ACCEPTED FROM OPEN CALL

AirSlice: A Network Slicing Framework for UAV Communications

Zhenhui Yuan and Gabriel-Miro Muntean

Network slicing technology delivers tailored network resources to individual applications on demand. The authors study the network demands of the most popular UAV applications. These applications are classified according to their quality of service requirements. A novel network slicing framework, AirSlice, is designed in the 5G domain to enable differentiated QoS support for UAV applications.

Abstract

Network slicing technology delivers tailored network resources to individual applications on demand. This article first studies the network demands of the most popular UAV applications. These applications are classified according to their quality of service (QoS) requirements. A novel network slicing framework, AirSlice, is designed in the 5G domain to enable differentiated QoS support for UAV applications. Proof-of-concept deployment of AirSlice is presented. AirSlice intends to inspire new businesses, new operations, and new experiences for both operators and UAV service providers.

Motivation and Background

Gartner (https://www.gartner.com/en/newsroom/press-releases/2017-02-09-gartner-says-almost-3-million-personal-and-commercial-UAVs-will-be-shipped-in-2017) and the Federal Aviation Administration (FAA) (https://www.faa.gov/uas/) predict that the global market for unmanned aerial vehicles (UAVs) will reach US$11.2 billion in revenue and 7 million in shipments by 2020. Advancements in communication, computing, embedded systems, and aerial engineering are boosting UAVs' capabilities in terms of higher payload, longer flight times, redundant sensors, and artificial intelligence support. These onboard technologies enable UAVs to become key platforms for emerging services, especially in the emerging 5G networks. The authors have already studied diverse opportunities and challenges of using mobile networking technologies with UAVs in a general 5G network context [1, 2]. Other researchers have considered using UAVs to support network services in specific circumstances such as those following natural disasters [3].

The 5G wireless networks are not only set to offer increased capacity, lower latency, and faster speeds. They are designed to provide support for a wide range of rich applications, consumer and business-oriented alike, in diverse markets related to transportation, energy, healthcare, and so on. To achieve this, a novel architecture addressing the requirements of the latest increasingly complex and highly diverse services is needed. A potential solution is employing network slices, which create support for multiple logical networks to exist over a shared physical network infrastructure. This infrastructure can be supported by different coexisting operators and network providers, each offering their resources. Diverse services should run in this ecosystem employing the end-to-end virtual network that best addresses their requirements. This emerging network slicing approach aims to transform the whole 5G network resource allocation in support of differentiated services. In a UAV network scenario, for instance, different network slices can be allocated to diverse UAV services based on their quality of service (QoS) requirements.

The International Telecommunication Union — Radiocommunication Standardization Sector (ITU-R) [4] and Third Generation Partnership Project (3GPP) TS23.501 [5] specify three classic 5G services: enhanced mobile broadband (eMBB), massive machine-type communications (mMTC), and ultra-reliable and low-latency communications (uRLLC) (or mission-critical communications). UAVs, used as emerging 5G devices, are currently attracting significant attention, mostly coming from telecom operators and manufacturers. In the context of the above list of 5G network services, a natural classification of UAV applications and their associated traffic is as follows:

- High-bandwidth-demanding applications, including high definition (HD) video streaming from UAVs (maps to eMBB).
- Long-term connectivity with small payload applications, including unmanned traffic management (UTM) data (corresponds to mMTC). mMTC refers to high density of devices in a certain area, for example, up to 300,000 mobile users in a cell [6].
- Applications involving short-term ultra-reliable communications with small payload, like control of UAVs in the beyond visual line of sight (BVLoS) (can be associated with uRLLC).
- Best effort applications which include the ones that generate and exchange all other traffic that does not require any special treatment such as infotainment data.

Infotainment data refers to information and entertainment data and in the context of UAVs includes audio-visual data, GPS positioning, Internet access, vehicle settings, and so on. This proposed traffic classification is in line with the 3GPP enhancements to support vehicle-to-everything (V2X) communications described in TR 22.886 [7].

This article introduces AirSlice, a novel network-slicing-based framework for UAV communications in a 5G network context. AirSlice provides support for service differentiation, enabling diverse treatment for four traffic classes, spanning from high

Digital Object Identifier: 10.1109/MCOM.001.1900325

Zhenhui Yuan is with Northumbria University; Gabriel-Miro Muntean is with Dublin City University.

UAV applications	Task	Data rate, uplink	Data rate, downlink	End-to-end delay (IoT data)	End-to-end delay (control data)	Positioning accuracy	Height
Package delivery	Auto-pilot flight; real-time HD video for tracking	25Mb/s	300 kb/s	200 ms	20–50 ms	0.5m	120m
Agriculture	Farmland monitoring	20Mb/s	300 kb/s	200 ms	20–50 ms	0.5m	120m
Surveillance; search and rescue	Video delivery	25Mb/s	300 kb/s	200 ms	20–50 ms	0.5m	120m
Live broadcasting	4K/8K video delivery	25 Mb/s (4K) 100 Mb/s (8K)	600 kb/s	200 ms	20–50 ms	0.5m	120m
Swarm	Light show; coordinated search and rescue	1 Mb/s	300 kb/s	300 ms	20–50 ms	0.5m	200m

Table 1. UAV applications network requirements. Targets for 2020 suggested by IMT-2020 (5G) Promotion Group in China.

bit rate data to ultra-reliable short-term traffic. The differentiated support is enabled by network slices according to the diverse UAV traffic classes' QoS requirements. Network slicing is a key technology that is included in the domain of 5G. In this article, by default, UAVs are connected to 5G backhaul. These UAVs are then allocated a certain amount of network resources based on the proposed allocation, or slicing, rules, which is the QoS bound.

The article is structured as follows. First, challenges related to 5G and aspects related to its readiness for UAV communications are discussed. We then present network slicing in the context of current 5G efforts and introduce the AirSlice framework for traffic differentiation support. Following that, we describe the AirSlice cloud demo. Finally, security aspects are discussed, before conclusions and future work directions are presented.

Is 5G Ready for UAVs?

This section discusses the potential capabilities and challenges of deploying 5G for UAV services. Table 1 presents network requirements of five typical UAV services as specified by the IMT-2020 5G Promoting Group in China (http://www.imt-2020.org.cn/en). Package delivery with UAVs involves multiple flight tasks such as auto-pilot flight, real-time HD video of surroundings, and real-time control in case of emergency. Agriculture service with UAV focuses on 24/7 monitoring of farmland. Surveillance as well as search and rescue service with UAsV need real-time video delivery and fully autonomous flight capability. Live broadcasting with UAVs requires ultra-high bandwidth to satisfy 4k/8k video transmission. Uplink traffic includes both flight data and onboard sensor data, while downlink traffic generally refers to time-critical control commands. All UAV services recommend a data rate of 25 Mb/s and delay of 200 ms for uplink and data rate of 300 kb/s and delay of 20 ms for downlink. UAV swarm service, such as light show and coordinated search and rescue, is not bandwidth sensitive, but rather latency critical. UAV swarm use recommends data rate of 1 Mb/s and delay of 300 ms for uplink and 300 kb/s and delay of 50 ms for downlink. There are additional aspects such as the fact that flight height for all UAVs is limited to maximum 120 m in China.

Network QoS

The reality is that UAVs (e.g., drones) and smart mobile devices (e.g., smartphones) would share the same 5G networks and use the same spectrum, which is limited. However, the two types of devices do not have the same type of traffic and do not fully share the same network requirements for support of their services. Smart mobile device applications require mostly high bandwidth support, whereas UAVs run services that require mostly high reliability and low latency communications.For most mobile customers, cheap bills plus high bandwidth are always preferred. However, drone customers might also be concerned about reliability and latency. However, there will be fierce competition between UAVs and existing mobile devices for access to the limited network resources. Their QoS awareness toward device-specific aspects and traffic requirements is critical for their success.

5G Coverage and Interference in the Sky

Existing mobile communication networks are designed and optimized for terrestrial communications. In order to enable the best ground coverage, antennas of a base station are tilted down in order to reduce inter-cell interference. Hence, there might be blind network zones in the sky. UAVs can produce more significant uplink interference in 5G networks than mobile phones because the free space propagation increases the interference energy received at neighboring cells.

Enabling Network Slicing

Standardization bodies, operators, network vendors, and service providers are all working actively to design and deploy customized network slicing policies. Some of these efforts are discussed next.

Beyond 3GPP Standardization

3GPP TR 23.799 [8] describes a QoS framework as a generic overall solution and outlines nine work tasks needed for QoS support. It describes how the QoS functionality is distributed between core network (CN), radio access network (RAN), and user equipment (UE). A high-level system architecture and high-level functions have also been introduced. For example, TR 23.799 has included the following high-level network slicing solutions:

- Slicing of a public land mobile network (PLMN) with no UE visibility
- Slice selection process at the CN
- Multiple slice sharing by a UE
- Life cycle management of a network slice instance
- Slice selection based on usage class
- Roaming architecture of network slicing

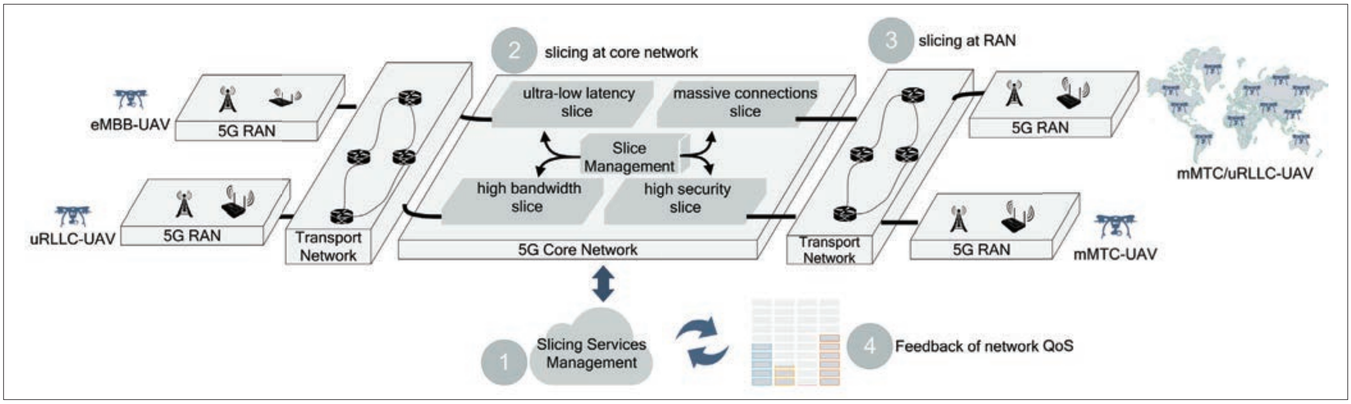

Figure 1. AirSlice architectural overview.

However, this 3GPP standard is not tailored for any applications, including those related to UAVs, leaving many issues open for future study.

The proposed AirSlice framework complements the work of 3GPP TR 23.799 R.14 by bridging the gap between the 3GPP standardized QoS framework and UAV-related services. In particular, based on UAV applications' specific requirements, AirSlice introduces support for QoS flow priority management, bandwidth and capacity estimation in the RAN, and slicing policy for multiple classes, including mission critical, bandwidth critical, and so on. This support is not included in the 3GPP TR 23.799 standard and enables QoS framework deployment in a UAV operational context.

Slicing Core and Access Networks

State-of-the-art network slice technologies concern both CNs and access networks. The key idea is to treat the core and access networks as a service. Taleb *et al.* [9] found the optimal number of virtual instances needed at the CN for a given service. The authors designed a new slicing technology that relates virtual instances of 4G and 5G to satisfy QoS of mobile traffic. Popovski *et al.* [10] proposed a communication-theoretic model that accounts for the requirements and characteristics of eMBB, mMTC, and URLLC services. Based on the communication models, the authors designed a heterogeneous non-orthogonal multiple access scheme to provide slicing for uplink traffic. Reliability diversity is the design principle that leverages the different reliability requirements across the services.

AirSlice: Framework and Architecture

This article introduces AirSlice, a new framework for 5G networks that creates and manages dedicated virtual networks for diverse UAV services. AirSlice coordinates network resources from the CN, transport network, and RAN. The major AirSlice goal is to identify the amount of network resources needed by particular UAV services and to allocate these resources. To resolve this challenge, first there is a need for a QoS model for the service to determine the amount of network resources required. Next, a resource allocation algorithm will translate the QoS requirements of the UAV service into practical network resources such as bandwidth. The AirSlice framework is described next.

AirSlice Life Cycle

Figure 1 illustrates the AirSlice architecture and indicates its life cycle. The architecture considers three major components, located at cloud, CN, and RAN. Figure 1 also shows how four fixed slices are mapped to three typical network slice categories: eMBB, uRLLC, and mMTC. In practice, the rules to configure new slices are flexible and can be determined by customers and/or network operators. For instance, such rules might consider billing cost, security levels, and so on.

The AirSlice life cycle includes the following four stages.

UAV Service-Oriented: QoS Management: The main contribution of AirSlice is the integration of UAV-specific QoS bound into the process of network management. Initially, it is important to configure the QoS bounds for different UAV services. AirSlice creates and classifies QoS profiles of UAV applications at the cloud. QoS network management recommendations are made by UAV service providers by considering not only traffic characteristics, but also high mobility and low reliability related aspects, which are specific to UAV communications. In practice, QoS is recommended by UAV service providers by jointly considering traffic characteristics and specific UAV tasks, as indicated in Table 1.

Mapping UAV-Specific Services QoS to Network Resources: Once QoS bounds of UAV services are configured, the next stage is to map these QoS bounds as the performance targets of network resource allocation techniques. For instance, by proper network management, a live broadcasting UAV shall receive 25 Mb/s uplink bandwidth for 4K video delivery. Network resources at the 5G CN are divided into slices according to the QoS profiles of UAV services. One or multiple slices could be mapped onto the same UAVs upon service request. For instance, a UAV with a search and rescue task is assigned both a uRLLC slice and an eMBB slice, for the purpose of real-time control and high definition video delivery, respectively.

Radio Resource Allocation: The core part of the AirSlice framework is the specific methodology for allocating, or slicing, network resources. In general, 5G network resources are distributed in access network, transport network, CN, and cloud. AirSlice performs radio resource allocation at the access network. Various solutions such as [11] can be employed to allocate bandwidth in

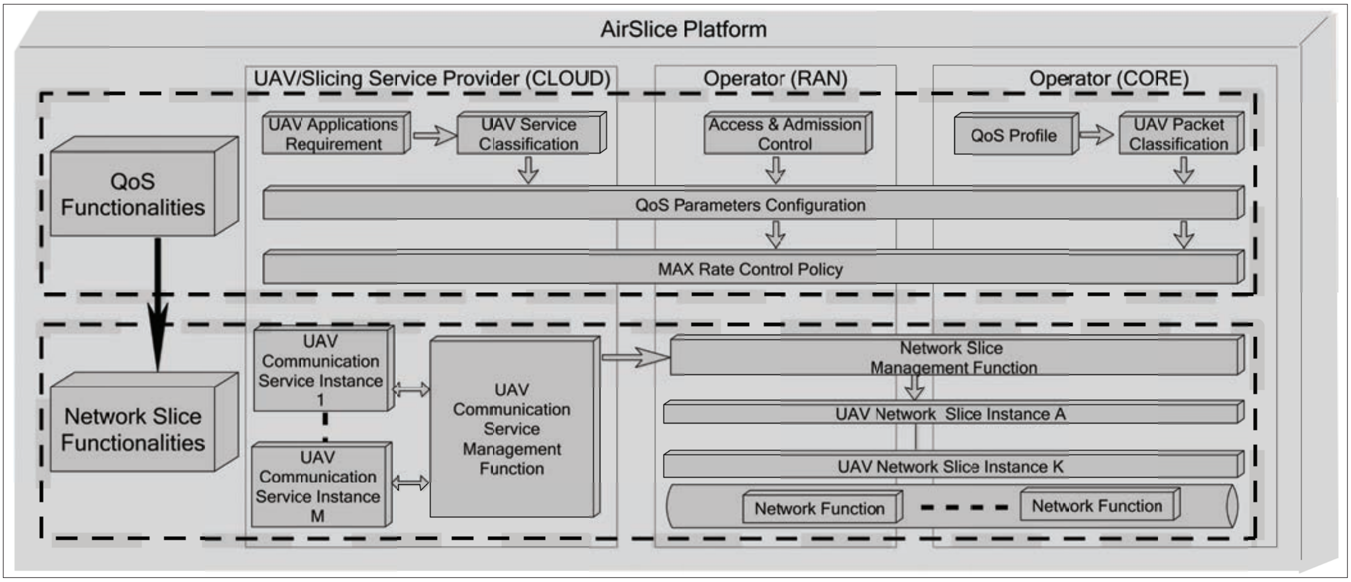

Figure 2. Block architecture of AirSlice.

accordance to QoS requirements and with various advantages and disadvantages.

Feedback: AirSlice monitors the achieved QoS levels of UAV applications and provides feedback to the QoS framework process at the cloud for dynamic adaptive adjustments.

A 3GPP system [12] is expected to provide QoS (e.g., reliability, latency, bandwidth) for various applications (e.g., emergency, medical) in 5G. 3GPP TR.22.864 [13] specifies that QoS is needed for equipment-to-equipment communications (e.g., UAV-to-UAV). A UAV-generated QoS flow may either be guaranteed flow bit rate (GBR) or non-guaranteed flow bit rate (non-GBR), depending on the UAV-QoS profile. The UAV-QoS profile of a certain flow is sent to the RAN and contains QoS parameters including 5G QoS identifier (5QI), downlink and uplink guaranteed flow bit rate (GFBR), downlink and uplink maximum flow bit rate (MFBR), downlink and uplink maximum packet loss rate, and so on. The proposed AirSlice framework uses per-flow instead of per-device differentiation in relation to the QoS granularity since any UAV might generate multiple flows that belong to different usage classes. For instance, a package delivery UAV sends back HD video traffic (eMBB class), while it receives real-time control commands (uRLLC class) to and from the delivery operation center.

AirSlice Block Architecture

A pre-condition for the deployment of AirSlice is that the operator maintains a white list of the types of UAV applications it can support for the customers. The customers can be third-party UAV service providers or the operator's own service departments. The customer can send requests to the operator for the support of a specific UAV service. These requests can include service requirements including QoS level, isolation, security, and so on. AirSlice makes decisions on whether the new service can be offered by checking the available 5G network resources in the specified geographical areas for the specified duration. AirSlice is designed and implemented with respect to 3GPP specifications on network slicing, 3GPP TR 23.799 and 3GPP TR 28.801. In order to support its core design ideas, AirSlice includes new modules, *QoS Functionalities* and *Network Slice Functionalities*, that are distributed between UAV service providers, slice service provider, and operators. Figure 2 illustrates the location of major AirSlice modules and discusses their role in the overall AirSlice context.

UAV/Slice Service Provider: Both the UAV service provider and slice service provider deploy AirSlice at the service level in the cloud. The UAV service provider creates QoS Functionalities by building QoS profiles for individual UAV applications. A QoS profile entry contains UAV application requirements and related service types. The network slice service provider configures network slice functionalities via the UAV Communication Service Management Function (UCSMF). The UCSMF receives the communication service related requirements from the UAV customers. The UCSMF translates the communication service requirement to network-slice-related requirements including network type, network capacity, QoS requirements, and so on. The UCSMF is an extention of the 3GPP Communication Service Management Function (CSMF). Multiple UAV communication service instances (UCSIs) are realized by the UCSMF module. A UCSI is an instance of a network slice that defines network characteristics and requirements of specific UAV services. Examples of network characteristics are ultra-low latency, high bandwidth, and so on. The UCSI extends the 3GPP service instance.

Operator (RAN and Core): Operators play a key role in the network slicing process, as it involves support of both RANs and CNs. RAN-side, QoS Functionalities are realized using an access and admission control scheme. The CN synchronizes QoS profiles with the UAV service providers and implements packet classifications per UAV flow. The operators are responsible for supporting QoS configurations and rate control policies during the life cycle of communications. The operator implements network slice support by using the Network Slice Management Function (NSMF) module. The NSMF manages and orchestrates instances of

NSMF determines the core network-related requirements and RAN-related requirements based on UAV customer requirements. The key function of NSMF is slice selection which handles UAV's attachment request by selecting an appropriate slice based on the UAV service type.

Figure 3. Real-life testbed with multi-access drones and diverse wireless networks: LoRa, WiFi, LTE, and 5G.

a network slice, the UAV network slice instance (UNSI). A set of UNSIs forms a logical network that meets the network characteristics required by the USI. Both the CN part and RAN part of a network slice can be shared by different UNSIs. The NSMF determines the CN-related requirements and RAN-related requirements based on UAV customer requirements. The key function of the NSMF is slice selection, which handles a UAV's attachment request by selecting an appropriate slice based on the UAV service type.

AirSlice Algorithm Implementation

A three-step algorithm for practical deployment of AirSlice is described next.

Step 1: Target a UAV service and derive its QoS requirements. Based on 3GPP TS 23.501, which specifies network statistics (e.g., packet loss rate, packet delay budget), identify QoS bounds to be observed. These QoS bounds are then used for network control policies such as adaptive scheduling, admission control, and queue management. In this article, the network QoS bounds are derived as indicated in Table 1.

Step 2: Compute the available bandwidth that a network can support for certain services under given channel conditions. IBM's Effective Capacity link model [14] is employed to model the wireless channel in terms of connection-level QoS metrics such as data rate, delay, and delay-violation probability. The channel model is then used to estimate the bandwidth for the QoS bounds.[1] The Effective Capacity link model employs the probability of non-empty buffer, marginal cumulative distribution function (e.g., Rayleigh-Ricean distribution), and a QoS exponent related to the Doppler spectrum. Based on this link model, derive the guarantees for each QoS metric (e.g., bit rate, delay) so that the network can tolerate a certain QoS-bound violation probability. This probability is dependent on the UAV service.

Step 3: Perform network slicing for different UAV services. This is achieved in two stages: one static and another one dynamic. During the static slicing stage, slice allocation is based on the bandwidth-delay product (BDP). As eMBB UAV services are associated with traffic with high BDP (e.g., high resolution video and/or large delivery latency), this stage selects network slices with high BDP. uRLLC UAV services produce lower BDP network traffic than that of eMBB services, so lower BDP slices will be statically allocated to them. During the dynamic slicing stage, BDP of the statically allocated network slice is changed in order to satisfy the QoS bounds as indicated in step 2. This is done by employing network adaptation technologies, such as adjusting physical layer modulations, to achieve different data rates.

AirSlice Cloud Instance

A network slice consists of a set of network functions, resources to run these network functions, policies, and configurations. It is not straightforward to test the feasibility of a certain network slicing scheme. However, a proof-of-concept AirSlice framework instance was deployed at the cloud. A graphic user interface was designed to visualize its operation. AirSlice was configured by taking input of UAV specifications, task description, access network, QoS bound, and so on. Key options and features include the following.

UAV Services Management: The system maintains profiles of typical UAV services such as search and rescue, video surveillance, UTM, swarm light show, remote control, delivery, and 3D mapping. It is allowed to load new services.

[1] The actual QoS metrics proved to be closely approximated by the QoS metrics predicted by IBM's Effective Capacity link model.

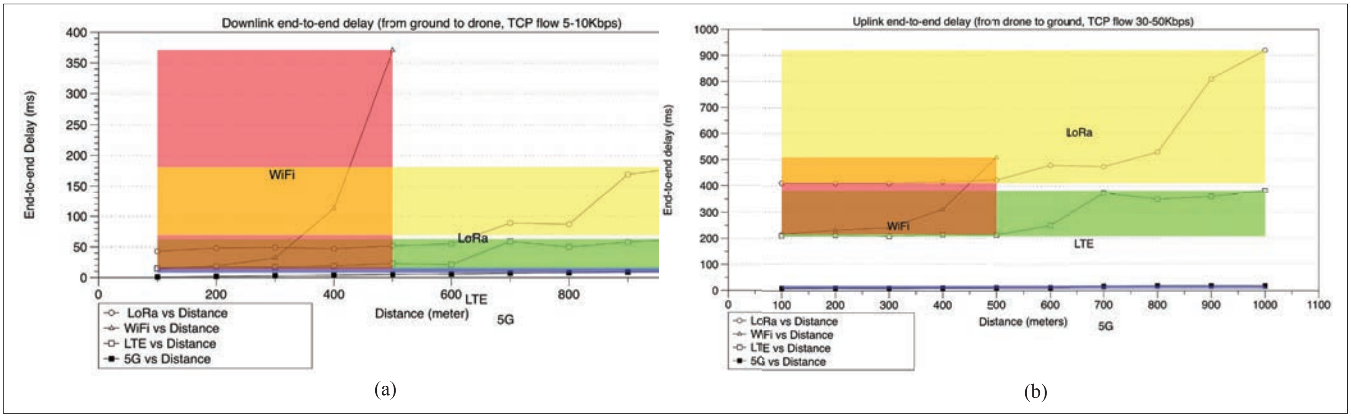

Figure 4. End-to-end delay for downlink and uplink TCP flows: a) downlink end-to-end delay (from ground to drone) of TCP flows; b) uplink end-to-end delay (from drone to ground) of TCP flows.

UAV Service Profile: The system allows configuration for the profile of a certain UAV service. The service profile includes parameters such as UAV type, network access technologies, number of active UAVs, flight speed, camera model, flight duration, and flight height.

Traffic Profile: Features of UAV-generated multimedia traffic can be input to indicate the traffic load. Default parameters include uplink telemetry throughput, image resolution, and video frame rate and resolution.

QoS Demands: The system supports setup of QoS requirement for certain UAV service in terms of throughput, latency, maximum latency variation, packet loss ratio, and network reliability.

Slice Recommendations: The system saves historical slice strategies for frequently used UAV services (e.g., UAV video surveillance). Whenever a new UAV service of the same type joins the network, the previously established slice is quickly assigned to the UAV. The system also supports modifications to network slices (e.g., adding, deleting, modifying) with minimal impact on active subscriber services.

AirSlice Real-Life Proof of Concept

Use Case Description

The proposed AirSlice is tested when slicing the radio resources for uRLLC services with QoS bound. Real-time control of UAVs typically employs a uRLLC service that requires ultra-reliable low-latency communications for both uplink and downlink. A drone light show is selected as the test use case as it involves regular downlink and uplink traffic. The downlink traffic, from the ground base station to UAVs, includes control commands for next movements (e.g., velocity, next position, emergency return). The uplink traffic, from UAVs to ground base station, includes rich flight data (GPS, battery life, velocity, flight log, etc). We evaluate the end-to-end latency between UAVs and ground base station in terms of communication distance. The latency bound recommended for downlink and uplink traffic in the drone light show use case are 50 ms and 300 ms, respectively.

Testbed Setup

Four different real networks, LoRa, WiFi, LTE, and 5G, were employed to mimic different levels of bandwidth capacity. These four network technologies are widely used by UAVs. LoRa has low bandwidth, but supports communication over large distance at low cost. WiFi is easily available and has enough bandwidth for most UAV traffic, but its communication distance is the shortest among the four access technologies. Both LTE and 5G can provide good bandwidth and coverage in cases where the budget is not limited. We demonstrate the remote control of a single UAV via the 5G network.

Figure 3 illustrates the test involving 100 light showing UAVs from RobSense. There are no other mobile devices than the 100 UAVs in the cell. Average latencies for the 100 UAVs are computed for both uplink and downlink communications. By default, each UAV is equipped with three radio modules, 5.8 GHz WiFi, 1.4 GHz LTE, and 433 MHz LoRa, which are connected separately to a ground WiFi base station (by Unify), an LTE femtocell (by Nokia), and a LoRa gateway (by RobSense). Additionally, we also test the latency of the real-time control of a single UAV via Huawei's 5G CPE (28 GHz) and China Mobile's 5G network in the city of Hangzhou, China.

Downlink traffic is sent from ground to the 100 UAVs at data rates between 5 and 10 kb/s, while uplink traffic is sent back to the ground at a data rate between 30 and 50 kb/s. The data rate is variable as the sampling rates of data sources are different; for instance, GPS packets are sent at 5 Hz and the battery life info is sent at 1 Hz. Both downlink and uplink traffic is delivered on top of TCP to guarantee delivery. Communication distance ranges from 100 m to 1000 m horizontally. The average flight height of the 100 UAVs is 90 m. The 100 UAVs were remotely controlled via a single ground control station (GCS), as shown in Fig. 3. The GCS shows real-time positions of the UAVs and includes flight options such as take off and landing. The GCS is then connected to the different wireless gateways, LoRa, WiFi, and LTE, via USB/Ethernet cables. These wireless gateways provide the last-mile network access.

Results Analysis

Figures 4a and 4b present the average communication latency variation with the distance for downlink and uplink, respectively. The plots are colored to indicate different virtual network slices. In both downlink and uplink, 5G can best

Future work may support offering Network Slice as a Service (NSaaS) by either operators or third-party organizations to customers. An operator might decide to slice the network using different criteria, such as service types, QoS level, mobility, and so on. Certain service types could be further classified to allow for finer granularity to UAV customers.

satisfy the communication latency (below 50 ms for downlink and 300 ms for uplink) across the whole communication distance from 100 m to 1000 m. LTE is within the latency bound for distances below 700 m, beyond which the latency increases greatly. This might be caused by interference from the neighboring cells; this is a common problem for existing operator networks. LoRa satisfies the downlink latency bound of 50 ms within a distance of 600 m, but finds it difficult to meet the uplink latency bound of 300 ms due to its limited bandwidth. WiFi satisfies both downlink and uplink latency bounds within 300 m distance only, but not beyond that, due to the poor signal coverage. In summary, communication distance, amount of traffic, and downlink and uplink QoS bound shall all be considered when slicing the network for a UAV service.

Conclusion and Future Work

The telecom industry has already launched commercialization of network slicing solutions. Huawei released the world's first 5G Slice Mall, comprising operators and industry partners, offering 5G network slicing services. Huawei demonstrated a network-slicing-based tele-medicine service developed with Deutsche Telekom and showcased a virtual reality game using 5G network slicing in partnership with Telefonica.

In such a context, this article is timely. It presents a novel AirSlice framework, which is designed to provide network slicing support for UAV communications in 5G networks. One of the major features of AirSlice is employment of network slicing to support QoS differentiation according to traffic classes. The article has presented standardization and research efforts as well as issues and challenges of design and deployment of AirSlice. It has also presented a proof-of-concept cloud deployment of AirSlice.

Future work may support offering network slice as a service (NSaaS) by either operators or third-party organizations to customers. An operator might decide to slice the network using different criteria, including service types, QoS level, mobility, and so on. Certain service types could be further classified to allow for finer granularity to UAV customers. For instance, eMBB traffic may be further classified into diverse QoS ranges. As specified in 3GPP TR 28.801 [15], typical network slice services could be characterized by other aspects including radio access technology, bandwidth, end-to-end latency, guaranteed and non-guaranteed QoS, security level, and so on. However, depending on the service offering, slice service providers impose limits on the network slice characteristics that will be exposed to the UAV customers. Finally, future work also considers building and testing a large-scale 5G network with many AirSlice-enabled UAVs, once commercial UAV-deployable 5G modules become available.

Acknowledgment

Gabriel-Miro Muntean would like to acknowledge the support of the Science Foundation Ireland (SFI) Research Centres Programme under Grant Numbers 12/RC/2289 (Insight Centre for Data Analytics) and 16/SP/3804 (ENABLE).

References

[1] Z. Yuan *et al.*, "Ultra-Reliable IoT Communications with UAVs: A Swarm Use Case," *IEEE Commun. Mag.*, vol. 56, no. 12, Dec 2018, pp. 90–96.
[2] Z. Yuan *et al.*, "ComProSe: Shaping Future Public Safety Communities with ProSe-Based UAVs," *IEEE Commun. Mag.*, vol. 55, no. 12, Dec 2017, pp. 165–71.
[3] N. Zhao *et al.*, "UAV-Assisted Emergency Networks in Disasters," *IEEE Wireless Commun.*, vol. 26, no. 1, Feb. 2019, pp. 45–51.
[4] ITU-R Rec. M.2410-0, "Minimum Requirements Related to Technical Performance for IMT2020 Radio Interface(s)," Nov. 2017.
[5] 3GPP TS 23.501, "System Architecture for the 5G System," Rel. 15.4.0, Dec. 2018.
[6] F. Boccardi *et al.*, "Five Disruptive Technology Directions for 5G," *IEEE Commun. Mag.*, 2014, pp. 74–80.
[7] 3GPP TR 22.886, "Study on Enhancement of 3GPP Support for 5G V2X Services," Rel. 15, Sophia Antipolis, France, Dec. 2016.
[8] 3GPP TR 23.799, "Study on Architecture for Next Generation System," v14.0.0, 16 Dec. 2016.
[9] M. Bagaa *et al.*, "Coalitional Game for the Creation of Efficient Virtual Core Network Slices in 5G Mobile Systems," *IEEE JSAC*, vol. 36, no. 3, Mar. 2018, pp. 469–84.
[10] P. Popovski *et al.*, "5G Wireless Network Slicing for eMBB, URLLC, and mMTC: A Communication-Theoretic View," *IEEE Access*, vol. 6, Sept, 2018, pp. 55,765–79.
[11] M. Asif-Ur-Rahman *et al.*, "Toward a Heterogeneous Mist, Fog, and Cloud-Based Framework for the Internet of Healthcare Things," *IEEE Internet of Things J.*, vol. 6, no. 3, June 2019, pp. 4049–62.
[12] 3GPP TS 23.107, "Quality of Service (QoS) Concept and Architecture," V15.0.0, June 2018.
[13] 3GPP TR 22.864, "Feasibility Study on New Services and Markets Technology Enablers — Network Operation," V14.1.0, Sept. 2016.
[14] D. Wu and R. Negi, "Effective Capacity: A Wireless Link Model for Support of Quality of Service," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, July 2003, pp. 630–43.
[15] 3GPP TR 28.801, "Study on Management and Orchestration of Network Slicing for Next Generation Network," V15.1.0, Jan 2018.

Biographies

Zhenhui Yuan received his Ph.D. degree in electronic engineering from Dublin City University (DCU), Ireland, in 2012. His research focuses on wireless communications with robotics. He was a postdoctoral researcher (2012-2014) on a project funded by Ericsson and Enterprise Ireland and a 3GPP researcher on 5G (2014–2015) at Huawei (Shanghai) Technologies Co. Ltd. In 2015, he became an associate professor at Hangzhou Dianzi University, China, and co-founder and chief scientist at RobSense (Hangzhou) Technology Co. Ltd. Currently, he is with Northumbria University, United Kingdom. He is an Associate Editor for *IEEE Access* and *IEEE Transactions on Consumer Electronics*. His research website is at www.zhenhuiyuan.com.

Gabriel-Miro Muntean is an associate professor with the School of Electronic Engineering, DCU, and co-director of the DCU Performance Engineering Laboratory. He has published over 350 papers in top-level international journals and conferences, authored three books and 18 book chapters, and edited seven additional books. His research interests include quality, performance, and energy issues related to rich media delivery, technology-enhanced learning, and other data communications over heterogeneous networks. He is an Associate Editor of *IEEE Transactions on Broadcasting*, and the Multimedia Communications Area Editor of *IEEE Communications Surveys & Tutorials*. He is the coordinator of the EU-funded project NEWTON (http://www.newtonproject.eu).