Understanding and Diagnosing Real-World Femtocell Performance Problems

Chunyi Peng*, Yuanjie Li†, Zhuoran Li*, Jie Zhao*, and Jiaqi Xu*
*Dept. of Computer Science Engineering, The Ohio State University, Columbus, Ohio, USA
†Dept. of Computer Science, University of California, Los Angeles, Los Angeles, California, USA
Email: chunyi@cse.ohio-state.edu, yuanjie.li@cs.ucla.edu, {li.5233, zhao.1359, xu.1629}@osu.edu

***Abstract*—Femtocells (small cells) augment the current mobile network by providing users short-range radio access at home and small-business settings. They have rapidly emerged as a promising scheme to alleviate capacity and coverage shortage by offloading traffic from the conventional Macrocells (large cells). Despite its increasing popularity, the real-world Femtocell performance has remained largely unexplored. In this paper, we conduct an in-depth study to assess Femtocell performance and diagnose identified issues in operational carrier networks. We focus on user-deployed Femtocells in a top-tier US mobile network. While the Femtocell generally works well, unanticipated performance degradations and even failures still occur. Contrary to conventional wisdom in the research community, we find that, *radio* link quality and interference is not the main bottleneck of Femtocells in many real-life usage scenarios. For instance, while Femtocell deployment at blind-zones with no radio coverage is desirable, not all deployments have succeeded; Compared with their Macrocell counterparts, Femtocells exhibit lower speed and larger speed variations, and induce larger delay for data services. Moreover, mobility support for femtocells is incomplete and no seamless migration is available under certain usage scenarios. We pinpoint their root causes, quantify the potential impacts, and share the learned lessons.**

I. Introduction

Femtocells are small, low-power base stations, typically designed for home or enterprise usage. They provide short-range radio access over the licensed frequency band, augmenting the conventional Macrocell[1]-based network infrastructure. They are plug-and-play, and allow for ad hoc deployment by home or business owners at areas of their interest, particularly at hotspots with high user demands or at coverage edges where performance significantly degrades. Figure 1 gives an illustrative scenario in a home setting with poor Macrocell coverage. Traffic is offloaded from Macrocells to Femtocells, thereby greatly boosting coverage and capacity.

Femtocells offer a cost-effective approach to improving spatial spectrum efficiency without upgrading the carrier network infrastructure. They leverage the available home/enterprise connectivity (wired Internet access) and supplements the existing radio access. Compared with WiFi and other wireless technologies, Femtocells work as an integral part of the cellular infrastructure and enable better data/voice services. Given such appealing features, Femtocells are deemed as an essential technique in Heterogeneous networks (HetNets) to tackle capacity crisis in the current cellular network and the upcoming 5G [1]–[4]. In reality, Femtocells have been widely deployed worldwide. In 2014, 5.7 million Femtocells were shipped, serving more than 47 global operators [5]. It is projected to reach $3 billion market with about 40% annual growth by 2020 [6]. In the US, Femtocells are even denser than Macrocells in some areas [7].

[1]We use Macrocells to denote macro-cells, micro-cells, and pico-cells together, which are deployed by the operators. They all have larger coverage areas than Femtocells.

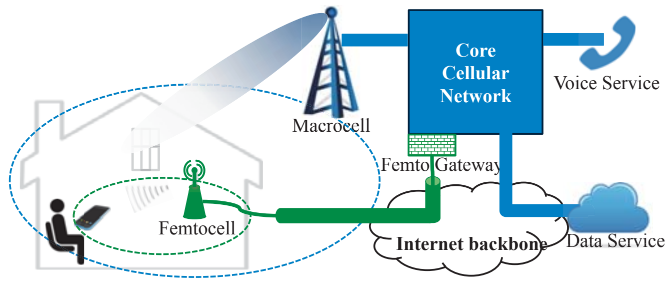

Figure 1: Femtocells augment the Macrocell-based mobile network to boost coverage and capacity. Example: the user deploys a femtocell at home with poor Macrocell coverage.

Femtocells have been actively studied by the research communtiy. Under the premise that Femtocells are operated over the same frequency as Macrocells, current research focuses on radio interference management, including co-channel interference alleviation [8]–[10], coordinated radio resource allocation [11]–[14], or radio coverage/capacity optimization [15]–[17]. Other efforts look into the key Femtocell operations, such as handoff [18]–[20], access control [21], and data forwarding [22]. All these studies either assume an ideal Macrocell-Femtocell framework (say, two-tier Macrocell-Femtocell architecture) for theoretical analysis or work with special Femtocells in the testbed settings. They focus on the Femtocell's radio access and idealized design. The issue of the Femtocell performance in reality is largely unaddressed.

In this work, we conduct an empirical study to assess and diagnose the Femtocell performance in real-life usage settings. We seek to answer three questions. *(1) How well do Femtocells perform in operational networks? Have they achieved their design goals of boosting coverage and capacity? (2) Are there performance glitches or even failures, largely unanticipated by the research community, observed in practice? What are their consequences and root causes? (3) What lessons can we learn* *to improve Femtocell design and operations?*

We conduct a 6-month measurement with a top-tier US carrier network in two metropolitan areas on the east and west coasts. Our study shows that Femtocells are largely successful to improve user experience in indoor settings, particularly for those with weak/no radio coverage. However, we uncover that the current operation and practice of Femtocells are not without problems. In certain common settings, users do suffer from call drops, data service disruption, unstable connectivity, and failure to camp on Femtocells, to name a few. We infer their root causes. To our surprise, the results show that radio interference is no longer the main performance bottleneck in practice. Instead, inappropriate forwarding architecture and control-plane functions for Femtocells should bear the main responsibility. The current practice largely inherits the legacy design for Macrocells (except its mobility support), but they do not work well in the Femtocell context. In addition, missing design components and inappropriate configurations contribute to other performance problems.

Our main contribution can be summarized as follows.

1) We conduct an in-depth study to examine the Femtocell performance in practice. To the best of our knowledge, this offers the first empirical assessment on commodity Femtocells in operational mobile networks.
2) We further identify performance glitches and failures in common settings. We pinpoint their root causes that have not been reported before.
3) We share the lessons and insights to design better Femtocells, and their implications on the emerging HetNets.

The structure of the rest paper is as follows. Section II introduces background on network architecture and procedures related to Femtocells. Section III gives an overall view of the Femtocell problems, followed by the empirical study of Femtocell deployment, performance and mobility support in Sections IV, V and VI. Section VII concludes our work.

II. Femtocell Primer

Network architecture. Figure 1 depicts the cellular network architecture using Macrocells and/or Femtocells. The network infrastructure is divided into two parts: the radio access subsystem and the core network. In the former, Macrocells (Femtocells) are deployed to offer wireless access to mobile devices within relatively large (small) geographic coverage area. The core network serves as a backhaul to the Internet and the telephony network.

Compared with public Macrocells, Femtocells are typically private and deployed by home/business owners. Hence, they do not have dedicated links to the core network. They have to traverse the public IP network (Internet backbone), which is usually outside the mobile carrier network. In order to integrate Femtocells as a coherent part of the mobile network infrastructure, an additional Femtocell gateway is launched as the portal to the core network. Moreover, a secure tunneling (via IPSec) is adopted to protect the communication between the Femtocell and the Femtocell gateway. This way, Femtocells

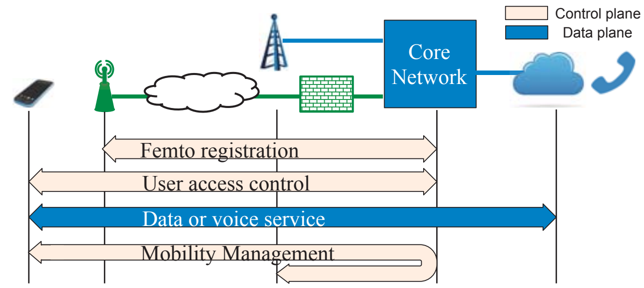

Figure 2: Main operations for Femtocells.

can be treated like Macrocells (except with smaller coverage) by simply offering alternative radio access in vicinity.

Main Procedures for Femtocells. Figure 2 describes main operation flows in the data plane and control plane. Basically, Femtocells follow the Macrocell's approach to provide network services. Inheriting the legacy design seems to be cost effective. They deliver user traffic (*e.g.*, voice or data) in the air and forward it to the core network through secure pipes over the IP backbone. Afterwards, the legacy core network mechanisms take over the rest delivery to the external Internet or the telephony network.

To facilitate traffic delivery on the data plane, the control plane provides a variety of essential signaling functions including Femtocell registration, user access control and mobility support. Femtocell registration is the first element at start up when the Femtocell is being installed. Once the customer plugs the Femtocell into the Internet and powers it on, the Femtocell first needs to register within the network. It reports its unique femtocell identifier, as well as the address where the Femtocell is deployed to the core network. With the femtocell successfully registered, femtocell configuration will be undertaken, such as checking and downloading the latest software, configuring radio parameters.

User access control is to determine whether the user is legitimate to use mobile network service. In the conventional Macrocell context, the device must first register itself with the core network. The core network checks its subscription information and grants the registered user appropriate network access. In the Femtocell context, access control is still performed in the core network, though Femtocells are typically open to a closed subscriber group (CSG), not like Macrocells open to all users.

Mobility support is a salient feature in cellular networks. It switches the serving cell to ensure ubiquitous radio access. This is realized through two procedures: handoff (with ongoing traffic) and cell re-selection (without traffic). Upon a handoff, the current serving cell negotiates with the target cell with the assistance of the core network. Afterwards, the target cell takes over the radio access, and the resulting seamless migration ensures minimal voice/data interruption. The cell re-selection procedure is initiated and performed by the device through directly choosing another serving cell. To learn whether to switch and which cell to switch to, measurements on neighboring cells (on their radio signal quality) are carried out at

Category	Type	Findings	
Deployment (§IV)	☺	**F1**: Femtocell deployment is desirable even only in order to enhance coverage.	Figure 3
	☹	**F2**: Not all Femtocell deployment can succeed due to GPS dependency.	
Data/Voice Performance (§V)	☺	**F3**: Femtocells boost data speed at `weak/dead`-zones, and even at `strong/OK` zones during busy hours.	Figure 4a,4b,4c
	☹	**F4**: Femtocells experience lower-than-expected and less reliable data speed.	Figure 4d
	☹	**F5**: Femtocells incur larger latency for data services in almost all the settings.	Figure 5, Table II
	☺	**F6**: Femtocells set up voice calls more quickly and successfully at weak/no coverage spots.	§V, Fig. 6
Mobility (§VI)	☹	**F7**: Mobility support for data in Femto→Macro is missing (except voice).	Figure 8,9
	☹	**F8**: Persistent Femtocell oscillations exist in real usage scenarios.	Figure 8
	☹	**F9**: Mobility support for voice and data in Macro→Femto and Femto→Femto cases is missing.	Table IV

Table I: Summary of main findings on real-world performance and problems. ☺: positive; ☹: negative.

the devices in both procedures. Once the serving cell switch succeeds, the device performs location update to report its new serving cell to the core network. In the Femtocell context, the same seamless mobility support requirement exist but in three distinct scenarios: from Macrocell to Femtocell (M→F), from Femtocell to Macrocell (F→M), and From Femtocell to Femtocell (F→F). The first two supports (M→F and F→M) are illustrated in Figure 2.

Related work. As described in Section I, most research efforts work on radio resource management [8]–[17] and few work on handoff [18]–[20], access control [21], and data forwarding [22]. All the above studies either work with theoretical modeling or special Femtocell equipments. Different from their work, we seek to investigate Femtocell performance in practice. We aim to identify possible problems and their root causes, and gain lessons to design and build better Femtocell-Macrocell networking architecture and operations.

III. Reality Check on Femtocell Performance

To understand usability and performance of Femtocells in reality, we conduct an empirical study of user-deployed Femtocells in the wild. We are motivated by three questions: (1) Are there any real demands to deploy Femtocells, given the fact that most areas may have been well covered by Macrocells? (2) Once deployed, do current Femtocells improve service performance, while replacing a remote Macrocell? (3) How does their mobility support perform? Will the users obtain seamless voice and data as they were in the Macrocell-based mobile network?

To answer these questions, we construct Femtocell cases in two US cities – Los Angeles on the west and Columbus, OH on the midwest – over a top-tier, US mobile carrier network. We use two commodity Femtocell models (2013 and 2014 versions) provided by the carrier. We use 3G Femtocells, since 4G Femtocells [23] are not supported by the US carriers yet[2]. We run experiments during two periods: from 11/01/2014 to 03/10/2015, and from 06/10/2015 - 07/20/2015. No obvious distinctions are observed over these two periods. Table I summarizes our major findings.

We next elaborate on each finding. We describe the detailed experiment settings and present measurement results. We validate benefits from Femtocell and also disclose problematic operations and their negative impacts, and discuss their root causes and lessons.

[2]For example, Verizon plans to release LTE Femtocells in 2015 [24], which are still currently unavailable.

IV. Should Femtocells be Deployed?

Our study yields a positive answer to this question. We measure Macrocell radio coverage in indoor settings and check whether there exist indoor zones that are not well covered.

Experiment settings. our measurement covers 21 home and campus buildings, including ten office buildings (eight lab buildings and two libraries), eight apartment buildings, one house, and two public buildings (the gym and the shopping center). All have 2–8 floors and about 60% of buildings have the basement floor. We use `OpenSignal` [25], a popular app to collect cellular network status and radio quality, including the serving network type, cell identity and serving radio signal strength. To measure co-located 4G and 3G Macrocell radio quality, we use multiple phones (each for 3G or 4G) simultaneously, while walking along some routes at a relative slow speed (about 0.5m/sec). We use several Android phone models: HTC one, Nexus 5 and Samsung Galaxy S3, S4 and Note 3. The results are similar across all the phone models.

F1. Coverage Enhancement. Figure 3a gives a snapshot of 3G and 4G radio quality per second while walking in three lab buildings. Since 3G and 4G use different physical-layer technologies, they adopt different radio quality metrics: Received Signal Strength Indicator (RSSI) for 3G and Reference Signal Received Power (RSRP) for 4G. The former is in a range of [-113, -70] dBm (out of coverage if -113dBm) and the latter is in [-135, -75] dBm and there is almost no coverage when the value is below -120 dBm. The out-of-coverage zones (also called dead zones) are marked with pink strips in Figure 3a. Figure 3b plots the cumulative distribution function (CDF) of radio quality measured in all the test buildings. It implies that while indoor coverage is largely acceptable, there indeed exists some coverage holes. 11.3% of test locations suffer from no 3G coverage while 6.1% have no 4G coverage.

Consider most phones support both 3G and 4G, the devices can automatically switch to another better technology if the currently serving one is too weak. As a result, the serving Macrocell at one location depends both co-located 3G and 4G radio signal strength. We thus define weak and strong thresholds and classify all the test spots into the following four coverage types: (1) `Strong`-zone (at least one strong): 3G, 4G or both have their measured radio quality better than their strong thresholds (*3G: -85dBm* and *4G: -95dBm*); (2) `OK`-zone

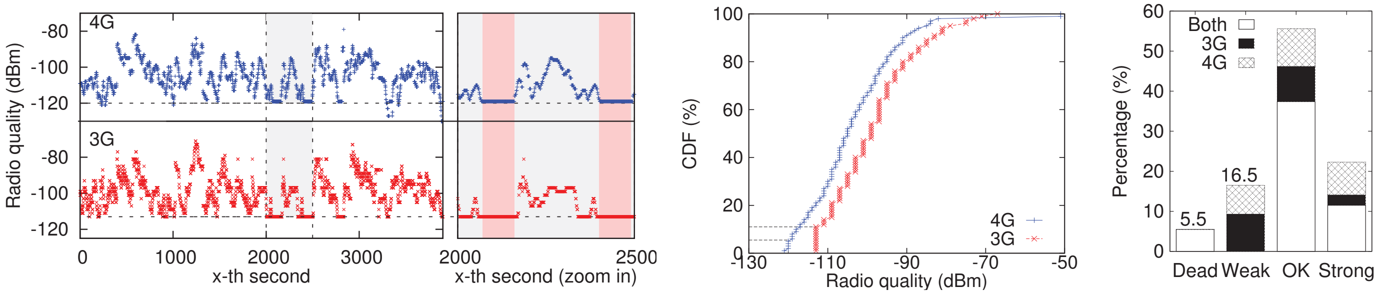

(a) 4G and 3G radio quality in three buildings

(b) CDF of 4G/3G radio quality

(c) Statistics

Figure 3: Macrocell radio quality in the indoor testbeds.

(at least one OK): at least one has its radio quality larger than its weak threshold (*3G: -105dBm and 4G: -110dBm*), but none is larger than its strong threshold; (3) `Weak`-zone (both weak): none of their radio quality is larger than their weak thresholds but is covered by at least one technology; and (4) `Dead`-zone (uncovered in both): both has no coverage. Thresholds are determined by our experience and match with a large amount of prior measurement [25]. Roughly, they correspond to different signal bars on the phone: 0 (`dead`), 1 (`weak`), 2-3 (`OK`), 3-4 (`strong`). Figure 3c plots the histogram of four coverage types at all the test spots. Clearly, it is seen that while the indoor coverage is largely acceptable, a portion of places (22% in our tests) suffer from `dead`/`weak` coverage: `dead`: 5.5%, `weak` (weaker 3G): 9.7%, `weak` (weaker 4G): 7.2%.

One thing worth noting is that the carrier should not take much blame for such indoor coverage. Radio waves experience severe multi-path fading and shadowing effects caused by complicated indoor environments and obstacles; The radio quality usually fades around the walls, elevator shafts, corners and staircases, in the ground and basement floor, *etc.*. This is extremely difficult, if not impossible, to guarantee full coverage through a *coarse-grained* Macrocell deployment. This does motivate a *finer-grained* Femtocell deployment in these areas, even solely for enhancing or extending coverage. We next show that Femtocells are also desirable to improve service performance, particularly during traffic congestion.

F2. Deployment failure. To our surprise, although Femtocells are highly desirable at `dead`/`weak` coverage spots, we indeed observe that not all Femtocell deployment can succeed at these places. The problem is rooted in the current Femtocell registration practice which reports its address solely depending on GPS signals, which are not available at all the spots. In one testbed with 20 `dead`/`weak` coverage spots, we find that 6 spots never obtain GPS information, and 2 spots occasionally receive valid GPS information. Registration fails in 30-40% cases with `dead`/`weak` coverage and succeeds at the all spots with `OK`/`strong` coverage. This is not hard to understand. GPS information is usually unavailable at certain complicated indoor environment, such as at the basement floor or around the corner/obstacles. This coincides with the `dead`/`weak` radio coverage. Ironically, such failures are exactly against one main purpose to boost coverage through deploying Femtocells.

However, it is unnecessary to count on GPS only in Femtocell registration. This is an operational issue. It is reasonable for the core cellular network to learn the locations of Femtocells which are deployed in an ad hoc manner. Otherwise, it may expose the existing infrastructure to unknown but significant interference. However, GPS should not be the only choice. Alternative options should be provided, especially based on Macrocells nearby. A Femtocell should be able to sense its neighboring Macrocells and learn its approximate location. This option instinctively matches with the actual need: Femtocell's relative position to Macrocells, rather than its absolute location, is more critical to manage Femtocells. In the worst case (*e.g.*, basement) without detecting any Macrocell nearby, Femtocells should be safe to deploy since it does no harm to the existing infrastructure. In the following experiments, we only consider the successful cases.

V. Can Femtocells Boost Service Performance?

We compare data and voice performance using Femtocells and 4G/3G Macrocells. We focus on data because voice requires less resource and there is no significant improvement.

Experiment settings. We deploy Femtocells at places with various radio `strong/OK/weak/dead` coverage. We run `Speedtest` [26] to collect downlink and uplink speed, as well as ping latency, at different hours of a day (8AM-22PM) while using 3G/4G Macrocells. We compare the results at same spots when 3G Femtocells in vicinity are in use. The Femtocell performance is independent of Macrocell coverage, and thus all the records are merged. We do not distinguish the results when Femtocells are deployed at campus and at home with different IP backhaul. The former yields a higher-speed access as large as 80–100Mbps, while the one in residence usually supports 10–25Mbps (both from `speedtest`). Compared with the Femtocell speed, the IP backhaul is not the bottleneck at most time and no clear distinction at campus/home is observed. For voice, we develop an autodialer app to measure voice performance.

F3. Data service speed enhancement. Figure 4 shows the boxplot (5th, 25th, 50th, 75th and 95th percentiles) of 4G/3G

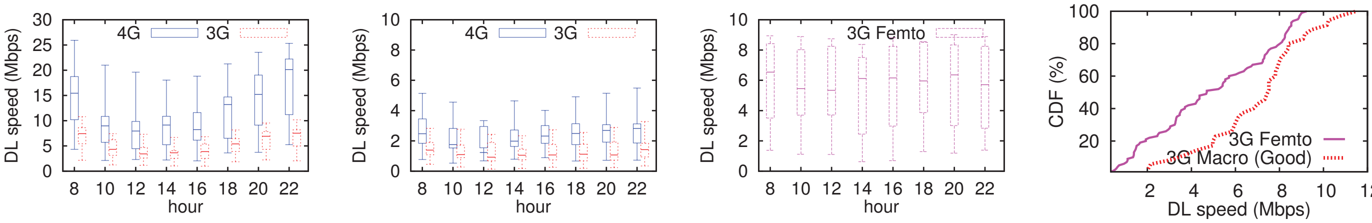

(a) Macrocell (`Strong`) (b) Macrocell (`Weak`) (c) 3G Femtocell (d) CDF of 3G downlink speed

Figure 4: Downlink data speed using 4G Macrocells, 3G Macrocells and 3G Femtocells in different scenarios. (a)(b)(c) show the boxplot of 5th, 25th, 50th, 75th, 95th percentiles at different hours of a day; (d) compares the CDF of downlink speed using 3G Femtocells and 3G Macrocells (8AM, 20-22PM) in `Good` coverage.

downlink speed with `strong` and `weak` radio coverage[3]. We omit the results for `OK` and `dead` coverage because the performance for `OK`-coverage is in between and no tests succeed without coverage. Clearly, we observe that Femtocells are able to boost throughput not only at `weak/dead` coverage areas, but also at `strong` coverage areas during congestion hours. First, the Macrocell speed is relative small (4G: <3Mbps, 3G: <1.5Mbps) due to weak radio quality. The median values are used if unspecified. Once Femtocells are in use, the data speed can significantly jump to 5.3 – 6.5Mbps, with about 100–234% gain over 4G and 300 – 490% gain over 3G. Note that the gain could be even larger if Femtocells support 4G. Second, Femtocells help even when Macrocells have `strong`-coverage. Though the peak speed via Macrocells is higher (4G: 15 – 20Mbps, 3G: 6.9 – 7.5 Mbps), the speed shrinks by 50-60% (4G: 7.8-9.1Mbps, 3G: 3.4–3.6Mbps) at busy hours (12PM-16PM). It is caused by heavy traffic load and huge dynamics. Femtocells can boost 3G throughput (65-90% growth) in this case. We do not observe the gain over 4G Macrocells, because 3G is relatively slower than 4G. It could be beneficial if 4G Femtocells are in use.

F4. Lower-than-expected and less reliable speed. To our surprise, we also observe that Femtocells do not fulfill its performance potential in replace of 3G Macrocells. In our tests, Femtocells yield lower peak speed than 3G Macrocells. Even using the same 3G technology, Femtocells support up to 9.2Mbps, smaller than the maximum speed observed in 3G Macrocells (11.5 Mpbs). Moreover, actual data speed vibrates more intensively in a larger scope (see 25th, 75th percentiles). We compare the CDF of downlink speed using 3G Femtocells and 3G Macrocells in good settings (with `Good` coverage during non-busy hours, around 8AM, 20PM-22PM) in Figure 4d. The speed via Femtocells is much lower than expected. Data speed using 3G Macrocells is larger than 6Mbps in 72% cases, whereas the speed via Femtocells is smaller than 6Mbps in 60% cases and is even smaller than 2Mbps in 21% tests.

We look into the problem and find that it is not due to radio interference. Femtocells are operated under a different frequency band from 3G Macrocells; We run spectrum sensing using a software-defined radio platform (USRP N210 + SBX [27]) and validates that no strong co-channel radio signal exists. In fact, we find that data speed is throttled due to the current Femtocell architecture. Data has to be redirected to the core network via the Internet before eventually being delivered to the destination. Without leveraging its direct path to the Internet, the current Femtocell practice offers lower-than-expected speed with extra overhead toward and at the core. This is also clearly observed in the next latency results.

For uplink speed, we have similar findings but the values are much smaller. For example, 4G LTE support at most 7Mbps for uplink while up to 25Mbps for downlink. 3G usually support up to 11Mbps for downlink and 2Mbps (mostly <1.5Mbps) for uplink.

F5. Larger latency in Femtocells. We then move to another data service performance metric: latency. Figure 5 shows the ping latency collected by `Speedtest` using 3G/4G Macrocells or Femtocells. It is consistently observed that

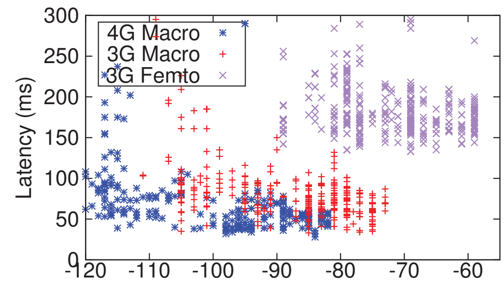

Figure 5: Radio quality (RSSI/RSRP) versus ping latency collected by SpeedTest via 4G Macrocell, 3G Macrocells and 3G Femtocells. Note that Femtocells have relatively larger RSSI in vicinity.

	Latency via SpeedTest			First-hop latency		
	Med	**Min**	**95th**	**Med**	**Min**	**95th**
4G Macro (`Strong`)	53	32	77	48	30	69
(`Weak`)	76	42	176	62	37	148
3G Macro (`Strong`)	69	36	94	59	31	81
(`Weak`)	87	38	181	79	37	153
3G Femto	176	132	246	136	94	208
IP backhaul (direct)	17.7	15.2	25.9	-	-	-

Table II: Comparison of data service latency (ms) using 4G Macrocells, 3G Macrocells and 3G Femtocells.

[3]Speedtest may not work at `weak` coverage sometimes; We take the successful tests into account, slightly inflating the achievable speed.

Femtocells induce larger delay almost in all the cases, regardless of its stronger radio quality. The median value via Femtocells is around 176ms, incurring extra 90-120ms delay compared with Macrocells (4G `strong/weak`: 53/76 ms, 3G `strong/weak`: 69/87 ms). Table II compares their median, minimal and 95th percentiles. We use the 95th percentile because it is inevitable to suffer a large latency in some rare cases due to networking environment dynamics. For example, $>$ 180ms is rare but still possible in 4G/3G Macrocells with extremely weak coverage.

To deduce its root cause, we analyze its latency breakdown and retrieve the first-hop latency via `traceroute`. We find that 4G/3G Macrocells and 3G Femtocells all have the same or similar (within the same pool, here, 172.26.196.xxx) first-hop destination. This is likely the gateway in the core network, which is responsible for delivering data packets to the external Internet. We then `ping` it and collect the first-hop delay in Table II. Clearly, additional latency via Femtocells is mainly attributed to the first hop. It implies that the current delivery path to the core network via IP backhaul is the main bottleneck. Without direct access on the data plane, offloading to Femtocells suffer from the traversal of the core network. We run ping tests using WiFi in order to emulate the latency via the direct path without traversing the core network. Latency can greatly shrinks to 15 – 25ms.

Lessons: In a nutshell, Femtocells indeed boost data speed in some cases (at `weak/dead`-zones, during congestion hours), but at the cost of a higher latency. It might hurt real-time data services or those with stringent timing requirements such as gaming and tactile Internet. The current offloading practice traverses the core network via Internet and thus induces unnecessary performance degradation and dynamics (throughput slump and extra latency). This calls for a new offloading scheme which makes full use of direct path to fulfill the Femtocell potential. In fact, two schemes have been proposed for Femtocells: local IP access (LIPA) and selected IP traffic offload (SIPTO) [28], [29]. Both allow traffic offloading away from the core network. However, LIPA is primarily for end devices to access their local network. For example, when a user has a femtocell at home or in the office, mobile devices can use LIPA to access other devices connected to the local home/office network such as TVs, computers and office servers. SIPTO allows selected IP traffic (*e.g.*, the one via Femtocells) to bypass the core network completely; However, it still requires the core network to control and determine what traffic can be bypassed. Both schemes have problems in mobility support and charging issues (bypassing traffic will not be charged) [30]. As a result, though both have been proposed around for years, they are not available in the market because vendors and operators do not buy into them.

F6. Limited voice improvement. We also compare voice performance using Macrocells and Femtocells. We develop an automatic dialer program to run two cases: call setup (dial and hang up once ringing) and 1-minute voice calls. We use the call setup time and the call drop rate as two main quality metrics.

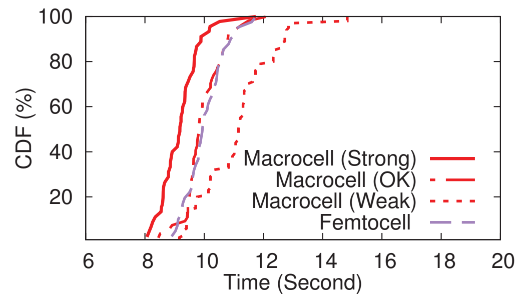

Figure 6: Dial-to-ring delay using Femtocells and Macrocells at `strong` and `weak` zones.

Our study shows that Femtocells can help to establish voice calls quickly and successfully at `weak/dead` zones but there is no significant improvement observed for alleviating call drops in our test cases. To assess call setup time, we measure dial-to-ring delay, an interval between dialing the number and receiving ringback, defined in the ITU standard [31]. Figure 6 plots their CDFs using Macrocells (at `strong/OK/weak` zones) and Femtocells. Note that 4G LTE users adopt circuit-switching Fallback (CSFB) to make voice calls [32] and thus all calls are still established via 3G circuit-switching technique. Femtocells reduce the call setup time at the `weak` coverage, from 10–12 seconds to 9-10 seconds. Moreover, we observe that not all the calls will not be established through at `weak` zone (certainly none at dead-zones) using Macrocells. While the ringback is heard at the caller, nothing rings at the callee. In our tests, each dialing lasts for 45 seconds and we find that 58 out of 150 runs (38.7%) at `weak` zones failed to reach the callee, whereas all the calls succeeded at `strong/OK` coverage. Regarding call drops, we did not observe call drops in our experiments even in the `weak` coverage, as long as the calls were established. At hence, there is no reduction on call drops in the tests.

There is no surprise to see little improvement on call setup and almost no help on call drops through Femtocells. First, the voice service has already adopts quality control and consumes much less resource (several or tens of kpbs). Once the call is established, resource is reserved and thus it is much less sensitive to radio quality and ongoing traffic, unless Femtocells are deployed at Macrocell's `dead`-zones. Second, call setup requires signaling exchange which is affected by radio quality. At `weak`-zones, it likely requires a longer setup time with a higher error rate in voice call signaling exchange.

VI. Do Femtocells Support Seamless Mobility?

We now examine mobility support for Femtocells. This is vital to ubiquitous services (*e.g.*, data and voice) in cellular networks, especially when Femtocell coverage is pretty small (in the range of 10–20 m). We consider slow mobility (in-building walking) only. When the user walks into/outside the Femtocell coverage, seamless service migration without disruption is highly desirable. This requirement is identical to that in the Macrocell-based network.

	Idle	Data	Voice	Data + Voice
F→M	√	×	√	√
M→F	√	×	×	×
F→F	√	×	×	×

Table III: Mobility support for Femtocells in operation. √: supported; ×: not supported. Data or voice get hurts in all × cases.

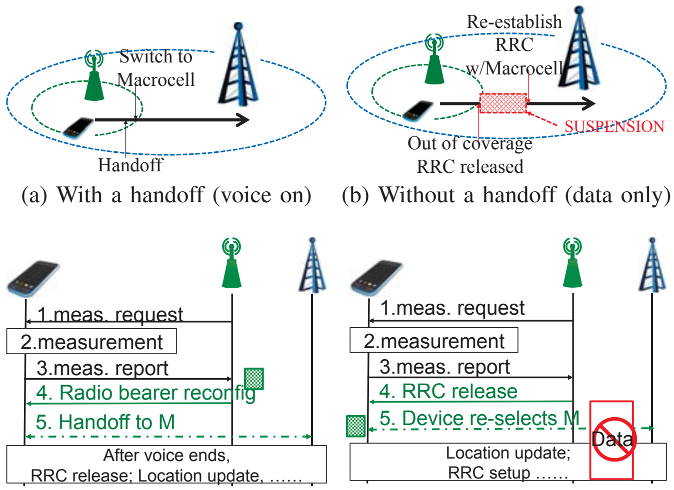

(a) With a handoff (voice on) (b) Without a handoff (data only)

(c) RRC events (voice on) (d) Data only

Figure 7: Events during F→M with/without seamless mobility support (a handoff).

However, our study finds that the current support for Femtocells is largely incomplete. Specifically, we discover that no handoff for data is provided in F→M mobility, whereas no handoff for any ongoing service (voice or data or both) is provided to migrate from an Macrocell/Femtocell to another Femtocell. Consequently, voice calls drop, and data sessions freeze or even abort. Table III summarizes the current practice of mobility support in three concrete mobility scenarios: F→M, M→F and F→F. Four traffic patterns are enumerated: idle (no traffic), data only, voice only, data+voice.

F7. F→M mobility support is missing (except voice). Our study shows that F→M handoff for voice is enabled so that the voice call will be migrated to the Macrocell to ensure call continuity but data not. Figure 7 illustrate the expected and actual events when a handoff is provided or not. With a handoff (left), the device is able to switch to the Macrocell before it releases its radio connection with the Femtocell (moving out of the edge of Femtocell coverage). In contrast, without a handoff (right), the device has to release its radio connection and thus data service gets disrupted. After a while (several seconds to tens of seconds), the device seeks for an available cell (here, Macrocell) and re-establish its radio connection. If such suspension is tolerable by the high-layer protocols (say, TCP or apps), the data service freezes otherwise it directly aborts. During this process, data throughput greatly shrinks and latency soars.

Figure 7d and Figure 7c reveal how to migrate to a Macrocell with voice (with a handoff) and data (without a handoff). The data-plus-voice has the same procedure as the voice-only one. In fact, both have the same measurement procedure triggered by mobility (Steps 1-3). When the device is about to leave (the Femtocell's radio strength smaller than some threshold), the serving Femtocell first asks the device to scan other suitable cells in its vicinity and learn the best candidate cell (here, a Macrocell). Afterwards, the device reports its measurement to the Femtocell. The distinction between voice and data occurs at Steps 4 and 5. To support ongoing voice or data, there exist an active radio bearer (aka to a link-layer connection) within cellular networks. For voice, the Femtocell then cooperates with the target Macrocell via the core network to invoke seamless roaming. In particular, it performs `radio bearer reconfiguration` and lets the device handoff to the Macrocell without disrupting the call. For data, instead, it releases the radio link connection (`RRC Release` in Figure 7d) and forces the device to go idle. Afterwards, the device itself scans the neighboring cells and re-selects the Macrocell. Note that, Step 5 for data starts a new connection, independent of the previous one. After switching to a new cell, the device has to notify the network (here, through `location update`) in both scenarios. The above signaling exchange are mandated by Radio Resource Control (RRC) [33], [34] and Mobility Management [35], specified in the 3GPP standard. The former is responsible for radio connection establishment and release (including radio bearer reconfiguration), while the latter is responsible for mobility-relevant control functions, such as handoff, cell re-selection and location update. For data, data packet delivery can be resumed until the the RRC establishment and location update complete after switching to a new cell. Clearly, data packet delivery suspends.

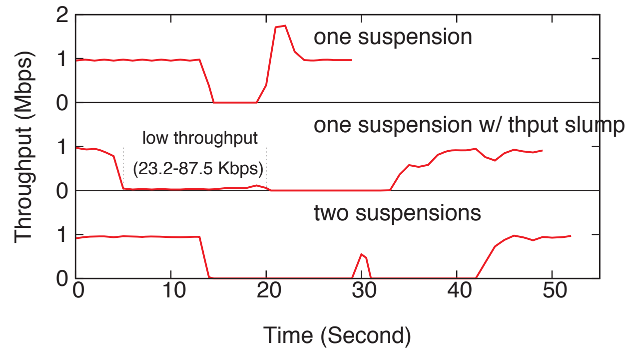

Figure 8: Three example traces in F→M experiments.

We run data services in F→M walking experiments. We first use constant-rate (say, 1Mbps) UDP traffic. Figure 8 gives three illustrative traces which represent three symptoms observed in our experiments: one suspension (top), one suspension with additional throughput slump ahead (middle) and two or more suspensions (bottom). First, data freezes when the radio bearer is down; at least one suspension is observed in all the tests. Second, prior to data suspension, data throughput may quickly shrinks to a small value (in this example, 23-87Kbps) due to lousy radio connection before it is released, in most tests. Third, two or more suspensions are observed in some tests. This is caused by RRC oscillation (explained in F8) in the current mobility support. Suspensions can happen many times at some place, *e.g.*, at the Femtocell edge with weak Macrocell coverage.

Figure 9 shows the CDF of the impact time, excluding the persistent-loop case (data always suspends). Overall, the suspension time lasts 2.6 – 45 seconds while the throughput slump may last up to 76 seconds. In addition to expected suspension caused by the RRC connection release-setup, we also observe the suspension contributed by location update (1.2–6.4 seconds). While it is ready for data transfer with an established radio connection in the data plane, the device has to wait for the signal of mobility management from the control-plane. We also test with popular apps: Webkit (web browsing), AndFtp (file downloading), Facebook, Skype and Youtube. Except Skype, all four use TCP. All respond slowly; Skype will abort when the suspension is too long and Youtube video streaming will freeze when the buffer runs out.

F8. Femtocell-Macrocell oscillation loop. We reveal one configuration problem in Femtocell-Macrocell mobility support: Femtocell oscillation loop will be persistently triggered in real-usage scenarios where its release threshold is not appropriately configured. We observe, as long as the Femtocell's signal strength is below its `RRC Release` threshold (−89 dBm, here) but still stronger than Macrocell, the data service is prone to the following loop: `RRC Release` → `re-selection` (Femtocell) → `RRC Setup` → `Data Transfer` → `RRC Release` ⋯. Each suspension lasts 3 seconds on average (in range of 2-6.7 seconds), followed by about 2.5-second delivery. As a result, throughput greatly shrinks and the data session almost freezes.

This loop is caused by two decisions involved. The first is to determine whether to release its connection with the Femtocell when being served by the Femtocell; The second is to determine whether to select the Femtocell when not being served by any cell. The criteria for both decisions are not coherent since they are designed for different purposes. Without a handoff, Femtocell releases its RRC connection as long as its radio quality degrades below certain threshold, regardless of how poor radio links the neighboring cells have. After the release, the device becomes idle and attempts to re-select the serving cell. However, it still chooses to stay in the Femtocell because its radio quality outperforms those neighboring cells. This causes unnecessary RRC release and setup and data suspension. We run a 2-hour experiment and this loop is repeatedly observed. Such mobility instability is caused by mobility management misconfiguration and more Macrocell cases are disclosed in our recent work [36].

F9. No M→F and F→F mobility support for voice and data. The mobility support for M→F and F→F is even worse. Handoff for both voice and data is missing. When moving from a Macrocell to a Femtocell, the device indeed initiates the measurement of neighboring cells. A list of other suitable nearby cells is configured by the Macrocell, and the device is instructed to measure their radio quality. However, the problem is that Femtocells are never on this list. Consequently, the device never monitors Femtocells. In F→F case, the original Femtocell even does nothing except watching their radio links broken. Without handoff to a Femtocell, voice/data service will be disrupted when the user moves outside the Macrocell (the original Femtocell) coverage. Note that this practice is against the original intention to boost the radio coverage via Femtocells in the area with weak/no Macrocell coverage.

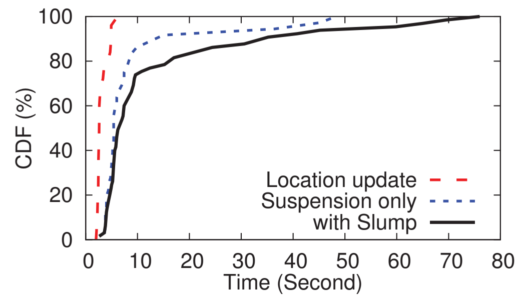

Figure 9: CDF of data suspension time during F→M.

	No	Weak	OK	Good
Voice drop (%)	100%	28.2%	0%	0%
Data suspension (second)	2.4 – 67	0 – 46.2	0	0

Table IV: Negative impacts on voice and data in M→ F.

We run M→F and F→F walking experiments. We observe that call drops and data suspends when the devices moves into no or extremely weak coverage, regardless of an available Femtocell in vicinity. In good/ok coverage, voice and data will be supported via the Macrocell (In F→F, it migrates to a Macrocell first). Only when the traffic ends and the radio connection is released, the device will choose the Femtocell over the Macrocell. Table IV shows call drops and data suspension with various Macrocell coverage. It matches with our previous measurement. The experimental results for popular apps are similar and omitted due to space limitation.

Lessons: The current practice does not offer appropriate mobility support from the network side, even when it is capable (at least handoff for data in the F→M case is possible, like voice support). The rationale possibly stems from different quality requirements on voice and data. Call is interactive and requires real-time continuity. In contrast, data are delivered in packets and can tolerate certain latency. Moreover, the transport and application layers provide additional mechanisms to resume data transmissions. The operator thus skips handoff for data to save network resource consumption. However, such a choice underestimates its negative impact on user experience and overlooks diverse requirements in data service (interactive and real-time applications).

Lacking M→F and F→F mobility support is a technical and operational problem. While inheriting the legacy Macrocell-Macrocell handoff mechanism, *all* Macrocells should be upgraded to take nearby Femtocells into account in their handoff strategy. To fully support ad-hoc Femtocell deployment, they thus need to be modified to sense their nearby Femtocells, or request Femtocell registration from the core network and keep it updated. Such upgrade incurs high operational cost. On the other hand, such new network infrastructure with ad-hoc deployment and diverse topology, warrants more research and engineering efforts in its mobility support.

VII. Conclusion

Femtocells are designated to augment the current 3G/4G cellular infrastructure. In this paper, we conduct reality check on their real-world performance. We confirm three baseline features for Femtocells, but also obtain three surprises. First, we confirm that Femtocell deployments are indeed justified and desirable. However, their deployment may not succeed due to location registration failures, which can be avoided. Second, the data and voice performances indeed improve with Femtocells, particularly in regions with weak or no coverage by Macrocells, or under cases with network congestion. However, given a Femtocell, its connectivity via the core network and its indirect data path limit the performance gain. As a result, the speed can be throttled and more delay is incurred. Third, mobility support is technically feasible, yet largely missing in practice. Moreover, mobility management has configuration issues, which expose design inconsistency between Femtocell and Macrocell handoffs.

Our study yields a few points we would like to share. First, the plug-and-play Femtocells offer a venue for adding a user-directed, ad-hoc network component into the existing, well-planned, well-designed cellular infrastructure. This nontrivial infrastructure shift should possibly call for new design approaches. As we have shown in this work, sticking to the legacy design may impede the performance gain, and even lead to errors or failures in the worst case. Second, network control and configuration would be equally important, together with data-plane issues (e.g., radio interference management, resource allocation). This is particularly important for the research community. Neglecting such design aspects may incur high price in performance penalty. Third, proper design and operations of control and management functions require concerted efforts at the end device, the network edge, and the core infrastructure. End devices alone may not deliver the desirable performance, e.g., in the mobility case. Finally, we believe our experiences on Femtocells may also shed lights on the ongoing research on HetNet, an important feature in LTE advanced and 5G. In a sense, Femtocells offer a simplistic case of HetNet. The lessons learned here may help to inspire more work to make HetNet a coherent component of future mobile networks.

Acknowledgments

This work was in part supported by NSF NeTS grants CNS-1421440, CNS-1423576, CNS-1526456, and CNS-1526985.

References

[1] Qualcomm. The 1000x data challenge. *https://www.qualcomm.com/invention/technologies/1000x*.

[2] NGMN. Ngmn 5g white paper. https://www.ngmn.org/work-programme/5g-initiative/.

[3] V. Chandrasekhar, J. G. Andrews, and A. Gatherer. Femtocell networks: a survey. *Communications Magazine, IEEE*, 46(9):59–67, 2008.

[4] H. Claussen, L. T. Ho, and L. G. Samuel. An overview of the femtocell concept. *Bell Labs Technical Journal*, 13(1):221–245, 2008.

[5] Global femtocell market to top $3b in 2020, April 2014. http://www.fiercewireless.com/tech/story/global-femtocell-market-top-3b-2020-report-says/2014-04-27.

[6] Global femtocell market (types, technology, applications and geography) - global analysis, trends & forecast through 2020, 2014.

[7] B. K. Horn. More femtocells than cell phone towers? *http://people.csail.mit.edu/bkph/other/Cellular/osm_blog*.

[8] V. Chandrasekhar and J. G. Andrews. Uplink capacity and interference avoidance for two-tier femtocell networks. *Wireless Communications, IEEE Transactions on*, 8(7):3498–3509, 2009.

[9] H.-S. Jo, C. Mun, J. Moon, and J.-G. Yook. Interference mitigation using uplink power control for two-tier femtocell networks. *Wireless Communications, IEEE Transactions on*, 8(10):4906–4910, 2009.

[10] J.-H. Yun and K. G. Shin. CTRL: A self-organizing femtocell management architecture for co-channel deployment. In *MobiCom*, 2010.

[11] S. Sen, T. Zhang, M. M. Buddhikot, S. Banerjee, D. Samardzija, and S. Walker. A dual technology femto cell architecture for robust communication using whitespaces. In *IEEE Dynamic Spectrum Access Networks (DYSPAN)*, 2012.

[12] K. Sundaresan and S. Rangarajan. Efficient resource management in ofdma femto cells. In *MobiHoc*, 2009.

[13] M. Y. Arslan, J. Yoon, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee. Fermi: A femtocell resource management system forinterference mitigation in ofdma networks. In *MobiCom*, 2011.

[14] J. Yoon, M. Y. Arslan, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee. A distributed resource management framework for interference mitigation in ofdma femtocell networks. In *MobiHoc*, 2012.

[15] S. Deb, P. Monogioudis, J. Miernik, and J. P. Seymour. Algorithms for Enhanced Inter-Cell Interference Coordination (eICIC) in LTE HetNets. *IEEE/ACM Trans. Netw.*, 22(1):137–150, February 2014.

[16] S. Ren, J. Park, and M. Van Der Schaar. Entry and spectrum sharing scheme selection in femtocell communications markets. *IEEE/ACM Trans. Netw.*, 21(1):218–232, February 2013.

[17] Y. Wang, J. M. Wang, and B. Bensaou. Regret-based learning for medium access in lte femtocell networks. In *MSWiM*, 2014.

[18] M. Z. Chowdhury, W. Ryu, E. Rhee, and Y. M. Jang. Handover between macrocell and femtocell for umts based networks. In *IEEE ICACT*, 2009.

[19] H. Zhang, X. Wen, B. Wang, W. Zheng, and Y. Sun. A novel handover mechanism between femtocell and macrocell for lte based networks. In *IEEE ICCSN'10*, 2010.

[20] T. Bai, Y. Wang, Y. Liu, and L. Zhang. A policy-based handover mechanism between femtocell and macrocell for lte based networks. In *Communication Technology (ICCT), 2011 IEEE 13th International Conference on*, pages 916–920. IEEE, 2011.

[21] A. Golaup, M. Mustapha, and L. B. Patanapongpibul. Femtocell access control strategy in umts and lte. *Communications Magazine, IEEE*, 47(9):117–123, 2009.

[22] T. Guo, N. Wang, R. Tafazolli, et al. Local mobility management for networked femtocells based on x2 traffic forwarding. *Vehicular Technology, IEEE Transactions on*, 62(1):326–340, 2013.

[23] Fujitsu. 4g femtocell for the enterprise environment. http://www.fujitsu.com/us/Images/FNC-4G-Femtocell.pdf.

[24] LightReading. Samsung snags verizon 4g small cell indoor deal. http://www.lightreading.com/mobile/small-cells/samsung-snags-verizon-4g-small-cell-indoor-deal/d/d-id/708029.

[25] Opensignal. http://opensignal.com/android/.

[26] Speedtest.net - Ookla. http://www.SpeedTest.net.

[27] Usrp n210. http://www.ettus.com/product/details/UN210-KIT.

[28] Radysis. Lipa: Local ip access via home node b. http://go.radisys.com/rs/radisys/images/paper-femto-lipa.pdf.

[29] 3GPP. TR23.829: Technical Report; Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO) (Release 10), 2011.

[30] R. Gupta and N. Rastogi. LTE Advanced: LIPA and SIPTO. https://www.aricent.com/pdf/Aricent_LIPA_SIPTO_Whitepaper.pdf.

[31] I. T. Union. Network grade of service parameters and target values for circuit-switched services in the evolving isdn, May 1999. Recommendation E.721, Telecommunication Standardization Sector of ITU.

[32] G. Tu, C. Peng, H. Wang, C. Li, and S. Lu. How Voice Calls Affect Data in Operational LTE Networks. In *MobiCom*, Oct. 2013.

[33] 3GPP. TS36.331: E-UTRA: Radio Resource Control (RRC).

[34] 3GPP. TS25.331: Radio Resource Control (RRC); Protocol specification.

[35] 3GPP. TS24.008: Core Network Protocols, 2012.

[36] Y. Li, J. Xu, C. Peng, and S. Lu. A First Look at Unstable Mobility Management in Cellular Networks. In *HotMobile*, Feb 2016.