

MIT Open Access Articles

LTE radio analytics made easy and accessible

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Swarun Kumar, Ezzeldin Hamed, Dina Katabi, and Li Erran Li. 2014. LTE radio analytics made easy and accessible. In Proceedings of the 2014 ACM conference on SIGCOMM (SIGCOMM '14). ACM, New York, NY, USA, 211-222.

Published Version: <http://dx.doi.org/10.1145/2619239.2626320>

Publisher: Association for Computing Machinery (ACM)

Permanent Link: <http://hdl.handle.net/1721.1/100518>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



LTE Radio Analytics Made Easy and Accessible

Swarun Kumar[†] Ezzeldin Hamed[†] Dina Katabi[†] Li Erran Li^{*}
[†]Massachusetts Institute of Technology {swarun, ezz, dk}@mit.edu
^{*}Bell Labs, Alcatel-Lucent erranli@research.bell-labs.com

ABSTRACT

Despite the rapid growth of next-generation cellular networks, researchers and end-users today have limited visibility into the performance and problems of these networks. As LTE deployments move towards femto and pico cells, even operators struggle to fully understand the propagation and interference patterns affecting their service, particularly indoors. This paper introduces LTEye, the first open platform to monitor and analyze LTE radio performance at a fine temporal and spatial granularity. LTEye accesses the LTE PHY layer without requiring private user information or provider support. It provides deep insights into the PHY-layer protocols deployed in these networks. LTEye’s analytics enable researchers and policy makers to uncover serious deficiencies in these networks due to inefficient spectrum utilization and inter-cell interference. In addition, LTEye extends synthetic aperture radar (SAR), widely used for radar and backscatter signals, to operate over cellular signals. This enables businesses and end-users to localize mobile users and capture the distribution of LTE performance across spatial locations in their facility. As a result, they can diagnose problems and better plan deployment of repeaters or femto cells. We implement LTEye on USRP software radios, and present empirical insights and analytics from multiple AT&T and Verizon base stations in our locality.

Categories and Subject Descriptors C.2.2 [Computer Systems Organization]: Computer-Communications Networks

Keywords LTE; Cellular; Analytics; PHY; Wireless

1. INTRODUCTION

Cellular service has become an integral part of our life. Yet as users and researchers, we have little visibility into the performance and real problems of these networks. Even the little information we have is primarily from trace analysis sanctioned by mobile operators [15, 16]. The lack of open and transparent access into the operation and inefficiencies of the cellular physical layer limits our ability as researchers to contribute effectively to the development of these networks. It also limits the ability of policy makers to independently verify operators’ claims of spectrum needs, and make informed decisions on licensed vs. unlicensed spectrum.

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGCOMM’14, August 17–22, 2014, Chicago, IL, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2836-4/14/08

<http://dx.doi.org/10.1145/2619239.2626320> ...\$15.00.

The need for increased visibility into the cellular PHY-layer is further emphasized by three recent trends. First, cellular deployment is moving towards small, femto, and pico cells [26], many of which will be deployed by a user to cover her home or small business. As a result, cellular operators no longer have full control over their LTE deployments, and struggle to understand the propagation and interference patterns affecting their service, particularly in indoor settings. Open, cheap, and ubiquitous radio monitoring can help deal with the challenging propagation patterns introduced by small cells. Second, the rise of LTE-based M2M applications motivates a more open access to LTE signal-based analytics. For example, Walmart, Home Depot, or Disneyland may leverage LTE signals and recent RF-based localization techniques to track how clients navigate their space and obtain business analytics. Also, oil and gas companies that deploy LTE-supported seismic sensors [7] can leverage access to LTE radio propagation to better plan their sensor network and debug connectivity problems. Third, the FCC plans to repurpose certain bands (e.g., 3.5 GHz) for multi-tier spectrum sharing, including LTE small cell deployments [10]. Operating in a shared spectrum naturally fits with an open model for signal monitoring and analysis, where all the entities sharing the spectrum can better understand the problems and cooperate to find solutions.

All of these reasons motivate a more open access to the cellular PHY layer, particularly LTE. In this paper, we ask the following question: Can we access the cellular PHY-layer of today’s LTE deployments, without support from mobile operators? Specifically, can we do this without requiring access to private user data or encrypted LTE channels? If such a service exists, it could enable better deployment of femto cells and repeaters, more businesses built over LTE networks, better informed spectrum policies, more efficient sharing of newly released bands, and an overall rise in transparency in an industry that is a major part of the world economy.

We introduce LTEye, an open platform for monitoring and analyzing the LTE PHY layer. LTEye is a passive sniffer that runs on off-the-shelf software radios (e.g. USRPs). It does not require provider support, and hence can serve end users, researchers, policy makers, or mobile broadband providers. LTEye extracts per-user analytics purely from the LTE control channels that contain meta-data on uplink and downlink transmissions, without accessing private data or system parameters from encrypted data channels. Specifically, it tracks individual users based on their temporal PHY-layer IDs, without requiring or exposing their private information. It then intelligently links these IDs across control messages to generate transmission records for each user. Each record reports the transmission’s resource utilization, modulation and coding rate, and frame loss rate. LTEye also records the wireless channel it perceives from base stations and mobile users within radio range. These channels are used to accurately monitor the 3-dimensional physical location of the users. LTEye maintains these records in a

database called LTEyeDB. It processes the records to generate two dimensions of fine-grained analytics: temporal analytics, to track LTE performance over time, and spatial analytics, to characterize LTE service across spatial locations.

We implemented LTEye on USRP software radios [8]. We deployed LTEye in four locations in our campus to compare the temporal performance of two major LTE providers: AT&T and Verizon. Our results revealed several inefficiencies in these networks. First, both providers deploy a Frequency Division Duplexing scheme, which uses independent equally sized frequency bands for uplink and downlink traffic. However, LTEye reported that for both providers, the average utilization of downlink resources (25.2% - AT&T, 58.2% - Verizon) far exceeded that of uplink resources (0.6% - AT&T, 2.6% - Verizon). While it is expected that the downlink is higher in demand, our results reveal that the LTE uplink is a remarkable 20 to 40 times less utilized than the downlink. LTEye’s analytics can therefore help policy makers encourage operators to adopt revised LTE standards that allow more prudent allocation of resources to the uplink and downlink¹, without relying on data from providers themselves to make the case.

Second, LTEye localized certain spots in our campus, where Verizon cellphones suffer poor link quality and often switch to 3G, despite reporting high signal power from the LTE base station. To investigate this, we moved our LTEye sniffers to these spots and found that they experienced high inter-cell interference (about 27 dB) from as many as five different base stations. To make matters worse, many of these base stations used overlapping channel estimation pilots, significantly impacting the decodability of these transmissions. These results help end-users identify poor placement of femto cells that cause such interference. Further, they benefit cellular providers themselves because they reveal interference problems that end-users face in indoors, hitherto inaccessible to providers. Interestingly, some of these PHY-layer inefficiencies may be unknown even to the operators as they are part of the PHY-layer implementations adopted by the base station vendors.²

LTEye also benefits researchers by providing deep insights into the PHY-layer protocols deployed by cellular providers. While the LTE standard spells out much of the PHY layer, the choice of rate adaptation algorithm is still left to individual operators. To gain insights into this algorithm, we analyzed LTEyeDB records of an AT&T base station in our locality. We found that even for static users with completely coherent channels and stable SNRs, the modulation and coding scheme changes significantly even between adjacent transmissions. More interestingly, the average modulation and coding of frames sent to a user changes, based not only on her wireless channels, but also on the network state as a whole. Specifically, if the network is scarcely utilized, the base station transmits to the user conservatively at low modulation on average, even if the wireless link is stable and supports much higher modulation. In contrast, as network utilization increases, under identical SNRs, the base station steadily increases its modulation to support more aggressive data rates to the user. Such analytics on the performance and design choices of today’s cellular operators help researchers design better LTE protocols.

LTEye enables businesses and network administrators to continuously monitor the spatial locations of mobile users, and build a geographic heatmap of LTE coverage and performance within their facility. However, accurately localizing mobile users purely based on their LTE signals is a challenging task. This is because past work on accurate indoor localization proposes two classes of solutions

that are ill-suited to LTE networks: First, localization using antenna arrays [30, 18] requires large bulky arrays, owing to the relatively low frequencies of LTE signals. Second, recent localization techniques using synthetic aperture radar (SAR) are less bulky, but are limited to signals transmitted and received by the same node (e.g. radar [12] or RFID backscatter [29] systems) and therefore do not apply to LTE signals. LTEye provides the best of both these solutions by extending SAR localization techniques to operate over communication signals as opposed to backscatter or radar signals. It also introduces a novel technique to handle errors due to multipath by identifying the shortest (or most direct) path.

Our evaluation of LTEye’s spatial analytics in large indoor environments reveals a median accuracy in 3D localization of mobile users of 61 cm in line-of-sight and 85 cm in non-line-of-sight settings. Further, we visualize the LTE performance of the mobile users across locations, helping building managers find optimal locations for relays or femtocells.

Contributions: This paper contributes the following:

- The paper presents LTEye, the first open platform to monitor and analyze per-user LTE PHY performance at fine temporal and spatial granularity.
- LTEye employs a new technique to identify and track individual users at the LTE PHY layer in a robust manner, without help from operators, and without requiring or exposing private user information.
- LTEye develops an innovative technique for accurate localization of users based on their LTE signals. This involves extending synthetic aperture radar (SAR) to operate over communication signals as opposed to backscatter and radar signals, and a novel technique for identifying the shortest and most direct path in the presence of multipath.
- LTEye’s evaluation on AT&T and Verizon LTE deployments reveal deep insights on the inefficiencies, utilization patterns, and differences between these providers. LTEye also provides heat maps to characterize LTE performance across indoor locations, without GPS support.

2. RELATED WORK

(a) LTE Sniffing Equipment: Devices such as Wavejudge, ThinkRF and IntelliJudge [25, 28] are wireless protocol sniffers to capture RF signals. They are mainly tools for wireless development and interoperability testing that provide visibility into the interaction between the PHY and protocol layers. Unlike LTEye, these devices need inputs from the cellular provider and do not perform localization or provide spatial analytics.

(b) Open LTE Implementations: There have been efforts in developing open source implementations of LTE protocols, notably OpenAirInterface [9], and OSLD [13]. These initiatives enable running LTE base stations on software radios; they do not extract spatial or temporal analytics.

(c) LTE Measurement Studies: Many recent LTE studies have been conducted using traces collected on participating smartphones or from inside LTE networks. Findings from these studies include: (1) The available bandwidth of LTE networks is highly variable and TCP is not able to fully utilize the bandwidth [16]; (2) LTE is significantly less power efficient than Wi-Fi [15]; (3) LTE latency is more consistent (less variable) than Wi-Fi [27]. Such studies focus on the higher layers of the stack, e.g., TCP throughput, transfer delay, and power usage. In contrast, LTEye focuses on the LTE radio layer; it provides fine-grained temporal and spatial information and does not require traces from the provider.

¹E.g. Asymmetric Carrier Aggregation [19] in LTE Advanced (3GPP Release 10) allows downlink resources to exceed the uplink.

²We confirmed this privately with some base station vendors.

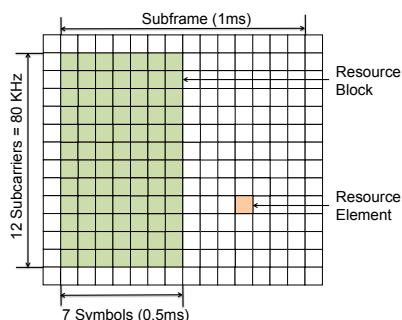


Figure 1: The Resource Grid is divided into multiple Resource Blocks.




Figure 2: Physical Channels are mapped to well defined Resource Elements in the Grid.

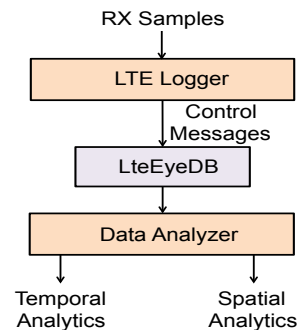


Figure 3: LTEEye's Architecture: Contains the LTE Logger and Data Analyzer.

(d) Cellular Location-Specific RF Measurements: Cellular operators need location-specific RF measurements to troubleshoot performance problems and plan future deployments. They typically obtain coarse location information by mapping a user to her serving cell. Operators then rely on drive tests to refine the spatial measurements. Drive tests are costly and constitute a big part of the network operating expenditure [17]. Further, they are increasingly inadequate as operators move toward femto cells, and need indoor coverage data. To reduce the cost and improve the spatial measurements, recent LTE releases propose mechanisms known as MDT [2]. MDT techniques localize a mobile phone either using in-network time measurement or by collecting location information using the phone's GPS. It is well-known however that in-network localization in cellular networks is not accurate (at hundreds of meters [22]) as time-delay measurements are only available for the serving cell of a mobile user. Even the E911 service using positioning reference signals can only guarantee 150m accuracy 95% of the time [21]. GPS measurements cannot be invoked often as they drain the user's battery. They also cannot capture indoor location.

(e) RF-based Localization Techniques: Our work is related to past work on RF-based localization. This problem has received much recent interest resulting in highly accurate systems. ArrayTrack [30] and PinPoint [18] are an antenna-array based indoor location systems that tracks wireless clients at fine granularity. Chen et. al.[14] build antenna arrays using software radios synchronized with a reference signal. PinIt [29] is an RFID localization system that combines SAR with a deployment of reference RFIDs to achieve highly accurate localization.

Our design builds on this past work but differs in that it introduces two innovative localization techniques. First, we extend SAR to operate over communication signals exchanged between a transmitter and a receiver. This contrasts with the current approach for SAR, which is limited to backscatter and radar signals, where the transmitter and receiver are a single node with no Carrier Frequency Offset (CFO) or Sampling Frequency offset (SFO). Second, we introduce a new technique that when combined with SAR or standard antenna arrays, estimates the delay difference between the various paths traversed by the signal to identify the shortest path. In particular, we measure these delays in time based on phase offsets in the frequency domain. Hence, LTEEye can resolve differences in delay below one time sample, unlike past work that estimates these delays via correlation in time [18].

3. LTE PRIMER

In this section, we briefly introduce LTE concepts relevant to this paper. LTE networks are divided into multiple geographical regions called cells. Each cell contains a cellular base station that serves multiple mobile users. We focus on Frequency Division Duplexing,

the mode of LTE most widely used by cellular operators. This LTE mode uses different dedicated carrier frequency bands for uplink and downlink transmissions. Hence, each base station uses a pair of frequency bands to communicate with users in its cell.

(a) Radio Resources. LTE's uplink and downlink transmissions are based on OFDM. While medium access and resource sharing is largely distributed in typical OFDM-based systems such as Wi-Fi, LTE centralizes much of resource allocation at the base station. Specifically, base stations divide radio resources into multiple frames over time, each containing ten subframes, spanning 1 ms each. Resources in each sub-frame are divided both along time and frequency as a 2-D time-frequency grid, as in Fig. 1. Each cell in the grid, called a resource element corresponds to one OFDM sub-carrier (15 KHz) over the duration of one OFDM symbol ($66.7 \mu s$).

The key task of an LTE base station is to allot both uplink and downlink resources between different users along both time and frequency. It allocates resources to users at the granularity of *resource blocks*, each spanning 0.5 ms (i.e., half a sub-frame) by 180 kHz (i.e., 12 sub-carriers). To combat frequency-selective fading, the assignment of resource blocks to users both on the uplink and downlink is not fixed, but hops from transmission to transmission.

(b) Physical Layer Channels. As LTE centralizes PHY-layer control, base stations need to transmit both data and control information to the users. To this end, LTE partitions network resources into well defined channels, each responsible for different types of information. These channels are mapped to well-known resource elements of the grid, as shown in Fig. 2.

Broadly, LTE base stations use four main channels on the downlink: (1) A data channel to send users their downlink data. (2) A control channel to allocate network resources. (3) A broadcast channel for new users to learn system parameters. (4) A hybrid-ARQ channel to send ACKs to the users. Similarly, the mobile users on the uplink are allocated data and control channels to transmit their uplink data and control messages. Table 1 describes these channels in greater detail.

We point out here that the downlink control channel bears rich information on the LTE PHY-layer. Specifically, it contains multiple *downlink control information messages* in which the base station allocates resource blocks to specific users for every transmission on either the uplink or downlink. In addition, the control information also specifies the modulation and coding to be used in these blocks. Upon hearing a control message, a user accesses the relevant data channel to send (receive) her uplink (downlink) transmission.

(c) PHY User Identifier. The LTE PHY refers to each mobile user using a temporary unique ID called the Cell Radio Network Temporary Identifier (C-RNTI). The C-RNTI reveals no private information about the user. It is local to the user's serving cell, and

Downlink LTE Channels	
Name	Description
PBCH	Physical Broadcast Channel: Carries general information about the cell, like number of antennas on the base station and total bandwidth.
PDCCH	Physical Downlink Control Channel: Sends downlink control messages, e.g. for resource allocation on uplink/downlink.
PDSCH	Physical Downlink Shared Channel: Holds downlink data meant for users in resource blocks indicated by the PDCCH.
PHICH	Physical Hybrid-ARQ Channel: Contains positive or negative acknowledgments for uplink data.
Uplink LTE Channels	
PUCCH	Physical Uplink Control Channel: Holds control information from users to base stations, e.g. ACKs, channel reports and uplink scheduling requests.
PUSCH	Physical Uplink Shared Channel: Mainly carries uplink data in resource blocks indicated by the PDCCH.

Table 1: **LTE Channels.** Details the name and function of PHY-layer channels, as defined in LTE standards [3].

is assigned when she enters the cell via a higher-layer connection establishment procedure[4].

A user continues to have the same C-RNTI as long as she is in the same cell and is not idle for more than the pre-configured timer value. The timer value is typically a few seconds to tens of seconds (12 sec in the measurement result of [15]). Hence, the C-RNTI assigned to a user may change quite often if it transmits sporadically.

4. LTEEYE

LTEeye is an open platform to analyze LTE radio performance. Its design aims to satisfy the following key attributes:

- *Provider-Independent:* LTEeye does not require any information or support from mobile providers. Thus, LTEeye allows end-users, policy makers and third parties to make a fair assessment and comparison of the service quality of different providers, without relying on information furnished by the providers themselves.
- *Off-the-Shelf:* LTEeye must be built on low-end off-the-shelf components, such as standard laptops and software radios. This makes LTEeye more accessible to end-users and easy to be deployed in large numbers.

4.1 System Architecture

LTEeye operates as a passive 2-antenna MIMO receiver. LTEeye’s architecture is a pipeline of two components: (1) the LTE Logger, and (2) the Data Analyzer. (see Fig. 3)

(a) **LTE Logger:** The LTE Logger sniffs on the LTE control channels to generate transmission records. The logger begins by listening to the broadcast channel to gather system parameters. It then sniffs the downlink control channel and performs LTE decoding, i.e., it demodulates the OFDM symbols, applies de-interleaving, de-scrambling and convolutional decoding to extract the actual bits of the downlink control messages.

The logger reads each of these control messages to populate LTEeyeDB, a database of LTE information records tagged with their transmission time. Each transmission record consists of several fields retrieved from the control messages, as listed in Table 2. Many of these fields help characterize network performance and utilization, both for the user over time, and for the network as a whole, if viewed in aggregate. In addition, notice two important fields in the records: (1) Each record is indexed by the user’s C-

RNTI (i.e. her PHY-layer user ID), which are key to link multiple records belonging to the same user. (2) Records maintain the uplink channels seen by the LTEeye sniffer over time. In §6, we show how these channels are essential to localize users.

(b) **Data Analyzer:** The data analyzer processes the records in LTEeyeDB to extract fine-grained analytics on both cellular base-stations and individual mobile users. It extracts two types of analytics: temporal analytics which describe LTE PHY metrics as functions of time (e.g., the per-user resource allocation over time), and spatial analytics which describe position-dependent RF metrics (e.g., the user location and the observed multipath effects). Sections §5 and §6 discuss both types of analytics in detail.

5. ENABLING TEMPORAL LTE ANALYTICS

In this section, we describe some of the challenges in obtaining temporal analytics, without access to private user information or system parameters in encrypted data channels.

Uniquely Identifying Users: To extract per-user temporal analytics, LTEeye needs to uniquely identify each user (i.e mobile device) served by the base station. We do not know the phone numbers of users, but can instead use the series of C-RNTIs (i.e. PHY-layer user IDs) assigned to them. One option is to require LTEeye to sniff the LTE channel continuously for an extended time to catch each user at the time she joins the cell and capture her C-RNTI assignment. Unfortunately this option has two limitations: First, the C-RNTI assignment is often transmitted from higher layers in the encrypted downlink data channel [4]. Second, low-end off-the-shelf equipment is unlikely to be able to continuously monitor and decode the LTE channel in real-time [6]. Instead, LTEeye sniffers should be able to periodically sniff the channel and obtain representative snapshots of the system. Hence, we need LTEeye to uniquely identify all users, including those who joined the cell even when LTEeye is not sniffing (i.e., LTEeye did not hear their C-RNTI assignment).

To address this problem, we observe that a user’s C-RNTI is used to scramble her control information on the downlink control channel. Specifically, recall from §3 that the control channel transmitted by the base station consists of multiple downlink control information messages, for different users. At the end of each message is a 16-bit sequence, which is the XOR of the checksum of the control message with the user’s C-RNTI. Traditionally, a user de-scrambles this sequence by her C-RNTI to retrieve the checksum and validate correctness of the control information. In contrast, LTEeye performs the *opposite* operation to retrieve the C-RNTI: It decodes each control message in the log including their corresponding scrambled checksums. For each of these packets, LTEeye reconstructs the expected checksum and XORs them with the scrambled checksum to recover the C-RNTI. Of course, it is important to verify if the control messages and C-RNTIs that are decoded are actually correct. To do this, LTEeye convolutionally re-encodes the retrieved control information message and compares it against the original coded control message to obtain the number of bit errors. It then discards control messages and C-RNTIs that report bit errors beyond a few bits.³ Hence, our solution enables LTEeye to map a C-RNTI to a user even if it was assigned when LTEeye is not sniffing the channel.

Tracking User IDs: LTEeye’s second challenge is to map C-RNTIs between logs. Recall from §3 that a user’s C-RNTI may be re-assigned if she is idle beyond a few seconds. Further, the same C-RNTI can map to multiple users over time. Hence, while the C-RNTI is a natural PHY-layer user ID, LTEeye must recognize when a user’s C-RNTI changes to find the list of C-RNTIs mapped to her.

³In our experiments, LTEeye correctly retrieved 99.5% of C-RNTIs across locations.

Field	Description	Format
Time	Transmission Time	32-bit timestamp
C-RNTI	PHY-layer ID	16-bit Sequence
UL/DL	Uplink or Downlink	0/1
nrb	Number of Resource Blocks	0 to N_{RB}
alloc	Bit-Map of Resource Blocks	N_{RB} bits
MCS	Modulation and Coding Scheme	5 bits
isAcked	Acknowledgment	± 1 (ACK/NACK)
UE-channel	Channel from user (if U/L)	$N_{sc}N_{rx}$ Complex Floats
BS-channel	Channel from Base Station	$N_{sc}N_{tx}N_{rx}$ Complex Floats
SNR, SINR	SNR and SINR of Base Station	Floating Point

Table 2: **Fields of LTEyeDB**. Where N_{RB} : Number of resource blocks, N_{tx} : Number of base-station antennas, N_{rx} : Number of sniffer antennas, N_{sc} : Number of OFDM sub-carriers.




Figure 4: **Mapping C-RNTI**. Mapping C-RNTI from old logs to the current log.

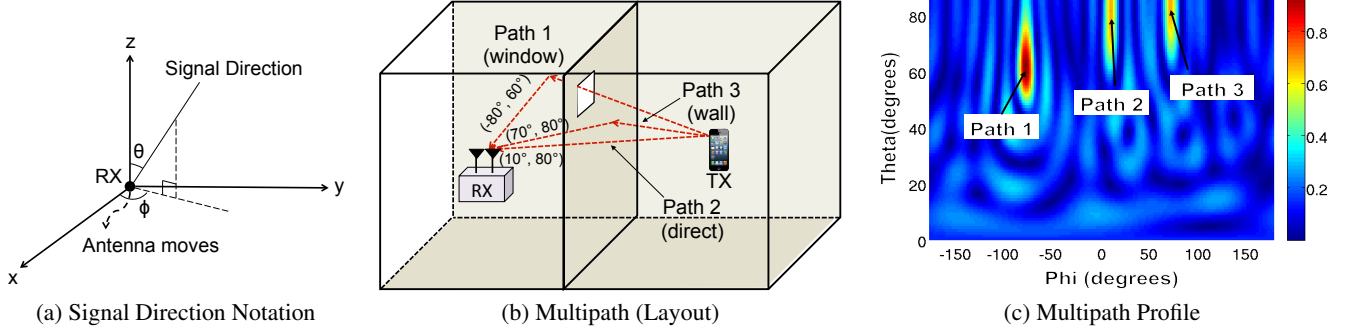


Figure 5: (a) Definition of (ϕ, θ) in 3-D space. (b) Depicts an example layout of a transmitter and receiver in two rooms separated by a wall. The signal has three main paths: Path 1 is the strongest, and reflects from the ceiling through a window. Path 2 is the direct path penetrating a wall. Path 3 is the weakest reflecting at the farthest wall. The figure labels (ϕ, θ) for each path. (c) Simulated multipath profile for (b) has peaks for each path at expected (ϕ, θ) . The height of the peaks (in shades of red) corresponds to the relative power of the corresponding paths.

LTEye addresses this challenge by formulating it as a matching problem. The goal of this problem is to map C-RNTIs in the current (most recent) log produced by the LTE logger with the C-RNTIs in prior logs as shown in the bi-partite graph in Fig. 4. In addition, the graph has two additional nodes: EXIT and NEW, which account for C-RNTIs in the old log that have “exit” the system, and C-RNTIs in the current log that are “new” to the cell, respectively. The weights in this graph must capture the similarity between the users associated with each pair of C-RNTIs. Specifically, we associate with each C-RNTI $\#i$ an RF fingerprint f_i that includes metrics such as the user’s location (extracted by our localization method described in §6), its SNR and multi-path characteristics. We can then assign a weight to each edge (i, j) in the graph by the similarity metric between these fingerprints $sim(f_i, f_j)$. In §6.3, we design effective RF fingerprints and similarity metrics based on spatial analytics.

Given the graph and weights, we can now solve this matching problem using the standard Hungarian Algorithm[20].⁴ The resulting matching either maps a C-RNTI to a user in a prior log, or identifies her as a new user.

Extracting System Parameters: LTEye needs to *reverse-engineer* several PHY parameters, otherwise available to users via encrypted data channels. For e.g., the LTE standard allows several possible formats for downlink control information, each spanning multiple lengths [5]. Exhaustively searching for each of the possible format-length pairs in all downlink control messages is highly expensive. Fortunately, operational LTE base stations use only a small subset of these formats for all users (only three possible formats for AT&T

and Verizon). As a result, LTEye learns the possible list of formats and lengths using the first few control messages to greatly reduce the search-space of formats for subsequent control information.

6. ENABLING SPATIAL LTE ANALYTICS

The core of LTEye’s spatial analytics is the ability to localize an LTE source. To this end, LTEye performs the following functions:

- *Extracts the multipath profile of an LTE signal:* First, LTEye extracts the multipath profile of a signal, which measures the power received along each spatial angle, i.e., along each signal path. Fig. 5(c) is an example multipath profile, where the signal traverses three paths causing three local maxima in the graph.
- *Identifies the direct path from the source:* Once LTEye has the multipath profile, it is natural to try to identify which path is in direct line of sight to the source (e.g., path 2 in Fig. 5(b)). Finding the direct path is an important step in localizing the source. Note that in some cases the direct path may be completely blocked and absent from the multipath profile. Our objective is to identify the direct path provided it exists in the multipath profile.
- *Localizes the source of the LTE signal:* Once LTEye finds the direct path to the source, it can localize the source to within a specific spatial direction. We can then deploy multiple LTEye sniffers to locate the source at the intersection of the direct paths as seen from these sniffers.

Past work on RF-based localization takes two approaches to build multipath profiles: The first approach, shown for Wi-Fi, uses an antenna array to steer its beam spatially and identify the power along each spatial direction [30]. However, LTE runs at much lower frequencies than Wi-Fi (700 MHz as opposed to 2.4 GHz), and

⁴We modify the algorithm to allow multiple C-RNTIs to map to NEW/EXIT simply by replicating these nodes. Edges at NEW/EXIT are weighted by a minimum threshold similarity.

hence an LTE antenna array will be 4 to 5 times more bulky than a comparable Wi-Fi array. The second approach uses synthetic aperture radar (SAR) [29], which uses a single movable antenna to emulate a virtual array of many antennas. As the antenna moves, it traces the locations of antenna elements in a virtual array.

SAR has traditionally been used in radar and RFID localization as it assumes backscatter signals where the transmitter and receiver are the same node and therefore have no carrier frequency offset (CFO) relative to each other. Hence, changes in the channel as the antenna moves are a function only of the antenna's location. In contrast, LTE signals are exchanged between an independent transmitter and receiver, with non-zero CFO. As the antenna moves, the channel changes both due to CFO and antenna movement. One option is to estimate and correct the CFO. Unfortunately this solution is fragile since any residual error in CFO estimation accumulates over the duration of movement and causes large localization errors. In the following section, we explain how we perform SAR over LTE signals without CFO estimation to realize the three functions in the beginning of this section.

6.1 SAR over LTE Signals Using Channel Ratios

LTEye operates as a 2-antenna MIMO receiver, with one static antenna and another movable antenna. The movable antenna may be mounted on a rotating arm attached to the device, which is the approach taken in our implementation (see Fig. 6(a)). The advantage of this approach is that it provides the 3-D spatial direction (i.e. both the azimuthal angle ϕ and polar angle θ) of the various signal paths as shown in Fig. 5(a). Alternatively, the antenna may slide back and forth on an arm fixed to the body of the device.

LTEye uses its MIMO capability to perform SAR over communication signals, but without frequency offset estimation. Our key idea is that instead of applying the SAR equations to the wireless channel of the moving antenna [12], we apply SAR equations to the *ratio* between the channel of the moving antenna to that of the static antenna. Taking the ratio of the two channels eliminates any effect of frequency offset since both MIMO antennas experience the same offset relative to the sender. However, since the ratio is between two antennas, one moving and the other static, it preserves how antenna displacement changes the channel of the moving antenna. This allows SAR to safely retrieve the multipath profile of the signal from the ratio, modulo frequency offsets.

Next, we mathematically show the validity of the above technique. Suppose that the receiver, placed at the origin, wants to measure the power of the signal $P(\theta, \phi)$ received from an independent transmitter along a spatial direction specified by the polar angle θ and azimuthal angle ϕ (see Fig. 5(a)). According to the SAR formulation, this quantity can be measured as: [11, 23]:

$$P(\theta, \phi) = |h(\theta, \phi)|^2, \text{ where } h(\theta, \phi) = \sum_t a_f(t, \theta, \phi) h(t) \quad (1)$$

Here, $h(t)$ is the wireless channel to the moving antenna at time t , assuming zero frequency offset between the transmitter and receiver. The quantity $a_f(t, \theta, \phi)$ captures the relative motion of the transmitter and receiver, and is independent of the wireless channels. For e.g., if the antenna moves along a straight line (i.e., linear SAR) $a_f(t, \theta, \phi)$ is defined as: $a_f(t, \theta, \phi) = e^{-j\frac{2\pi f}{c} x(t) \cos(\phi)}$, where $x(t)$ is the antenna location at time t , and f is the frequency of the signal [29]. Similarly, if the antenna rotates with radius r (i.e., circular SAR) $a_f(t, \theta, \phi) = e^{-j\frac{2\pi f}{c} r \cos(\phi - \phi_0(t))}$, where $\phi_0(t)$ is the angular position of the antenna at time t (see Fig. 6(b)).

Past work on SAR require both transmitters and receivers to share a common reference clock. For e.g., SAR devices on airplanes or satellites both transmit signals and receive their reflections to im-

age the topography of the ground[12]. Consequently, the measured channel $\tilde{h}(t)$ at the moving antenna is independent of frequency offset, i.e., $\tilde{h}(t) = h(t)$.

Unfortunately, when performing SAR between *independent* transmitters and receivers, the measured wireless channel $\tilde{h}(t)$ varies both due to position and due to carrier and sampling frequency offset between the transmitter and receiver, as well as any phase noise. In particular, we denote:

$$\tilde{h}(t) = h(t)e^{j\psi(t)} \quad (2)$$

Where $\psi(t)$ denotes the phase accumulated due to any carrier frequency offset, sampling frequency offset or phase noise between the reference clocks of the transmitter and receiver until time t . Thus, the key challenge to measure the power of the signal along any spatial direction, $P(\theta, \phi)$, as in the SAR Eqn. 1, is to eliminate this accumulated phase.

To resolve this challenge, LTEye is built on receivers that have at least two antennas: a static antenna, and a mobile antenna that moves along a known path. Let $\tilde{h}_1(t)$ and $\tilde{h}_2(t)$ denote the measured wireless channels from a given transmit antenna to a mobile and static antennas respectively (see Fig. 6(c)). As both antennas are connected to the same reference clock, they both accumulate the same phase $\psi(t)$ until time t . Hence, from Eqn. 2, we have:

$$\tilde{h}_1(t) = h_1(t)e^{j\psi(t)}, \quad \tilde{h}_2(t) = h_2(t)e^{j\psi(t)} \quad (3)$$

where $h_2(t) \approx h_2$ is relatively constant over a short duration since antenna-2 is static. Hence, the ratio of the wireless channels is: $\tilde{h}_r(t) = \frac{\tilde{h}_1(t)}{\tilde{h}_2(t)} = \frac{1}{h_2} h_1(t)$; that is, the channel ratio is a constant multiple of the moving antenna channel without the phase accumulation from frequency offset or phase noise. Thus, we can perform SAR as in Eqn. 1 by substituting the channel ratio $\tilde{h}_r(t)$ for the value of $h(t)$. Hence, LTEye allows a wireless receiver to perform SAR over LTE signals without frequency offset or phase noise estimation.

Finally, we make a few important observations:

- The above approach readily extends to OFDM / OFDMA. Specifically, let $h(\theta, \phi) = \sum_f \sum_t a_f(t, \theta, \phi) h_{r,f}(t) / h_{r,f}(0)$ in Eqn. 1, where the quantity $\tilde{h}_{r,f}(t)$ on subcarrier f of the OFDM signal is the ratio of the frequency (Fourier) Domain channels $\tilde{h}_1(t)$ and $\tilde{h}_2(t)$ measured on that subcarrier.⁵
- Our solution is resilient to movement of the transmitter that can be neglected relative to the movement of the rotating antenna. LTEye's rotation speed is chosen to deal with typical dynamism in indoor settings, e.g. walking speeds. However, one can easily detect and exclude fast moving transmitters by checking if the channel of the static antenna $\tilde{h}_2(t)$ is coherent over a rotation of the moving antenna, using the coherence metric in §8.3.
- While this paper applies our solution to LTE, our technique can be extended to apply to Wi-Fi and other such technologies.

6.2 Identifying the Direct Signal Path

In this section, we describe how LTEye separates the direct path from the reflected paths in a multipath profile reported by SAR so as to localize the transmitter. Intuitively, the direct path is the shortest path among all paths traversed by the signal (even if the direct path is completely blocked, the shortest path is the path closest to the direct path). Thus, one may identify the direct path (or the path closest to the direct path) by measuring the delay difference between the various paths in the multipath profile of the signal.

Directly measuring time delays (e.g. by correlating with known pilot signals), however is not sufficiently accurate. Specifically,

⁵Note that this is robust to frequency hopping by LTE users.




Figure 6: (a) Prototype of the rotating antenna on a LTEEye sniffer. (b) Circular Synthetic Aperture Radar (c) Depicts $\tilde{h}_1(t)$ and $\tilde{h}_2(t)$, the measured wireless channels from a transmit antenna to the mobile and static antennas respectively

LTE receivers have a channel bandwidth of 10 MHz. However, electromagnetic waves travel at the speed of light. Hence an error of even one time sample for a 10 MHz sampling rate (i.e., each time sample spans $0.1\mu s$) translates to an error in path lengths of 30 m.

Below we explain how LTEEye can measure *sub-sample delay differences* between the signal paths. The key idea is to exploit that delay in time translates into phase rotation in the frequency domain. Since LTE signals use OFDM, a time delay of the signal translates into phase rotation in the OFDM subcarriers. Yet, different OFDM subcarriers rotate at different speeds – i.e., higher OFDM frequencies rotate faster than lower frequencies. In fact, the phase rotation of a particular OFDM subcarrier f_i as a result of a delay τ is $\psi_i = 2\pi f_i \tau$. Thus, for each subcarrier, the difference in delay between two paths, p and q , for a particular subcarrier is:

$$\tau = \frac{\psi_{p,i} - \psi_{q,i}}{2\pi f_i} \quad (4)$$

One may also average across subcarriers to improve robustness to noise. Multiple sub-carriers can also resolve ambiguity if $|\psi_{p,i} - \psi_{q,i}| > 2\pi$ by correcting discontinuities in $\psi_{p,i} - \psi_{q,i}$ across frequencies. Thus, to identify the shortest path in a multipath profile, LTEEye does the following:

- First, it computes the phase of the channel for each subcarrier for each path separately. This can be done using the fact that the SAR formulation defines $h(\theta, \phi)$, which provides not just the power, but also the phase of the channel component along each spatial direction (see Eqn. 1). Hence, we can simply measure the phase of this path as $\psi = \arg[h(\theta, \phi)]$, for each OFDM subcarrier.
- Second, it computes the delay difference between each pair of paths using the Eqn. 4 above. It then identifies the shortest path as the one with least delay.

Once the direct path is found, the source is localized along this path. Complete localization can be performed using multiple LTEEye receivers and intersecting their direct paths. If the direct paths do not intersect, the best estimate is the point minimizing total distance (or equivalently, delay) to all LTEEye receivers. This point can be found using a simple geometric optimization omitted for brevity.

Next, we show how the multipath profile reported by SAR allows us to compute unique RF fingerprints for the users.

6.3 Measuring RF Fingerprints with Multipath Profiles

As described in §5, LTEEye tracks the different C-RNTIs (PHY-layer user IDs) assigned to mobile users between logs. To this end, LTEEye employs RF fingerprints [31, 29] to map C-RNTIs between logs. LTEEye defines a user’s fingerprint as the set of observed multipath profiles at each LTEEye sniffer (See §6.1). The key advantage

of this fingerprint is that it captures the user’s location, multipath, and SNR, as perceived from LTEEye sniffers. To measure similarity of two fingerprints, we employ dynamic time warping (DTW[24]), a technique that has recently been applied to capture similarity of two multipath profiles[29]. Given any two multipath profiles, DTW returns a cost function that varies inversely with their similarity. Hence, LTEEye defines the similarity metric between two RF fingerprints as the inverse of the total DTW cost function[29] between each pair of profiles in the fingerprints.

One might wonder if LTEEye’s fingerprint matching algorithm scales, given that a cell may serve a large number of users. Fortunately, while LTE cells can serve hundreds of users, we observed that only a small fraction of these users (about 4%) are re-assigned C-RNTIs between two logs.⁶ Further, while capturing C-RNTI re-assignments is essential to track individual users over time, it does not alter aggregate radio analytics in a statistically significant way, given that 96% of users in a cell retain their C-RNTI between logs.

7. IMPLEMENTATION

We implemented LTEEye on USRP N210 software radios[8] and WBX daughter-boards. The USRPs receive in the 700MHz frequency range at a bandwidth of 10 MHz on uplink and downlink channels corresponding to either AT&T (734-744MHz) or Verizon (746-756MHz). We implement an OFDMA receiver for LTE signals that interfaces directly with the USRP Hardware driver (UHD).

To obtain temporal analytics, we decode the downlink control channel in a pipeline of two modules: The first module in C++ performs synchronization, channel estimation and QPSK demodulation. It logs the demodulated soft bits received over one second into a file, and repeats this process every three seconds. The second module in MEX (C++) and Matlab reads the file and performs descrambling, de-interleaving and convolutional decoding. It validates the downlink control information messages by only admitting those passing the convolutional decoder with very high confidence, as described in §5. It then processes the control messages to get per-user LTEEyeDB records for uplink and downlink traffic, as in §4.1.

To obtain spatial analytics, we build prototype LTEEye sniffers, each containing two USRPs connected to an external clock. We mount the antenna of one of the USRPs on a light-weight rotating arm fabricated by a 3D printer, with an adjustable radius of 15-30 inches, as shown in Fig. 6. The arm is driven by an off-the-shelf stepper motor rotate at 30-120 rotations per minute.⁷ We use an Arduino UNO board to rotate the stepper motor accurately at a constant speed and provide real-time feedback on the position

⁶This is on average 1-2 users for AT&T and 3-4 users for Verizon.

⁷We plan to support higher torque in future iterations of the device.

Metric	AT&T		Verizon	
	Mean	Std-dev	Mean	Std-dev
Number of Users	23.37	7.90	87.66	44.75
% Downlink Utilization	25.20	17.35	58.18	20.58
% Uplink Utilization	0.59	0.47	2.60	1.06
Downlink MCS (bits/RE)	4.56	0.26	5.23	0.30
Uplink MCS (bits/RE)	3.25	0.22	3.61	0.18

Table 3: **Aggregate Statistics.** Tabulates mean and standard deviation of statistics over four locations for AT&T and Verizon.

of the rotating antenna. We implement a C++ module to use these positions and channel measurements, as in §6, to localize the users.

We evaluate LTEye’s spatial analytics using five LTEye sniffers in multiple indoor testbeds, in both line-of-sight and non-line-of-sight settings. We employ ten Samsung Galaxy Note LTE smart phones as users. Unless specified otherwise, each user communicates over LTE with a mix of varying traffic patterns including browsing activity, file transfer, Skype calls, and video streaming.

8. RESULTS ON RADIO ANALYTICS

In this section, we perform an extensive evaluation of LTEye’s temporal and spatial analytics:

- We compare temporal analytics of two LTE providers and highlight their PHY-layer inefficiencies in §8.1 and §8.2.
- We provide insights on the LTE rate adaptation algorithm in §8.3.
- We apply LTEye’s spatial analytics to two applications: detecting cheaters in an exam hall in §8.4 and visualizing a spatial heatmap of LTE performance in §8.5.
- We perform micro-benchmarks to evaluate the accuracy of LTEye’s localization and RF fingerprints in §8.6 and §8.7.

8.1 Comparing Temporal Analytics of Providers

End-users can deploy LTEye sniffers to compare the providers in their locality in terms of usage patterns, quality of service and congestion. In this experiment, we compare aggregate temporal analytics of two providers in our campus: AT&T and Verizon.

Setup. We place LTEye sniffer in four locations in the MIT campus, each listening to the AT&T and Verizon base stations that serve that location. We populate LTEyeDB over the duration of a representative weekday from 9:00am to 9:00pm. To reduce processing overhead, LTEye’s logger collects traces for a duration of one second, every three seconds. It validates these traces by only accepting control information with low bit error rate as reported by the convolutional decoder. We then measure the following metrics for each one-second trace: (1) Number of Active Users; (2) Mean Utilization of the Uplink and Downlink; (3) Mean Link Quality measured as the number of bits per resource element (bits/RE) in the uplink and downlink. We average each of these quantities over one minute intervals and plot them over time of day at a representative location (Fig. 7). We also estimate the mean value of these metrics across locations over one week to infer aggregate trends (Table 3).

Number of Active Users. Fig. 7(a) measures the number of active mobile users in a representative AT&T and Verizon cell over different times of the day. We observe that for both providers, the number of users in the morning increases steadily, and peaks at around 12:00 pm, after which the number begins to decrease. The increase in activity at around noon may be attributed to a greater number of subscribers who access LTE outdoors as they leave for lunch. Across locations, we observe that Verizon has a greater number of active users on average at 87.7, while AT&T has 23.4 active users through the day (see Table 3).

Network Utilization. Fig. 7(b) plots the utilization of a representative AT&T and Verizon cell, over different times of the day. Specifically, we measure the percentage of resource elements used by uplink and downlink traffic. Two trends emerge: First, both providers often achieve high downlink utilization (over 80%) through the day. AT&T achieves such high utilization sporadically through the day (for 2% of the day), while Verizon is heavily utilized for a more significant fraction of the day (for 18% of the day). Second, the utilization of the uplink is significantly lower than the utilization of the downlink, both for AT&T and Verizon.

Specifically, the mean downlink utilization (25.2% - AT&T, 58.2% - Verizon), far exceeds uplink utilization (0.6% - AT&T, 2.6% - Verizon) even when averaged across locations. While it is expected that the downlink is higher in demand, our results reveal that the LTE uplink is an unprecedented 20 to 40 times less utilized than the downlink. This exposes the practical limits of Frequency Division Duplexing mode of the LTE standard used by both AT&T and Verizon, which provides independent equally sized uplink and downlink frequency bands. Hence, our results can help policy makers encourage operators to adopt revised LTE Advanced standards that permit unequal allocation of resources to the uplink and downlink (e.g., via asymmetric carrier aggregation[19]) without relying on data from providers themselves to make the case.

Link Quality. Fig. 7(c) measures the average quality of channels in the network for a representative AT&T And Verizon cell measured over different times of the day. In particular, we measure the mean number of bits transmitted per resource element (bits/RE), capturing the modulation and coding scheme (MCS) on the uplink and downlink. Our results show that the mean quality on the downlink (5.2 - Verizon, 4.6 - AT&T) exceeds that of the uplink (3.6 - Verizon, 3.3 - AT&T). As mentioned earlier, this is because users are limited in transmit power and number of MIMO antennas, when compared to the base station. Further, the mean link quality of Verizon is marginally higher when compared to AT&T.

8.2 Identifying PHY-Layer Problems and Inefficiency

Cellular Providers and independent researchers can use LTEye to diagnose problems and inefficiencies at the LTE-PHY layer. In this experiment, we identify such deficiencies by analyzing the LTEyeDB database populated for both AT&T and Verizon. In particular, we consider the traces gathered over four locations in our campus served by different base stations over one week, as explained in §8.1. Interestingly, many of these PHY-layer inefficiencies may be unknown even to the operators as they are part of the PHY-layer implementations adopted by the base station vendors.⁸

Unnecessary Control Overhead. As explained in §3, the LTE resource grid on the downlink is divided into three main PHY-layer channels: the broadcast channel, the control channel and the data channel. The control channel occupies the first 1-3 symbols of each LTE sub-frame resulting in a control overhead ranging from 7% to 21% of all downlink resources. Ideally, an LTE base station should adapt the number of control symbols used in each sub-frame depending on the amount of control traffic that is required to be sent.

In practice, we discovered that the control overhead of AT&T base-stations was 10%, while that of Verizon base-stations was 21%. This is because, unlike AT&T, Verizon always uses three control symbols in each sub-frame, regardless of the amount of control traffic they contain. One might wonder if this is because Verizon’s control channels are significantly more utilized, warranting the additional overhead. Fig. 8(a) plots the control overhead of both AT&T and Verizon, as well as how much of this overhead is

⁸We confirmed this privately with some base station vendors.

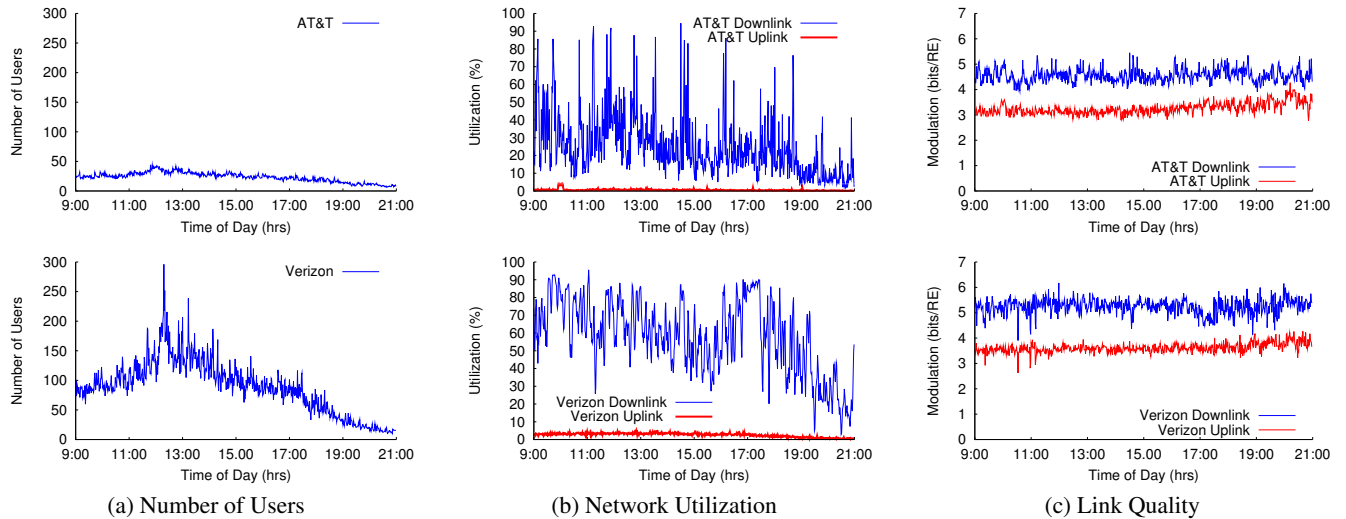


Figure 7: The plots show the following metrics for AT&T (above) and Verizon (below), measured every minute over a typical working day from a representative base station: (a) Number of Active Users per second; (b) Percentage of Utilized Resource Elements on the Uplink (red) and Downlink (blue); (c) Mean Number of Bits per Resource Element on the Uplink (red) and Downlink (blue).




Figure 8: **PHY-Layer Inefficiencies:** (a) Measures the utilization of the control channel for AT&T and Verizon, across different number of symbols allotted to the control channel. (b) Plots the percentage of data allotted to different downlink resource blocks for AT&T and Verizon. (c) CDF of measured Signal-to-Noise Ratio (SNR) and Signal to Interference Plus Noise Ratio (SINR) at spots of high inter-cell interference.

utilized. Clearly, the overhead of Verizon is significantly larger, despite having only a marginally higher control traffic (7.5%) than AT&T (6.1%). As a result, we estimate that Verizon can gain as much as 10% of additional downlink resources for data, just by adapting its control overhead to control-traffic demand.

Inefficient Resource Allocation. In this experiment, we analyze the downlink resource allocation mechanism of LTE base-stations during periods of high network utilization (over 80%). Fig. 8(b) plots the percentage of downlink data transmitted in each resource block across users, measured for both AT&T and Verizon. For AT&T, the graph remains flat across resource blocks, indicating that on average, a user is equally likely to get any of the downlink resource blocks. For Verizon however, we observe a peculiar dip around resource blocks 22-27. To investigate this, we noticed that Verizon avoids these resource blocks completely on subframes 0 and 5. This is because a few symbols in these resource blocks are dedicated for the broadcast channel. As a result, Verizon avoids allocating these resource blocks completely, leading the remaining symbols in these resource blocks to lie completely unused, even during peak hours of demand.

Inter-Cell Interference. One of the key benefits of LTEye is to provide insight into why users obtain poor performance. During the course of our experiments, LTEye localized users at certain spots that achieved poor link quality on the downlink (the lowest QPSK

rate), but high quality on the uplink. To investigate this, we moved our LTEye sniffers and testbed mobile phones to these locations. Surprisingly, our phones at these spots reported very high RSSI⁹ from the base station, yet often switched to 3G. We then used our LTEye sniffers at these spots to measure the signal-to-noise ratio (SNR) as well as the signal-to-interference plus noise ratio (SINR) on the downlink. Fig. 8(c) reports the CDF of these quantities across these locations. The figure demonstrates that these locations suffer from significant interference, with a mean SINR of 1.3 dB and a minimum of -1.2 dB, despite a high mean SNR of 29 dB. We realized the source of about 27 dB of interference is from neighboring Verizon base stations sharing same downlink spectrum. Specifically, our sniffer could sense as many as five distinct base station cell IDs at a single location. This is problematic as it affects pilot reference signals (known as cell-specific reference signals) that are critical for channel estimation. Base stations transmit these pilots in one of three different subsets of resource elements depending on their cell ID.¹⁰ Given that five base stations are observed at a given location, some of these pilots will inevitably collide, significantly impacting the decodability of signals from those base-

⁹The phone reported a receive signal strength (RSSI) of around -85 to -95 dBm, where the noise floor is at -120 dBm.

¹⁰For 2-antenna base stations, reference signals on both antennas occupy one of three subsets of REs, based on cell ID modulo 3 [3].




Figure 9: **Insights into Rate Adaptation:** (a) Trace of Downlink MCS for a user over time; (b) Channel Coherence metric of a USRP placed at the user’s location over time; (c) Mean downlink MCS (bits/resource element) across demand, with and without another high-demand user.

stations. Hence, these observations emphasize the need for effective placement and power control of base stations and small cells. They further highlight the importance of careful assignment of cell IDs to neighboring cells, to avoid interference between their pilots.

8.3 Insights into LTE Rate Adaptation

While the LTE standard describes much of the PHY-layer protocol and procedures, the rate adaptation algorithm is still left to the choice of individual operators. In this section, we show how LTEye can shed light on some interesting aspects of this algorithm for an AT&T base-station in our locality.

We consider a single user device that downloads a random bitfile over a 1 Mbps TCP link from a server (we control TCP rate by throttling the bandwidth at the server using `tc[1]`). We conduct our experiment during periods of low network utilization, where the mobile device in our testbed is the sole active user of the network. The user device is placed at a static high SNR location, which should support the highest modulation and coding on the downlink.

Fig. 9(a) plots the user’s modulation and coding scheme for one second on the downlink. Surprisingly, the graphs indicate that the rate adaptation algorithm hops over a wide range of modulation, during the experiment. One possible explanation is that the base station is responding to loss of downlink packets. However, the control channel indicated no packet loss on the downlink over the entire experiment. A second explanation is that the wireless channel is not actually coherent over the duration of one second, even though the phone is static. To investigate this, we place a USRP at the user’s location and estimate a channel coherence metric¹¹ capturing the base station’s SNR over two seconds, assuming the channel was estimated only once at time 0 (see Fig. 9(b)). We observe that the channel is indeed coherent throughout the experiment. Hence, the rate adaptation algorithm is fairly complex, involving aggressive modulation exploration.

Next, we repeat the above experiment for different downlink demands (i.e. TCP throughput) under identical SNR, and plot the mean downlink modulation (in bits per resource element) as shown by the blue line in Fig. 9(c). Interestingly, as the downlink demand increases, the observed modulation also increases as well. In other words, the base station avoids transmitting packets to the user at high modulation (i.e. avoids risking higher loss probability), unless it is forced to, since the user demands high throughput.

Finally, we repeat the above experiment, this time adding a second user device (User-2) to the network, while first user (User-1) downloads a file at different TCP throughputs, as before. We allow User-2 to download a large bitfile containing random strings via six simultaneous TCP connections so as to saturate the LTE down-

link demand. Interestingly, we now observe that packets to User-1 are sent at higher modulation, across demands (see the red line in Fig. 9(c)). To understand why this is the case, note that by sending data to User-1 at higher modulation, the base station consumes less downlink resources per bit for User-1’s data. This relieves more network resources that the base station can now assign to User-2 to serve its high demand. Therefore, base stations transmit packets at conservative modulation, only when this does not impact overall network throughput.

8.4 Detecting Cheaters in a Large Exam Hall

LTEye can enable new applications customized to the need of a particular community of users. For example, many modern exams follow an open book/material policy and allow students access to computers during the exams. However, students are asked to abstain from using the Internet to chat and collude online. Enforcing this policy over Wi-Fi is relatively easy by monitoring the Wi-Fi channels, turning the access point off, or even jamming the signal. However cheaters can still use their cellular service to chat with an accomplice. In this experiment, we demonstrate how LTEye’s spatial analytics can help localize such cheaters in a large exam hall that accommodates up to 300 students.

Setup. We consider a large $24\text{m} \times 17\text{m}$ exam hall as shown in Fig. 10(a) that seats up to 300 students. The exam hall has multiple chairs on a platform that slopes upwards from the podium. We place two LTEye sniffers on ledges close to the walls, as shown by the blue squares in the figure. The sniffers localize the 3-D location of ten LTE cellphones, corresponding to ten active cheaters accessing the Internet, placed among twenty randomly chosen locations (the red circles in the figure).

Results. Fig. 10(b) plots a CDF of the error in each dimension of the estimated 3D-location of each cellphone. We observe a mean error in localization of 34 cm along each dimension and 61 cm in 3D displacement between the measured and actual location. Note that our errors are in 3-D space unlike past work [30, 29] and the experiments were performed in a large $24\text{m} \times 17\text{m}$ area. LTEye identified the cheater’s seat among the 300 seats in the room with 95% accuracy. Note that a random guess has an accuracy of just $1/300$. Even when LTEye makes an error, it reports a seat adjacent to the cheater, which is sufficient to visually identify the cheater in most proctored situations. Note that the cellphones are predominantly in line-of-sight to the LTEye sniffers, due to the nature of the exam hall. In §8.5, we estimate the error in localization for non-line-of-sight scenarios as well.

8.5 Visualizing LTE Performance over Space

As end-users, we have limited visibility into how LTE performance varies in different parts of our home or work place. In this

¹¹Given an initial channel $h(0)$ and current CFO-adjusted channel $h(t)$, channel coherence metric is $10 \log_{10} |h(t)|^2 / |h(t) - h(0)|^2$

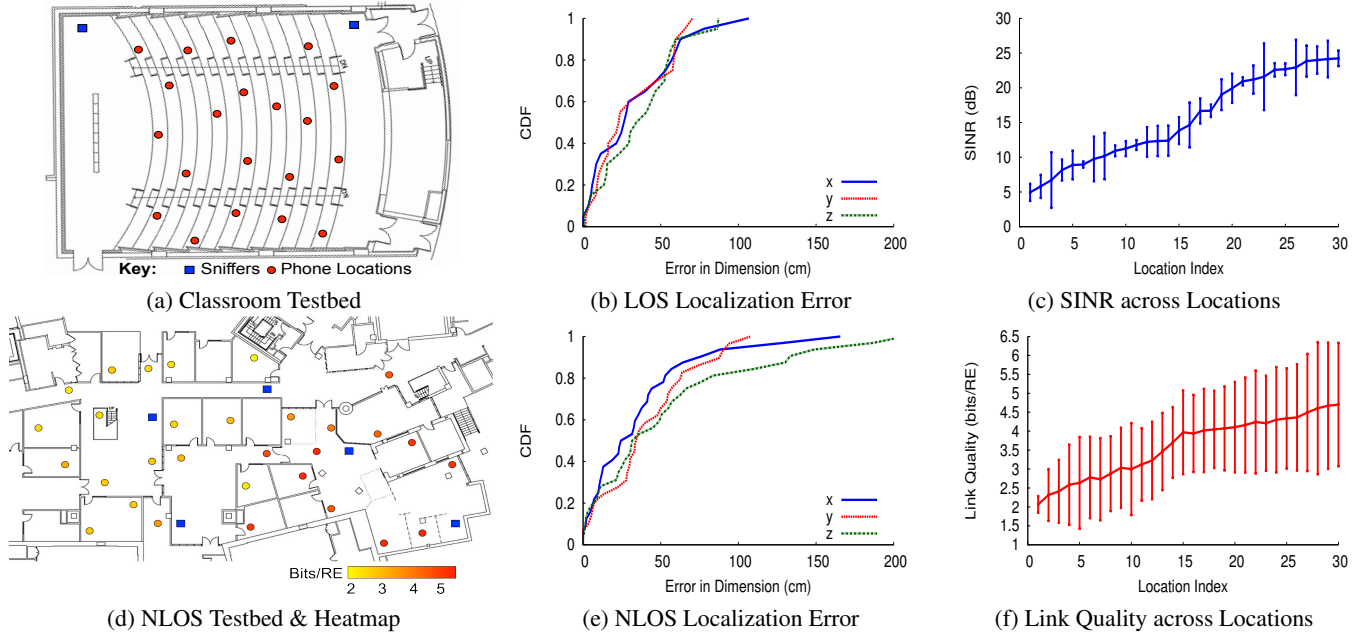


Figure 10: (a),(d): Depicts our two testbeds a large exam hall and a floor of a large building. LTEye sniffers are denoted as blue squares and candidate phone locations as red circles. Further, the red circles in (d) are colored in shades of red, based on observed link quality from the base station; (b),(e): CDF of the error in 3D localization on each dimension in line-of-sight and non-line-of-sight scenarios respectively; (c),(f): Plots the measured SINR using USRPs and observed link quality from the downlink control channel across phone locations.




Figure 11: **Accuracy of θ and ϕ** : Plots CDF of error in measured spatial angles in line-of-sight (LOS) and non-line-of-sight (NLOS).

experiment, we address this issue by synergizing LTEye’s temporal and spatial analytics to visualize the performance of a cellular provider across spatial locations.

Setup. We deploy five LTEye sniffers on a 60m×34m floor of a large building, denoted by the blue squares in Fig. 10(d). We place ten phones in each of thirty randomly chosen locations shown as red circles in the figure. Note that several phones are in non-line-of-sight relative to all LTEye sniffers. We emphasize that for each client device, 3D-localization is performed using the spatial angles from at most three LTEye sniffers. The localization error can be further improved by incorporating spatial angles from additional LTEye sniffers.

Results. Fig. 10(e) plots the CDF of localization error along each of the three dimensions for phones that are in non-line-of-sight relative to all LTEye sniffers. Our results show a mean error in localization of 43.7 cm along each dimension and 84.6 cm in net 3-D displacement. Our algorithm to identify the direct line-of-sight path from §6.2 is crucial to localize phones in non-line-of-sight. Of course, the algorithm hinges on the fact that the line of sight path is, at the very least, observable in the multipath profile produced by SAR (See §6.1), even if it is not the most dominant path.

Our experiments revealed that the line-of-sight path was always observed in the multipath profile of every phone in our large indoor testbed, including those furthest away from each LTEye sniffer in Fig. 10(d). Of course, while this may not generalize to every environment, our observations show the benefits of better penetration of signals in the 700 MHz frequency range through walls and obstacles, compared to Wi-Fi signals at 2.4 GHz or 5 GHz, and the higher transmit power of LTE devices in general.

Fig. 10(c) measures the mean and variance of Signal to Interference plus Noise Ratio (SINR) observed by a USRP placed in each of the thirty phone locations. The locations are sorted by mean SINR for ease of visualization. Fig. 10(f) plots the mean and variance of link quality for each phone (as bits per resource element) at the same locations, measured from LTEyeDB based on the downlink control channels. We observe that the link quality and SNR follow similar trends, showing that LTEye can effectively characterize the performance of the LTE network across spatial locations.

Fig. 10(d) visualizes the spatial distribution of link quality across phone locations, denoting positions of high quality as circles with darker shades of red, and low quality with lighter shades of red. The figure indicates that the link quality is strongest at locations to the

bottom right, and weakest along locations to the top left. In fact, we found that placing an LTE relay at the top-left part of the floor significantly improves LTE service across the floor.

8.6 Measuring Accuracy of Observed Spatial Angles

In this experiment, we measure the accuracy of the polar angle θ and azimuthal angle ϕ , that are key primitives to LTEye's localization algorithm in §6, across spatial locations.

Setup. We consider an LTEye sniffer placed in one of five possible locations (denoted by blue squares) in a floor of a large building, as in Fig. 10(d). The LTEye sniffers are elevated on ledges close to the walls.¹² We place ten phones in several randomly chosen locations in both line-of-sight and non-line-of-sight, spanning the full range of spatial angles. We find the ground truth of the spatial angles by noting the actual 3D positions of the phones and the sniffer on a scaled high-resolution building floorplan.

Results. Fig. 11(a)-(d) plot the CDF of error in θ and ϕ in line-of-sight (LOS) and non-line-of-sight (NLOS) locations. The figures show a low median error in both ϕ (LOS: 6.9°, NLOS: 7.8°) and θ (LOS: 7.2°, NLOS: 9.9°) across locations. Note that the accuracy can be further improved with multiple LTEye sniffers, particularly, in cases where the deviation in angles is large. Note that our algorithm to find the path of minimum delay was crucial for the accuracy of spatial angles in non-line-of-sight.

8.7 Tracking C-RNTIs using RF Fingerprints

In this section, we evaluate LTEye's RF fingerprinting to map C-RNTIs assigned to the same phone. We consider the setup in §8.6 above and populate LTEyeDB across several experiments spanning ten minutes each. We track the correct C-RNTI mapping by constantly listening on the uplink from multiple USRPs placed close to each phone to recognize the high power signals that are sent during connection establishment. We also measure the inferred C-RNTI mapping from RF-fingerprints (See §5 and §6.3).

Results. We measure two quantities: (1) Precision: The percentage of new C-RNTIs which were correctly mapped to old C-RNTIs. (2) Recall: The percentage of correctly retrieved C-RNTIs-mappings among all actual C-RNTI mappings. Our algorithm achieves a high mean precision of $98.4 \pm 1.3\%$ and mean recall of $96.7 \pm 1.4\%$, demonstrating the effectiveness of LTEye's C-RNTI matching algorithm. Note that we leveraged RF fingerprints to track C-RNTIs of users in §8.4 and §8.5 above.

9. CONCLUSION

We presented LTEye, the first open platform to provide fine-grained temporal and spatial analytics on LTE radio performance, without private user information or provider support. LTEye employs a novel extension of synthetic aperture radar to communication signals to accurately localize mobile users, despite the presence of multipath. We empirically evaluate LTEye on software radios and provide deep insights on the LTE PHY and highlight shortcomings such as inter-cell interference and inefficient spectrum utilization.

Acknowledgments: We are grateful to John Romanishin for fabricating the rotating arm used in our experiments. We thank Vimalkumar Jeyakumar, Srinivas N.G., Arthur Berger, the NETMIT group and our reviewers for their insightful comments. This work is funded by NSF. We thank members of the MIT Center for Wireless Networks and Mobile Computing: Amazon, Cisco, Google, Intel, Mediatek, Microsoft, ST Microelectronics and Telefonica for their interest and support.

¹²LTEye cannot tell apart up from down as the rotating antenna path is symmetric. Placing sniffers on ledges removes this ambiguity.

10. REFERENCES

- [1] Linux Advanced Routing and Traffic Control. <http://lartc.org>.
- [2] 3gpp. Radio Measurement Collection for Minimization of Drive Tests (MDT). http://www.etsi.org/deliver/etsi_ts/137300_137399/137320/11.03.00_60/ts_137320v110300p.pdf.
- [3] 3rd Generation Partnership Project. E-UTRA, Physical Channels and Modulation, TS 36.211, v8.9.0. 2010.
- [4] 3rd Generation Partnership Project. E-UTRA, Radio Resource Control, TS 36.331, v8.9.0. Apr 2010.
- [5] 3rd Generation Partnership Project. E-UTRA, v8.8.0, Multiplexing and Channel Coding, TS 36.212. 2010.
- [6] B. Ayvazian. LTE Deployment Strategies: Network Overlay vs. Single RAN. *White Paper: UBM Tech*, 2013.
- [7] B. Conley. M2M Implications of the 4G LTE Buildout. *Remote Magazine*, 2013.
- [8] USRP N210. <http://www.ettus.com>. Ettus Inc.
- [9] Eurecom. OpenAirInterface. <http://www.openairinterface.org/>.
- [10] Federal Communications Commission. Enabling Innovative Small Cell Use In 3.5 GHz Band NPRM & Order. FCC 12-148, 2012.
- [11] R. Fenby. Limitations on Directional Patterns of Phase-Compensated Circular Arrays. *Radio and Electronic Engineer*, 1965.
- [12] P. J. Fitch. *Synthetic Aperture Radar*. Springer-Verlag, 1988.
- [13] FlexNets Group. Open-Source Long-Term Evolution (LTE) Deployment (OSLD). <https://sites.google.com/site/osldproject/>.
- [14] C. Hsieh-Chung et al. Determining RF Angle of Arrival using COTS Antenna Arrays: A Field Evaluation. In *MILCOM*, 2012.
- [15] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. *MobiSys*, 2012.
- [16] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. *SIGCOMM*, 2013.
- [17] Huawei. SingleSON Whitepaper. <http://www.huawei.com/en/static/SingleSON-75714-1-127574.pdf>, 2012.
- [18] K. Joshi, S. Hong, and S. Katti. PinPoint: Localizing Interfering Radios. *NSDI*, 2013.
- [19] Y. Kakishima, T. Kawamura, Y. Kishiyama, H. Taoka, and T. Nakamura. Experimental Evaluation on Throughput Performance of Asymmetric Carrier Aggregation in LTE-Advanced. In *VTC Spring*, 2011.
- [20] H. W. Kuhn. The Hungarian Method for the Assignment Problem. *NRL Quarterly*, 1955.
- [21] C. Laurendeau and M. Barbeau. Centroid Localization of Uncooperative Nodes in Wireless Networks Using a Relative Span Weighting Method. *EURASIP Journal on Wireless Communications and Networking*, 2010.
- [22] V. Loscri et al. Carrier Independent Localization Techniques for GSM Terminals. In *PIMRC*, 2008.
- [23] S. J. Orfanidis. Electromagnetic Waves and Antennas. <http://www.ece.rutgers.edu/orfanidi/ewa/>.
- [24] S. Salvador and P. Chan. Toward Accurate Dynamic Time Warping in Linear Time and Space. *IDA*, 2007.
- [25] Sanjole. The WaveJudge Wireless Test System. <http://www.sanjole.com/our-products/wavejudge-test-system/>.
- [26] M. Solutions. White Paper: Femtocells - The Gateway to the Home. 2010.
- [27] J. Sommers and P. Barford. Cell vs. WiFi: On the Performance of Metro Area Mobile Connections. *IMC*, 2012.
- [28] ThinkRF. Product Concept Brief for Distributed and Remote LTE Analysis. <http://thinkrfrf.com/>.
- [29] J. Wang and D. Katabi. Dude, Where's My Card?: RFID Positioning That Works with Multipath and Non-line of Sight. *SIGCOMM*, 2013.
- [30] J. Xiong and K. Jamieson. ArrayTrack: A Fine-grained Indoor Location System. *NSDI*, 2013.
- [31] J. Xiong and K. Jamieson. SecureArray: Improving WiFi Security with Fine-grained Physical-layer Information. *MobiCom*, 2013.