis. | 41.9 km | 75.5 km | 89.6 km | 76.8 km | 18.8 km |
| V_{UE} (MB) | 37.2 | 198.7 | 1204.3 | 387.2 | 73.9 |
| V_{OP} (MB) | 37.2 | 199.6 | 1249.7 | 389.8 | 74.3 |
| Gap (MB) | 0.0 | 0.9 | 48.2 | 2.6 | 0.4 |
| Gap Ratio | 0.0% | 0.4% | 3.6% | 0.6% | 0.5% |

Table 12: Accounting gap for driving commuters during March 18-29, 2013.

Real Mobile Users: We also study the impact of mobility on data accounting for five commuters who drive in the LA area (driving distance ranges from 18.8 km to 89.6 km) and have data plans with OP-I, OP-II and OP-III. We record the data usage measured at the mobile devices and the one accounted by operator for two weeks (March 18-29, 2013) in Table 12. In this field trial, we only focus the accounting gap observed during daily commute. Since these five commuters have WiFi access in office or at home, we thus configure their mobile devices to disable cellular networks (packet-switched services only) and use WiFi networks whenever WiFi networks are available (i.e., commuters arrive at offices or homes). To quantify the accounting gap on real commuters, except Traffic-Monitor and our traffic capture tool, we retain those applications on their devices already and do not install other new applications. Note that, we do not claim that these samples well represent the daily usage for average commuters in all scenarios. Instead, we show that mobile accounting gap exists for real mobile users. Among these users, the most popular applications are Gmail and messaging services, e.g., Whatsapp, LINE or Skype (not for video/voice services here). We see that the accounting gap of User 3 is 48.2 MB (3.6%), whereas that of Users 1, 2, 4 and 5 is below 2.6 MB and 0.6%. User 3 takes his children to preschool and occasionally plays cartoon movies to his kids via PPS during his daily commute. This scenario may be quite common for family commuters but not for those who remain single.

Application Source Rate: We now discuss how application source rate affects accounting gap in the test routes. Due to space limit, we present only the results of Routes 1 and 2 with three different source rates. Figure 12 plots the accounting gap of these two routes. We observe that the accounting gap increases with the application source rate. For example, in OP-II, the overall accounting differences of Route 1 are 0.05 MB, 2.6 MB and 9.6 MB for the 200kbps, 400kbps and 800kbps sources, respectively.

We also observe that the source rate may affect the number of




Figure 12: Accounting gap with source rate.

intra-system handoff occurrences. On Route 7 (19.2 km), we observe only one intra-system handoff for both OP-I and OP-II, but 76 times for OP-III, when the source rate is 200kbps. However, when the source rate decreases to 1.6kbps, the numbers of handoffs for OP-I and OP-II increase to 14 and 15, respectively. Similar results are observed on our all Samsung and HTC phones. We think that this phenomenon may be induced by the operator's network deployment or its packet forwarding mechanism similar to Mobile IP [25], which is used to reduce handoffs and packet losses during an ongoing data session. Unfortunately, due to lack of information on the practices by OP-I and OP-II, we are unable to verify our conjecture.

Mobility Speed: We now examine how mobility speed affects the accounting gap incurred by handoff. There are two findings. First, low mobility speed incurs more inter-system handoffs. We perform the DL-UDP-200kbps experiments on our test routes in the absence of no-coverage areas but in the presence of hybrid networks, e.g., Route 1 or 10, at three speeds: low, medium and high. In local routes, the low, medium, and high speeds are 24 km/h, 40 km/h and 56 km/h, respectively. In freeway routes, they are 72 km/h, 88 km/h and 104 km/h, respectively. We find that the speed will not directly affect the accounting gap caused by handoff. Note that our speed does not exceed the speed limit for a roaming terminal that cellular networks support (e.g., 560 km/h over LTE [26]). We observe that lower mobility speed leads to more inter-system handoff occurrences.

To understand the root cause, we use an example (in OP-II) to illustrate what happens at both low and high speeds. In Route 10, we observe 6 and 2 handoff occurrences for low mobility speed and high mobility speed, respectively. To illustrate the impact of mobility speed on handoff occurrence, we look into the first 0.35 km of Route 10, when the first inter-system handoff is observed during driving at low speed mobility. Figure 13 shows the RSSI changes over time and distance on the test route. The dash line specifies when and where inter-system handoffs occur at low and high speeds, respectively. When a mobile device stays longer in the zone with weak RSSI (i.e., < -106 dBm), an inter-system handoff is more likely to be observed. For example, the RSSI within the distance range of [0.08, 0.16] on the route is around -108 dBm to -113 dBm. The mobile device stays in this area for 13 seconds and 4 seconds at low and high speeds, respectively. An inter-system handoff occurs at the distance of 0.18 km when the mobility speed is low. The RSSI improves to -93 dBm at the distance of 0.19km due to this handoff. In contrast, the mobile device does not have any inter-system handoff triggered at the same spot at high mobility speed. However, its RSSI quickly improves to -106 dBm at the distance of 0.19 km. The mobile carrier does not initiate an inter-system handoff unless the mobile device stays in the zone with weak signals longer than a pre-specified time threshold. We have not fully analyzed the precise, triggering threshold for




Figure 13: Handoff occurrence with mobility speed.

an inter-system handoff. Based on our experimental estimate, the time threshold is about 5 to 15 seconds and varies with operators.

Second, we observe that the 3G radio access technology adopted by OP-III is not as dependable at high mobility speed as that used by OP-I and OP-II. When we remain static, the transmission rate of EVDO is around 452 kbps with the RSSI being -95 dBm. We then configure phones to use 3G networks only and perform the DL-UDP-200kbps experiment on Route 13 at the speed of 104 km/h. We find that the average transmission rate of 3G EVDO for OP-III is reduced to 160 kbps (with the RSSI being between -75 dBm and -95 dBm). Accounting gap is observed and mostly due to the SC events. In contrast, the average transmission rate is close to 200 kbps for both OP-I and OP-II, and there is no visible accounting gap (i.e., the gap ratio below 1%).

Vehicle traffic: In order to gauge the impact of vehicle traffic on the accounting gap, we perform the DL-UDP-200kbps experiments on Route 13, which traverses the LA downtown area. The experiments are conducted during both rush hours (6pm-7pm, weekdays) and non-rush hours (9pm-10pm, weekdays), during the two-week period (from October 1st to October 12th, 2012). We do not observe any inter-system handoff occurrence for both OP-I and OP-II. However, in OP-III, we observe that two inter-system handoffs were triggered during non-rush hours when traffic jam was observed on October 5th, and six inter-system handoffs occurred during rush hours on Oct-1st, Oct-2nd and Oct-12th. Since the RSSI is strong (i.e., from -75 dBm to -85 dBm) at the spot where inter-system handoffs are triggered, there is no reason for OP-III to migrate users to another system (e.g., 3G EDVO) due to low RSSI. Therefore, we presume that this migration is mainly due to its system capacity limit. However, an inter-system handoff is not always observed within OP-III even during rush hours, though heavy vehicle traffic implies more roaming users and may potentially consume more capacity. We speculate that vehicle traffic may not have strong correlations with the inter-system handoff occurrence. Not all vehicles during rush hours use OP-III networks and request radio resources for data transfer or voice calls. Consequently, heavy vehicle traffic does not imply that all carriers always experience resource shortage.

Gray Zone: We observe gray zones similar to the spots shown in [4] on some test routes. On Route 4, there is 0.5-km road nearby a shopping mall. When we drive through it, we do not receive any packets, and an accounting gap thus occurs within OP-I. However, the interesting thing is that the RSSI of the road is from -76 dBm

to -87 dBm and no handoff occurs. Moreover, this is also observed on Route 11 for OP-III. Although gray zones are not commonly observed (i.e., there are only two zones within 232 km for all three operators), they may affect a large number of users. The first one is nearby a shopping mall in city, and the second one is on the freeway with the heaviest traffic in LA. Both zones observe many users on a daily basis.

9. POSSIBLE SOLUTIONS

In this section, we discuss several possible solutions to mitigating the accounting gap caused by handoff or insufficient coverage. Each solution has its pros and cons. Which one is more appropriate highly depends on the usage scenario.

Suspend Accounting: Stop accounting incoming packets until handoff is completed. The merit of this approach is that, users completely eliminate the accounting gap caused by handoff. However, it does not reduce the accounting gap incurred by insufficient coverage. The perceived quality for data transfer is worse due to longer data transmission suspension time.

Report Unsent Packets: The RAN records the volume of packets that are not successfully sent to the phone, and reports this volume to internal routers (i.e., accounting elements). Then internal routers deduct the unsent data volume from the user's data usage. The merit of this proposal is to minimize the accounting gap caused by both handoff and insufficient coverage. However, its prerequisite is that reliable delivery mechanism, e.g., acknowledgement mode [16], is enabled in RAN. Without the feedback from the mobile device, RAN does not know whether packets are delivered or not. Therefore, if the unreliable delivery mechanism, e.g., unacknowledged transmission mode [16], is used, RAN cannot provide such information to internal routers.

Client-Based: Each application server shall effectively control its sending rate when delivering data to the mobile device depending on the mobile device's feedback. In the absence of feedback from the device, the application server shall immediately decrease its rate. Accounting gap is then mitigated (i.e., little data will be sent to the mobile during handoff). However, it is not quite practical to expect all applications to implement efficient rate control mechanisms.

Proxy-Based: Setup a proxy for mobile devices. All packets for such devices shall be relayed to this proxy and then forwarded to mobile devices. The proxy can dynamically enable/disable packet forwarding to mobile devices through cellular networks. The merit of this proposal is that users are able to control packet forwarding at the proxy depending on the network status (e.g., when approaching handoffs or insufficient coverage areas), and reduce the accounting gap without modifying current applications. Unfortunately, deployment and operations of the proxy raise concerns. In current practice, carriers distribute user data traffic to various deployed middle-box (e.g., http proxy or NAT) machines. If our proxy-based solution can be combined with these middle-box functionalities at the deployed servers, the maintenance cost will be reduced.

10. RELATED WORK

In 3G/4G cellular networks, subscribers are charged based on the traffic volume of mobile users. A recent work has studied the discrepancy of the 3G accounting system [23]. It reported several scenarios where users are charged for what they do not receive. It identifies the root cause as open-loop operations in 3G accounting without taking proper feedback from users. In contrast, our work focuses on data accounting under the mobility scenario,

where users handover from one network to another. It complements the scenarios studied in [23].

Mobility-related performance issues in cellular networks have been reported in several earlier studies. [27] offers comprehensive mobility performance assessment of a commercial HSPA network. Regarding handoffs, it notes that the triggers and the consequences of handoffs are not predictable and favorable in many cases. It shows that in nearly 30% of all handoffs, selecting a base station with poorer signal quality has happened. [21] states that the maximum time for handover is 114 to 140 seconds and the average time is 20 to 30 seconds. The instability and unpredictability of handovers on performance is observed, but accounting-related issues are not addressed.

11. CONCLUSION

With rapid deployment of 3G/4G cellular infrastructure, more users have data access while roaming around. Unlike the wired internet, most operators adopt usage-based charging, rather than flat-rate charging. Users are thus concerned about whether their data usage is correctly accounted for or not. In this work, we conduct experiments on three US carriers to study how accounting works for users on the go.

Our study shows that roaming users do not receive packets but are charged by operators during inter-system handoffs (across 2G, 3G, and 4G systems) and when driving through no-signal zones. This is mainly because packet drops during handoff events are not taken into consideration during the standardized accounting operations. The problem is that, data accounting is not halted when handoff is performed and buffered packets are dropped. Consequently, mobile users pay for what they do not receive during inter-system handoffs. Despite being operator specific and route dependent, the accounting gap is largely predictable since handoffs and insufficient coverage can be traced and gauged over time. Our ongoing effort seeks to build an accounting map that logs the observed gap volume for roaming users. Once constructed, users can prefetch it and act accordingly to minimize the potential overcharging.

12. ACKNOWLEDGMENTS

We greatly appreciate the insightful and timely comments from our shepherd, Dr. Matt Welsh, and the anonymous reviewers for their constructive feedback. We also thank all participants in the mobile accounting experiments.

13. REFERENCES

- [1] AndFTP - LYSESOFT, v2.9.8. <http://www.lysesoft.com>.
- [2] At&t coverage viewer. <http://www.wireless.att.com/coverageviewer/>.
- [3] Cisco visual networking index: Global mobile data traffic forecast update, 2011-2016. <http://www.cisco.com>.
- [4] Dead Cell Zones. <http://www.deadcellzones.com/>.
- [5] Gmail. <https://play.google.com/store/apps/details?id=com.google.android.gm>.
- [6] Open Signal. <http://opensignal.com/coverage-maps/US>.
- [7] PPS. <http://www.pps.tv/>.
- [8] Speedtest.net - Ookla. <https://play.google.com/store/apps/details?id=org.zwanoo.android.speedtest>.
- [9] Traffic Monitor - RadioOpt GmbH. <https://play.google.com/store/apps/details?id=com.radioopt.widget>.
- [10] TrafficStats. <http://developer.android.com>.
- [11] WebKit, Default Web Browser in Android. <http://developer.android.com/reference/android/webkit/package-summary.html>.

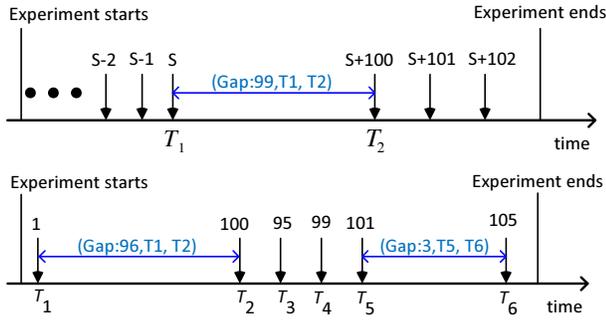


Figure 14: Accounting gap processing with packets arrival in order and out of order.

- [12] 4g americas - 3gpp white paper for release 7 and release 8, 2008. http://www.4gamericas.org/documents/3GPPRel-7andRel-8_White_Paper07-08-08.pdf.
- [13] 3GPP. TS23.060: GPRS; Service description; Stage 2, Dec. 2006.
- [14] 3GPP. TS32.298: Telecommunication management; Charging management; Charging Data Record (CDR) parameter description, Sep. 2007.
- [15] 3GPP. TS27.007: AT command set for User Equipment (UE), Dec. 2011.
- [16] 3GPP. TS25.322: Radio Link Control (RLC) protocol specification, Sep. 2012.
- [17] 3GPP. TS43.129: GSM/EDGE Radio Access Network; Packet-switched handover for GERAN A/Gb mode, Sep. 2012.
- [18] CNET. Competitive wireless carriers take on at&t and verizon, 2012. http://news.cnet.com/8301-1035_3-57505803-94/competitive-wireless-carriers-take-on-at-t-and-verizon/.
- [19] H. Holma and A. Toskala. *WCDMA FOR UMTS - HSPA EVOLUTION AND LTE*. Wiley, 2007.
- [20] H. Holma and A. Toskala. *LTE for UMTS: Evolution to LTE-Advanced*. Wiley, 2011.
- [21] X. Liu, M. Seshadri, A. Sridharan, H. Zang, and S. Machiraju. Experiences in a 3g network: Interplay between the wireless channel and applications. In *ACM Mobicom*, 2008.
- [22] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das. Understanding traffic dynamics in cellular data networks. In *INFOCOM'11*, April 2011.
- [23] C. Peng, G. hua Tu, C. yu Li, and S. Lu. Can We Pay for What We Get in 3G Data Access? In *ACM MOBICOM*, 2012.
- [24] C. Peng, C.-y. Li, G.-H. Tu, S. Lu, and L. Zhang. Mobile data charging: new attacks and countermeasures. In *ACM CCS*, pages 195–204, 2012.
- [25] RFC2002: IP Mobility Support, 1996.
- [26] I. T. S. Sesia and M. Baker. *LTE - The UMTS Long Term Evolution: From Theory to Practice*. Wiley, 2009.
- [27] F. P. Tso, J. Teng, W. Jia, and D. Xuan. *Mobility: A double-edged sword for hspa networks*. In *ACM MobiHoc*, 2010.
- [28] A. D. V. Vanghi and B. Vojcic. *The cdma2000 System for Mobile Communications: 3G Wireless Evolution*. Pearson Education, 2004.
- [29] I. C. Yi-Bing Lin. *Wireless and mobile network architectures*. Wiley, 2001.

APPENDIX

A. Event Detection: We first extract time windows for handoffs and non-handoffs. After that, we classify them according to RSSI values or handoff types. We induce the occurrence of handoff based on the change of TYPE or CID in network traces. Say, one network trace has $TYPE_X$ and CID_X at the x -th timestamp. Therefore,

$$HO \Leftarrow TYPE_X \neq TYPE_{X+1} \mid CID_X \neq CID_{X+1},$$

For instance, Table 3 shows two examples: one handoff from 2G EDGE to 3G UMTS and the other handoff within 3G UMTS. Note that the network trace only indicates when a handoff completes. In fact, the handoff starts before that and its impact lasts longer. So we define the time window of an handoff as the one between the timestamp (t_{start}) that the last packet is received before handoff and the timestamp (t_{end}) that the first packet is received after handoff. Given the timestamp of handoff occurrence, we look for these two time (t_{start} and t_{end}) in its packet delivery log. This time window is also called handoff suspension time.

For the remaining time, we use RSSI thresholds to distinguish non-handoff events: NH/SC, NH/WC and NH/NC, defined in Table 4. The handoff classification will be described in Section 6.1.

A1. Handoff type detection: Broadly, handoff can be classified into two categories: intra-system handoff and inter-system handoff. When an intra-system handoff finishes, the mobile device still uses the same RAN and CN, but different base stations. In contrast, the mobile device switches to different RAN and CN after an inter-system handoff is completed. Say, one network trace has $TYPE_X$ and CID_X at the x -th timestamp. Therefore,

$$\begin{aligned} HO &\Leftarrow TYPE_X \neq TYPE_{X+1} \mid CID_X \neq CID_{X+1}, \\ \text{Intra-HO} &\Leftarrow HO \& TYPE_X \equiv TYPE_{X+1}, \\ \text{Inter-HO} &\Leftarrow HO \& TYPE_X \neq TYPE_{X+1}. \end{aligned}$$

The equivalence of two network types will be extended in Section 6. For instance, Table 3 shows two examples: one is an inter-system handoff from 2G EDGE to 3G UMTS and the other is an intra-system handoff within 3G UMTS.

B. Gap Calculation: Now the key question is to learn how much packets are not received in each specific time window. These packet loss finally contribute to the accounting gap as discussed before. The challenge we have is that packet delivery is logged in an event-driven way. Constrained by this, we can only induce the packet loss in the time period starting and ending with packet reception. As shown in Figure 14 (above), we receive packet S at T_1 and packet $S + 100$ at T_2 . We can induce the packet loss in $[T_1, T_2]$, but not in a smaller time window part of $[T_1, T_2]$. We will address it in the following mapping (Step C).

Now we investigate how to calculate the gap from the packet delivery log. Say, the phone receives a packet with sequence number S_X at time T_X . If there is no packet loss between two continuous timestamps, that is, $S_{X+1} - S_X \equiv 1$, the gap in $[T_X, T_{X+1}] = 0$. In other words, the gap in $[T_X, T_{X+1}]$ is calculated as $S_{X+1} - S_X - 1$. For example, the packet loss in the above plot of Figure 14 is 99 in $[T_1, T_2]$. But the actual packet delivery can be out-of-order. As shown in the bottom plot, Packets 95 and 99 may arrive after packet 100. To handle packet out-of-delivery problem, we update the highest sequence number $S_{highest}$ in the process and look forward to check if any packets have been received latter. We only calculate the volume gap if the sequence number is larger than $S_{highest}$. For those packets that will be received later, we exclude them in the accounting gap. Take the example in Figure 14 (below), we receive Packets 1, 100, 95, 99, 101 and 105 in order. When Packet 100 is examined, we retrieve the accounting gap in $[T_1, T_2]$ as 97 packets, i.e., Packets 2-94 and 96-98. We then set $S_{highest}$ to 100. Thus, we will not calculate accounting gap when we examine Packets 95 and 99. Note that the packet log might fail to record the last packet reception, for example, Route 12 ends in no coverage zone and phones do not receive any packet at the end. In this case, we determine the last accounting gap using the V_{OP} and V_{UE} .

C. Event and Gap Mapping: We now need to map the accounting gap within each event. By compare each gap time period and event time period, we are able to find out which event happens during gap time period. In most cases, the gap time window is much smaller than the event one, and we will sum up all the gaps for one single event. In case more than one event crosses the same gap time window, e.g., NC and SC, we treat this accounting gap into the category NC+SC, instead of NC or SC. Our measurement shows that less than 60KB accounting gap are contributed by multiple events, no more than 0.3% of the total gap. Due to space limit, we do not discuss the accounting gap caused by multiple events in this paper.