2020 IEEE 21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)

SkyCell: A Prototyping Platform for 5G Aerial Base Stations

Ludovico Ferranti*†, Leonardo Bonati*, Salvatore D'Oro*, Tommaso Melodia*
*Institute for the Wireless Internet of Things, Northeastern University, Boston, MA
†Sapienza University of Rome, Italy
Email: {ferranti.l, bonati.l, s.doro, melodia}@northeastern.edu

***Abstract*—In this paper we propose SkyCell, a prototyping platform for 5G autonomous aerial base stations. While the majority of work on the topic focuses on theoretical and rarely implemented solutions, SkyCell practically demonstrates the feasibility of an aerial base station where wireless backhaul, autonomous mobility and 5G functionalities are integrated within a unified framework. We showcase the advantages of Unmanned Aerial Vehicles for 5G applications, discuss the design challenges, and ultimately propose a prototyping framework to develop aerial cellular base stations. Experimental results demonstrate that SkyCell not only supports heterogeneous data traffic demand and services, but also enables the implementation of autonomous flight control algorithms while improving metrics such as network throughput (up to 35%) and user fairness (up to 39%).**

I. Introduction

In the last few years it became more evident that traditional and inflexible approaches for cellular networking that only rely on ground-based infrastructure are not capable to keep the pace with the ever increasing traffic demand [1, 2]. Thanks to their unique features, such as rapid deployment, high mobility and accurate positioning, Unmanned Aerial Vehicles (UAVs) have been identified as perfect candidates to either assist 5G cellular networks, or to compensate for their temporary failure due to disasters and emergencies [3–5].

The advent of Software-Defined Radios (SDRs), together with their small form factor and high flexibility, have made it possible to equip UAVs with lightweight transceivers and antennas, thus paving the way to the new concept of Unmanned Aerial Base Stations (UABSs) [6]. Empowered by high mobility, reconfigurability of RF front-end components, and connected to the Internet through a wireless backhaul, UAVs can serve as 5G pico- and femto-cells that can be easily deployed in real-time according to the traffic demand (Fig. 1).

Even though UABSs have attracted attention from both industry and academia due to their undeniable advantages in the 5G scene [7], a unified framework combining flight control and networking functionalities is currently still missing. In fact, the majority of works on 5G UABSs covers these two core aspects separately, thus not fully leveraging their potential and functionalities [8]. The reasons behind the lack of such integrated solutions are manifold. UAVs are resource-constrained devices in terms of battery, payload, and computational power. Thus, design choices concerning them should comply with these limitations. For instance, hardware components should be lightweight and energy-efficient but, at the same time, powerful enough to perform convoluted signal processing operations, and support Quality of Service (QoS)-aware mobility algorithms.

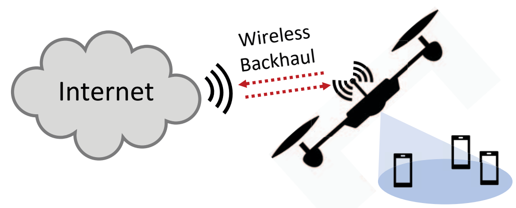

Fig. 1: SkyCell application scenario.

The goal of this paper is to tackle all of the above challenges and take a further step toward the design, development and prototyping of UABSs. Specifically, in this paper we make the following contributions:

- We present SkyCell, a prototyping platform for 5G autonomous UABSs. SkyCell seamlessly integrates UAV flight control directives and procedures with networking capabilities to provide improved cellular network service to users. As an example, SkyCell allows network operators to deploy UABSs that autonomously adapt their trajectories to optimize relevant QoS metrics, such as network throughput and Signal-to-Interference-plus-Noise Ratio (SINR);
- We develop SkyCell prototype to demonstrate the effectiveness of the proposed unified framework by using Commercial Off-The-Shelf (COTS) hardware and open-source software.
- We make our prototype compliant with the existing cellular architecture by leveraging the srsLTE software to jointly instantiate cellular core network and base station, and provide services to mobile subscribers;
- We demonstrate the benefits of using UABSs by assessing SkyCell performance in a variety of real-world scenarios

978-1-7281-7374-0/20/$31.00 ©2020 IEEE
DOI 10.1109/WoWMoM49955.2020.00062

through extensive experiments in an outdoor UAV facility. Our experiments show that QoS-driven mobility schemes improve network metrics such as network throughput (up to 35%) and user fairness (up to 39%).

The remainder of this paper is organized as follows. In Section II, we illustrate SkyCell architecture and system design. In Section III, we describe the developed SkyCell prototype, while we experimentally evaluate its performance in Section IV. Finally, we discuss related work in Section V, and draw our conclusions in Section VI.

II. SkyCell Architecture

The ability of SkyCell to access functions that allow both the control of a cellular base station, motion control of the drone, and sensor readings is a sought-after feature in platforms for aerial networks. One of SkyCell main strengths is to bring together within a unified framework the communication and flight management aspects of UABSs. Fig. 2 shows the core components of SkyCell architecture: The measurements *Database*, the *Mobility Module*, and the *Network Module*.

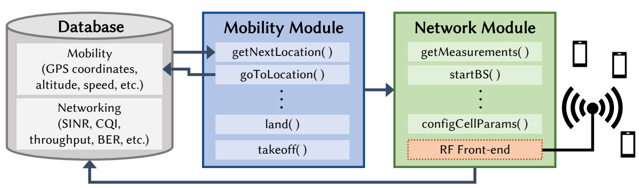

Fig. 2: SkyCell system architecture.

Database: It stores all the information regarding the UAV sensor readings, e.g., GPS coordinates, altitude, and speed, together with network-related information, such as SINR, CQI, throughput, and bit error rate (BER). A history of the measured metrics is saved in SkyCell database together with their acquisition timestamp. The availability of such information enables the swift implementation of optimization and reinforcement learning algorithms that require current and past information of the network state. We will show how this database can be utilized to automatically adjust the trajectory of the UABS in Section IV.

Mobility Module: It provides all the basic UAV functionalities, such as *takeOff()*, *land()*, and *goToLocation()*, together with the function *getNextLocation()* used to compute the next position of the UAV. This function reads information from the database (e.g., past and current metric readings) and can be programmed by the network operator to implement the desired optimization or machine learning algorithm, as well as reading a list of stored GPS coordinates. Moreover, this module computes the current UAV information (e.g., GPS location and sensor readings), stores it in the database, and sends it to the network module.

Network Module: It provides APIs to instantiate cellular base station (i.e., *startBS()*) and core network (i.e., *startCN()*) applications on the drone, as well as to provide radio access to User Equipments (UEs) through its RF front-end component. The network module periodically stores network measurements (e.g., throughput, SINR, CQI, etc.) in the database through the function *recordUEsMetric()*.

The role of this information is twofold: (i) It is accessed by the mobility module when executing *getNextLocation()* (e.g., to compute the next UAV location as shown in Algorithm 1), and (ii) it allows the network module to adjust the UABS parameters and configuration in real-time (e.g., modulation scheme, number of resource blocks, etc.) through the *configCellParams()* function. For example, if a UE moves, SkyCell will notice a degradation in the computed SINR. To react to this negative effect, SkyCell can jointly instruct the network module to use a more robust modulation scheme, and the mobility module to move to a different location.

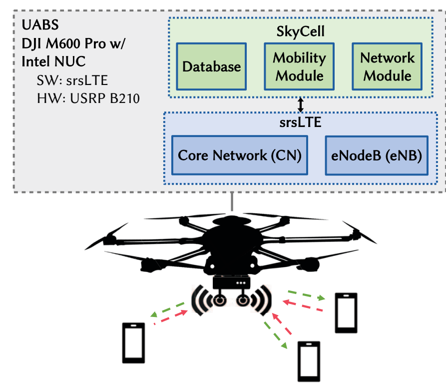

Fig. 3: SkyCell stack.

Ultimately, the consistent interactions between all these modules make SkyCell API library a suitable platform to prototype and test advanced machine learning and artificial intelligence algorithms, as well as traditional optimization techniques for aerial cellular networks. Algorithm 1 provides a practical example of SkyCell APIs and their utilization in a practical exploration-exploitation scenario whose details will be given in Section IV.

III. SkyCell Prototype

SkyCell stack and prototype implementation are shown in Figs. 3 and 4. We leveraged a DJI Matrice 600 (M600) Pro UAV equipped with an Intel NUC 7i7DNKE Mini PC and an Ettus Research USRP B210 with VERT2450 omnidirectional antennas. Our prototype is shown in Figs. 5a and 5b.

The DJI M600 Pro UAV is a commercial hexacopter for industrial application with a $1.67\,\mathrm{m}$ wingspan and a payload capacity of $6\,\mathrm{kg}$. It is equipped with a variety of sensors, such as three inertial measurement unit sensors with magnetometer, accelerometer, and gyroscope, and three GPS sensors for accurate UAV localization. The drone is maneuvered through the DJI Onboard SDK installed on the Intel NUC board that communicates with the DJI A3 Pro flight controller through a set of APIs. This enables a swift and real-time flying mission update by reading data from the sensors and setting specific GPS coordinates the drone needs to visit.

The Intel NUC Mini PC is powered directly by the UAV batteries through a DC-DC upconverter, and runs the DJI Onboard SDK that interfaces with the A3 flight controller through a JTAG-USB cable. The NUC runs Ubuntu 18.04 LTS and provides service to mobile subscribers by jointly implementing a cellular base station and core network. Due to the lack of research-oriented experimental 5G platforms, these functionalities have been implemented through srsLTE [9]. Because of the similarities between LTE and 5G numerologies, though, we are confident that our findings remain valid for 5G scenarios, nonetheless.

The srsLTE software provides an LTE-compliant protocol stack implementation that includes evolved packet core and base station (eNB) applications. We employed this tool to instantiate an eNB on the USRP B210 SDR mounted on the UAV that acts as radio front-end and interfaces to the NUC through a USB 3.0 connection. VERT2450 antennas have a 3 dBi gain. The eNB, then, serves three commercial off-the-shelf cellular phones (Samsung Galaxy S5), and provides them Internet connectivity. This is achieved through a wireless backhaul implemented by establishing a Wi-Fi link between the NUC and a ground-based Wi-Fi router connected to the Internet.

The *getNextLocation()* function is executed by the NUC, which reads statistics relative to the over-the-air link between eNB and users, e.g., SINR and downlink throughput, to compute the next UAV location. These metrics are either directly measured by the UABS, or sent as user feedback as per standard LTE control channel procedures.

For illustrative purposes only, network throughput maximization, and SINR maximization have been chosen as policies to compute the next UAV location in the experiments described in Section IV.

IV. Experimental Evaluation

To demonstrate SkyCell capabilities we ran extensive experiments at the Unmanned Aircraft Systems (UAS) Drone Lab at Northeastern University's Kostas Research Institute in Burlington, MA. These experiments have been performed in the $45 \times 60 \times 16\,\mathrm{m}$ UAS Drone Lab outdoor drone cage whose purpose is to foster UAV testing and prototyping activities.

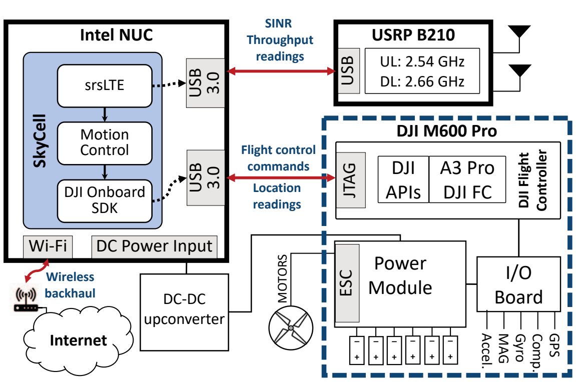

Fig. 4: SkyCell prototype implementation.

Algorithm 1 Exploration-Exploitation used in Section IV

Input: a UABS V implementing SkyCell APIs;
Output: An exploration path of the UABS;
$V.startCN()$; ▷ Start the Core Network
$V.startBS()$; ▷ Start the eNB
$V.takeOff()$;
while convergence criterion not met **do**
  $p_{curr} \leftarrow V.getCurrentPosition()$;
  $m_{p_{curr}} \leftarrow V.recordUEsMetric(p_{curr}, m_{curr})$; ▷ Record UE metrics (e.g., SINR, throughput)
  $p_{next} \leftarrow V.getNextLocation(p_{curr}, m_{p_{curr}})$;
  $V.flyTo(p_{next})$; ▷ Compute next location
$V.returnHome()$;
$V.land()$;
$V.stopCN()$; ▷ Stop the Core Network
$V.stopBS()$; ▷ Stop the eNB

function GETNEXTLOCATION($p_{curr}, m_{p_{curr}}$)
  $\mathbf{p}_{exp} \leftarrow V.getNeighborhood(p_{curr})$; ▷ Set of locations to explore
  for $p \in \mathbf{p}_{exp}$ **do**
    $V.flyTo(p)$;
    $m_p \leftarrow V.recordUEsMetric(p)$;
  $p_{next} \leftarrow \arg\max_{p \in \mathbf{p}_{exp} \cup p_{curr}} \{m_p\}$;
  return p_{next}; ▷ The best next location

The experimental setup, shown in Fig. 5c, consists of a SkyCell UABS (Figs. 5a and 5b) and three COTS Android smartphones serving as UEs. In order to assess SkyCell performance under different network conditions, we considered three different traffic schemes:

- *Edge Video Streaming*: Each UE requests a video to SkyCell UABS through the Samba (SMB) networking protocol. The video is then transmitted to the UE through the established LTE link;
- *Online Video Streaming*: Each UE requests a video from YouTube through the mobile app installed on the smartphone;
- *Speedtest*: Each UE runs the speedtest hosted by Google Measurement Lab.

It is worth noticing that differently from the online video streaming application where contents are hosted on remote servers accessed through the Internet via SkyCell UABSs, edge video streaming leverages locally available videos cached on the UABS itself that are ultimately transmitted to UEs over the LTE link. Moreover, despite both online video streaming and speedtest applications require Internet access to retrieve the requested data, the kind of traffic they generate is very diverse, and thus has a different impact on the overall performance of SkyCell. As an example, in Figs. 6 and 7 we show a snapshot of the instantaneous throughput provided by SkyCell to the three UEs relative to the online video streaming and speedtest applications, respectively. While online video streaming results are characterized by a periodic and bursty traffic pattern that only sporadically occupies the wireless channel (Fig. 6), the speedtest application is a more demanding process (Fig. 7) as it generates a continuous data stream which requires unceasing allocation of network resources.

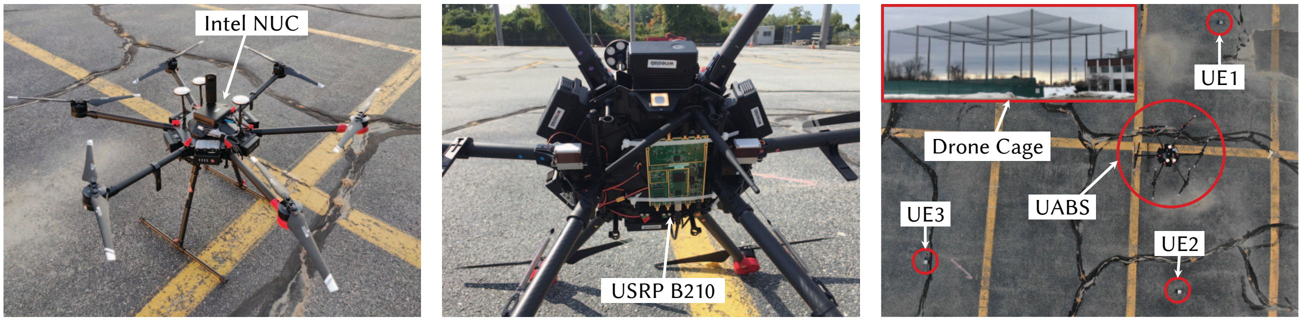

(a) (b) (c)

Fig. 5: Experimental setup: (a)-(b) DJI Matrice 600 Pro UAV equipped with Intel NUC compute board and Ettus Research USRP B210; (c) UAS Drone Lab outdoor drone cage and initial position of UABS and UEs.

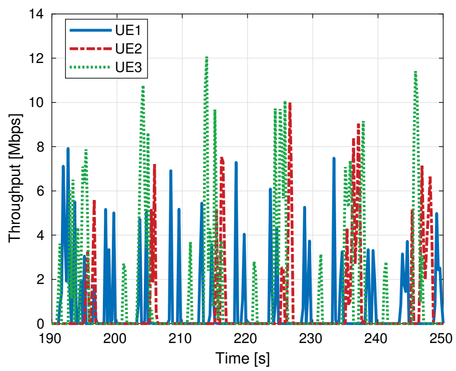

Fig. 6: Online video streaming traffic with bursty behavior.

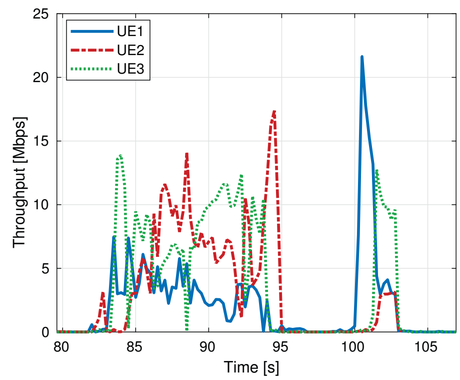

Fig. 7: Speedtest traffic with continuous behavior.

A. *Edge Video Streaming*

One of the unique features of SkyCell is its ability to combine networking and flight control capabilities to provide improved overall performance to UEs. To showcase this feature, and to test SkyCell autonomous and adaptive flight capabilities, we implemented an exploration-exploitation search algorithm (i.e., Algorithm 1) that periodically monitors network and QoS parameters while exploring the surrounding environment and makes decisions accordingly. The idea is the following: The drone takes off and *explores* the environment by visiting four different positions with respect to its current position (e.g., north, east, south, and west). For each of these positions it measures and records the value of relevant metrics, such as throughput, SINR, and CQI, among others. SkyCell, then, *exploits* these metrics to select the next UAV location such that a desired metric is maximized, e.g., the aggregate user downlink throughput.

This procedure is, then, periodically iterated until a convergence condition is met, e.g., no further improvement can be obtained by changing the UAV location.

In Fig. 8, we focus on the edge video streaming application and compare SkyCell performance under different metrics criteria. Specifically, we consider two different kinds of experiments: Throughput-based, and SINR-based.

• *Throughput-based (TB)*: in this scheme SkyCell UABS aims at finding a profitable hovering point such that the overall throughput of the network is maximized;

• *SINR-based (SB)*: SkyCell UABS computes an efficient trajectory to maximize the aggregate SINR of the UEs.

Figs. 8b and 8d show the path traveled by the UAV while serving the cellular subscribers in the *TB* and *SB* cases, respectively. From these figures, we notice that different metrics have significant impact on the trajectory of the SkyCell UABS. Similarly, Figs. 8a and 8c show the throughput measured by the three UEs. Both figures provide interesting insights. First, the take-off procedure causes a performance degradation under both schemes. This comes at no surprise as take-off: (i) consumes more power than hovering procedures, thus subtracting it from the both the NUC board and the USRP, i.e., from the eNB, and (ii) generates a steep altitude variation with a consequent increase in the Doppler shift, which eventually deteriorates the measured throughput due to a temporary synchronization loss between the eNB and the UEs. Second, it is worth noticing that despite the *SB* approach shows higher peak throughput values (Fig. 8c) than the *TB* scheme (Fig. 8a), the average per-UE throughput is higher in the *TB* scheme (approximately 14%).

B. *Online Video Streaming*

We now consider two different positioning schemes, the autonomous scheme presented in Section II, and the *waypoint* scheme. In the latter, SkyCell eNB is iteratively positioned above each UE in Fig. 5c after the initial take-off.

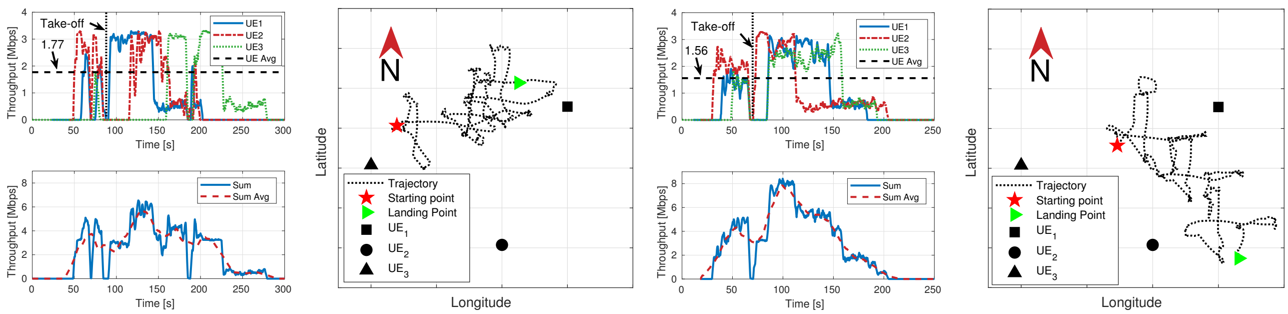

(a) Throughput under the TB autonomous flight scheme.

(b) Flight trajectory under the TB autonomous flight scheme.

(c) Throughput under the SB autonomous flight scheme.

(d) Flight trajectory under the SB autonomous flight scheme.

Fig. 8: Comparison of TB (a-b) and SB (c-d) autonomous flight schemes for Edge video streaming applications.

Fig. 9 shows the average throughput of the three UEs, as well as the corresponding Jain's fairness index in the video streaming case. As expected, our results demonstrate that the autonomous scheme provides better throughput and is more fair if compared to the waypoint scheme. In fact, the waypoint scheme deterministically selects the positioning of the UAV completely disregarding the impact on the resulting performance. This highlights that a proper choice for the UABS positioning is paramount.

C. *Speedtest*

Finally, in Fig. 10 we present the average and maximum overall throughput of the system when considering the waypoint mobility scheme and the case in which the SkyCell UABS is positioned on the ground equidistant from the three UEs, i.e., *ground* scheme. Fig. 10 clearly shows the superior performance of the mobile UABS, which outperforms the ground-based eNB in terms of both average and maximum throughput.

V. Related Work

The application of UAVs to cellular networks has gained momentum in the recent literature. Despite being beneficial to extend network connectivity and coverage, the integration of UAV mobility and cellular networking functionalities comes with a variety of unique challenges that call for specific features and design choices [10]. How to determine the minimum number of UAVs required to cover a specific area is investigated in [11, 12], while in-depth coverage analysis for urban and rural areas are given in [7] and [13], respectively. A theoretical framework to perform energy-efficient radio resource management and to find the optimal coverage radius for UAV-assisted mmWave cellular networks is presented in [14]. Iellamo et al., instead, proposed a clustering algorithm to compute the position of UABS such that the network capacity is maximized [15].

Since the quality of cellular network services strongly depends on the location of aerial base stations, the community has focused on devising solutions to efficiently control UAV trajectories in the presence of obstacles [16], on improving the quality of video-streaming services [17], and on offloading ground-based base stations [18]. Kalantari et al. proposed an algorithm to optimize the position of UABSs based on the availability of wireless backhaul links [19]. Similarly, a mathematical model for UAV-aided mmWave backhaul management in 5G systems is given in [20].

Although these works have provided significant advancement to the literature, they are supported by simulation results only and—unlike SkyCell—they do not consider practical implementation aspects and challenges of real systems.

First efforts to close the gap between theoretical works and practical implementations are presented in [21] and [22], which analyze the feasibility of UAV-based LTE relays. A new radio access network control mechanism that extends

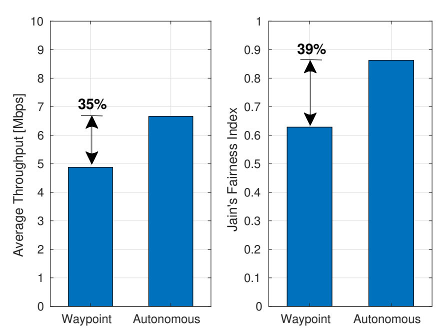

Fig. 9: Average online video streaming throughput and Jain's fairness index comparison with waypoint and autonomous schemes.

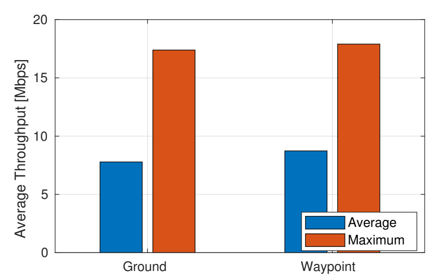

Fig. 10: Average and maximum speedtest throughput.

the ground-designed cellular infrastructure to support UAV-based UEs is proposed by Bertizzolo et al. in [23]. An exparimental evaluation of location-aided mmWave backhaul link management for UABSs is, instead, carried out in [24], while [25] prototypes a framework for automatic optimization of arbitrary network functions in ad-hoc swarms of drones.

The work that gets closer to SkyCell is that in [26]. In this work, D'Alterio et al. propose a UABS prototype implementing a gradient-based optimization algorithm that adapts the placement of the UABS such that a target metric is maximized (e.g., the throughput of a UE). Although the prototype in [26] clearly demonstrates the benefits and effectiveness of aerial cellular networks, it considers a very simple application scenario with a single UE and no backhaul from the UABS to the internet. These important limitations prevent a thorough assessment of UABS-enabled cellular networks. These are, instead, addressed in SkyCell, where we carry out an exhaustive demonstration of UABSs in real-world scenarios with wireless backhaul connectivity and multiple UEs.

VI. Conclusions

In this paper we presented SkyCell, an SDR-equipped prototyping platform for Unmanned Aerial Base Stations (UABS) that jointly integrates flight control and cellular networking functionalities. We discussed SkyCell's system architecture and presented a prototype implementation with COTS hardware, LTE-compliant open-source software, and a Wi-Fi backhaul. To demonstrate the effectiveness of SkyCell in real-world network deployments, we performed an extensive experimental campaign showing that autonomous and QoS-driven UABS deployments improve network throughput (up to 35%) and user fairness (up to 39%) if compared to static flight control algorithms. Future work will aim to improve performance, integrate a mmWave backhaul, and experiment with a higher number or UABSs and users.

VII. Acknowledgements

This work was supported in part by the US National Science Foundation under Grant CNS-1618727 and in part by the US Office of Naval Research under Grants N00014-19-1-2409 and N00014-20-1-2132.

References

[1] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "Open, programmable, and virtualized 5G networks: State-of-the-art and the road ahead," *arXiv:2005.10027 [cs.NI]*, May 2020.

[2] L. Bonati, S. D'Oro, L. Bertizzolo, E. Demirors, Z. Guan, S. Basagni, and T. Melodia, "CellOS: Zero-touch softwarized open cellular networks," *Computer Networks (COMNET). Preprint available at arXiv:2004.03418*, 2020.

[3] L. Gupta, R. Jain, and G. Vaszkun, "Survey of important issues in UAV communication networks," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 2, pp. 1123–1152, Nov. 2016.

[4] M. Deruyck, J. Wyckmans, L. Martens, and W. Joseph, "Emergency ad-hoc networks by using drone mounted base stations for a disaster scenario," in *Proc. of IEEE WiMob*, New York, NY, USA, Oct. 2016, pp. 1–7.

[5] L. Ferranti, S. D'Oro, L. Bonati, E. Demirors, F. Cuomo, and T. Melodia, "HIRO-NET: Self-organized robotic mesh networking for Internet sharing in disaster scenarios," in *Proc. of IEEE WoWMoM*, Washington, D.C., USA, June 2019.

[6] M. Alzenad, A. El-Keyi, F. Lagum, and H. Yanikomeroglu, "3-d placement of an unmanned aerial vehicle base station (uav-bs) for energy-efficient maximal coverage," *IEEE Wireless Communications Letters*, vol. 6, no. 4, pp. 434–437, 2017.

[7] H. C. Nguyen, R. Amorim, J. Wigard, I. Z. Kovacs, and P. Mogensen, "Using LTE networks for UAV command and control link: A rural-area coverage analysis," in *Proc. of IEEE VTC-Fall*, Toronto, ON, Canada, Sept. 2017, pp. 1–6.

[8] R. Amorim, H. Nguyen, P. Mogensen, I. Z. Kovács, J. Wigard, and T. B. Sørensen, "Radio channel modeling for uav communication over cellular networks," *IEEE Wireless Communications Letters*, vol. 6, no. 4, pp. 514–517, 2017.

[9] I. Gomez-Miguelez, A. Garcia-Saavedra, P. Sutton, P. Serrano, C. Cano, and D. Leith, "srsLTE: An open-source platform for LTE evolution and experimentation," in *Proc. of ACM WiNTECH*, New York City, NY, USA, Oct. 2016.

[10] Y. Zeng, R. Zhang, and T. J. Lim, "Wireless communications with unmanned aerial vehicles: Opportunities and challenges," *IEEE Communications Magazine*, vol. 54, no. 5, pp. 36–42, May 2016.

[11] J. Lyu, Y. Zeng, R. Zhang, and T. J. Lim, "Placement optimization of UAV-mounted mobile base stations," *IEEE Communications Letters*, vol. 21, no. 3, pp. 604–607, Nov. 2017.

[12] H. Wang, H. Zhao, W. Wu, J. Xiong, D. Ma, and J. Wei, "Deployment algorithms of flying base stations: 5G and beyond with UAVs," *IEEE Internet of Things Journal*, 2019.

[13] E. Greenberg and P. Levy, "Channel characteristics of UAV to ground links over multipath urban environments," in *Proc. of IEEE COMCAS*, Tel-Aviv, Israel, Jan. 2017, pp. 1–4.

[14] J. Chakareski, S. Naqvi, N. Mastronarde, J. Xu, F. Afghah, and A. Razi, "An energy efficient framework for UAV-assisted millimeter wave 5G heterogeneous cellular networks," *IEEE Trans. on Green Communications and Networking*, vol. 3, no. 1, pp. 37–44, Mar. 2019.

[15] S. Iellamo, J. J. Lehtomaki, and Z. Khan, "Placement of 5G drone base stations by data field clustering," in *Proc. of IEEE VTC Spring*, Sydney, NSW, Australia, Nov. 2017, pp. 1–5.

[16] R. de Silva and S. Rajasinghege, "Optimal desired trajectories of UAVs in private UAV networks," in *Proc. of IEEE ATC*, Ho Chi Minh City, Vietnam, Dec. 2018, pp. 310–314.

[17] L. Ferranti, F. Cuomo, S. Colonnese, and T. Melodia, "Drone cellular networks: Enhancing the quality of experience of video streaming applications," *Elsevier Ad Hoc Networks Journal*, vol. 78, pp. 1–12, 2018.

[18] F. Cheng, S. Zhang, Z. Li, Y. Chen, N. Zhao, F. R. Yu, and V. C. M. Leung, "UAV trajectory optimization for data offloading at the edge of multiple cells," *IEEE Trans. on Vehicular Technololgy*, vol. 67, no. 7, pp. 6732–6736, July 2018.

[19] E. Kalantari, M. Z. Shakir, H. Yanikomeroglu, and A. Yongacoglu, "Backhaul-aware robust 3D drone placement in 5G+ wireless networks," in *Proc. of IEEE ICC Workshops*, Paris, France, July 2017, pp. 109–114.

[20] M. Gapeyenko, V. Petrov, D. Moltchanov, S. Andreev, N. Himayat, and Y. Koucheryavy, "Flexible and reliable UAV-assisted backhaul operation in 5G mmWave cellular networks," *IEEE JSAC*, vol. 36, no. 11, pp. 2486–2496, Nov. 2018.

[21] R. Gangula, O. Esrafilian, D. Gesbert, C. Roux, F. Kaltenberger, and R. Knopp, "Flying rebots: First results on an autonomous UAV-based LTE relay using OpenAirinterface," in *Proc. of IEEE SPAWC*, Kalamata, Greece, June 2018.

[22] P. Ladosz, H. Oh, and W.-H. Chen, "Optimal positioning of communication relay unmanned aerial vehicles in urban environments," in *Proc. of IEEE ICUAS*, Arlington, VA, USA, July 2012, pp. 1140–1147.

[23] L. Bertizzolo, X. Tran, Tuyen, B. Amento, B. Balasubramanian, R. Jana, H. Purdy, Z. Yu, and T. Melodia, "Live and let live:flying UAVs without affecting terrestrial UEs," in *Proc. of ACM HotMobile*, Austin, TX, USA, Mar. 2020.

[24] L. Bertizzolo, M. Polese, L. Bonati, A. Gosain, M. Zorzi, and T. Melodia, "mmBAC: Location-aided mmWave backhaul management for UAV-based aerial cells," in *Proc. of ACM mmNets*, Los Cabos, Mexico, Oct. 2019.

[25] L. Bertizzolo, S. D'Oro, L. Ferranti, L. Bonati, E. Demirors, Z. Guan, T. Melodia, and S. Pudlewski, "SwarmControl: An automated distributed control framework for self-optimizing drone networks," in *Proc. of IEEE INFOCOM*, Toronto, ON, Canada, July 2020.

[26] F. D'Alterio, L. Ferranti, L. Bonati, F. Cuomo, and T. Melodia, "Quality aware aerial-to-ground 5G cells through open-source software," in *Proc. of IEEE GLOBECOM*, Waikoloa, HI, USA, Dec. 2019.