March 2022

AERIAL EDGE ORCHESTRATOR

Sudha Katgeri

Lubna Rasheed

Jasna Jayaraj

Recommended Citation

Katgeri, Sudha; Rasheed, Lubna; and Jayaraj, Jasna, "AERIAL EDGE ORCHESTRATOR", Technical Disclosure Commons, (March 14, 2022)
https://www.tdcommons.org/dpubs_series/4967

AERIAL EDGE ORCHESTRATOR

AUTHORS:
Sudha Katgeri
Lubna Rasheed
Jasna Jayaraj

ABSTRACT

Techniques are presented herein that support an optimized approach to deploying and re-deploying aerial edge nodes to dynamically meet fluctuating network needs, thus ensuring a better utilization and reuse of network resources. Aspects of the presented techniques include providing an Unmanned Aerial Edge Node (UA-EN) that may be managed by an Aerial Edge Orchestrator (AEO) in which the AEO can serve as a 'single pane of glass' for the deployment, management, and teardown of UA-ENs. The AEO may employ an automated workflow to provision UA-ENs, over a secure communication channel, to direct them, as needed, to congested locations. The AEO may also ensure the teardown and return, recharging, and reuse of UA-ENs through automated workflows and computations. UA-ENs themselves have the built-in intelligence to detect and avoid obstacles in real time (e.g., by overriding a provisioned flight path during flight), are able to establish mesh-based communication with enterprise wireless deployments, and may switch between satellite, Long-Term Evolution (LTE), and wireless backhauls.

DETAILED DESCRIPTION

Networks are traditionally designed with a maximum possible capacity for a particular coverage area rather than being based on usage, invariably resulting in a capacity that exceeds usage. In most real-world deployments, especially in outdoor locations and in public spaces, network utilization fluctuates over time, even during the span of a single day. Such fluctuations result in the fixed network being significantly underutilized during off-peak hours.

As one example, an open-air stadium with a capacity of 10,000 does not need to deploy a network that is capable of constantly servicing 10,000 people. As another example, a hiking trail is typically busier during evenings and weekends as opposed to midday on a work day and therefore it does not need to accommodate 100 connections at all times during the day.

Additionally, user density within a deployment can vary over time dependent on location and events associated with a location. For example, in a conference center an area that is hosting a popular TED Talk will see denser demand than other areas of the center. Similarly, a scheduled popular event in a park will see more users than the rest of the park.

Existing solutions to resolve density demand variation focus on optimizing the usage of deployed fixed network resources. Techniques presented herein discuss augmenting the service edge with dynamic deployment of a new type of edge node, namely an Unmanned Aerial Edge Node (UA-EN), that is capable of communicating over Long-Term Evolution (LTE) and enterprise wireless networks and which is managed and deployed, as needed, by an Aerial Edge Orchestrator (AEO).

Figure 1, below, illustrates elements of an exemplary architecture according to aspects of the techniques presented herein.

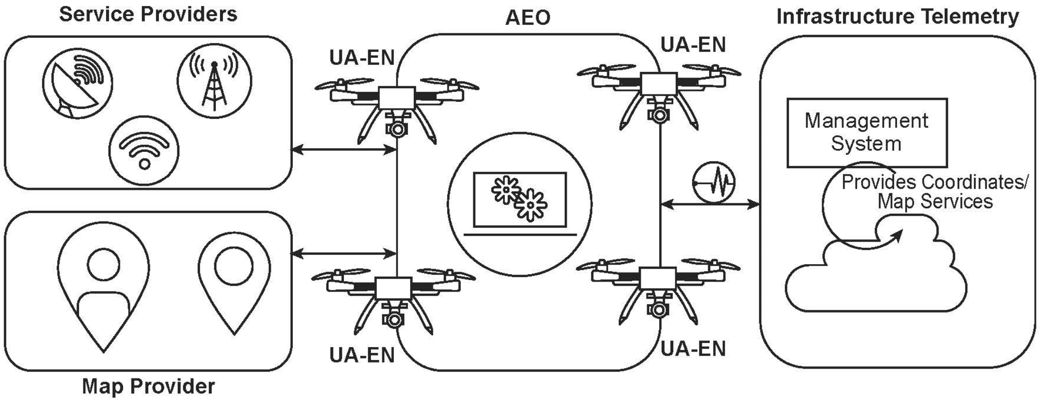

Figure 1: Exemplary Architecture

The different elements that are depicted in Figure 1, above, will be described in the following narrative. As the name suggests, a UA-EN may be primarily responsible for providing last-mile access to end users in areas of congestion as directed by an AEO. The hardware and software features of a UA-EN include a number of elements.

A radio-frequency identification (RFID) asset tag or a serial number uniquely inventories each UA-EN. Furthermore, such an identifier may be used to authenticate a UA-EN over an LTE network or an Institute of Electrical and Electronics Engineers (IEEE)

802.11 standard-based network. An LTE or a 3rd Generation Partnership Project (3GPP) fifth generation (5G) chipset can be provided to support communication from a UA-EN to an AEO during flight between locations and when Wi-Fi access is not available. As well, an 802.11 a, b, g, n, ac, ax, 6, 6E, or 7 chipset (or any other future 802.11 variant) can be provided to serve as a wireless backhaul to communicate with infrastructure access points through, for example, the Adaptive Wireless Path Protocol (AWPP). A Manufacturing Installed Certificate (MIC) on a UA-EN can serve to establish an encrypted communication channel to an AEO.

An onboard Global Positioning System (GPS) facility can provide location coordinates and help to track a UA-EN during flight. Additionally, an obstacle sensor can be used to prevent in-flight collisions as the UA-EN flies towards a destination. Further, a camera with a built-in microphone can be utilized to capture and store real-time images and video during flight. Also, a modular slot may be used to embed an Ethernet card to plug in external input devices (such as a microphone, an additional camera, sensors, etc.). Further, a Universal Serial Bus (USB) port can be used to extend local storage. Finally, the UA-EN includes a battery that, in some instances, can be powered by a solar panel (to continue charging the battery, when it is feasible).

In addition to providing edge access, the UA-EN can support a number of functionalities. For example, the UA-EN can include intelligence that may be responsible for quick real-time decisions to ensure a smooth flight. In the case of an obstacle sensor being triggered, the UA-EN can respond by avoiding the obstacle. Further, the UA-EN can employ an onboard camera to assess the obstacle and generate maneuverability code that will override the trajectory that was programmed by an AEO. Such a modification can be governed and optimized with the intent that the new trajectory can align with the trajectory that was provisioned by the AEO.

During operation, the UA-EN may also be responsible for triggering a protocol (e.g., AWPP) and establishing mesh-based communication with an infrastructure access point. Further, the UA-EN may hand-off backhaul communication from an LTE network to a wireless network and leverage a payload (e.g., AWPP or Control and Provisioning of Wireless Access Points (CAPWAP)) to communicate any non-Radio Frequency (RF) telemetry to an AEO through a monitoring application.

It is important to note that some existing 3GPP solutions involving unmanned aerial system (UAS) functionality focus on the UAD leveraging a cellular infrastructure for communication within the system and with UAS traffic management (UTM). However, under aspects of the techniques presented herein, an AEO and the UA-EN can access multiple network infrastructures including, but not limited to, cellular facilities. For example, the UA-EN can be provisioned with onboard radios to leverage wireless, cellular, and satellite infrastructures for end user data forwarding and for communicating with an AEO. This enables UA-ENs to extend an existing infrastructure as well as augment an existing infrastructure with alternate networks to service end users in case of congestion on one network type. Additionally, an AEO is also integrated with the different providers to communicate with and track UA-ENs.

Under aspects of the techniques presented herein, an AEO is the control node which is responsible for deploying and managing UA-ENs. Such a capability may encompass a number of elements. For example, the AEO may integrate with a network monitoring application and network assurance telemetry receivers to monitor enterprise telemetry key performance indicators (KPIs) and establish a secure communication channel with an authenticated UA-EN in which the secure communication channel can be utilized for provisioning as well as receiving telemetry over any network infrastructure, including wireless, satellite and/or cellular network (e.g., LTE, 5G, 6G, etc.).

The AEO determines the UA-EN that is best suited for deployment for a specific circumstance. Among other things, the AEO may consider the current battery charge percentage of a UA-EN (which may be referred to as C), the time that will be taken for a UA-EN to cover the distance between the current location of each UA-EN and a destination location (which may be referred to as t1), the rate of battery depletion during flight (which may be referred to as fbdr1/hr), and the rate of battery depletion while hovering and providing mesh neighbor relationship and servicing users (which may be referred to as obdr/hr). By employing the above, the service battery life (SBR) for each UA-EN may be calculated as: `SBR = [C - (t1*(fbdr1/hr))] / obdr` in which, by default, a higher numerical SBR for a UA-EN is considered to be most optimal.

If historic data (such as peak hours and corresponding demand (e.g., between 5 p.m. and 7 p.m. on weekends), the duration of specific events that are scheduled in a location (e.g., a movie night every Friday during the summer), etc.) is available where an increase in demand may be predicted, an automated workflow encompassing the sequential deployment and teardown of UA-ENs may be programmed to ensure uninterrupted service.

The AEO can further communicate with satellite image providers to retrieve and process satellite images to compute the trajectory for the UA-ENs. Such a trajectory computation involves identifying each obstacle (e.g., buildings, trees, etc.), estimating their dimensions, and generating code that provides commands (such as roll, pitch, yaw, turn, stop, etc.) and provisioning the same to the UA-EN that was selected, as described above.

Accordingly, aspects of the techniques presented herein support an intelligent UA-EN which can self-compute to correct, in real time, its flight path around obstacles. A UA-EN may apply Artificial intelligence (AI) and machine learning (ML) analysis to sensor data, images, and live video feeds from the on-board camera(s). The updated path and obstacle details are relayed to an AEO so subsequent UA-EN deployments may leverage a more optimized flight path. In addition to communicating with an AEO, UA-ENs may establish a peer-to-peer mesh network over a wireless network and continuously share their health metrics, location, networks being serviced, etc. All of the communications that are exchanged between UA-ENs may be synchronized to an AEO to achieve a coordinated, updated, and self-managed system.

Once a UA-EN has switched over from LTE to wireless backhaul, it may rely on streaming telemetry from a monitoring application to receive non-RF parameters such as battery life, images, videos, hover height, etc.

The AEO may tear down a UA-EN (based on, for example, network usage KPIs and/or the battery life of the UA-EN) either for redeployment to another location or for return to a dock for charging, as applicable. A teardown decision may be based on a drop in the total associated client count, congestion markers (such as drops, retransmissions, etc.) not being observed for an extended period of time, the time of the day (vis-à-vis, for example, historic data of non-peak hours, a venue’s closing, etc.), and the battery life of the UA-EN.

If network KPIs indicate a drop in usage, a trajectory may be recalculated to either redeploy the same UA-ENs to a different location (as was described above) where a need exists or return the UA-ENs to a dock (as will be described below).

Since UA-ENs are battery powered, they require close monitoring to ensure uninterrupted service and to ensure that a UA-EN can return to its dock for recharging. In determining when a UA-EN ‘return to dock’ workflow needs to be triggered, the AEO may consider the rate of battery life depletion during flight back to a docking station (which may be referred to as fbdr2/hr), the duration of flight back to a docking station (which may be referred to as t2), and an oversight threshold (which may be referred to as N) that is user configurable with a minimum allowed value of 15%. By employing the above, the Minimum Battery Life of a UA-EN to return to dock/re-home may be calculated using the following formula:

```
Minimum Battery Life (to dock/re-home) = (t2*(fbdr2/hr)) + (t1*(fbdr1/hr) + N
```

A UA-EN will be deployed only when its current battery life exceeds the Minimum battery Life to return to dock/re-home.

The AEO can manage a number of different workflows. For example, a first workflow encompasses priming a UA-EN. When it first comes out of a box, a UA-EN can leverage Bluetooth, LTE, and manufacturing-provisioned service set identifier (SSID) beacons to initiate communication with an AEO. The serial number of the UA-EN and/or a certificate exchange may be used to authenticate the UA-EN and establish secure communication. After a secure channel is established, the AEO may enable the UA-EN’s GPS and capture the docked location of the UA-EN as a home position. The UA-EN’s RF chipset may be programmed by the AEO with a bridge-group name, a data rate, a channel, etc. for the UA-EN to establish an 802.11 wireless mesh over a (e.g., AWPP) protocol.

A second workflow encompasses the deployment of a UA-EN. During operation, the AEO can subscribe to RF network congestion parameters such as client association failures, call admission control (CAC) failures, quality of service (QoS) drops, etc. as provided by a monitoring application. Additionally, the AEO continuously monitors the battery lifetime of each UA-EN. The deployment of UA-ENs is triggered based on such criteria.

After a deployment trigger is identified, a trajectory is calculated for all of the UA-ENs as described above and an optimal UA-EN is selected for deployment. The trajectory is provisioned to the UA-EN along with the motion command set that is needed for the

UA-EN to trace the trajectory. After it is deployed, the UA-EN travels along the predetermined path executing a set of spin, roll, yaw, stop, pitch, etc. commands as provisioned by the AEO. However, the UA-EN monitors an obstacle sensor and an onboard camera to react to unexpected obstacles and maneuver around the obstacles by overriding a portion of the trajectory code.

Once the UA-EN is within "hearing" distance of an infrastructure access point it may leverage a (e.g., AWPP) protocol to establish a wireless mesh connection and then initiate a switchover from LTE to wireless backhaul. Following such a switchover, the UA-EN may stream telemetry (both RF-related and non-RF-related) back to the AEO through the infrastructure access point and the monitoring application (and on to the AEO).

A third workflow encompasses the teardown of a UA-EN. Within this workflow, an AEO constantly monitors the network KPIs and when it observes decongestion markers over a period of time, it will initiate a teardown workflow.

A first teardown step encompasses calculating a trajectory – from a UA-EN to either its new destination or to a docking station. In addition, if a battery depletion notification for a currently-deployed UA-EN is the trigger, the AEO may reinitiate the deployment workflow to deploy another UA-EN (e.g., a new UA-EN) to ensure continued service.

The trajectory that was calculated for the original UA-EN is then provisioned along with instructions to the UA-EN to switchover from wireless to LTE and schedule deployment. The UA-EN switches over to LTE (with potentially a blacklist of an RF link for a fixed duration to avoid a flip-flop and packet loss), re-establishes LTE communication with the AEO, and then traces the trajectory as directed.

Figure 2, below, depicts elements of an exemplary sequence diagram according to aspects of the techniques presented herein and reflective of the narrative that was presented above.

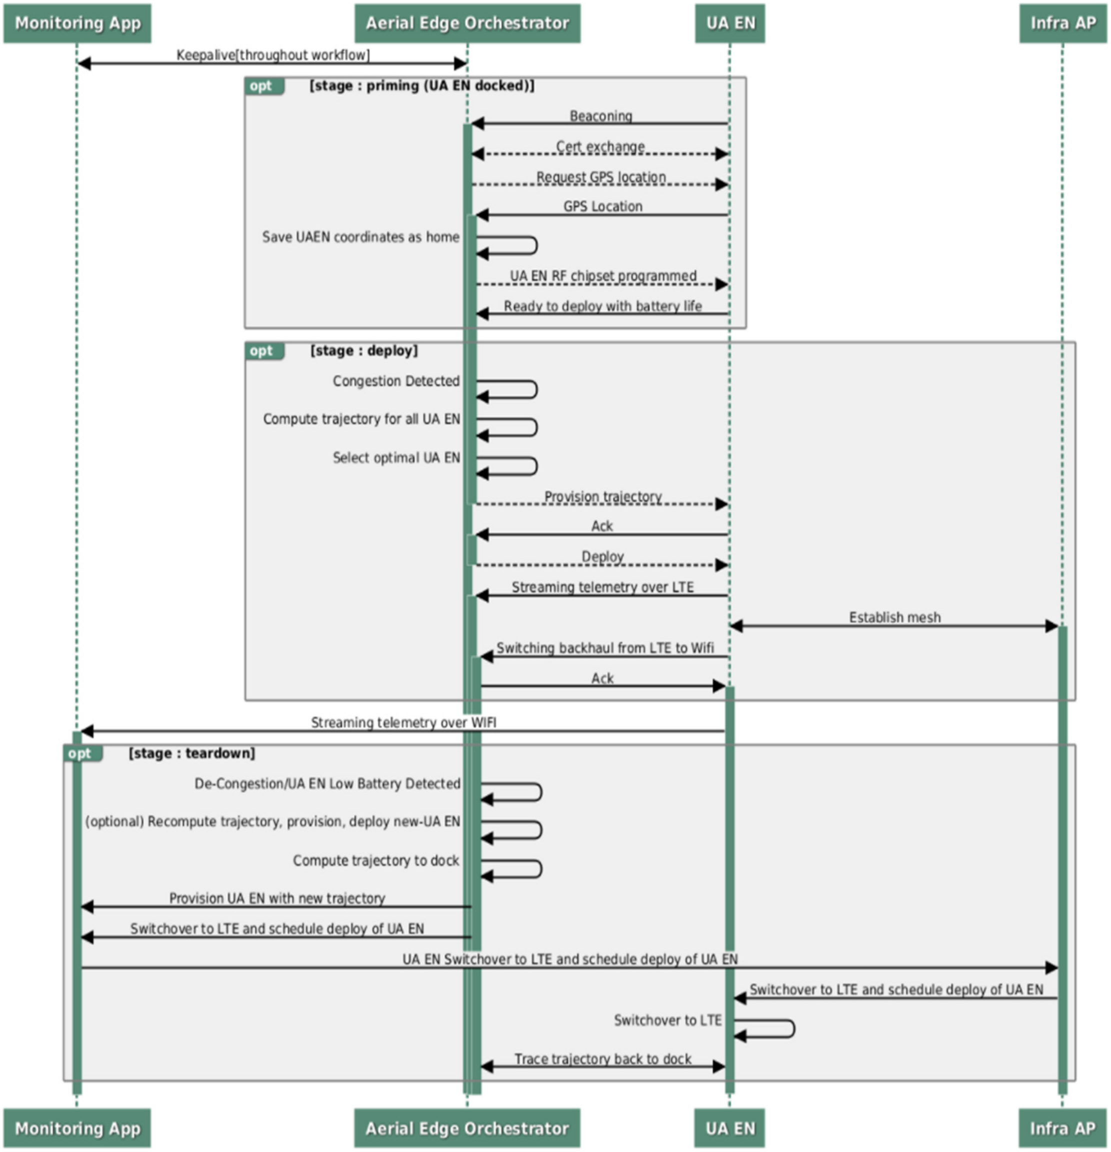

Figure 2: Exemplary Sequence Diagram

It is important to note that under elements of existing solutions battery life is monitored to determine when a UA-EN should be re-homed. In contrast, under aspects of the techniques presented herein, besides battery life, network demand is considered to determine if a UA-EN needs to discontinue service. Further, the presented techniques support an automated orchestration to ensure that subsequent UA-ENs continue to be deployed to meet network demand and service-level agreements (SLAs) for any planned or unplanned service outage of a UA-EN. The mesh network between UA-ENs is further leveraged to triangulate the location of any UA-EN that experiences an unplanned outage due to hardware failure or some environmental event. An AEO may provide application programming interfaces (APIs) that fixed network providers can utilize to understand the demand and usage for a given location in order to maintain, append, or redistribute their deployments.

In summary, techniques have been presented that support an optimized approach to deploying and re-deploying of aerial edge nodes to dynamically meet fluctuating network needs, thus ensuring a better utilization and reuse of network resources. Aspects of the presented techniques include a UA-EN that may be managed by an AEO in which the AEO serves as a 'single pane of glass' for the deployment, management, and teardown of UA-ENs. The AEO may employ an automated workflow to provision UA-ENs over a secure communication channel to direct them, as needed, to congested locations. Further, the AEO may also ensure the teardown and return, recharging, and reuse of UA-ENs through automated workflows and computations. UA-ENs themselves have the built-in intelligence to detect and avoid obstacles in real time (e.g., by overriding a provisioned flight path during flight), are able to establish mesh-based communication with enterprise wireless deployments, and may switch between satellite, LTE, and wireless backhauls to service congested areas dynamically.