Ayon Chakraborty

Eugene Chai

Karthikeyan Sundaresan Sampath Rangarajan NEC Laboratories America, Inc.

Amir Khojastepour

ABSTRACT

We envision a flexible, dynamic airborne LTE infrastructure built upon Unmanned Autonomous Vehicles (UAVs) that will provide on-demand, on-time, network access, anywhere. In this paper, we design, implement and evaluate SKYRAN, a self-organizing UAV-based LTE RAN (Radio Access Network) that is a key component of this UAV LTE infrastructure network. SKyRAN determines the UAV's operating position in 3D airspace so as to optimize connectivity to all the UEs on the ground. It realizes this by overcoming various challenges in constructing and maintaining radio environment maps to UEs that guide the UAV's position in real-time. SKyRAN is designed to be *scalable* in that it can be quickly deployed to provide efficient connectivity even over a larger area. It is adaptive in that it reacts to changes in the terrain and UE mobility, to maximize LTE coverage performance while minimizing operating overhead. We implement SkyRAN on a DJI Matrice 600 Pro drone and evaluate it over a 90 000 m² operating area. Our testbed results indicate that SKyRAN can place the UAV in the optimal location with about 30 secs of a measurement flight. On an average, SKyRAN achieves a throughput of $0.9 - 0.95 \times$ of optimal, which is about $1.5 - 2 \times$ over other popular baseline schemes.

KEYWORDS

LTE, RAN, UAV, 5G, Localization, Radio Environment Map, Tactical Communications

INTRODUCTION 1

We envision an autonomous, cellular network built upon untethered, low-altitude Unmanned Aerial Vehicles (UAVs) that are rapidly deployable to provide uninterrupted, on-demand LTE network access, anywhere. This UAV network can be flown into desired regions and be deployed either in conjunction with existing fixed infrastructure to augment capacity and coverage (e.g., surge in traffic demands, stadiums, concerts) or independently into remote or challenging environments, such as rural locations, mountainous terrains or disaster areas, to provide temporary connectivity in areas that will otherwise be unreachable by fixed infrastructure. The operational endurance of each such LTE node is heavily dependent on the UAV platform. This endurance can range from several minutes to

CoNEXT '18. December 4-7. 2018. Heraklion. Greece

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6080-7/18/12...\$15.00

https://doi.org/10.1145/3281411.3281437



Figure 1: Visualization of the UAV positioning problem. 20 UEs are deployed across a 250 m×250 m area in Manhattan, New York (details in §5). The deployment reflects a natural setting where UEs are concentrated in few pockets of locations/roads. Figure (a) shows the average throughput per UE corresponding to the UAV's position in the airspace at a fixed altitude. Note that favorable UAV positions are not abundant. Only \approx 5 % positions result in an average throughput higher than 26 Mbps which is about 52 % higher than the median. This highlights the benefits of optimized UAV positioning.

hours [3] with battery-powered UAVs [4, 11], and up to several days with gasoline-powered UAVs [5].

Mobile operators have started to evaluate the benefits of such UAV-based LTE networks: AT&T [2] deployed a UAV-based LTE base station (BS) that is tethered to a ground station to provide LTE service in the aftermath of hurricane Maria in Puerto Rico [7]; Verizon [12] has also trialed and evaluated an LTE BS mounted on a fixed-wing aircraft that is designed for a large-area coverage. However, these are highly specific and limited, early-stage instantiations of UAV-based on-demand LTE platforms. In this paper, we adopt a broader, more unified view of this problem, and ask the question: How can we design and build a cost-effective, autonomous UAV RAN solution that can be rapidly deployed in any situation to provide on-demand mobile connectivity infrastructure with complete flexibility to track and serve dynamic connectivity demands?

Need for UAV Positioning. The goal for each UAV in such a network is to position itself in an appropriate location, in 3D space, that will offer optimal performance¹ to all UEs within the area of operation. Fig. 1 demonstrates the importance of careful UAV positioning in delivering substantial performance gains. While offline approaches (path-loss model based) to predict UAV-UE channel or those leveraging the UE locations to directly place the UAV, can be employed, we demonstrate their inability to capture practical,

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 ACM 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.

¹Optimality can be defined as a function of throughput, SNR, outage, etc. of a desired set of clients/UEs jointly; e.g. maximizing their min or weighted throughput.

heterogeneous terrains, as well as adapt to UE dynamics in §2.2. Hence, realizing this objective in practice requires the UAV to probe and profile the radio channel between itself (at a multitude of locations in 3D air-space) and every UE location on the ground (i.e. construct a UE-specific RF map) in *real-time*, before solving a joint optimization objective across all UEs to determine its optimal position. Solutions today largely focus on the optimization problem, while assuming that all the required RF information is known and UEs being static (or their locations known) [15, 26]. Indeed, the bigger challenge facing practical deployment of these networks, lies in the accurate estimation of real-time information (RF maps) critical to determining the optimal UAV position even in the face of UE dynamics. This in turn forms the focus of this work.

Challenges. The finer the granularity of the air-space from which the RF channel is probed (longer probing duration/overhead), the greater is the accuracy of the estimated RF maps and the resulting optimality of the UAV position. However, the UAV's motion during probing can cause the path loss (Tx-Rx power in dB) to the UEs to vary rapidly and significantly (over 20 dB in our experiments, Fig. 7), resulting in a highly sub-optimal connectivity for the UEs. Hence, it is critical for the UAV to minimize its probing overhead (time in motion) and maximize the time spent in delivering optimized coverage to its UEs from a stationary position. Tackling this tradeoff requires us to carefully address three key challenges:

(1) Scalable Radio Environment Map (REM) Estimation. The larger the coverage area $(O(N^2), N$ being points in 1D) of the ground terrain, the larger is the corresponding airspace $(O(N^3))$ to be probed. Further, the higher the complexity of the terrain (buildings, foliage, etc.), the higher the heterogeneity (variations and features) in the REM (i.e. RF map). Extrapolating simple path loss models is no longer sufficient to capture such heterogeneity, and requires a finer granularity of probing (large number of UAV probing positions) in the 3D space, adding to more overhead.

(2) Adapting to UE Dynamics. Note that REMs are estimated with respect to specific UE locations. Hence, appreciable UE mobility can contribute to changes in their REM and in turn to the optimal UAV position. However, frequent repositioning of the UAV to cater to individual UE mobility would incur significant overhead and might not be practical. On the other hand, being overly conservative about repositioning the UAV might lead to degraded performance. Thus, it is important to strike a balance between frequent UAV position updates and deteriorated connectivity.

(3) UE Location-awareness. Given that a REM corresponds to a fixed UE location on ground, knowledge about the UE location will facilitate reuse of previously gathered data or estimated REM (e.g., for a UE that moves to a location or neighborhood, where REM is already available). The lack of UE location information at the UAV, leaves little room for balancing the tradeoff between probing overhead and estimation accuracy. Our experiments show that a UE localization error of about 10 m or more can have a significant impact on UAV positioning and hence system performance (§2.4, fig. 9). Existing UE localization algorithms [16, 28, 34, 36] that are designed for static BSs (macro-cells) deliver localization accuracies in the order of several tens to hundred meters. This is not sufficient, especially for small cells that cater to 1 Km² area or less. Further,

the noisy measurements resulting from the UAV's mobility makes it challenging to obtain the desired localization accuracy.

SKYRAN. Towards addressing these challenges, we present SKYRAN– a first-of-its-kind system that automates and optimizes the entire process of UAV-driven LTE RAN deployments, namely REM construction and UAV placement. SKYRAN adopts a *measurement-based* approach to construct and maintain REMs: it leverages channel data collected from measurement flights for UE localization and an REM estimation algorithm that is both scalable to large coverage areas and adaptive to UE dynamics. Briefly, SKYRAN's operation involves three key design elements (seen in Fig. 2):

(1) UAV-Optimized UE Localization. SKYRAN leverages the mobility of the UAV to create a synthetic aperture array, and the synchronous nature of LTE transmissions to range the UE organically (using just LTE) from this array (multiple locations) and eventually localize it. At the start of an epoch, it executes a short random flight trajectory during which it records LTE's PHY-layer Synchronization Reference Signals (SRS) to each UE. From the latter, it infers the respective signal time-of-flight (ToF) and hence range to each UE. Such signals being LTE standards compliant, are supported on all LTE UE devices. To derive the UE's location, the UAV's GPS and ToF data are then used in a multilateration algorithm that is robust to measurement inaccuracies arising from terrain obstacles and UAV mobility.

(2) Spatial Filtering for Scalability. Armed with the UE locations, SKyRAN then computes and executes an intelligent flight trajectory for probing all the UEs simultaneously. The trajectory is computed by identifying unexplored regions in the airspace with high signal gradients (variations) and evaluating their contribution to increased REM accuracies for the UEs against the added cost to probe them. This allows SkyRAN to prioritize its probing overhead for spatial regions that have a larger impact on the estimation of REM. This provides scalability to larger areas of operation than would be otherwise possible with exhaustive measurement approaches. The RF data from probing is then interpolated to obtain an REM for each UE, which is finally used to determine the UAV's optimal position. (3) Temporal Aggregation for Adaptability. Instead of reacting to individual UE mobility, SKyRAN triggers a new epoch only when when collective UE dynamics (mobility of multiple UEs) impacts the aggregate system performance (by a configured margin). Thereupon, it executes its probing step to refine and update its REMs as needed. Further, REMs constructed for locations in prior epochs are reused when UEs visit these locations or their neighborhood in future. Using such prior information, SKYRAN is able to adapt and optimize its probing flight trajectory to further minimize overhead even in the face of UE dynamics.

We realize SKYRAN with a custom designed payload consisting of two Small Board Computers (SBCs), one to execute the software EPC and the other to run an OpenAirInterface eNodeB, a USRP B210 and an LTE antenna, a LTE power amplifier and duplexer. The payload is mounted on a DJI M600Pro drone, which serves as our UAV. Our onboard custom flight control software, built on top of the DJI OnBoard SDK, enables SKYRAN to operate autonomously on the drone.

In our real-world flight experiments, given a fixed overhead (distance flown), SKYRAN reconstructs an REM of within 2 dB of

SkyRAN: A Self-Organizing LTE RAN in the Sky



Figure 2: Probing the 3D space to obtain best operating point. SKyRAN chooses an optimal height for the UAV to operate in and constructs an optimal measurement trajectory in the 2D plane to estimate the REMs for a set of UE locations. Such REMs are stored and historical data is used in case UEs reappear in similar locations where REMs have been estimated priorly.

the optimal, which helps SKYRAN optimize the UAV's position. On an average, SkyRAN achieves a throughput of 0.9 - 0.95× of optimal, which is about 2× that of a baseline scheme employing a uniform search trajectory for probing, and about 1.5× over non-REM based approaches that employ just the UE locations.

A demo of SkyRAN is available at http://www.nec-labs.com/ skyran. We also supplement our evaluation with large-scale tracedriven simulations that capture more UEs, heterogeneous terrains, UE dynamics, etc. by employing real-world LIDAR datasets for accurate terrain construction.

BACKGROUND AND MOTIVATION 2

In this section, we discuss the deployment of the UAV-based LTE network along with the challenges and design principles chosen to address them.

Deployment Model 2.1

We envision a holistic UAV LTE infrastructure network consisting of UAV RANs that can either operate independently of other fixed networks, deployed in hard-to-reach terrain that is outside the range of existing fixed infrastructure, or used as hotspots to augment the capacity of fixed networks for targeted areas such as sports stadiums, concert halls, parades or in response to local, high attendance events.

Each SkyRAN UAV accomplishes this task of autonomous coverage operation using its three components (as shown in Fig. 13): (a) an LTE Evolved Node B (eNodeB), which directly connects to UEs on the ground; (b) an Evolved Packet Core (EPC) that provides LTE core functionality (i.e. UE authentication and registration, routing, RAN state management etc) to the UEs in the UAV's area of operations; and (c) a flight control core that provides autonomous flight capability, and runs all of the SkyRAN algorithms.

The focus of our work is the design, implementation and evaluation of a single SKYRAN UAV in this UAV LTE network. Our design is inherently scalable and can be extended to operate in a multi-UAV scenario, as briefly discussed in §8.





Figure 3: Positioning based Figure 4: Estimated RF map only on UE locations leads to error w.r.t. ground truth suboptimal placement.

ground-truth RF map.

from a exhaustive UAV probing flights over four different terrains.



Figure 5: Different flight Figure 6: RF map error detrajectories overlaid on a creases as increasing proportion of the area of operations is probed.

20 30 40 50

2.2 A Measurement-Based Approach

A measurement-based approach is necessary for a SKYRAN UAV to capture an accurate picture of the radio-frequency characteristics. Without measurements, the UAV can only choose an operating position either randomly, or with respect to the UE locations. While, random UAV positioning offers no guarantee on performance, a UAV positioning that is based solely on UE locations can also be highly sub-optimal. As an example, Fig. 3 shows a real-world RF map with the average RF signal strength to all three UEs from every point in the operating area. This RF map is obtained from a UAV flying in an exhaustive measurement trajectory over the testing area. A UAV operating at the centroid of these three UEs will achieve \approx 30 – 50% lower throughput (see Fig. 21 in §4.5 for further details) than a UAV at the optimal position derived from the RF map measurement. The degradation is more pronounced in complex terrains (i.e. with more natural or man-made obstructions), highlighting the limitations of a geographical approach (based on UE locations) to UAV positioning. Hence, there is a need to characterize the radio environment through a measurement-based approach.

Use of Radio Environment Maps 2.3

SkyRAN adopts a measurement-based approach to REM construction, which in turn determines the optimal operating position of the UAV. It can be argued that there are simpler alternatives to radio environment maps (REMs): one can either construct the REM using a free-space path-loss model without measurements, or use a throughput map instead of a REM. When compared to throughput maps, REMs offer a lower-level, higher fidelity view of the actual



Figure 7: Variation in Figure 8: Reducing operatpathloss from the UAV ing altitude decreases path-LTE eNodeB to a UE loss. during flight.



Figure 9: Localization accuracy impacts the eventual UAV positioning in SKYRAN that affects aggregate performance. Note that SKYRAN achieves a relative throughput of about $0.9 - 0.95 \times$ when the localization accuracy is within 5 m. Performance loss of about 10% is observed with 10 m error that goes beyond 50% with 20 m or more error.

channel conditions between the UE and UAV (i.e. without incorporating MAC-layer artifacts like rate adaptation), and enables the UE to better approximate the *true* RF characteristics of its operating environment. Furthermore, the cost associated with constructing a measurement-based throughput map is similar to that of REMs.

Similarly, simple pathloss models (e.g. free space path loss) cannot capture the terrain variations that have a significant impact on the actual channel characteristics, while detailed ray-tracing model [18], requires fine-grained topology information (i.e. LiDAR terrain scans) which are not easily available. As an example to demonstrate the advantage of REMs, we consider four different real-world terrains (details in §3.5), each with 3 UEs. In each terrain environment, we obtain the fine-grained ground truth RF characteristics by exhaustively measuring the RSS between each point in space and the UEs on the ground via a detailed flight trajectory. We then compare two different approaches to approximate this ground truth RSS data: (a) with a REM constructed using a SKYRAN UAV, and (b) an approximate path-loss map using free-space path-loss models given the UE locations. Fig. 4 shows that the error of the path-loss map, w.r.t. ground truth, is up to 4× greater than that of the data-driven map (10 vs. 4dB in Terrain-4). The SkyRAN UAV, using its REM, selects a better operating position with higher overall SNR and throughput.

The construction of REMs is also inherently distributed in nature, and can be easily scaled to operate in multi-UAV networks in a future extension of SKYRAN.

2.4 Leveraging UE Locations

To understand the importance of UE location, first consider the ground-truth REM, collected via exhaustive measurements, as shown in Fig. 5. Two different flight trajectories are shown here, (a) an

exhaustive search path that begins at one corner and systematically explores across the measurement area, and (b) another that prioritizes its measurement trajectory towards the location of the UEs.

Fig. 6 shows the median error of the REM constructed by these two flight trajectories, as a function of the proportion of the total area explored. Observe that a UE location-aware trajectory that prioritizes measurements towards areas most affected by UE activity returns useful RF information at a higher rate, and thus achieves a more accurate REM in a shorter time: with only 15% of the area measured, the REM constructed by the location-aware trajectory differs from the optimal by 5dB, while the error of the map constructed by the naive approach is about 12.5× greater at 16dB.

However, existing LTE UE positioning algorithms such as Enhanced Cell-ID [6], Assisted Global Navigation Satellite Systems (A-GNSS) [1], Uplink/Downlink Time Difference of Arrival (TDoA) [17] techniques are designed for static eNodeBs and assume features such as clock-synchronization across macro cells that are not suitable for SkvRAN UAV RANs. Further, RF fingerprinting techniques without eNB synchronization achieve an accuracy of approximately 50–100 m. This is an order of magnitude greater than the 5–10 m accuracy desired for accurate REM construction, the lack of which affects optimized UAV positioning and hence performance by over 50% (see Fig. 9).

2.5 Efficient REM Construction and Updates

SKYRAN uses a measurement-based REM construction technique. It also periodically repeats part (or all) of the measurements to refresh the map in response to appreciable UE mobility and environmental changes.

Update Overhead Scales with Operating Area Size. The larger the operating area, the larger the 3D space that the UAV has to probe in order to estimate the REM. Given that real-world operating areas are expected to be large (up to several 10s of Km², it is imperative that SkyRAN algorithm must readily scale to support larger operating environments.

Suboptimal LTE Performance During Probing. The REM estimation duration has to be limited because the performance of the LTE network degrades during channel measurement. For example, Fig. 7 shows that during a 50m flight segment, the pathloss between the UAV and UE on the ground can vary from 77dB to 95dB. High LTE overhead, in the form of PHY layer signaling (e.g. Channel Quality Indicator messages), is needed for LTE to track this rapidly changing channel, thus reducing the throughput over the network. **Limited UAV Battery Power.** Our DJI M600Pro platform consumes more battery during forward motion, and we expect that other UAV platforms with similar designs and operating procedures may have similar power-drain characteristics. Hence, the shorter the duration of the measurement flight, the longer the UAV LTE endurance when providing LTE service.

UE Dynamics. Appreciable UE mobility leads to changes in REM and in turn to the optimal UAV position. While frequent repositioning of the UAV to cater to individual UE mobility incurs significant overhead, not being adaptive can lead to degraded performance. Thus, SKYRAN needs to balance adaptability and overhead in the presence of UE dynamics.

3 SKYRAN: DESIGN

SKYRAN is designed for *scalable* and *adaptable* UAV RAN operation. In this section, we give an overview of the design and operational procedures of SKYRAN, followed by a description of the algorithms employed.

3.1 SKYRAN Overview

Epoch-based Operation. SKYRAN operation follows multiple *time epochs.* Fig. 10 illustrates the operation of a SKYRAN UAV in a single time epoch, as it is providing LTE coverage to an area.

SKYRAN Operation. The SKYRAN UAV is launched with the boundaries of the operation area as its sole parameter. In each time epoch, the UAV first undertakes a **UE localization** flight, which is a short random flight trajectory within the operation area. This flight is used to collect LTE SRS data from the UEs so that their location can be determined by the SKYRAN UAV. Once the UE positions are found, the SKYRAN UAV computes and executes an optimal measurement flight trajectory to **estimate the REM**. The SKYRAN UAV then **computes the optimal position** from this REM, and positions itself at this position to **provide LTE service**. This position is retained until a new epoch is dynamically triggered in response to appreciable UE dynamics (discussed in Section §3.5), upon which, the sequence of operations in Fig. 10 is repeated.

3.2 UE Localization

SKYRAN uses UE location information to plan flight trajectories that maximize the accuracy of the REMs. UE locations allow SKYRAN to coordinate measurement flights across large areas to best cover the UEs at known positions (scalability), and monitor UEs for mobility to guide update and future measurement flights to keep the REMs updated (adaptability). It is thus a key feature in SKYRAN.

Localization Overview. SKYRAN uses multilateration in principle to determine the location of each UE in the operating area. We note that all of these UEs are associated with the LTE eNodeB on the SKYRAN UAV.

Steps 1–4 of Fig. 10 illustrates the process of UE localization. The SKYRAN UAV moves along a short, random flight trajectory and records the GPS positions of the UAV (Step 1) and the uplink LTE Synchronization Reference Signals (SRS) (Step 2) from each UE. The SKYRAN UAV computes the time-of-flight (ToF) of the signal from the UE to the eNodeB using these SRS data (Step 3). The ToF and UAV GPS data is then used by the multilateration algorithm to determine the UE location.

3.2.1 *Raw Data Collection (Steps 1 and 2).* **UAV GPS Position.** The GPS positions of the UAV are sampled at 50 Hz during the UAV flight. Each position reading is timestamped using the global system clock, so that it can be time-aligned with the SRS data.

LTE SRS Data. The uplink Synchronization Reference Signal (SRS) [9] is a known PHY-layer signal sent from the UE and received at the eNodeB, and is used by the eNodeB to measure the state of the uplink LTE channel from the UE. This signal occupies an LTE OFDM symbol with *N* subcarriers, which we refer to as the *SRS symbol*. The SKYRAN UAV receives uplink frequency-domain SRS symbols from each UE once every 10 ms (or 100 Hz). Each report is timestamped by the eNodeB using same system clock.

3.2.2 *ToF Computation (Step 3).* LTE subframes are transmitted in a time-deterministic manner. Hence, the time-of-flight of the RF signal from the UE to the eNodeB can be inferred from the delay between the time the SRS symbol is transmitted by the UE, to the time it is received at the eNodeB. The UE transmission time is synchronized with the eNodeB, and thus is known at the SKYRAN UAV.

A received SRS symbol with its *N* OFDM sample values can be expressed as $\mathbf{s} = [s^{(1)}, \ldots, s^{(N)}]$. Let the known SRS symbol transmitted by the UE be represented similarly as $\mathbf{h} = [h^{(1)}, \ldots, h^{(N)}]$. Note that \mathbf{s} and \mathbf{h} are frequency-domain values. Using the cross-correlation property of the Discrete Fourier Transform (DFT), we can find the delay of the SRS symbol from the magnitude peak of the sequence

$$ifft(s \odot h^*) = ifft([s^{(1)}h^{*(1)}, \dots, s^{(N)}h^{*(N)}])$$
(1)

where \odot is the element-wise multiplication operator, ifft(·) is the inverse fast fourier transform function, and (·)* is the complex conjugate function.

However, (1) only gives the delay to the nearest time-domain sample offset, which has a low resolution. For example, for a 10 MHz LTE band (that is actually sampled at 15.36 MHz), each sample difference corresponds to a real-world distance of 19.5 m. In order to get a higher delay measurement resolution, we upsample both s_i and h: to get a delay resolution to within 1/K of a sample, we upsample the SRS sequence by K times. Since s_i is a frequency domain SRS symbol, this upsampling process is accomplished by zero-padding the sequence:

$$\hat{\mathbf{s}} = [s^{(1)}, \dots, s^{(N/2-1)}, \\ 0^{(1)}, \dots, 0^{(N(K-1))}, s^{(N/2)}, \dots, s^{(N)}]$$
(2)

where $0^{(1)}, \ldots, 0^{(N(K-1))}$ is a sequence of N(K-1) zeros inserted into the middle of the SRS symbol \mathbf{s}_i . The transmitted SRS symbol \mathbf{h} is upsampled in the same manner to obtain $\hat{\mathbf{h}}$. Let $\hat{\mathbf{y}} = ifft(\hat{\mathbf{s}}_i \odot \hat{\mathbf{h}}^*)$ as in (1), then the delay offset of the received SRS symbol is

$$t = \frac{1}{K} \max pos(\hat{\mathbf{y}}) = \frac{1}{K} \max pos([|y^{(1)}|, \dots, |y^{(N)}|])$$
(3)

where $\max pos(\cdot)$ returns the index of the element in sequence $\hat{\mathbf{y}}$ with the largest absolute value, and $|\cdot|$ is the magnitude operator. Note that the larger the value of K, the lower the SNR of the peak of the correlation². This limits the practical accuracy of this upsampling method for large K. In our implementation, we select K = 4 to strike a good balance and yields good ToF accuracy.

Due to the higher SRS reporting rate, there are multiple ToF values, $t_i^{(1)}, \ldots, t_i^{(M)}$, between any two consecutive UAV GPS reports g_i and g_{i+1} . We average these M ToF values $\bar{t}_i = (\sum_{k=0}^M t_i^{(k)})/M$ and assign the mean value \bar{t}_i to UAV GPS report g_i to generate a stream of GPS-ToF tuples (g_i, \bar{t}_i) at the same UAV GPS reporting rate of 50 Hz. This continuous sequence of GPS-ToF tuples are then used in the multilateration algorithm to resolve the position of the UE.

 $^{^2 {\}rm The}$ magnitude of the inverse FFT is scaled by 1/(KN) while the noise is unchanged

CoNEXT '18, December 4-7, 2018, Heraklion, Greece



Figure 10: Overview of the SKYRAN UAV RAN operation when providing LTE service.

3.2.3 UE Localization (Step 4). The ToF can be mapped to physical distance via multiplication by the propagation distance of the RF signal per sample time. However, ToF processing onboard the UAV incurs a constant processing delay, which manifests as a constant ToF/distance offset that must be eliminated. SKYRAN compensates for this offset by directly incorporating it as an unknown in the equation for ranging distance between the UE and the UAV. Such a series of equations (from different UAV positions) can then be solved for the UE's location by re-casting it into a least-squares formulation (for better robustness to noisy UAV measurements), and adopting a gradient descent based iterative solution. In the interest of space, we omit the mathematical details behind the offset incorporated multi-lateration formulation.

3.3 Estimating Radio Environment Maps

After the UEs have been localized, SKYRAN estimates the Radio Environment Maps (REMs) for the UEs in its coverage area, and uses them to estimate the optimal placement of the UAV in the 3D airspace. However the sheer scale of measurement overhead required, assuming a brute-force approach, to sample the entire 3D airspace is prohibitive. Also such REMs need to be created across UEs. SKYRAN intelligently *prunes* the airspace to filter out *informative* points to conduct measurements in, that enriches the REMs with minimal overhead.

Steps 5–7 of Fig. 10 shows the three key steps involved in the REM estimation: SKYRAN first determines the *optimal operating altitude (Step 5)*. It then *finds the REM measurement trajectory (Step 6)*, and uses the data collected over that trajectory to *update the REMs (Step 7)* for individual UEs. Finally, SKYRAN *estimates the optimal UAV position* and moves to provide LTE service at that position.

Quantizing Space. The UAV GPS coordinates in SKYRAN are observed to have an accuracy of 1 m-5 m. SKYRAN thus quantizes its operating area into $1 \text{ m} \times 1 \text{ m}$ grid cells.

3.3.1 Finding Optimal Altitude (Step 5). A comprehensive RF channel profile requires a REM at every altitude to each UE. Notwithstanding its excessive overhead, REM at different altitudes tend to reveal dependent RF information capturing the underlying terrain characteristics. Hence, SKYRAN strikes a balance between overhead and accurate REM construction – it identifies a target operating altitude, for which it constructs an accurate REM.

It finds this target altitude using a key insight related to RF path loss: *there exists an optimal altitude where UAV-to-UE path loss is minimized*. Increasing the altitude further results in increased path loss (greater distance) whereas decreasing the altitude magnifies the impact of shadowing effects due terrain obstacles (e.g., buildings and trees) increasing path loss. This effect is shown in Fig. 8 that reports UAV-to-UE path loss as a function of the UAV's altitude. We note that while [14] presents an analytical way of determining the optimal altitude of the UAV for maximizing coverage, it requires detailed terrain information unavailable in a practical setup like SKYRAN.

SKYRAN's Approach: In the first epoch, the SKYRAN UAV positions itself directly above the centroid of the locations of the active UEs at an altitude of 120 m (the maximum allowable altitude under FAA regulations). Next the UAV starts decreasing its altitude while tracking the decreasing path loss till the altitude with minimum path loss is found. In all subsequent epochs, SKYRAN estimates the relevant REMs specific to that target altitude. Constructing REMs at this target altitude has shown to yield $0.9-0.95 \times$ of the optimal performance in our experiments. Note that this target altitude is not updated every epoch but only when appreciable network dynamics is detected (discussed in §3.5).

3.3.2 Computing Flight Trajectory (Step 6). A flight trajectory is a path taken by the SKYRAN UAV through its operating area to obtain RF measurements that will feed into REM construction. We quantize this path into points that are 1 m apart. This quantized path is computed over the following four steps:

Step 6.1: Aggregating REMs. A REM for a UE is a 2D grid of $1 \text{ m} \times 1 \text{ m}$ cells in the airspace, covering the operating area of SxyRAN UAV (at the target altitude). Each grid cell shows the signal-to-noise ratio (SNR) from that cell to the UE's location. The *aggregate REM* is the end-result of a grid cell-wise sum of the *per-UE REMs* in the current epoch. Fig. 11 illustrates this aggregation step.

Step 6.2: Gradient Map (Spatial Filtering). SKYRAN computes the SNR *gradient map* from the *aggregate* REM. The gradient of each grid cell is the greatest difference between its SNR and the SNR of its directly adjacent, neighboring cells. The gradient map is thus a 2D map showing the SNR's gradient of each cell within the operating area of the SKYRAN UAV. Cells with higher values of gradients denote areas of higher SNR fluctuations. The goal of the SKYRAN UAV is to bias its measurement efforts to such high gradient grid cells. This enables us to more accurately capture the fine grained variations in SNR (critical for eventual REM estimation), while limiting overhead even in large terrains.

Step 6.3: Location Clusters. SKYRAN partitions the grid cells into high and low gradient cells. The cells with gradients greater than the median of the gradient map are the high-gradient cells, with the other cells considered low-gradient. Measurement efforts are only focused on the high-gradient grid cells. In spite of this partition, it

is still infeasible for a SKYRAN UAV to visit each grid cell. SKYRAN thus applies K-means clustering over the high-gradient cells to spatially group them into K clusters, each with its own cluster head.

Step 6.4: Trajectory and Information Gain. SKYRAN constructs a flight trajectory through the *K* cluster heads by solving a traveling salesman problem with these *K* cluster heads as nodes to visit. The length of the trajectory and its spatial coverage depends on *K*. SKYRAN constructs paths for each $K \in \{K_{min}, \ldots, K_{max}\}$ and selects the best trajectory, i.e. the path with the highest *information*-*to-cost ratio*.

Trajectory Information. Each existing (i.e. non-new) UE is associated with a set of flight trajectories from previous epochs. Note that a new UE does not have any flight trajectory history, and is thus assigned an empty set. For each UE, each measurement flight obtains new information about state of the channel from itself to other points in the operating area. To quantify this information, we define the information gain that a new trajectory provides to a UE as the shortest distance between the new trajectory and all the historical trajectories in the set assigned to the UE. The information gain for a new UE (i.e. with an empty trajectory set) is high. For the sake of mathematical tractability, we assigned a large fixed value I_{max} to it. The average information gain is the the mean information gains over all UEs in the current epoch. SKyRAN's formulation for information helps measurements in relatively unexplored areas. Information-to-Cost Ratio. The cost of a trajectory is its length. The information-to-cost ratio is thus the ratio of the average information gain of the trajectory to its length. SKyRAN selects the new trajectory with the highest information-to-cost ratio, thus maximizing the value of measurement data collected over the flight trajectory.



Figure 11: Schematic of the SKYRAN trajectory construction algorithm.

3.3.3 Update REM (Step 7). Measurement Update. During the measurement flight, the LTE eNodeB PHY reports the SNR to each UE at 100 Hz. The SKYRAN UAV also reads its GPS position (from the UAV flight controller) at 50 Hz. For each per-UE REM, the SNR of the a grid cell along the flight trajectory is assigned the average of all SNR readings taken within that grid cell.

Interpolation. SKyRAN uses a relatively lightweight interpolation technique called *Inverse Distance Weighting*, *IDW*³ to estimate the

SNRs at grid cells that are not directly on the measurement flight trajectory. To estimate the value of SNR at a given cell, IDW uses the weighted mean of the SNR values of its neighboring cells, where the weight is determined by square of the inverse distances between the center of the cell and the center of its neighboring cells.

3.4 Estimate Optimal UAV Position

SKYRAN UAV positions itself at the grid call that satisifies the *maxmin* SNR across all UEs to provide LTE service during the epoch. This choice of a *max-min* metric ensures a minimal QoS to all the UEs present in the SKYRAN's network and is a common objective for coverage problems. To achieve this, SKYRAN first constructs a *min-SNR* map, where each $1 \text{ m} \times 1 \text{ m}$ grid cell is assigned the minimum SNR values of the corresponding grid cells over all other per-UE REMs. The optimal UAV position is then selected as the min-SNR grid cell with the maximum value. Note that SKYRAN is equally applicable to other performance objectives (e.g. weighted UE SNR) as well. SKYRAN then moves itself to this optimal position to provide LTE service to the UEs until the next epoch, where the steps in Fig. 10 are repeated all over.

3.5 Adapting to UE Dynamics

Dynamic epoch to balance overhead and performance. Changes in UE locations due to mobility can affect their REM, thereby necessitating a reposition of the UAV for optimal connectivity. However, UAV repositioning requires a new epoch, where the UEs' REMs are re-estimated by probing more locations along a computed flight trajectory - an overhead that is prohibitive at a frequent scale. The shorter the epoch duration (i.e. frequent epochs), the more responsive the UAV's positioning can be to individual UE movements. However, this not only contributes to excessive probing overhead, but might also result in the UAV constantly adapting its position without being able to deliver optimized connectivity from a given position. On the other hand, a longer epoch duration is not favorable either as substantial UE mobility may result in significant changes to the network topology, thereby resulting in heavily degraded connectivity. Given that SkyRAN's objective is to provide optimized connectivity to all the UEs jointly, it strikes a balance between overhead and performance by dynamically triggering epochs that respond to aggregate UE dynamics. Specifically, a new epoch is triggered when the aggregate performance drops beyond a predefined threshold that can be set by the operator (say, 10%). Fig. 12 shows the degradation of overall throughput performance as a function of the epoch duration. Initially, the UAV is positioned optimally but as time advances, a fraction of the UEs move along certain predefined routes (scripted to closely mimic human mobility) without any change to UAV position. It can be observed that even a small threshold (e.g. 10% loss) allows for a reasonable epoch duration (10 minutes), thereby allowing SKyRAN to balance overhead and adaptability.

Temporal aggregation of REMs for minimizing overhead. SKYRAN's approach of REM construction and update is designed to implicitly optimize overhead in the presence of UE dynamics. Note that a REM is constructed for the *position* of a UE, rather than for the UE itself, as the latter can move around. Thus, when a UE moves by a large distance (over *R*) with respect to its current position, SKYRAN

³Although sophisticated and more computationally intensive interpolation techniques like Gaussian Process Regression or Ordinary Kriging have been used to interpolate radio maps but it has been shown to offer marginal improvement over IDW (see [30]).

Ayon Chakraborty et al.



Figure 12: Degradation of throughput performance with time where the UAV does not reposition itself while a fraction the UEs move. The results give a ballpark on the length of an epoch, say 10 mins while the relative throughput is within 80% of optimal.

can initialize the UE, at its new position, with an existing REM that is measured for a position within R of its latest position. This allows SKYRAN to leverage REMs from prior epochs that are spatially relevant for the UE. This contributes to efficient REM updates with minimal overhead even in the face of UE mobility.

Only when a UE has moved to a position, where no prior REM exists (within range *R*), SKYRAN initializes a new REM using a free-space path-loss (FSPL) model: each cell in the new REM will represent the SNR from the UAV to the UE, at its new position, as predicted by the FSPL model, rather than through measurements. However, this new REM will be updated with measurement data as they come in during successive epochs that are relevant to the given UE. The factor *R* is chosen depending on the scale and the terrain complexity of the coverage area. A higher value of *R* trades off REM's accuracy for scalability. In our current setup, we choose a *R* of 10 m from Fig. 9, as it provides minimal performance loss (<10%), while allowing for reduced overhead through REM reuse.

4 IMPLEMENTATION AND EVALUATION

4.1 Implementation

We implement SxyRAN with a custom-built LTE system mounted under a DJI Matrice 600 Pro (M600Pro) drone. Fig. 13 shows a picture of our SxyRAN LTE UAV.

Hardware. The SKYRAN drone is built using two single-board computers (SBCs), one with a Intel Core i7 and the other with a J1900 CPU. A USRP B210 with a GPS clock serves as the LTE RF frontend. This is connected to a LTE duplexer, power amplifier (PA) and low-noise amplifier (LNA) that provides about 18dB gain on both the LTE downlink and uplink channels. This is all connected to a 5dBi LTE antenna. All the hardware is mounted to the drone using a custom-built carbon fiber frame with dampers to minimize vibration transfer from the drone. The backhaul (command and control) network in our current implementation uses the AT&T LTE network via a separate LTE phone tether. This design choice is made for the sake of convenience as the backhaul design is outside the scope of our work. This LTE backhaul link can be replaced with other technologies such as mmWave, WiFi, or LTE-U etc.



Figure 13: A snapshot of our SKyRAN LTE UAV in operation and an UE connected to its network.

Software. We use OpenAirInterface [10] eNodeB and EPC software on the SBCs. The i7 SBC executes the LTE eNodeB, flight control software and SKYRAN algorithms, and the J1900 SBC runs the LTE EPC. The flight control software is built using the DJI OnBoard SDK, and enables fully hands-free autonomous drone flight. When launched, the SKYRAN drone autonomously executes the localization, REM measurement and optimal placement without any direct user input (e.g. through the remote controller).

Operating Range. With the LTE PA and LNA, the SKYRAN LTE system has a real-world operating range of over 300 m. This operating range is achieved even when the UE is in a NLOS situation, with buildings and/or tall trees between the UE and the drone. The flight time of the SKYRAN drone is up to 30 minutes. The maximum operating altitude of the drone is 120 m above ground level, as per FAA regulations.

SKYRAN Demonstration. A demonstration of the SKYRAN UAV in operation, providing LTE service within our test area can be found at http://www.nec-labs.com/skyran. This demonstration flight shows the SKYRAN UAV conducing the UE localization flight, followed by a measurement trajectory to estimate the the REM. The SKYRAN UAV then moves to the position that provides optimized LTE service for all UEs jointly.

4.2 Testbed Operation

We deploy and test $S\kappa vRAN$ in a 90 000 m² area (shown in Fig. 15) surrounding (and including) our campus building. Our performance results reported in this section are based on 35 S $\kappa vRAN$ test flights within this test area.

<u>Testbed Limitation</u>: We evaluate SKYRAN in our testbed using only static UEs. A larger scale performance study of SKYRAN over multiple epochs, and larger, more varied environments is presented in §5.

Smartphone UEs. We deploy seven Moto G5 LTE smartphones as UEs that connect to the SKYRAN LTE network. The UE locations are selected to ensure that all UEs experience both LOS and NLOS channels to the UAV over the course of a measurement flight. The UEs will thus experience highly varying channel conditions during the SKYRAN operation, as shown in Fig 14.





Figure 14: SNR distribution Figure 15: Coverage area of same SKYRAN flight.

of selected UEs during the SKYRAN. The yellow trajectory shows the tracks taken by SKYRAN's UAV to collect ground truth measurements.



0.8 0 5 10 15 Ranging Error (m) 0 20 25

Figure 16: Trajectory of the Figure 17: CDF showing dis-UNIFORM and CENTROID af- tribution of accuracy for ToF ter the UE localization flight. based ranging for UEs 1,5 and 7 for a 20 m long flight trajectory.

Algorithms. In our testbed experiments, we compare the performance of SKyRAN with two other UAV placement algorithms: Uniform and Centroid.

UNIFORM does not use UE location information and REMs, and instead adopts a zigzag trajectory across the test area, starting from one corner of the test area boundary, to measure the channel state uniformly as shown in Fig. 16.

CENTROID uses UE locations, but does not use REMs. Instead, it moves directly to the centroid of the UE locations to provide UAV LTE service, and is illustrated in Fig. 16.

Ground Truth Channel State. For a given set of UE positions (based on SKYRAN's localization), we fly the SKYRAN in a zigzag manner to collect detailed measurements from each UE to UAV positions throughout the test area. Fig. 15 shows the trajectory of this ground-truth measurement flight. This allows us to construct a detailed ground-truth REM, and determine the true optimal UAV operating point. We compare the UAV position obtained through SKYRAN, centroid and uniform with this ground-truth location. While we do this at multiple altitudes, for easier exposition of different schemes, we present results for UAV positioning at a given altitude.

4.3 UE Localization Accuracy

As mentioned in §3, the measurement trajectory is preceded by a short flight to localize the UEs. In the following section, we present



Figure 18: Distribution of lo- Figure 19: Median localizacalization errors for UEs 1,6 tion error as a function of and 7 for a 20 m long flight. flight trajectory budget.

the localization performance for UEs located in three different types of environments: UE 1 is located in the parking lot region (open space); UE 6 is right beside a large office building; and UE 7 is within a heavily forested portion of the test area, with 35 m high trees.

Ranging Accuracy. The ranging accuracy of course depends on accurate estimation of the time of flight (ToF). Depending on the environment, the ToF could be very noisy (standard deviation as high as 25 ns). In LoS, however, ToF is less noisy (standard deviation \approx 5 ns) The median ranging error is about 4 – 5 m over a 20 m localization flight as shown in Fig 17, and this accuracy is not affected by the UE location. This is achieved with an SRS upsampling of K = 4that improves the ranging accuracy with a minimal reduction in SNR. We consider an LTE bandwidth of 10 MHz in our experiments. We have found, through further experiments, that this accuracy holds when UEs are randomly placed at other locations in the test area.

Localization Accuracy. Fig. 18 shows the CDF of the UE locations determined by the SKYRAN UAV using the ToF measurements from a 20 m flight. We can see that within our 300×300 m test area, the median accuracy of SkyRAN is within 5 - 7 m.

UE Mobility. SKYRAN is primarily geared towards nomadic or relatively static users. However, with fast moving users (e.g., moving at car speeds), the localization accuracies for such UEs are deteriorated by as much as 3-4×. However such errors do not significantly impact the positioning of the UAV as such. SKyRAN handles and optimizes on aggregate clusters of users rather than individual user's mobility.

SkyRAN UE localization is an order of magnitude more accurate than the expected 50 - 100 m accuracy expected of existing LTE localization techniques (using macrocells). SKYRAN achieves this accuracy using only one eNodeB, and without complex time synchronization requirements. Furthermore, this is achieved only using a very short flight of 20 m. Fig. 19 shows that with our localization technique, longer flights are not needed to improve accuracy any further.

4.4 REM Construction Efficiency/Accuracy

The more accurately and efficiently that SKyRAN can estimate the REM in each epoch, the better that it can adapt its operating position, and the less time required in subsequent epochs to adapt the REMs to UE mobility.

Fig. 20 shows that the accuracy of the estimated REM using both SkyRAN and UNIFORM, w.r.t. the ground truth, increases as the amount of time spent (overhead) estimating it increases. This



0.6 0.4 0.2 0 6 Number of UEs

ror decreases with increasing measurement time.

Figure 20: Estimated REM er- Figure 21: Centroid-based placement of the UAV offers only 0.4×-0.6× of the optimal throughout performance.

measurement flight is conducted using the known location of the UEs in the testbed, as shown in Fig. 15. The longer the length of the measurement trajectory, the greater the number of channel measurements the UAV collects. However, observe that SkyRAN reduces the estimated REM error to its lower bound of 3 dB in a mere 82 s. The REM estimated by UNIFORM, on the other hand, reaches only 7dB even after 120 s.

SKYRAN is thus able to achieve better REM accuracy in each epoch which (a) reduces the overhead spent to position itself in an area of operations (scalability) and (b) reduces the overhead needed to track UE mobility across epochs (adaptability).

4.5 **UAV Placement**

Placement with only UE Locations. CENTROID is a simpler 4.5.1 version of SkyRAN that uses UE locations alone for UAV positioning. However, its performance is far below that of SKYRAN.

Fig. 21 shows the average LTE throughput of UEs using the CEN-TROID placement algorithm. Recall that only UE positions, and no measurement-based REMs are used by CENTROID. Observe that without a measurement-based REM to guide positioning decisions, CENTROID only achieves 60 % of the optimal ground-truth throughput, and only with a larger number of UEs. When the number of UEs is small, the achieved throughput can be as low as 40 % of optimal, along with higher variance. This is because as the number of UEs is small, the optimal UAV position is sensitive to environmental and terrain obstacles. With larger numbers of UEs, the effects of these obstructions will be "averaged out".

We will show later in this section that with only a short 2.5 min probing flight, SkyRAN can achieve over 90 % of the optimal throughput in all these cases. The low comparative performance of the CENTROID placement algorithm reflects the importance of measurementbased REMs in UAV positioning.

4.5.2 Placement with UE Locations and REMs. We evaluate SKYRAN and the UNIFORM algorithm in two topologies with different UE distributions shown in Fig. 22: Topology A with uniform UE distribution, and Topology B with a clustered placement of the UEs.

SKYRAN using UE locations and REMs outperforms UNI-FORM. Fig. 23 shows the throughput (relative to optimal) obtained by SKYRAN and UNIFORM for varying levels of measurement budget in these two topologies. UAV moves at 30 km/h during measurements in this experiment. Observe that SkyRAN is able to achieve a significant fraction of the optimal throughput even under small budgets. By being judicious about the spatial regions explored,



Figure 22: Different UE topologies formed in our SKyRAN testbed.



Figure 23: Relative throughput w.r.t. the optimal position for varying measurement budgets.



Figure 24: Median REM accuracy for a total measurement budget of 1000 meters.

SKYRAN delivers close to 2x throughput gain over UNIFORM for smaller budgets.

Its gains are further enhanced in the clustered UE topology B. Here, UNIFORM struggles to reach 70% optimality even with a 1000m budget, because its measurement flight does not prioritize the UE cluster. In contrast, SKYRAN automatically localizes the UE cluster, and biases its flight trajectory to maximize information collection rate from the UE cluster. Thus, SKYRAN is able to deliver close to 95% optimality, and incurs less than half the overhead (400m budget) to achieve UNIFORM's performance at 1000m. The corresponding REM error of under 3 dB for topology B in Fig. 24, clearly highlights the role of efficient REM construction behind SkyRAN's superior performance.

Comparing also with CENTROID's performance in Fig. 21 that maxes out at 50-60% optimality, we find that UE location (CENTROID) and RF measurements (UNIFORM) alone are unable to deliver satisfactory in isolation. This highlights the merits of SkyRAN's approach in leveraging UE location to conduct intelligent measurements and thereby construct efficient REMs.

SCALE UP STUDY 5

We augment the testbed experiments with scale-up simulation studies to demonstrate the performance of SkyRAN in larger terrains,

Ayon Chakraborty et al.



Figure 25: UAV hovering in the 3D airspace. The terrain represents the NYC dataset. The attenuation of the signal from the UAV's eNodeB to the UE is computed using the traced ray from the UAV to the UE on the ground. The ray could be obstructed by one or more building structures adding to NLOS path loss.

and with significantly more UEs. Our experiments are geared to explore the scalability and adaptability of the SKyRAN system.

5.1 **Simulation Setup**

LiDAR Topology. We use publicly available LiDAR datasets [8] to capture fine-grained details about various landscapes to simulate different terrain characteristics. In particular, we use data from three distinct types of terrains:

(i) RURAL is a 250 m×250 m rural area that consists of mostly open spaces, trees and a few small buildings;

(ii) NYC is a 250 m×250 m section of heavily urbanized area of downtown Manhattan; and

(iii) LARGE, which is a $1 \text{ km} \times 1 \text{ km}$ area of a semi-urban township in Wisconsin.

We pre-process the point-clouds to obtain a spatial granularity of 1 m. Fig. 25 shows the point cloud for the NYC dataset.

REM Generation. We model the channel between a UAV (in 3D space) and a UE on the ground using terrain-aware ray-tracing. We general realistic REMs by ray tracing with fine grained LiDAR terrain information. For each UE location on the ground, we construct an REM via a detailed ray tracing model to every point in 3D space (see Fig. 25). We use the LiDAR data to determine the portion of each ray that is obstructed by terrain features, and the portion that experiences only free space attenuation. We thus build an accurate model of the attenuation experienced by each direct signal path between the UAV and the UE, and the collection of rays over the 3D space forms the REM for the specific UE ground position.

SKYRAN In Dense Urban and Dynamic 5.2 Environments

SKYRAN aims to build an REM between every possible UE position and its 3D operating airspace above. In this section, we show that



Figure

26: w.r.t. the optimal position.



overhead to improve the overhead to improve the relative throughput to $0.9 \times$ relative throughput to $0.9 \times$ w.r.t. the optimal position for different types of simulated terrains.

this objective is scalable to more complex terrains. We consider six UEs (similar to the number in our testbed) randomly placed throughout the NYC topology, where UE-to-UAV links are subject to greater degrees of blockages from the high-rise buildings.

Fig. 26 shows that over 50 random simulation instances, when the UEs are stationary (i.e. STATIC), SKYRAN takes a similar amount of time (≈ 100 secs) to reach 90 % optimal throughput, which is similar to our testbed results (At our UAV speed of 30 km/h, this corresponds to 833m of flight distance) This is true even though the NYC topology has an order of magnitude more buildings (and thus, blockages) than our testbed.

When UEs are mobile, multiple epochs are required to estimate the REM. Using the same NYC topology and six UEs, we consider the situation where in each epoch, half of the UEs are randomly moved to different positions. This models a highly mobile UE environment. Observe from Fig. 26 that SKyRAN takes a combined six minutes of flight time across multiple epochs to reach 90 % optimal throughput, and is almost only half as long as required by UNIFORM.

Similar performances of SkyRAN vs UNIFORM can be seen in the delay required to obtain accurate REM estimation, as shown in Fig. 28. This shows that SKYRAN is scalable and adaptable even to dense urban and highly dynamic environments.

5.3 **SKYRAN** with Limited Resources

Recall that in practice, LTE performance suffers during UE localization and channel measurement flights. Hence, it is helpful to limit the amount of time spent in probing to an upper bound. Here, we use the same mobility model as before, along with a total flight distance of 5000 m over all epochs. Fig. 29 shows that while SKYRAN does not have any advantage over UNIFORM in the flat RURAL terrain, SKYRAN achieves 1.4× the throughput of UNIFORM in NYC and LARGE. SKYRAN maintains such gains over UNIFORM in the LARGE terrain even though it is 16× the area of NYC. Equally impressive SKYRAN gains over UNIFORM w.r.t. REM accuracy is also achieved, as shown in Fig. 30.

5.4 SKYRAN with Varying Number of UEs

We increase the number of active UEs from 2 to 10 while moving half of them in each epoch. The performance of SKYRAN improves roughly linearly till we have 8 active UEs (see Fig. 31). More the number of UEs, greater is the information gathered in parallel for constructing their respective REMs, and to figure out an optimal location for operation.





28: Figure overhead to improve the put w.r.t. optimal UAV posimedian REM accuracy to within 5 dB.



racy for a total measurement put w.r.t. optimal UAV posibudget of 5000 m of flight tion for a total measurement across epochs.



tion for a total measurement budget of 5000 m of flight across epochs.



Number of UEs per epoch Figure 30: Median REM accu- Figure 31: Relative through-

budget of 5000 m of flight

across epochs.

RELATED WORK 6

LTE Localization: Accuracy for network-side localization in today's LTE networks using state-of-the-art techniques [16, 28, 34, 36] can vary from a few tens to about a couple of hundred meters. While our LTE localization scheme is based on ranging from ToF information, [28, 36] uses data-driven techniques obtained through war-driving or crowd-sourcing to build location classification models. [34] uses a combination of Timing Advance information and RSRP to build a SVM-based classifier model that has an accuracy of 40 m-50 m. To the best of our knowledge this is the first work that reports LTE localization accuracies with a mobile base station. Second, the mobility of the base station is used to our advantage to improve median localization accuracy to sub-10 m.

UAV Placement: Recently UAV based deployment of communication infrastructure has caught much attention, alike from the industry [20, 27] and academia [18, 21-24, 38]. Given the characteristics of the coverage area (e.g, terrain) or the underlying RF environment several works [15, 25, 26] address the problem of 3-D placement of the drone in the operational airspace. Different types of objectives have been addressed while estimating an optimal placement. For instance [15, 26] maximizes the network usage in terms of the number of users on the ground that are covered by the cellular service. [25]'s placement objective takes into account the backhaul link from the drone to other relay drones as well. [14] provides fundamental results about the operating altitude of the UAV to maximize coverage. While the above works analytically models the optimal position of the UAV, a few works [19, 29, 35] focus on a measurement based strategy to estimate the optimal position.

RF/Propagation Modeling: Some recent works have reported field trials and measurements [13, 27, 38] to model RF pathloss and interference from LTE enabled UEs on UAVs associated to macrocell base stations. However these works focus on connectivity with respect to a single base station as opposed to build REMs. It brings completely a different set of constraints and challenges with the base station on the UAV catering to multiple UEs on the ground.

While most of the above work is analytical or based purely on simulation, SKYRAN takes a step ahead to realize a real LTE endto-end deployment.

DISCUSSION 7

Although SKyRAN is self-contained for single UAV operation, however, in order to scale such a network multiple UAVs need to work in unison to form a connected mesh. This involves multiple research challenges where SKYRAN addresses only the access part of the network. SKYLITE [37] paints a bigger picture highlighting an end-to-end system design of a UAV-based LTE network. Sky-CORE [31] redesigns the LTE evolved packet core (EPC) into a single light-weight, self-contained entity that is co-located with each UAV. SKYHAUL [37] optimizes the relative position and orientation of the UAVs for optimized backhaul connectivity. Nevertheless, SKYRAN gives a technological primitive to perform low cost REM estimation from which our entire SKYLITE eco-system is benefited.

A few recent works [32, 33, 39] have proposed analytical or numerical approaches for multi-UAV scenarios, however, in all cases it solely depends on an accurate estimation of the path loss model. SKYRAN is designed for scalability and adaptability, and directly supports multi-UAV deployments. In a multi-UAV scenario, the REM are cooperatively constructed and shared amongst multiple SKYRAN UAVs, thus increasing the scalability of SKYRAN to support even larger operating areas with more UEs. The adaptability feature of SkyRAN ensures that it will still be able to monitor and react to changes in UE mobility and the environment.

Placement objective: SKYRAN places the UAV based on a minmax throughput criteria, a fairly common practice in cellular networks. However, the usage of REMs makes our system flexible enough to incorporate other objective measures, e.g., increasing total coverage or maximizing users in the network and so on [15, 26].

CONCLUSION 8

SKYRAN is the first of its kind system with an end-to-end implementation of an aerial LTE base station catering to multiple UEs on the ground. This paper focus on the associated set of challenges in the LTE RAN related to optimally positioning the UAV in the 3D airspace. We show how a brute-force approach to scan the entire 3D airspace is infeasible to scale for a larger coverage area or with substantial client dynamics. Our solution is UE location-aware. Knowing the location of the UEs, SKYRAN intelligently prunes the 3D airspace for collecting measurements based on which it constructs Radio Environment Maps (REMs) for individual UEs. The estimated REMs serve as a basis for determining the optimal placement of the UAV to operate at.

CoNEXT '18, December 4-7, 2018, Heraklion, Greece

REFERENCES

- [1] Assisted GNSS. https://en.wikipedia.org/wiki/Assisted_GPS.
- [2] AT&T Cell on Wings. http://about.att.com/innovationblog/cows_fly.
- [3] DJI Phantom Drone. https://www.dji.com/phantom.
- [4] DJI Wind 8 Octocopter. https://www.quadrocopter.com/DJI-Wind-8-Octocopter_ p_1841.html.
- [5] Engineers design drones that can stay aloft for five days. http://news.mit.edu/ 2017/drones-stay-aloft-five-days-0607.
- [6] Enhanced Cell-ID. https://en.wikipedia.org/wiki/E-CellID.
- [7] Hurricane Maria. https://en.wikipedia.org/wiki/Hurricane_Maria.
- [8] LiDaR Traces. ftp://rockyftp.cr.usgs.gov/vdelivery/Datasets/Staged/NED/LPC/ projects/.
- [9] LTE SRS. http://www.sharetechnote.com/html/Handbook_LTE_SRS.html.
- [10] Open Air Interface. http://www.openairinterface.org/.
- [11] Vapor 55 Helicopter UAV. http://www.pulseaero.com/uas-products/vapor-55.
- [12] Verizon Cell on Wings. https://newatlas.com/verizon-drones-internet-trials/ 45818/.
- [13] AL-HOURANI, A., KANDEEPAN, S., AND JAMALIPOUR, A. Modeling air-to-ground path loss for low altitude platforms in urban environments. In *Global Communications Conference (GLOBECOM), 2014 IEEE* (2014), IEEE, pp. 2898–2904.
- [14] AL-HOURANI, A., KANDEEPAN, S., AND LARDNER, S. Optimal LAP altitude for maximum coverage. IEEE Wireless Communications Letters 3, 6 (2014), 569–572.
- [15] BOR-YALINIZ, R. I., EL-KEYI, A., AND YANIKOMEROGLU, H. Efficient 3-d placement of an aerial base station in next generation cellular networks. In *Communications* (ICC), 2016 IEEE International Conference on (2016), IEEE, pp. 1–5.
- [16] CHAKRABORTY, A., ORTIZ, L. E., AND DAS, S. R. Network-side positioning of cellular-band devices with minimal effort. In *Computer Communications (INFO-COM), 2015 IEEE Conference on* (2015), IEEE, pp. 2767–2775.
- [17] CHAN, Y.-T., AND HO, K. A simple and efficient estimator for hyperbolic location. IEEE Transactions on signal processing 42, 8 (1994), 1905–1915.
- [18] DHEKNE, A., GOWDA, M., AND CHOUDHURY, R. R. Extending cell tower coverage through drones. In Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications (2017), ACM, pp. 7–12.
- [19] DORLING, K., HEINRICHS, J., MESSIER, G. G., AND MAGIEROWSKI, S. Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics:* Systems 47, 1 (2017), 70–85.
- [20] E.T. AL, X. L. Mobile networks connected drones: Field trials, simulations, and design insights. arXiv preprint arXiv:1801.10508 (2018).
- [21] FOTOUHI, A. Towards intelligent flying base stations in future wireless network. In A World of Wireless, Mobile and Multimedia Networks (WoWMOM), 2017 IEEE 18th International Symposium on (2017), IEEE, pp. 1–3.
- [22] FOTOUHI, A., DING, M., AND HASSAN, M. Dronecells: Improving 5g spectral efficiency using drone-mounted flying base stations. arXiv preprint arXiv:1707.02041 (2017).
- [23] FOTOUHI, A., DING, M., AND HASSAN, M. Service on demand: Drone base stations cruising in the cellular network. arXiv preprint arXiv:1710.09504 (2017).
- [24] GANGULA, R., ESRAFILIAN, O., GESBERT, D., ROUX, C., KALTENBERGER, F., AND KNOPP, R. Flying Rebots: First Results on an Autonomous UAV-Based LTE Relay Using Open Airinterface. In 2018 IEEE 19th International Workshop on Signal

Processing Advances in Wireless Communications (SPAWC) (2018), IEEE, pp. 1–5.] KALANTARI, E., SHAKIR, M. Z., YANIKOMEROGLU, H., AND YONGACOGLU, A.

- [25] KALANTARI, E., SHAKIR, M. Z., YANIKOMEROGLU, H., AND YONGACOGLU, A. Backhaul-aware robust 3d drone placement in 5g+ wireless networks. In Communications Workshops (ICC Workshops), 2017 IEEE International Conference on (2017), IEEE, pp. 109–114.
- [26] KALANTARI, E., YANIKOMEROGLU, H., AND YONGACOGLU, A. On the number and 3d placement of drone base stations in wireless cellular networks. In Vehicular Technology Conference (VTC-Fall), 2016 IEEE 84th (2016), IEEE, pp. 1–6.
- [27] LIN, X., YAJNANARAYANA, V., MURUGANATHAN, S. D., GAO, S., ASPLUND, H., MAAT-TANEN, H.-L., EULER, S., WANG, Y.-P. E., ET AL. The sky is not the limit: Lte for unmanned aerial vehicles. arXiv preprint arXiv:1707.07534 (2017).
- [28] MARGOLIES, R., BECKER, R., BYERS, S., DEB, S., JANA, R., URBANEK, S., AND VOLIN-SKY, C. Can you find me now? Evaluation of network-based localization in a 4G LTE network. In *INFOCOM 2017-IEEE Conference on Computer Communications*, *IEEE* (2017), IEEE, pp. 1–9.
- [29] MODARES, J., GHANEI, F., MASTRONARDE, N., AND DANTU, K. UB-ANC planner: Energy efficient coverage path planning with multiple drones. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on* (2017), IEEE, pp. 6182– 6189.
- [30] MOLINARI, M., FIDA, M.-R., MARINA, M. K., AND PESCAPE, A. Spatial interpolation based cellular coverage prediction with crowdsourced measurements. In Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data (2015), ACM, pp. 33–38.
- [31] MORADI, M., SUNDARESAN, K., CHAI, E., RANGARAJAN, S., AND MORLEY MAO, Z. SkyCore: Moving Core to the Edge for Untethered and Reliable UAV-based LTE networks. ACM Mobicom (2018).
- [32] MOZAFFARI, M., SAAD, W., BENNIS, M., AND DEBBAH, M. Drone small cells in the clouds: Design, deployment and performance analysis. In *Global Communications Conference (GLOBECOM)*, 2015 IEEE (2015), IEEE, pp. 1–6.
- [33] MOZAFFARI, M., SAAD, W., BENNIS, M., AND DEBBAH, M. Efficient deployment of multiple unmanned aerial vehicles for optimal wireless coverage. *IEEE Communications Letters 20*, 8 (2016), 1647–1650.
- [34] NI, L., WANG, Y., TANG, H., YIN, Z., AND SHEN, Y. Accurate Localization Using LTE Signaling Data. In Computer and Information Technology (CIT), 2017 IEEE International Conference on (2017), IEEE, pp. 268–273.
- [35] NIU, S., ZHANG, J., ZHANG, F., AND LI, H. A method of UAVs route optimization based on the structure of the highway network. *International Journal of Distributed Sensor Networks* 11, 12 (2015), 359657.
- [36] RAY, A., DEB, S., AND MONOGIOUDIS, P. Localization of LTE measurement records with missing information. In Computer Communications, IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on (2016), IEEE, pp. 1–9.
- [37] SUNDARESAN, K., CHAI, E., CHAKRABORTY, A., AND RANGARAJAN, S. SkyLiTE: Endto-End Design of Low-Altitutde UAV Networks for Providing LTE Connectivity. arXiv preprint arXiv:1802.06042 (2018).
- [38] VAN DER BERGH, B., CHIUMENTO, A., AND POLLIN, S. Lte in the sky: trading off propagation benefits with interference costs for aerial nodes. *IEEE Communications Magazine* 54, 5 (2016), 44–50.
- [39] WU, Q., ZENG, Y., AND ZHANG, R. Joint trajectory and communication design for uav-enabled multiple access. arXiv preprint arXiv:1704.01765 (2017).