FFI-RAPPORT

16/01132

LINE EW-UAS Link and Network Performance

passive performance analysis from the November 2015 flights

–

Erlend Larsen
Øivind Kure
Lars Landmark
Jostein Sander

LINE EW-UAS Link and Network Performance
passive performance analysis from the November 2015 flights

Erlend Larsen
Øivind Kure
Lars Landmark
Jostein Sander

Norwegian Defence Research Establishment (FFI) 13 December 2016

Keywords
Antenner
Mobilitet
Nettverk
Radiokommunikasjon
UAV (Unmanned Aerial Vehichles)

FFI-rapport
FFI-RAPPORT 16/01132

Project number
136701

ISBN
P: 978-82-464-2754-6
E: 978-82-464-2755-3

Approved by
Jan erik Voldhaug, *Research Manager*
Anders Eggen, *Director*

Summary

The communication capacity needs in the Norwegian Armed Forces (NAF) will increase in the coming years. Radios with high capacity often have strict requirements for Line-of-Sight (LoS). Unmanned Aerial Vehicles (UAVs) can mitigate this increased LoS requirement, and UAVs with varying levels of autonomy will be developed. Autonomy for UAVs tailored for communication purposes requires an intimate understanding of the behavior of networks with UAVs. Range limitations, signal propagation behavior, antenna radiation patterns, Medium Access Control (MAC) technology, mobility patterns and network dynamics will, both by themselves and combined, be factors that can make a large impact on the performance of UAV-enabled networks. Therefore it is necessary to experiment with real UAVs that communicate using real radio equipment in real usage conditions.

This report presents the passive network performance analysis of the sensor network during the experiment Light Navigation-radar ESM Electronic Warfare Unmanned Aircraft System (LINE EW-UAS) on November 5, 2015 at Ørland. In this experiment, two Penguin B fixed wing UAVs flew west of Ørland airbase. The UAVs had sensors to geo-locate a ship's X-band navigation radar, and by using military broadband radios they transmitted sensor information to the Ground Control Station (GCS). This information was further automatically relayed to an operational analysis cell at FFI Kjeller. This report focuses on the radio network between the two UAVs and the GCS. Specifically, we address the following questions:

- What is the performance of the employed radios in UAV-to-ground communication?
- Can we confirm that there are antenna radiation challenges in UAV-to-ground communication?
- Are there other lessons learned from this first FFI communications experiment on UAV-to-ground communication?

Based on passive packet logging of the application data, the results from the experiment were divided into phases, based on the set radio link bit rate (1024 kbps and 128 kbps) and the mobility patterns of the two UAVs. The number of sent, received, forwarded and duplicated packets are discussed for each phase, and by using the mobility patterns we are able to evaluate the link performance with regard to the platform position and attitude.

The results show a difference between the links UAV1-GCS and UAV2-GCS, with much better performance for UAV2-GCS compared to the former. Measures for the performance are in this report packet loss, forwards, duplicates and delay. Further, the Kongsberg Defence and Communications (KDC) TacLAN radios produced a lot of duplicate packets, probably caused by a bug in the radios' MAC implementation. We confirmed that there are antenna radiation challenges in the UAV-to-ground communication. In this respect we also see benefits of forwarding using a dynamic ad hoc network. There were also both pros and cons when using only user application traffic to measure the link and network performance.

Range, coverage and LoS challenges in various terrain and environments should be tested in further work, in addition to link rate experiments.

Sammendrag

Behovet for kommunikasjonskapasitet i Forsvaret øker i årene fremover. Radioer med høy kapasitet har ofte sterke krav til frisikt. Unmanned Aerial Vehicles (UAV-er) kan avhjelpe dette økte behovet for frisikt-kommunikasjon, og UAVer med varierende grad av autonomi kommer til å bli utviklet. Å støtte utvikling av autonomi for UAVer skreddersydd for kommunikasjonsformål krever en inngående kjennskap til hvordan nettverk med UAV-er oppfører seg. Faktorer som rekkeviddebegrensninger, signalpropagasjon, antennestråling, teknologi for medium aksess kontroll, mobilitetsmønstre og nettverksdynamikk vil, både alene og kombinert, kunne ha stor innvirkning på ytelsen til UAV-støttede nettverk. Derfor er det nødvendig å eksperimentere med faktiske UAV-er som bruker virkelig radioutstyr i reell bruk.

Denne rapporten presenterer passiv nettverksytelsesanalyse av sensornettverket som ble etablert og brukt i eksperimentet Light Navigation-radar ESM Electronic Warfare Unmanned Aircraft System (LINE EW-UAS) 5. november 2015 på Ørland. I dette eksperimentet fløy to Penguin B fixed-wing UAV-er vest for Ørland flystasjon. UAV-ene hadde sensorer ombord for å geo-lokalisere et skips X-bånd navigasjonsradar, og ved bruk av militære bredbåndsradioer sendte de sensorinformasjon til bakkestasjonen. Informasjonen ble videresendt automatisk til en analysecelle på FFI Kjeller. Denne rapporten fokuserer på radionettverket mellom de to UAV-ene og bakkestasjonen. Spesifikt tar vi for oss følgende spørsmål:

- Hva er ytelsen til radiosystemet når det brukes i UAV-til-bakke-kommunikasjon?
- Kan vi bekrefte at det eksisterer antennestrålingsutfordringer i forbindelse med UAV-til-bakke-kommunikasjon?
- Finnes det andre erfaringer fra dette første FFI-kommunikasjonseksperimentet på UAV-til-bakke-kommunikasjon?

Basert på passiv pakkelogging av applikasjonsdata deler vi resultatene fra eksperimentet inn i faser etter radiolink-datarate (1024 kbps og 128 kbps) og mobilitetsmønstrene til de to UAV-ene. Antallet sendte, mottatte, videresendte og duplikerte pakker evalueres for hver fase, og mobilitetsmønstrene blir utnyttet for å evaluere linkytelsen med hensyn på plattformenes posisjon og stilling i rommet.

Resultatene viser en forskjell mellom linkene UAV1-Ground Control Station (GCS) and UAV2-GCS, med mye bedre ytelse for UAV2-GCS sammenlignet med den andre. Mål på ytelse i denne forbindelse er pakketap, videresending, duplikater og forsinkelse. Videre produserte Kongsberg Defence and Communications (KDC) TacLAN-radioene mange duplikatpakker, sannsynligvis forårsaket av en feil i radioenes implementasjon av medium aksess kontroll. Vi fikk bekreftet at det er antennestrålingsutfordringer i kommunikasjonen mellom UAV og bakke. I dette henseendet ser vi også fordelen av videresending gjennom bruk av et dynamisk ad hoc-nettverk. Vi identifiserte også fordeler og ulemper ved å kun bruke brukerapplikasjonstrafikk for å måle nettverksytelsen.

Av videre arbeid ser vi behov for å undersøke rekkevidde, dekning og Line-of-Sight (LoS)-utfordringer i forskjellig terreng og omgivelser, i tillegg til kapasitetsytelsesmålinger.

Contents

1 Introduction

The need for communication capacity in the Norwegian Armed Forces (NAF) will increase in the coming years, both due to the requirement for closer cooperation between units, and due to the use of additional sensors. The use of broadband radios on higher frequencies can support this increased demand, but such radios have stronger requirements for Line-of-Sight (LoS) than current narrow-band systems. Elevated relays, e.g., Unmanned Aerial Vehicles (UAVs), could alleviate this.

UAVs are highly suited platforms for sensor elevation. However, UAVs can also be used to extend the communication range of radio systems. The later years' technological development increases the possibility of UAVs that require little human input to be in a position where they can improve ground-based networks. Taking advantage of already deployed sensor platform UAVs as communication relays may increase the robustness of ground-based networks.

Utilizing an elevated relay as part of a larger network entails complexity that is currently researched under the concept of Flying Ad Hoc Network (FANET) [1]. FANETs are a new type of ad hoc network, represented as a special subset of Mobile Ad Hoc Networks (MANETs) and Vehicle Ad Hoc Networks (VANETs), addressing challenges that are especially prominent for flying network nodes. One major challenge is antenna structure for nodes operating in a 3D environment. Current omni-directional antennas common in MANETs turn out to have limitations in 3D, since the antenna radiation pattern for omni-directional antennas can often be compared to a donut, with weak spots along the axis of the antenna. Increased understanding of FANET challenges could lead to better use of UAVs as communication relays, not least as the technological development drives towards more autonomous behavior for the UAVs.

This report presents passive link and network monitoring results from the Light Navigation-radar ESM Electronic Warfare Unmanned Aircraft System (LINE EW-UAS) experiment [2] performed at Ørland November 5^{th} 2015. Through this work we want to gain a better understanding of UAV-ground communications challenges in general, spanning from hardware integration via radio choices and stability performance to achieved throughput. The monitored traffic is the network traffic generated by the LINE EW-UAS application, in combination with control of sequence numbers and timestamps for application traffic traversing the UAV radios.

Specifically, we address the following questions through this report:

- What is the performance of the employed radios in UAV-ground communication?
- Can we confirm that there are antenna radiation challenges in UAV-ground communication?
- Are there other lessons learned from this first experiment on UAV-ground communication?

1.1 Background

The LINE EW-UAS experiment was a combined effort of eight different research groups from five different departments at the Norwegian Defence Resesarch Establishment (FFI). It was established to investigate advantages and challenges of utilizing unmanned systems for use in Electronic

Warfare (EW). The overall goal was to fly multiple sensors using UAVs to detect ship-based navigation radars, comparing these results with Automatic Identification System (AIS)-data tracks, and then potentially forward anomalies to an operational headquarters for further action.

The research program on Communications Infrastructure at FFI had the responsibility of supporting the experiment with communication knowledge and resources. The research program is also investigating the potential use of autonomous UAVs for improving military ground-based mobile tactical networks. Thus, the experiment represented a major potential to learn more about the behavior of a UAV-ground link, and also UAV-UAV communication.

Through the research program on Communications Infrastructure, FFI is also collaborating internationally with Fraunhofer, Germany, through the Coalition Networks for Secure Information Sharing (CoNSIS)-project. Through this project, the Norwegian Ministry of Defence (MoD) finances the Norwegian industry for Research and Development (R&D) in communications. A one-year project with Kongsberg Defence and Communications (KDC) allowed the use of their Tactical Local Area Network (TacLAN) radio collection [3], specifically the SR600 and WM600 broadband radios with possibilities to implement logging support mechanisms.

The analysis of the communication performance was not part of the primary objectives, and could only be realized through passive monitoring and programmable functions in the communications equipment implemented in other projects.

1.2 Report structure

The rest of the report is structured as follows. First, the experiment components are described in Chapter 2. Second, the method for gaining insights into the network performance is described in Chapter 3. Third, the nodes' movements and the phases of the experiment are explained in Chapter 4. Then, the communication results are presented in Chapter 5. Chapter 6 concludes the report and presents plans for future work.

2 Experiment components

The overall experiment, described extensively in [2], involved a multitude of different components. In this chapter, we give an overview of the relevant components of the experiment for the purpose of this report. This means that we limit the scope to components relevant to evaluate the radio link and network performance between the UAVs and the ground segment. For practical purposes, the ground segment is referred to as the Ground Control Station (GCS)[1].

2.1 Node and system setup

The main components of the experiment were two UAVs and the GCS. The two UAVs were each loaded with the following hardware:

- A Light Navigation-radar ESM (LINE) Electronic Support Measures (ESM) sensor.
- Hardware to process the signals from the LINE sensor and generate Cooperative ESM Operations (CESMO) [4] messages (Sensor Processor).
- A KDC SR600 broadband radio.
- Antennas for sensor and radio communications.

The standard monopole whip antenna for the UAV-mounted SR600 was replaced with a custom antenna designed and produced at FFI (Figure 2.1). There was not enough space in the payload compartment in the Penguin B UAV from UAV Factory to use the SR600 radio with the default antenna. The antennas (two UAVs) and their performance are described in detail in [2, Chapter 5.3.3.2]. In particular it is documented that the communication antennas of the two UAVs performed differently. The measurements of the reflection coefficient of both antennas (Figure 2.1) show that around 323 MHz, antenna 1 reaches -20 dB, while antenna 2 reaches -35 dB at minimum reflection. However, the peak absorption does not by itself indicate less gain for the UAV1 antenna. Instead, the antennas are vulnerable for unwanted modifications to the cone shape of the upper part of the antenna. If the cone form is modified, the antenna center frequency may move lower or higher than the 323 MHz it is supposed to have. An accidental bend of one or several of the cone legs during handling of the vehicle could alter the center frequency to an extent where the antenna gain in the 5 MHz band with center 323 MHz is reduced considerably.

The GCS consisted of a CESMO message receiver PC, a hub, a PC running Wireshark, and a KDC WM600 broadband radio with the same configuration as the SR600 radios.

The radios were configured to transmit on a center frequency of 323.000 MHz on a 5 MHz channel. The SR600 radios transmitted at 1 W power, while the WM600 radio transmitted at 5 W power. All radios transmitted on the same channel and formed an ad hoc network using the Open Shortest Path First version 2 (OSPFv2) [5] routing protocol. The SR600 and WM600 radios were compatible over the air and were set up to communicate in ad hoc mode. I.e., they could relay traffic from each other to reach beyond the range of one link. The network architecture can be seen in Figure 2.2.

[1] In lack of a more specific notion for the payload data reception node and adjacent logging unit on ground.

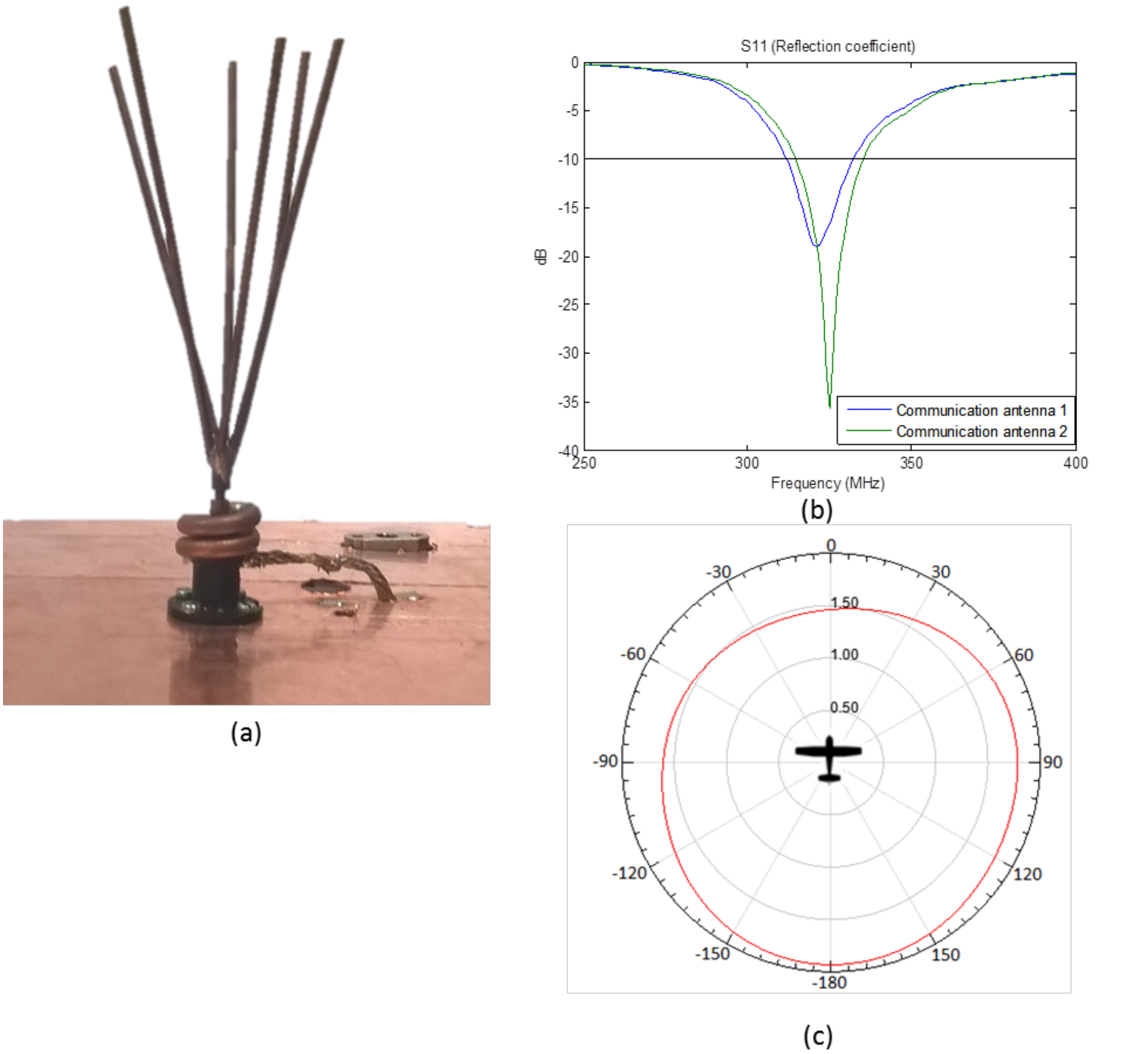

Figure 2.1 *(a) Picture of the finished antenna. (b) Measurements of the reflection coefficient of both antennas. (c) simulated radiation pattern of the antenna mounted on the UAV. From Grimstvedt et al. [2].*

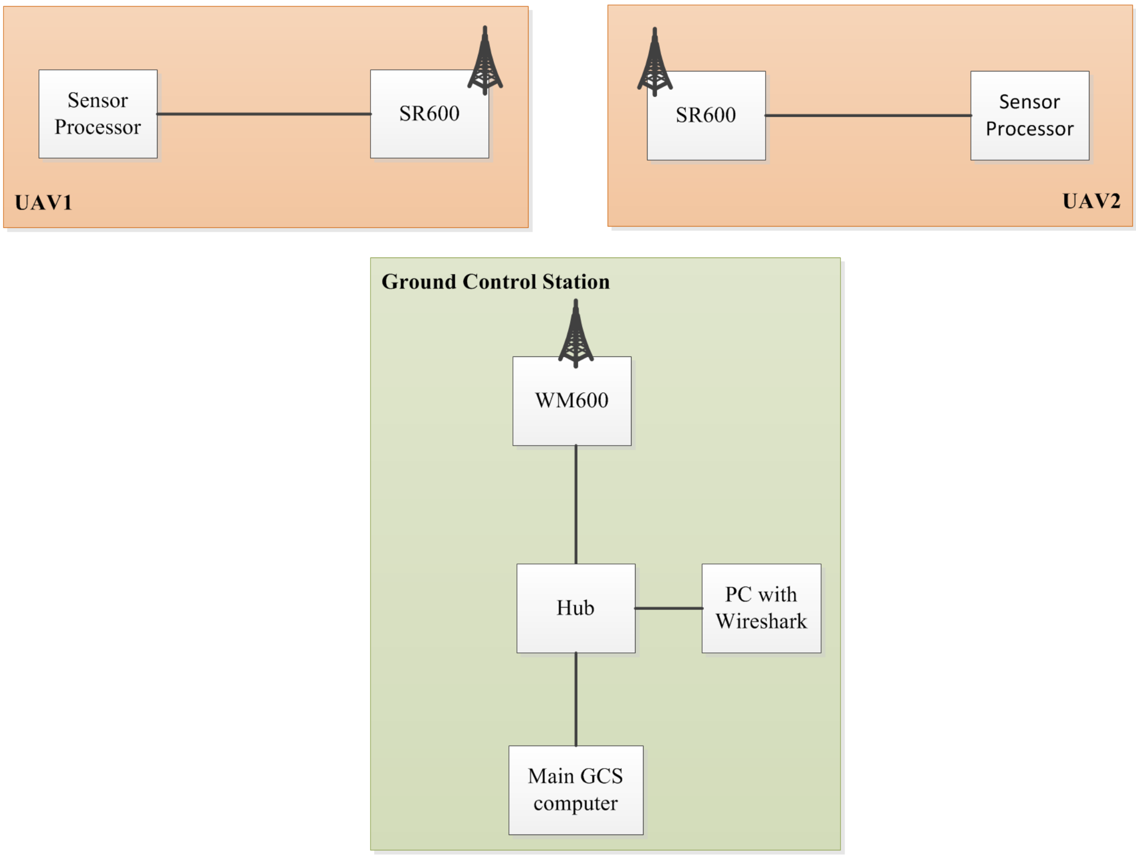

Figure 2.2 *The network architecture.*

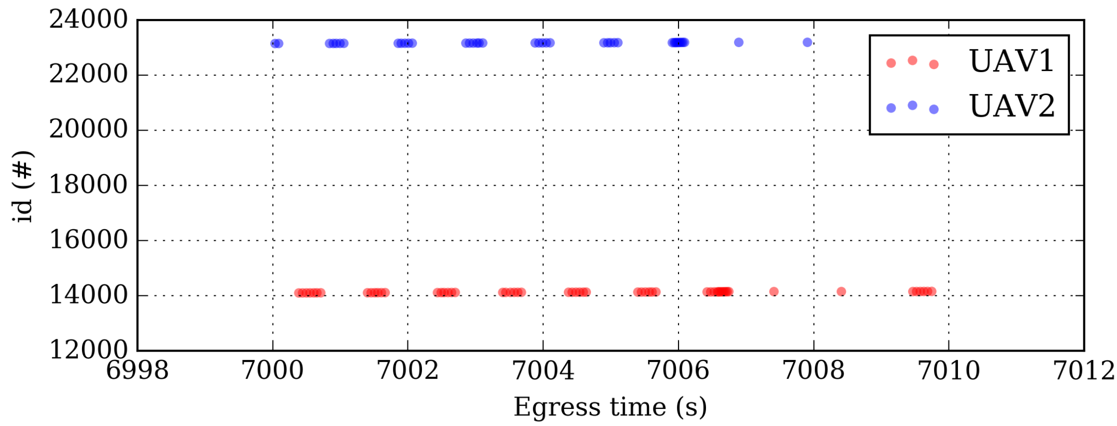

Figure 2.3 *Coordinated application traffic transmission showing the id of received packets over egress time.*

The radio Medium Access Control (MAC) protocol was based on Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), and adhered in most ways to the IEEE 802.11 standard [6]. However, as the radios had no requirement to be compatible on air with other IEEE 802.11 equipment, and due to the lower channel frequency, some implementation differences had to be expected.

The radios' data rate is static and was configured via a web interface. All other radios were able to read the data, regardless of their own rate setting. The available rates spanned from 128 kbps to 2500 kbps. The rates used in our experiment were 1024 kbps and 128 kbps.

The radio interface queue size was also configurable, albeit through command line interaction. Without manually configuring the interface queue size, the default queue size depended on the radio data rate. At 1024 kbps, the interface queue size was 64 packets, and at 128 kbps it was 8 packets.

2.2 Traffic generation

The analysis in this report is a passive[2] link and network performance analysis. See Chapter 3 for further description. The analysis depends on measurements of the traffic generated by the sensor application on-board the UAVs. The sensor application was not a typical Constant Bit Rate (CBR) application with a predefined separation between each packet. Instead it produced packets depending on the sensor input, making it less predictable in terms of resources needed for communication. When examining the results, it is important to know the behavior of the application to understand its impact on the results, and further interpretation of these. For simplicity, we refer to the application as the CESMO-application, since it used the CESMO-message format [4] to send messages.

[2] We acknowledge that we do add information during the measurement (12 bytes per packet). However, no extra packets are generated, meaning that for the application there is miniscule impact. The passiveness also limits our ability to test the network in cases where the application does not generate traffic.

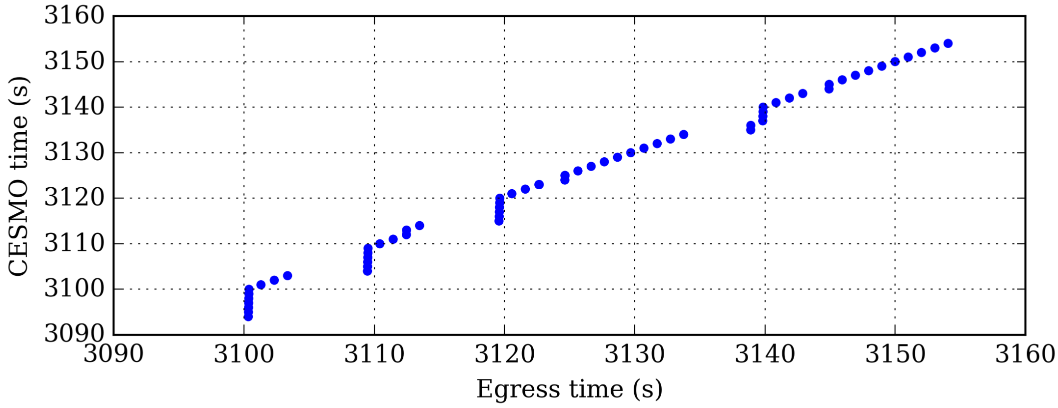

(a) CESMO time for position messages.

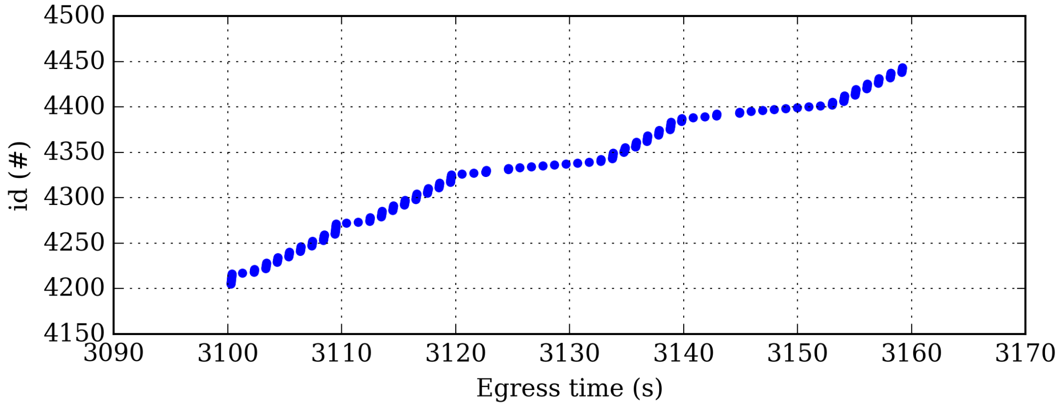

(b) Id of all received packets.

Figure 2.4 *CESMO time and id over egress time between time 3100 s and 3160 s from UAV2.*

The CESMO-application on-board the UAVs sent data upon receiving requests from the GCS. These requests were sent alternately to the two UAVs with 0.5 s intervals between each. The application was further designed to limit its transmission within two bounds, time and rate. The application transmitted data for 0.5 s. Thus, it was expected that when the first application received a new request, the other was finished transmitting its 0.5 s of data. The other bound was the rate limit. The application was restricted not to deliver more than a total of 32 kbps of data. With two UAVs, this gave a rate limit of 64 kbps. The application held back any remaining data while waiting for its next half second of transmission time. The coordinated application traffic transmission at a 128 kbps radio link rate can be observed in Figure 2.3. Note that the reason for not seeing an id increase is the scale of the Y-axis and the different starting points of the packet id for UAV1 and UAV2.

The application also generated position update messages once a second. These were buffered locally until receiving a request. The reception of such a request could delay for multiple seconds, depending on the link performance, as requests sent from the GCS could be lost. This can be observed in Figure 2.4a. Here, the application has queued several consecutive position update messages over several seconds, before transmitting them all at once. Looking at Figure 2.4b within the same time frame, we see that the application has been sending other traffic while holding back the position messages.

The application generates sensor status messages containing the position of the sensor that can be easily identified by size. All status messages are 160 bytes. Sensor data messages on the other hand, can vary in size, but are commonly around 600 bytes.

During pre-flight tests at Ørland, the sensor application did not work as intended on UAV1. It was receiving a lot more noise compared to UAV2. Therefore, the operators had to increase the threshold for the sensor raw data on UAV1 compared to UAV2 [2]. However, this new threshold was set to make both sensors perform similarly. Thus, we did not expect any differences between the amount of traffic from each UAV. Regardless, both UAVs produced enough data to analyze the communication network performance.

3 Method

The report examines network data mainly logged at a laptop in the sensor network. The data was logged using Wireshark [7] running on a Windows 7 Enterprise Operating System (OS). The traffic was fed both to the Wireshark-PC and the ground node using a Local Area Network (LAN) hub (Figure 2.2). This way the Wireshark-PC logged all payload data passing between the UAVs and the ground node.

The UAVs were deployed to detect and position navigation radars using an FFI-developed ESM-sensor. The ESM-application transmitted sensor data to the GCS using the radio network, but only after a request from the GCS. Such a request was generated once each second by the GCS. No additional traffic to test capacity or second-to-second connectivity was generated. Thus, the measurements were performed on the application traffic, mainly through examining the Wireshark log. From the log, the IP Time-to-Live (TTL) field can be used to tell how many router hops the packet has passed, providing that the start value is known. With a total of three routers in the network, the TTL can be used to learn which routers the packet has been forwarded through, i.e., the forward path of the packet. In our setup, the default TTL was 64. Each radio was a router. Thus, a packet transmitted successfully from the UAV to the GCS would experience a TTL reduction of two, due to two radios forwarding the packet. A packet received at the GCS with a lower TTL than 62 would have been forwarded via the other UAV.

The application that produced the payload traffic was a request-response application. The application in the UAVs only sent response traffic upon requests from the ground node. This reply could consist of one or several packets, depending on whether the application had created and buffered any data since the last request, or if there was still a backlog of packets not tranmitted during the last transmit period. A broken link could cause a lost request packet. The consequence would be a lack of transmitted packet(s) from the UAVs, creating periods of time without transmissions across the connection, i.e., lack of connection measurements.

The CESMO-application produced packets in conformance with the CESMO protocol. This enabled us to extract the reported GPS position as well as an application-internal timestamp. The timestamp was synchronized with GPS time, but fixed to be updated once per second. The timestamp tells us when the GPS position was recorded. The application often queued these position reports, sending instead sensor data in its alotted transmission window. Using the timestamp we can see at what position the UAV was in at the time when the position report was generated.

In addition to the Wireshark log of application traffic, telemetry data for the two UAVs were recorded by the UAV control station. This telemetry data was recorded by the vehicle command and control link, which was a more robust Very High Frequency (30-300 MHz) (VHF) Direct Sequence Spread Spectrum (DSSS) link. The telemetry data show among other information the position and velocity of the two vehicles. Using the take-off velocity (approximately 10 m/s) as a synchronization point between the telemetry data and the Wireshark log, the position of the UAVs for each received traffic packet was identified.

While the traffic measurement was passive (no extra data packet generation), the IP header of the application packets was manipulated in two ways at the radio in the UAVs:

1. The packets were assigned a sequential IP identification number.

2. The packets were assigned an ingress timestamp[3].

Being able to mark all packets with a sequential IP identification number brings several advantages. Comparing the IP identification number of the first and last packet gives us the number of sent, or attempted sent, packets. The sequential IP identification number makes it possible to evaluate how many of the packets received by the UAV radio that are subsequently lost before reception at the ground node. However, it is not able to separate between loss due to transmission errors and loss due to congestion (queue tail drop). In addition this would enable detection of packet reordering.

The UAV radios were modified to add a timestamp [8] by means of the IPv4 Option header field, an additional information field in the IPv4 header. The extension amounted to 12 bytes. Care was taken not to change packets that would increase in size beyond Maximum Transmission Unit (MTU) and become fragmented because of the extra 12 bytes in the IP header. The timestamp was in milliseconds, starting at 0 for every new day. However, there was no time synchronization in this experiment, and the radios were reset to Unix epoch (1. January 1970 00:00:00) at each power-up. Thus, as long as the power-up occurred no more than 24 hours ago, the timestamp would reflect the time since the last radio power-up. The added timestamp makes it possible to evaluate the experienced delay for each packet while forwarded in the radio network.

The process of extending the packet header with timestamp information was performed using the *Click Modular Router* framework [9, 10], or *Click* for short. Click is a software architecture for building flexible and configurable routers. In Click, packet processing modules implement simple router functions like packet classification, queueing, scheduling, and interfacing with network devices. Through the multi-national CoNSIS-project KDC enabled Click-scripting on their TacLAN radios.

In the report we operate with two notions: *Ingress time* and *Egress time*. Ingress time is the time when a packet is processed by each of the UAV radios, before it is queued for transmission from the radio into the radio network. Egress time is the time when the packet exits the radio network, logged by Wireshark at the GCS.

The work in this report depends on several log files that are not synchronized at recording. Thus, to identify the time and position of events, the egress time (Wireshark log timestamp) has been used as reference, and the events of the other log files have been mapped to this time reference. We define a common time reference point, called time zero, defined as the time point of the Wireshark-PC receiving the first packet from the first UAV (UAV1). We define the take-off as the time when the ground speed of the vehicle was higher than 10 m/s. UAV1 took off at 205.874 seconds and UAV2 at 2088.661 seconds. UAV2 started its first operational circle at 2720 s. Thus, at this time, both UAVs were in similar motion pattern at the same distance from the GCS. The rate change from 1024 kbps to 128 kbps had to be done manually from the ground. Therefore, the rate was not changed at the two UAVs simultaneously. Through evaluating the behavior of the time delta between ingress and egress times, we determined that for UAV1 the rate was reduced to 128 kbps at time 6810 and for UAV2 this shift occurred at time 6830. For simplicity, the change is considered implemented on both vehicles at time 6830.

The radios were set to run at 1024 kbps. This rate was chosen for two reasons:

[3]This is not a CESMO-timestamp, but a timestamp generated from the SR600 radio clock.

- KDC has experienced extensively with this bit rate, e.g., in [11][4], and have great confidence in it.
- The bit rate is in the same order of magnitude as comparable Ultra High Frequency (300-3000 MHz) (UHF) radios.

The default radio interface queue size for 1024 kbps was 64 packets. This is a small queue size, should the link experience temporary packet loss. In many cases, a small queue size is proper in a tactical network, since delayed traffic often loses its importance. However, in our case, the objective was to identify the link and network performance. Therefore, to increase the likelihood that the identified packet loss was caused by transmission loss instead of queue overflow, the radio interface queue size was set to 1000 packets.

[4]The bit rate 920kbps in [11] is the same as our 1024 kbps rate.

4 Experiment events

In this chapter, we address major events during the experiment, that may help us to achieve a better understanding of the behavior of UAV-ground links. These events include the radio link rate change and specific UAV mobility patterns. We thereafter define phases based on these events that enable us to understand the results better, as they are achieved in a given context.

4.1 Radio link rate

The radios were set to run at 1024 kbps, and ran using this rate for most of the experiment. However, at around 6810 s, ($\frac{4}{5}$ of the experiment time), the rate was reduced to 128 kbps. This choice was made since there was reason to believe that the perceived low performance of UAV1 could be caused by a bad link or by interference from the radio onto the sensor. The link rate was changed using an ssh-connection from the Wireshark-PC at the GCS to each of the two SR600 radios.

When the rate was changed to 128 kbps, this caused a re-installation of the queue limit of 8 packets.[5] As is apparent from the results, this led to more packet loss for the low link rate selection. This increased packet loss is suspected to be caused only by packet loss from the queue.

Regardless of the reason for the increased packet loss, the packet loss uncertainty is only the case for the low bit rate phase of the experiment. For the first $\frac{4}{5}$ of the experiment, the queue size was 1000 packets and the rate was 1024 kbps, making us confident that the packet loss by and all is due to transmission loss.

We identified the change from 1024 kbps to 128 kbps by looking at the lower bound of the variation in difference of arrival at the GCS[6] (Figure 4.1). Closer inspection reveals that the rate change occurred at 6810 s for UAV1 while it took place at 6830 s for UAV2.

4.2 UAV movements

The experiment ran as follows. First, UAV1 taxied out on the runway, backtracked and took off from RWY15[7]. It flew first in a racetrack pattern 500 m west of the airstrip, one lap. Then it flew north into the northern end of the Grandvika bay near Frigård for a circling pattern in a small circle. After one or two circles it moved further west and started circling in a large circle west of the bay, between Grandholman and Hoøyskjæret. Figure 4.2 shows the experiment area.

While UAV1 circled between Grandholman and Hoøyskjæret, UAV2 took off and moved first in the same way as UAV1 to the small circle in the bay near Frigård. Next, UAV2 moved into the

[5] This automatic queue limit depended on the set rate for the radio at boot. If the rate was 128 kbps at boot, the queue was reset to 8 packets. Conversely, a 1024 kbps rate at boot enforced a 64 packet queue upon a rate change.

[6] The difference in arrival between two consecutive packets from the same UAV at the GCS is the egress time delta.

[7] Using RWY15 means that the aircraft took off with its nose in the S-SE direction (approximately 150°)

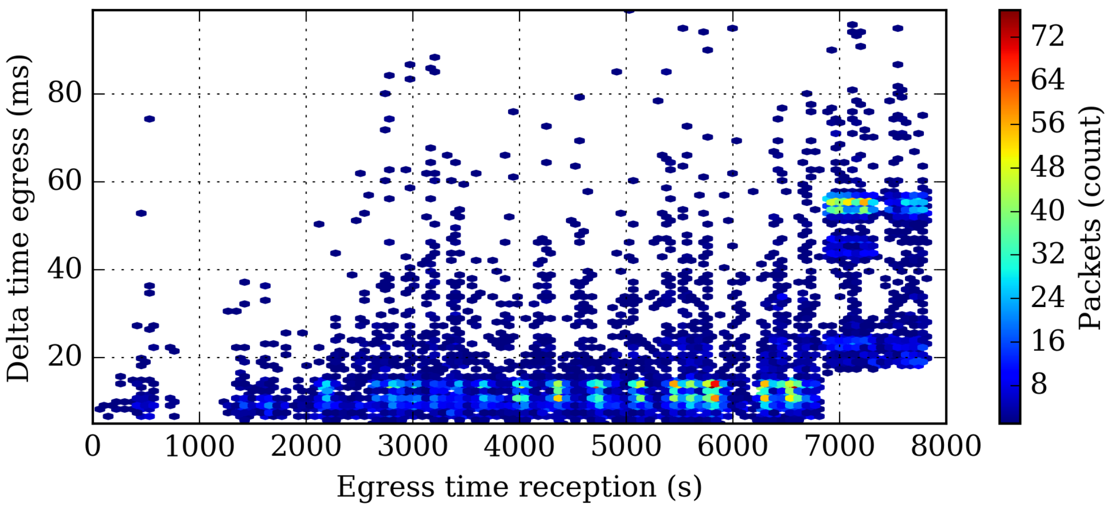

Figure 4.1 *Lower bound (<100 ms) of the egress time delta between packets over egress time for UAV1.*

Figure 4.2 *Map of the experiment area west of Ørland airbase.*

Figure 4.3 *Top-down track view of UAV1.*

large circle where UAV1 was circling. At first, the two UAVs were separated by about a quarter of a circle circumference. To get the UAVs diametrically across from each other, the pilot sent UAV2 on a larger North circle bend on one occasion.

After a while circling the large circle, the pilot put the UAVs in a racetrack pattern by extending the southern part of the circle southward. They continued in this pattern until the data logging was switched off at the end of the experiment. In the racetrack pattern, the two UAVs were not positioned diametrically across from each other.

Figure 4.3 shows the telemetry track of UAV1 plotted in Google Earth. Looking at the vertical view of the UAV's track (Figure 4.4), we see that the UAV reaches it operational height of 400 m (1312 ft) after multiple laps in the small circle[8] just west of Frigård. The UAV also makes some laps in the large circle at 400 m before it ascends to 500 m (1640 ft) altitude.

Figure 4.5 shows the track of UAV2 plotted in Google Earth. Looking at the same time at the side track view (Figure 4.6), we see how UAV2 follows the same height increase pattern of UAV1, but ascends more quickly to around 380 m in the small circle west of Frigård. However, it only reaches its operational altitude of 460 m (1509 ft) after it has begun circling in the large circle.

The distances between each of the UAVs and the GCS are presented in Figure 4.7. We see that after the transportation to "target", the distance varied between 4 km and 6 km, both in the large circle and in the racetrack. We also see that UAV2 started its flight and reached the stable large circle much later than UAV1.

[8]The spiral between -2800 and -1800 meters on the x-axis.

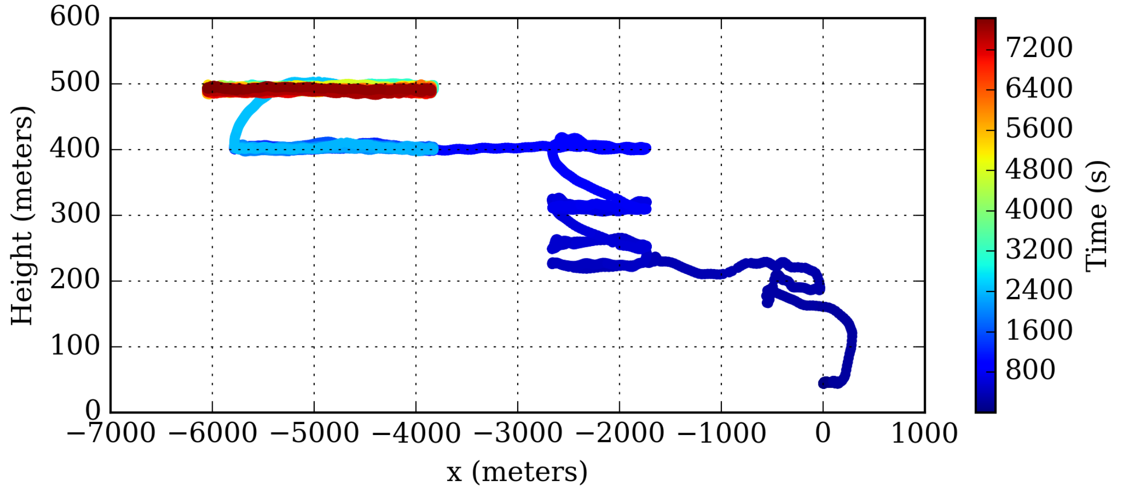

Figure 4.4 *Side track view of UAV1. Blue is early, green is in the middle, and red is late in the experiment.*

Figure 4.5 *Top-down track view of UAV2.*

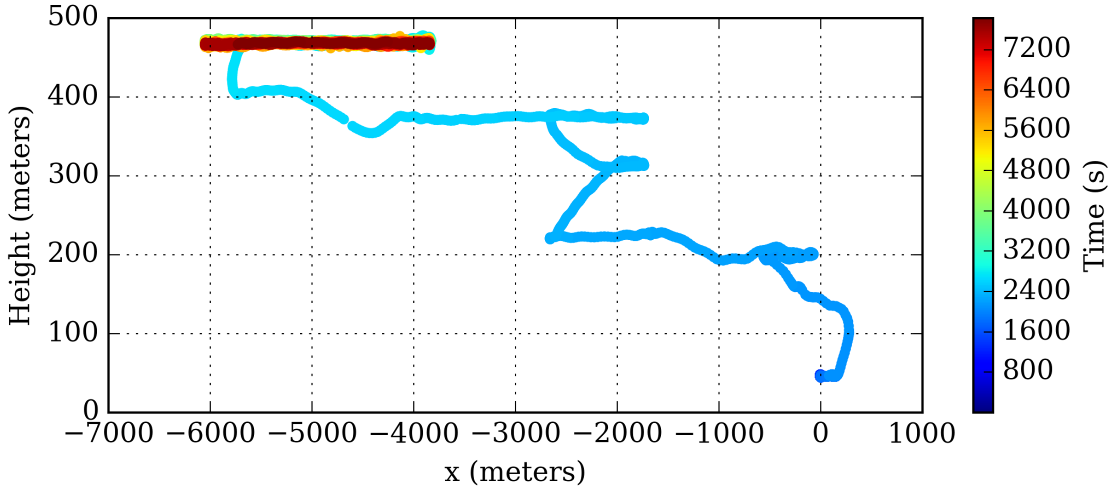

Figure 4.6 *Side track view of UAV2. Blue is early, green is in the middle, and red is late in the experiment.*

4.3 Phases

To better understand the experiment results, we have chosen to structure the results according to phases (Figure 4.8), based on the radio link bit rate setting and the UAVs' movement patterns. The defined phases are the basis for the results presented in Chapter 5.

We define two main phases, to evaluate the link during the high and low bit rate link setting. While the high bit rate setting was valid from *Time 0*[9], we want to evaluate the link when it was utilized for transportation of sensor data. Thus, we set the starting point for Phase H (High capacity) at the moment both UAVs are performing their first operational circle (2720 s). The end of Phase H is defined to be at the time when the first UAV (UAV1) changes bit rate setting from 1024 kbps to 128 kbps (at 6810 s). The start of Phase L (Low capacity) is defined to be at the time when the last UAV (UAV2) changes bit rate setting to 128 kbps (at 6830 s). The end of Phase L is defined as the end of the payload traffic logged part of the experiment (at 7829 s).

In addition to the two main phases H and L, we have defined phases where the UAVs are moving in repetitive patterns. Using these phases, we can analyze the effect of the UAV's position and attitude with regards to the direct link between the UAV and the GCS.

The UAVs operated mainly in two patterns at constant altitudes: large circle and racetrack. UAV1 moved in a large circle, first at 400 m (from 985 s to 2430 s), and then climbed to 500 m, where it circled from 2480 s to 3730 s. It then flew a racetrack pattern (the large oval circle) from 3730 s to the end of the experiment. Since the bit rate was changed at 6810 s, this is also defined as the end of this phase.

UAV2 flew in a large circle together with UAV1 at 460 m altitude (2720 s - 3830 s) and changed its motion pattern to racetrack together with UAV1, flying this from 3830 s to the end of the

[9] *Time 0* is the time when the first packet was received from UAV1 at the GCS. At this moment, both UAVs were on the ground.

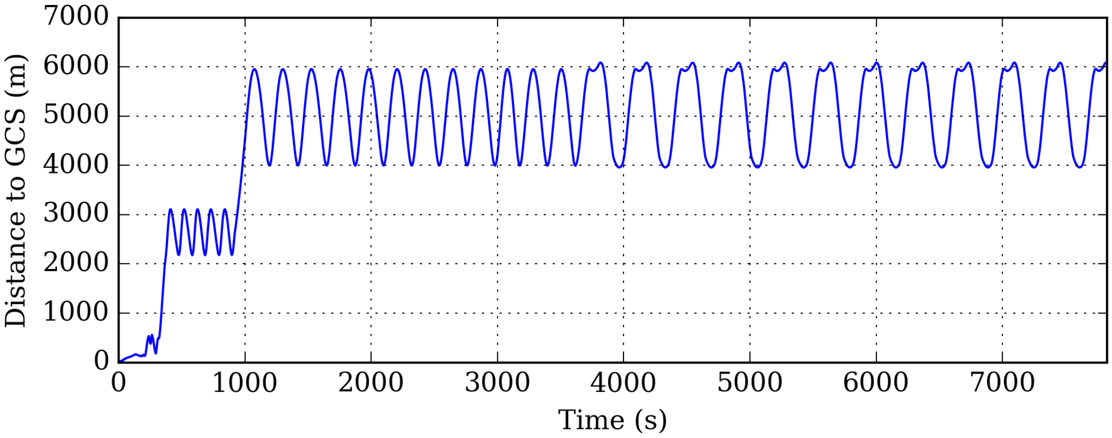

(a) UAV1

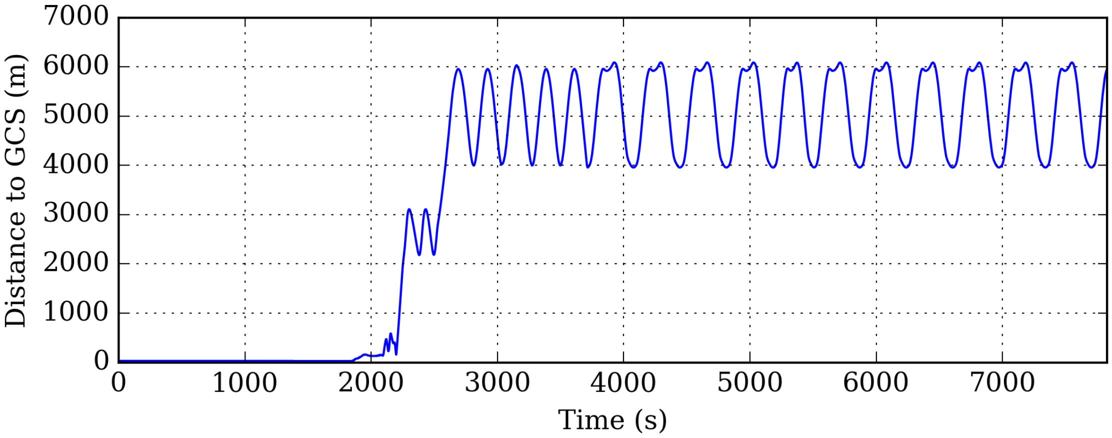

(b) UAV2

Figure 4.7 *Distance between the UAVs and the GCS during the experiment.*

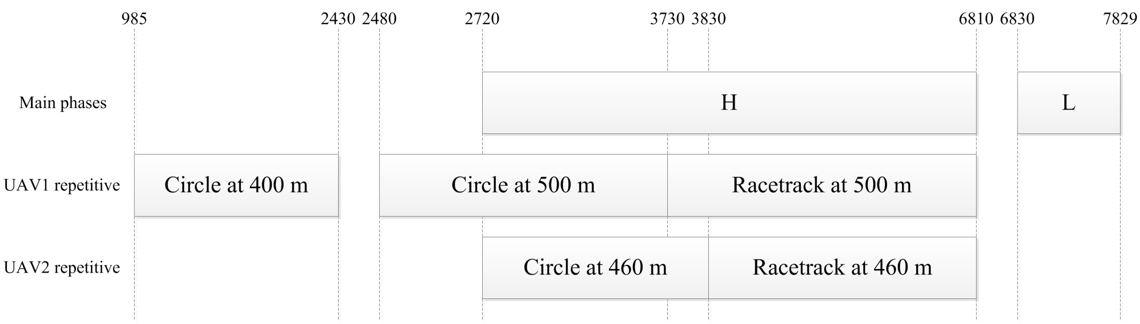

Figure 4.8 *Main and repetitive pattern phases.*

experiment. For UAV2, we also define the bit rate change time as the end of the racetrack pattern phase. Thus, all repetitive pattern phases took place with a radio bit rate of 1024 kbps.

5 Results

This chapter presents the communications results from the LINE EW-UAS experiment. The structure is based on the phases defined in Chapter 4.3. First, the overall performance for each of the main phases is addressed. This gives us a general understanding of the performance of the radios in the UAV-ground communication context. Second, we investigate repetitive flight patterns to see if we can verify antenna radiation challenges.

5.1 Overall performance

5.1.1 Phase H, 1024 kbps

In the high data rate phase, both UAVs are operational, flying in the large circle and thereafter in the racetrack pattern. Since the two UAVs have similar movement patterns, we would expect similar performance from the UAVs, with regards to packets sent, packet loss, duplicates and latency.

5.1.1.1 Packets sent

Beginning with the number of sent packets[10] (Figure 5.1), we observe that there is great difference between the two UAVs. UAV1 has produced only 53.4 % of the packets produced by UAV2. As mentioned in Chapter 2.2, this could be a consequence of the threshold adjustment, but there are also other possible explanations:

- UAV1 has not been able to receive the requests for data.
- Route loss has prevented packets from the application to be accepted by the radio.
- The sensor in UAV1 has produced less data than UAV2.

The request from the GCS triggers the position message and sensor data transmission. The consequence of each lost GCS request to UAV1 is that the data, which is continuously produced by the application, is not sent. Instead, all data are buffered by the application, and sent upon receiving a new request. If the lower number of packets sent was due to loss of requests, we would be able to see this as lack of CESMO status messages received from UAV1. However, the packet size distribution for UAV1 in Phase H (Figure 5.2) shows that while the number of 600 bytes packets is much lower for UAV1 than UAV2, the number of packets with size 160 is almost the same for both UAVs (Actually a little higher for UAV1 than UAV2). The 160 bytes packets are CESMO sensor status messages. Therefore, we conclude that the loss of requests is not the reason for the lower packet production of UAV1.

A second explanation could be that a loss of route from UAV1 to the GCS has prevented Click from incrementing its IP identification counter when the application has produced new packets. The

[10]The number of sent packets is defined as the delta of the IP identification value between the first and the last packet received, plus one.

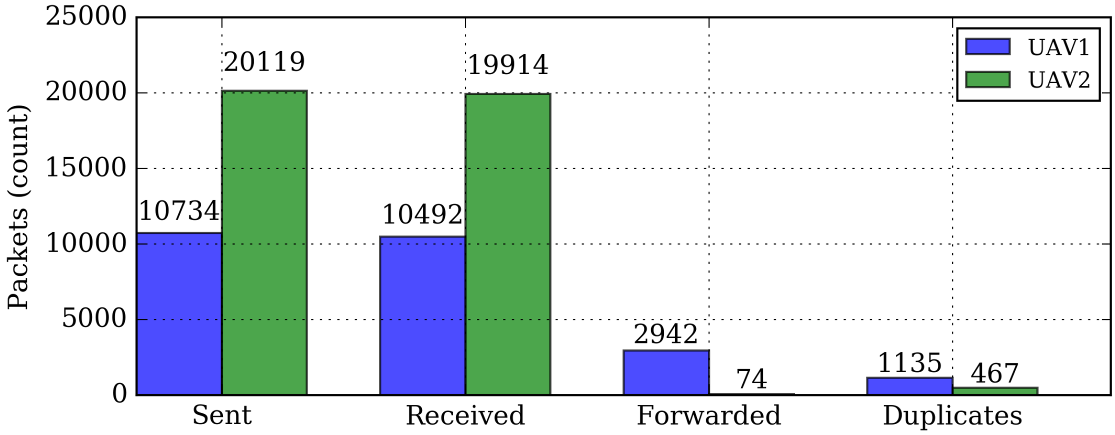

Figure 5.1 *Packets sent, received forwarded and duplicated for UAV1 and UAV2 in Phase H, as logged at the GCS.*

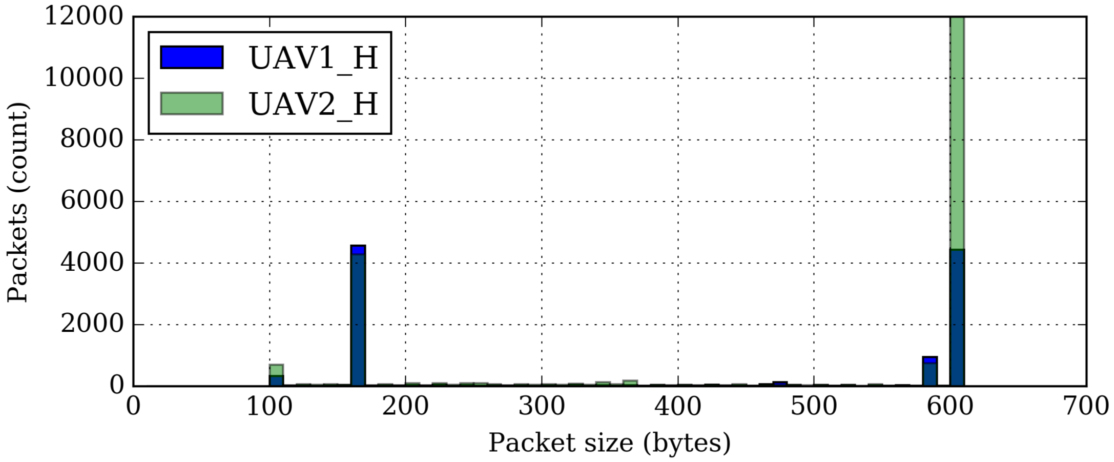

Figure 5.2 *Packet size distribution for the UAVs in Phase H.*

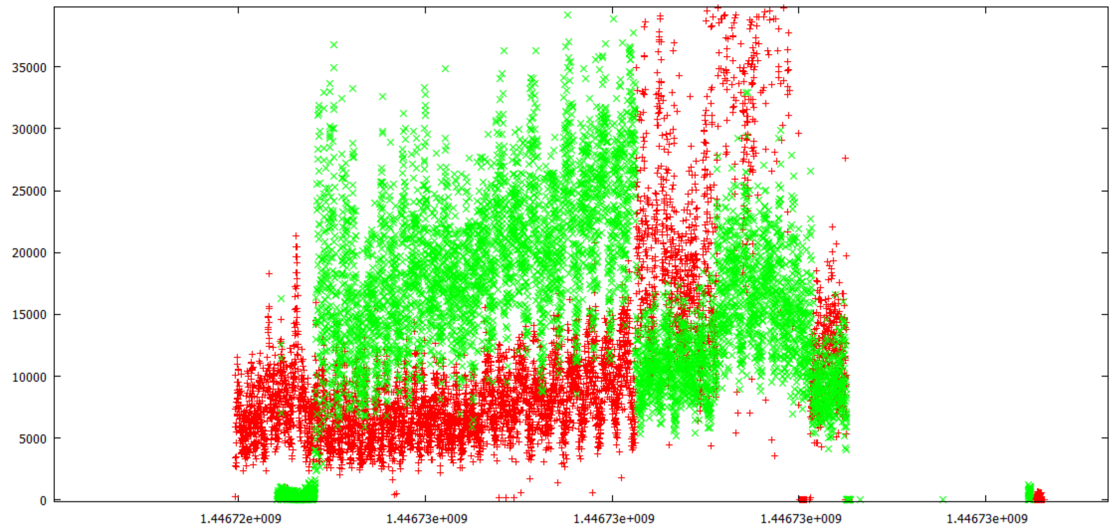

Figure 5.3 *On-board recorded sensor data for both UAVs for the duration of the experiment. UAV1 in red and UAV2 in green color.*

radio may have responded to the application with an Internet Control Message Protocol (ICMP) message of "Destination unreachable" and discarded the packet before it was treated by Click. This could also explain the lack of burst loss documented later in the report. However, these ICMP messages would only have been transmitted between the radios, and not reach the logging PC which was placed on the wired side of the WM600 radio at the GCS.

A third possibility is that the sensor in UAV1 has been unable to collect as much data as UAV2, due to the threshold adjustment. Thus, while the requests have been received, they have not resulted in as many packets sent from UAV1 as from UAV2. This is indicated by Figure 5.2, since the number of packets at 600 bytes is only a third for UAV1 compared with UAV2. Figure 5.3 shows the on-board recorded sensor data for the experiment for both UAVs. These data appear to confirm our number of packets sent in that it seems UAV1 (red) have produced only half of the packets of UAV2 (green).

5.1.1.2 Packet loss

The packet loss ratio for the communication between the UAVs and the GCS was for UAV1 2.3 %, while UAV2's loss ratio was 1.0 %. In other words, UAV1's packet loss ratio was over double that of UAV2. While the packet loss is low for both UAVs, the difference between UAV1 and UAV2 is interesting. Could the higher packet loss have the same cause as that which led to the much lower number of packets sent for UAV1, for instance a problem with the communication antenna of UAV1? To answer this question, we need to investigate the results further.

Our data shows that the packet loss has not happened in bursts. By comparing the value difference between the IP identification field of two packets (called id_step), we can evaluate whether the radio sent the two packets consecutively ($id\_step = 1$), or if there have been lost packets in between them ($id\_step > 1$). Figure 5.4 shows the id_step between packets received from UAV1 in

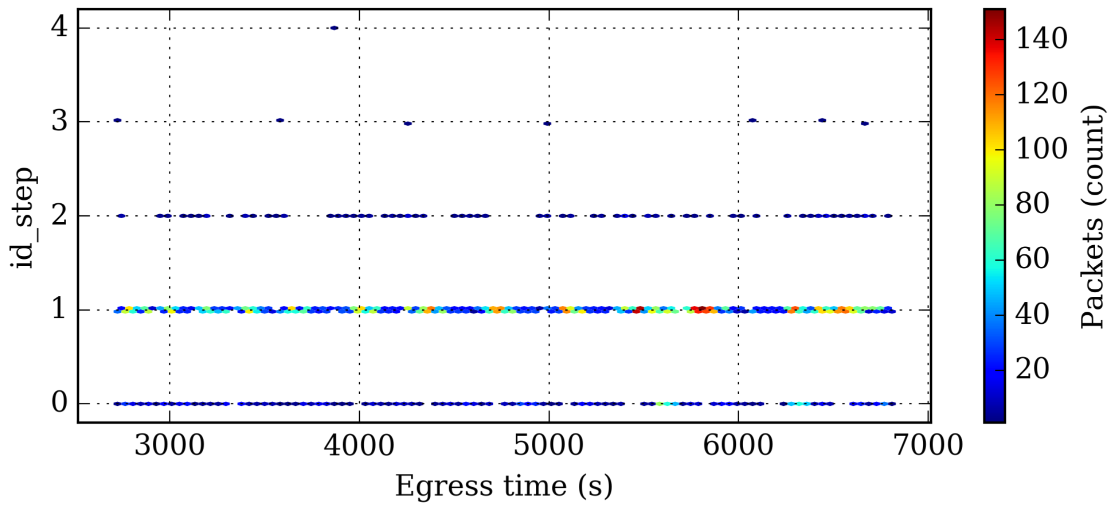

Figure 5.4 *Packets' id_step over egress time for UAV1 in Phase H.*

Phase H. If there were major bursts of packet loss, this should have shown as very high id_step increments. The highest increment is 4 and most of the packet loss occurs as individual packets lost ($id\_step = 2$), spread throughout Phase H. Note also that the $id\_step = 0$ indicates that the received packet is a duplicate of the previously received packet.

5.1.1.3 Packets forwarded

The TacLAN radio system establishes and maintains a dynamic ad hoc network. In our experiment, this meant that if any of the UAVs had problems reaching the GCS directly, a route via the other UAV would be established. Thus, the traffic could reach the destination, even though there was no direct link to the GCS. For a forwarding event to occur, the routing protocol on the radio in the UAV has to experience a time-out of the direct link to the GCS, while still having both a link to the other UAV and knowing that the other UAV can reach the GCS.

The results for Phase H show that the number of forwarded packets (Figure 5.1) was magnitudes larger for UAV1 than for UAV2 (2942 vs. 74, or a ratio of 27.4 % vs. 0.4 %).

First, since the ad hoc routing protocol only reroutes over a relay when the direct link is lost[11], the 27.4 % of UAV1's packets would not have been received at all, if it had not been for the ad hoc network functionality and UAV2 being nearby. In that case, the packet loss for UAV1 would have increased from 2.3 % to 29.7 %.

Second, it is interesting that we see the same tendency in the performance difference between the two UAVs, with UAV1 performing worse than UAV2. It is difficult to directly determine the reason for the difference in the number of rerouted packets between the two UAVs, but a difference in antenna gain performance between the antennas on the two UAVs could be the culprit.

The rerouting did not occur as one burst for UAV1. Instead, it was spread out over the entire phase duration. However, as we will see, it seems to correlate with the repetitive patterns. UAV2

[11] There was no Link Layer Notification (LLN) enabled.

only experienced one instance of rerouting, counting for all the 74 packets, between time 2781 s and 2798 s.

5.1.1.4 Duplicate packets

The wireless medium typically has a much higher Bit Error Rate (BER) than wired networks. Therefore, it is usual to confirm a unicast transmission with an Acknowledgment (ACK)-message from the receiver to the sender. If the original data transmission is not received, or the ACK-message fails to reach the sender, the original packet is retransmitted. In cases where the original packet is received, but the ACK-message is not, a retransmission may cause the receiver to receive a duplicate transmission. Normally, duplicates that occur due to MAC retransmissions should be filtered by the receiving MAC-layer.

Even though duplicate packets may occur, the number of duplicates (Figure 5.1) seen in our experiment is surprisingly high. For UAV1, the number of duplicates amounts to 10.6 % of the packets sent. While the number is lower for UAV2, it is still substantial (2.3 %). When looking at the actual rate, the duplicate packet rate was 0.28 Packets per second (pps) for UAV1 and 0.11 pps for UAV2. Similarly to the packet loss, the duplicate packet events also occur throughout the simulation (Figure 5.4, $id\_step = 0$).

The much higher number of duplicates for UAV1 than for UAV2 is difficult to explain. A plausible explanation is that it is caused by worse link conditions for ACK-messages from the GCS to UAV1 than from the GCS to UAV2.

Duplicate packets should be detected and filtered by the MAC layer, and not be delivered to the network layer. The number of duplicate packets points towards a radio MAC-layer implementation bug. If the original packet fails to reach the receiver, a retransmission is in order. However, if the data transmission reaches the receiver, and instead it is the resulting ACK-message that is lost on its way to the data sender, the current implementation seems to be unable to avoid duplicates. KDC has been notified of this bug, and they have since addressed the issue.

We consider the impact that the duplicates have had on the experiment as minuscule, since most packets have only been forwarded over one hop. Although UAV1 has 27.4 % forwarded packets, which could cause a portion of the duplicates to have a cascading transmission effect, the distribution of the duplicate packets with respect to TTL is not gravitating towards forwarded packets. However, in a situation where the traffic is forwarded over numerous hops, the consequence could be a disastrous cascade of duplicate packets.

5.1.1.5 Summary Phase H

There is a clear difference between the performance of traffic originating from UAV1 and UAV2 for the Phase H results. UAV1 has transmitted only half the packets of UAV2. The traffic originating from UAV1 has suffered over twice the loss, a quarter of the packets were forwarded via UAV2 and the number of duplicate packets was more than double that of UAV2. Although we were aware that the antennas have some performance differences, per [2, Chapter 5.3.3.2], the difference caught us

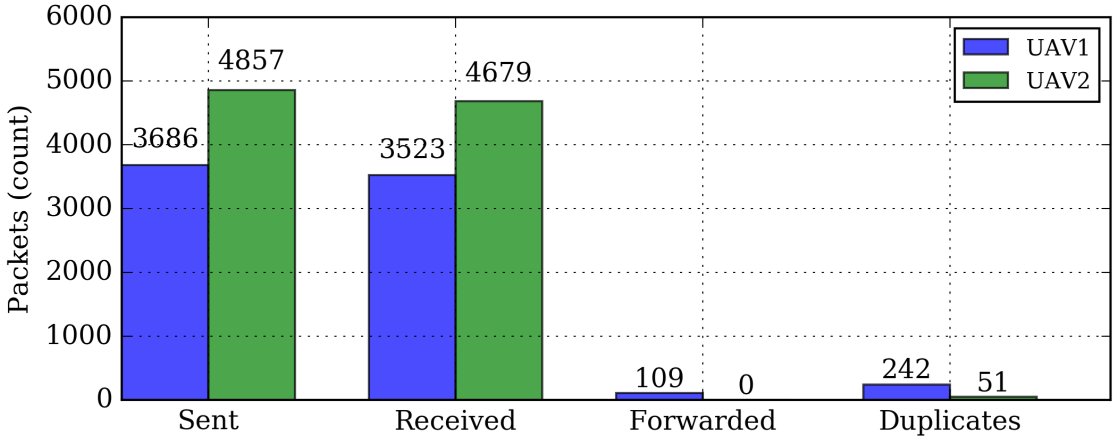

Figure 5.5 *Packets sent, received forwarded and duplicated for UAV1 and UAV2 in Phase L, as logged at the GCS.*

by surprise. The two UAVs were attempted fitted with the same equipment, and thus expected to perform equally. Except for the difference in the number of packets sent, the reason for the differing performance is not clear. The higher loss and forwards for UAV1 may be explained by differences in the antennas.

5.1.2 Phase L, 128 kbps

5.1.2.1 Results

We now turn to the results for the low data rate phase (Phase L) (Figure 5.5). UAV1 again produced only 75.9 % of the packets produced by UAV2. This is an increase from Phase H. The total number of produced packets was reduced from the Phase H, but the packet rate production for both UAVs combined increased from 7.5 pps in Phase H to 8.6 pps in Phase L. The reason for this increase is not obvious.

The distribution of the packet sizes in Phase L (Figure 5.6) for the received packets, shows that the increase in the number of packets sent occurs mainly for the 600 bytes packets. The distribution is more even than in Phase H. In L, UAV1 provided two thirds of the 600 bytes packet volume created by UAV2, compared to just one third in Phase H.

The packet loss in Phase L (Figure 5.5) increased for both UAVs. The traffic originating from UAV1 experienced a 4.4 % packet loss, while UAV2-originating traffic had a 3.7 % packet loss. The total loss increased, but the difference in loss between UAV1 and UAV2 was reduced.

The reason for the higher loss may be explained by queue loss. The difference between Phases H and L is the link rate of 1024 kbps and 128 kbps, respectively. With the reduced link rate, there is an increased risk of congestion. We recall also that when the rate was reduced, the queue size was also reduced from 1000 packets to only 8. Therefore a packet burst from the application computer could lead to packet loss even with an average load well below the link rate. Unfortunately, the log

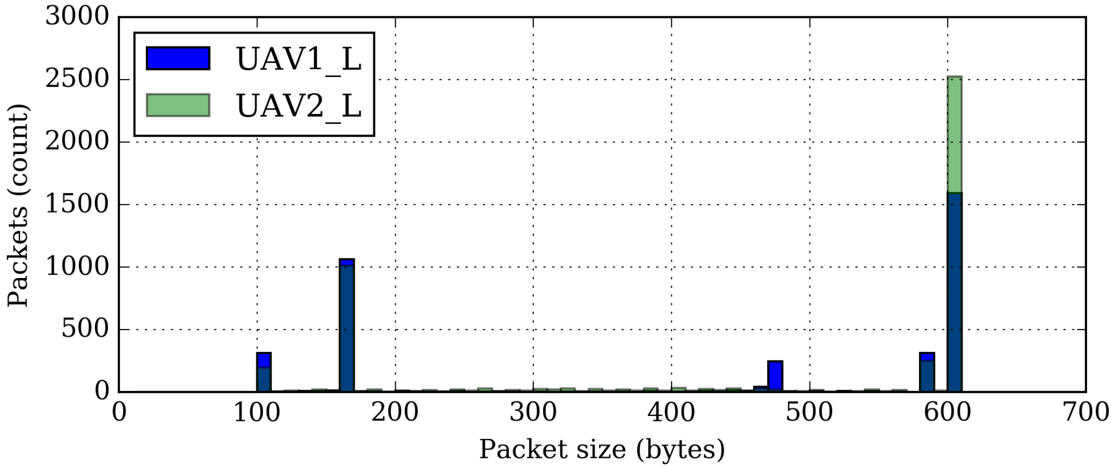

Figure 5.6 Packet size distribution for the UAVs in Phase L.

data post-processing gives us no possibility to discriminate between loss caused by queue drop and loss caused by transmission failure.

The number of forwarded packets (Figure 5.5) was reduced tremendously in Phase L. Recalling that UAV1 sent a quarter of its packets via UAV2 in Phase H, the ratio is 3.0 % of its packets in Phase L.

The lower number of forwarded packets may be interpreted in terms of improved link quality between UAV1 and the GCS, as a consequence of the lower radio bit rate setting.

The duplicates ratio of UAV1 (Figure 5.5) was 6.6 % in Phase L. This is much lower than that of Phase H (10.6 %). However, the number of duplicates is 4.7 times larger for UAV1 than UAV2. The duplicate packet rate was about the same for UAV1 at 0.24 pps (down from 0.28 pps in Phase H). For UAV2 we see a drop to 0.05 pps (down from 0.11 pps). The cause of this is uncertain.

5.1.2.2 Summary Phase L

The difference between the two UAVs was still pronounced in Phase L, compared with Phase H. The difference in the number of sent packets between the two UAVs was smaller than in Phase H, but the reason for this is unknown. It might be due to adjustments done by the operators while the vehicles were in flight. The packet loss was higher for both UAVs, which we suspect is due to queue loss. The number of forwarded packets also indicates that the link improved for UAV1, since the number of packets forwarded via UAV2 was reduced to only 3.0 %. This would also explain why the two UAVs are closer in terms of packet loss than in Phase H, separated now by only 0.7 % points (1.2 % points in Phase H).

5.2 Repetitive motion patterns

The results from the main phases showed that the traffic originating from UAV1 throughout both phases experienced more packet loss, was to a larger degree forwarded and contained more duplicate

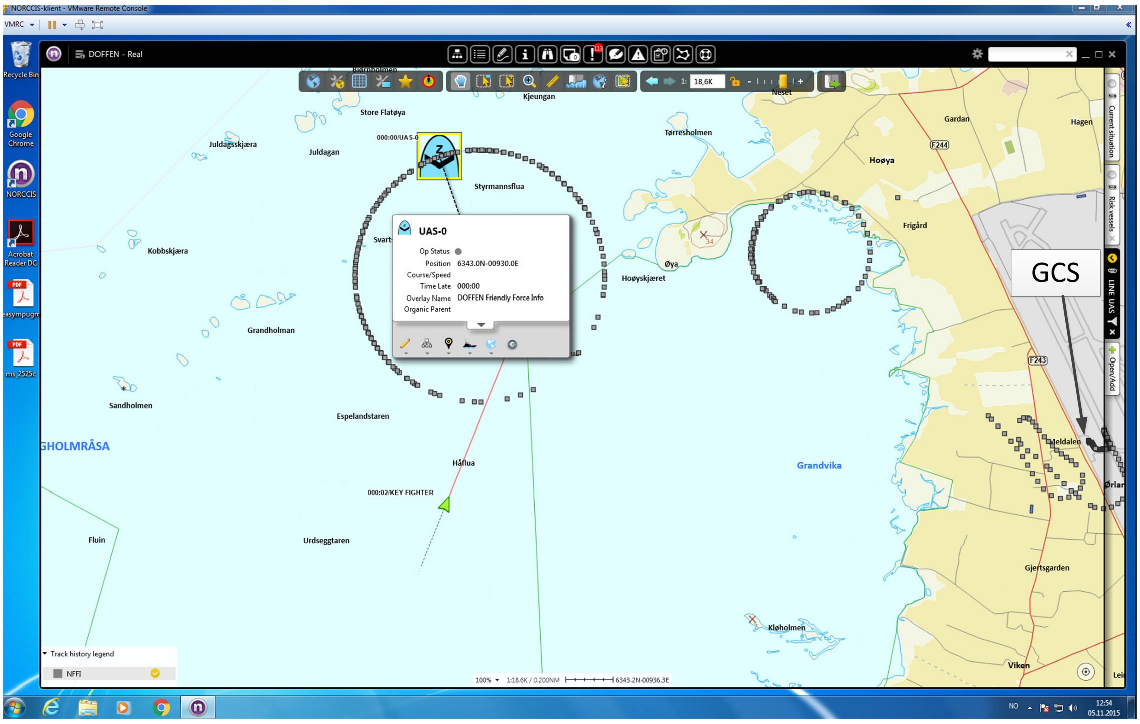

Figure 5.7 NorCCIS positions plot for UAV1. The GCS is at the right, by the airstrip. Time estimated to 1700 s into the experiment.

packets than UAV2. However, we were not able to identify the reason for the worse performance.

Based on the positions received and monitored at FFI, it appeared that the link between UAV1 and GCS was unreliable in several phases of the flight. Figure 5.7 shows a plot of the early stages of the experiment (around time 1700 s), before UAV2 had taken off. In the southeastern sector of the large circle, the number of position updates is much lower than in the rest of the circle. We can also see that position updates are lacking, both from the transit between just west of the airport and the small circle, and also from this first circle over to the large circle. However, as we will see, the observed lack of incoming position plots did not give a complete impression of the received data.

5.2.1 Circling pattern

The UAVs circled around the same center position at three different heights. UAV1 circled first at 400 m for 1445 seconds, and afterwards ascended to 500 m, circling here for 1250 seconds. UAV2 started circling at 460 m 240 s after UAV1 began its 500 m circling, and circled here for 1110 seconds.

The main statistics (Figure 5.8) show that there are differences between the repetitive phases, both between the two UAV1 circle phases and between UAV1 and UAV2. For UAV1 the packet loss ratio was very high when circling at 400 m. While the loss ratio was reduced to merely a fifth when the altitude was increased to 500 m, the loss ratio was still close to double that of UAV2 (Table 5.1).

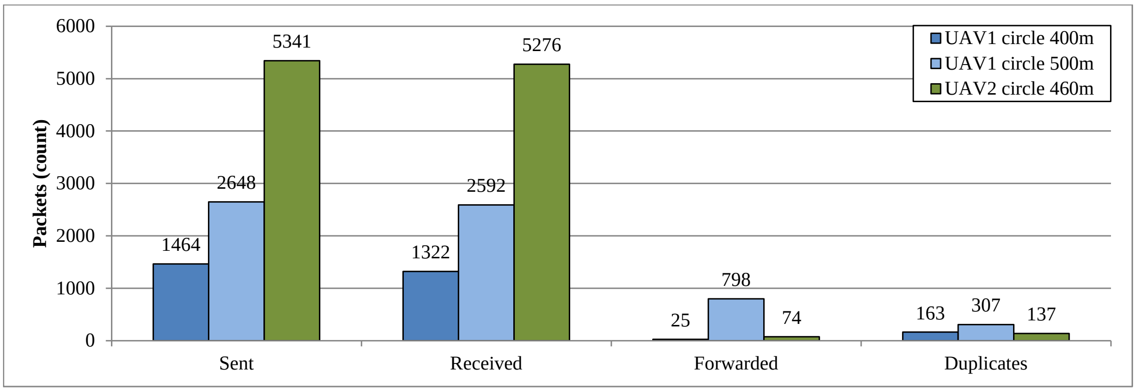

Figure 5.8 *Packets sent, received forwarded and duplicated for UAV1 and UAV2 in the repetitive circling phases, as logged at the GCS.*

Table 5.1 *Packet loss ratio for circling phases.*

Phase	Packet loss ratio
UAV1 circle 400 m	9.7 %
UAV1 circle 500 m	2.1 %
UAV2 circle 460 m	1.2 %

5.2.1.1 UAV1 circle at 400 m altitude

The first repetitive phase we examine is UAV1 circling at 400 m altitude between 985 s and 2430 s. The data reception varied a lot with the egress time (Figure 5.9). Note that the first packet reception did not happen until at 1250 s, even though UAV1 started the circling at 985 s. Further, we see that there are several gaps, some with small "islands" of packet receptions where the id seems to increase rapidly (e.g., time 1600 s). This could potentially be caused by buffered packets, by on-air packet loss or a combination of the two. In addition, the packet rate is much higher around time 2200 s than elsewhere during the phase, due to more packets being produced by the application. The application produces data at a variable rate.

To evaluate how the packet reception maps with the movement of the vehicle, we have plotted the packet reception with regards to the egress time and the Y-axis position (Figure 5.10). The GCS is placed in (0,0). Positive Y-axis position means that the vehicle was north of the GCS when the packet was received, and the unit on the Y-axis is meters from the GCS on the north-south axis. From the reception of the first packet (at around 1250 s), we see that the lack of received packets maps relatively good with the UAV reaching the southernmost position and starting to travel north again. We also see that towards the end of the circling at 400 m altitude, some packets have been forwarded via UAV2 (red crosses). This late in the experiment, UAV2 had taken off and was circling near Frigård.

5.2.1.2 UAV1 circle at 500 m altitude

We continue with UAV1's circling at 500 m altitude, which took place between 2480 s and 3730 s. The egress time synchronized data for packets from UAV1 in this phase (Figure 5.11) shows that for

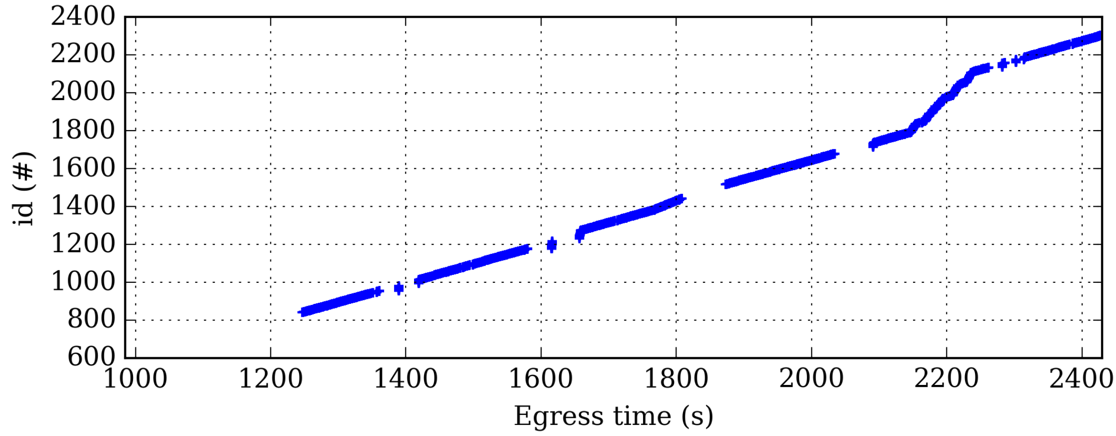

Figure 5.9 *IP identification for packet receptions over egress time. UAV1 circling at 400 m altitude.*

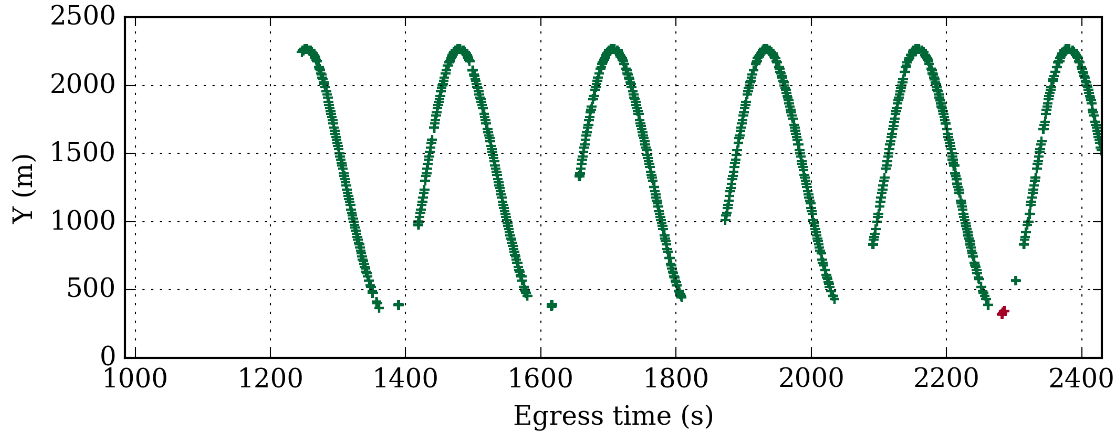

Figure 5.10 *Packet reception over egress time colored by TTL (red=forwarded). UAV1 circling at 400 m altitude.*

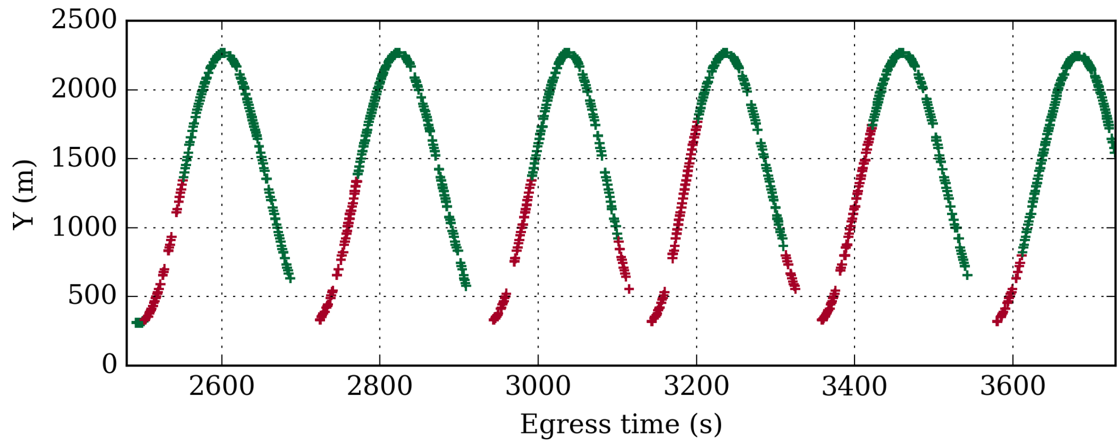

Figure 5.11 *Packet reception over egress time colored by TTL (red=forwarded). UAV1 circling at 500 m altitude.*

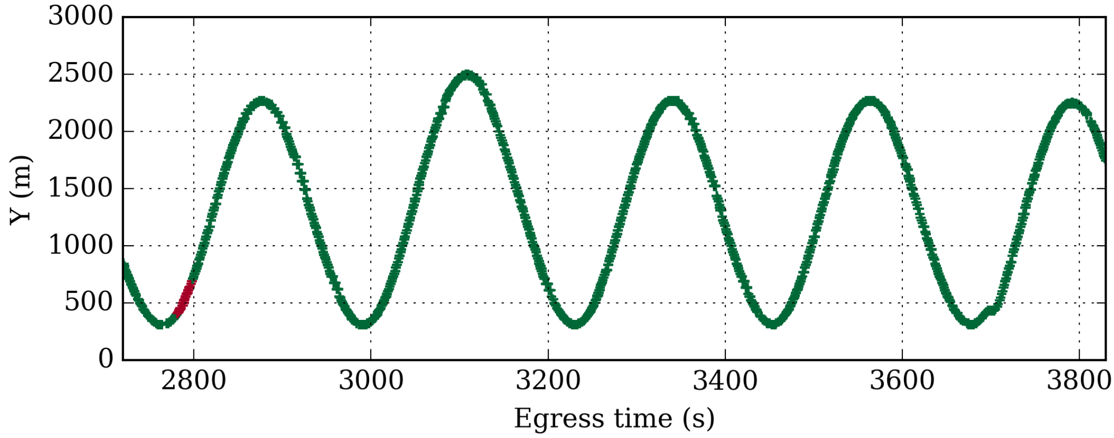

Figure 5.12 *Packet reception over egress time colored by TTL (red=forwarded). UAV2 circling at 460 m altitude.*

the most part the line is continuous throughout the circling. UAV2 has moved into the area, and we see that it forwards traffic for UAV1 in each and every circle. There is only a small gap around 2500 s where there has been an interrupt. Although the packet reception has been better at this altitude, more packets have been forwarded via UAV2, for larger sectors of the circle. Comparing with UAV1 circling at 400 m (Figure 5.10), the loss starts earlier and last for longer. There is some variation for each circle concerning where the direct link is broken, and where it returns. However, for all circles, the southern bend is vulnerable, and the link returns approximately half-way towards the northern bend.

5.2.1.3 UAV2 circle at 460 m altitude

UAV2 circled at 460 m between 2720 s and 3830 s. At this time, UAV1 was in the same circle, so any direct communication problems could be helped through rerouting via UAV1. Figure 5.12 shows that in fact UAV1 helps UAV2 just after its first southern bend (before 2800 s). However, except for this one place, causing a forward of the only 74 packets forwarded by UAV1 (Figure 5.8), the packets flow directly between UAV2 and the GCS throughout the circling phase.

5.2.2 Racetrack pattern

After making stable circles at 500 m and 460 m for UAV1 and UAV2 respectively, it was concluded that the pattern should be changed to a racetrack pattern to get better reception for the sensors. At 3730 s UAV1 started its racetrack pattern, while UAV2 waited until 3830 s before it began moving in the same pattern.

The main statistics (Figure 5.13) show that there are differences between the racetrack phases of UAV1 and UAV2. Even though UAV2 started its racetrack pattern 100 seconds later than UAV1, UAV1 only sent 56.8 % of UAV2's packets during this phase. The packet loss ratio (Table 5.2) shows that UAV1 continues to lose a lot more packets than UAV2, and while UAV2 forwarded

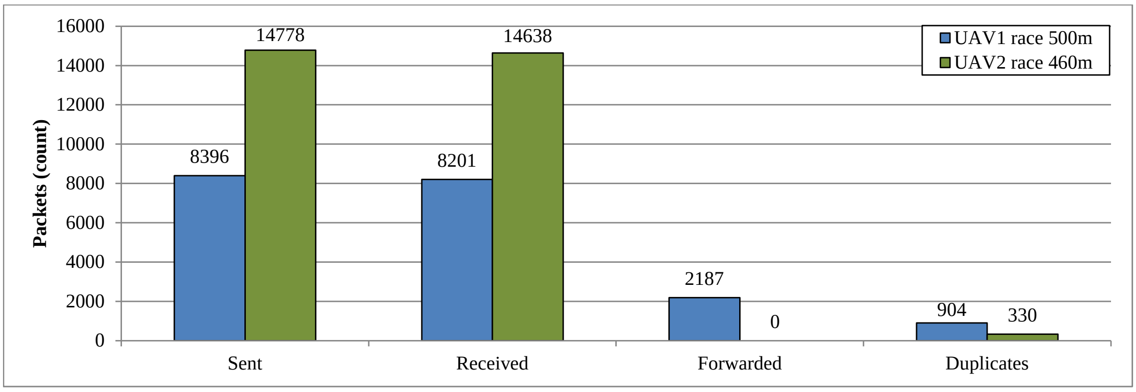

Figure 5.13 *Packets sent, received forwarded and duplicated for UAV1 and UAV2 in the repetitive racetrack phases, as logged at the GCS.*

Table 5.2 *Packet loss ratio for racetrack phases.*

Phase	Packet loss ratio
UAV1 race 500 m	2.3 %
UAV2 race 460 m	0.9 %

packets for UAV1, UAV1 has not done the opposite for UAV2. The ratio of forwarded packets for UAV1 in the racetrack pattern was 26.0 %, compared to 30.1 % while in the circle at 500 m.

5.2.2.1 UAV1 racetrack at 500 m altitude

We see that the packet reception (Figure 5.14) was stable for UAV1 during the entire phase. Even so, from the figure it appears that over a third of the packets have been relayed via UAV2. However, as the statistics show that the ratio of relayed to sent packets for UAV1 is 26.0 %, the green areas must be more densely packed with packet transmissions than the red areas. Regardless, the direct link is unavailable for substantial periods of time. The periodicity of these direct-link breaks is not as clear as in the circle phases. The link is mostly broken as the UAV hits the southern bend, but sometimes the link breaks when the UAV is abeam the GCS (Y=0) heading southwards, other times it only breaks just before the bend (at 4900 s and 5650 s). Likewise we see that some laps in the racetrack are very green (direct link), while others have a lot of red.

5.2.2.2 UAV2 in racetrack at 460 m altitude

Interestingly, UAV2 does not forward any of its packets through UAV1 in the racetrack phase. As Figure 5.15 shows, the direct link to the GCS is stable throughout the phase. There are some points where the line seems a little sparse, e.g., at 5600 s, and the packet loss ratio of 0.9 % confirms the visual impression of Figure 5.15 that UAV2 has also lost packets.

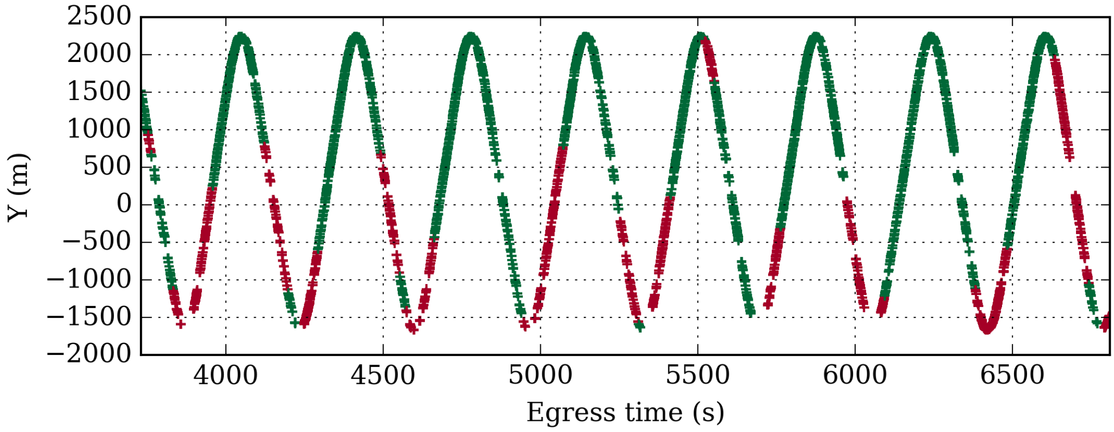

Figure 5.14 *Packet reception over egress time colored by TTL (red=forwarded). UAV1 in racetrack at 500 m altitude.*

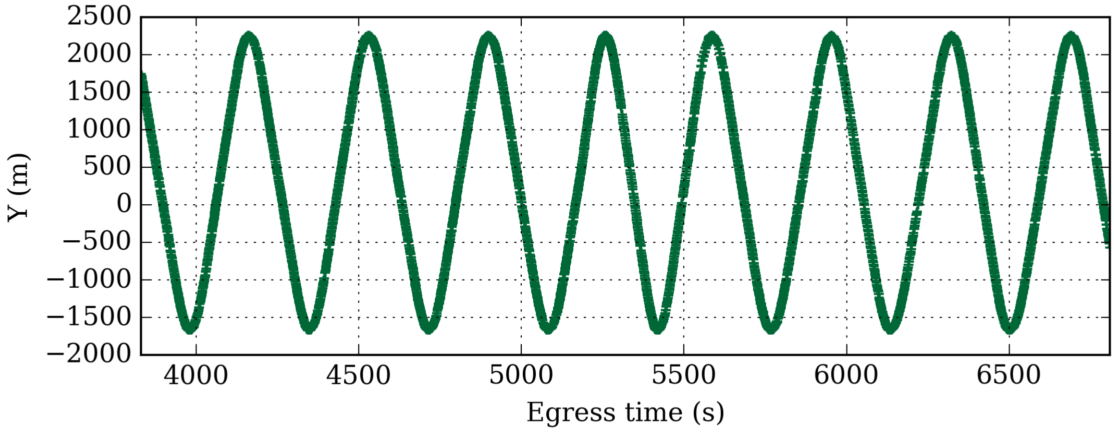

Figure 5.15 *Packet reception over egress time colored by TTL (red=forwarded). UAV2 in racetrack at 460 m altitude.*

5.2.3 Summary repetitive phases

UAV1 and UAV2 have very different performance when looking at the results scoped according to the repetitive phases. Although it is unfortunate that the link between UAV1 and the GCS was so unreliable, we were able to establish the following lessons learned:

- Creating a robust network using ad hoc technology may sustain communications during unforeseen events, such as a poor link.
- There are clear signs that there are challenges using an "omni-directional" antenna in a fixed-wing UAV, due to the movement of the vehicle.

Looking closer at the repetitive link breaks, it seems that the link is least reliable when the UAV is closer to the GCS. While this seems counter-intuitive, it can be explained by the antenna radiation pattern and a combination of the mounted antenna position and the vehicle itself. The nose wheel was positioned right in front of the antenna, breaking the LoS when the vehicle flew straight towards the GCS. However, it was not sufficiently explained how the vehicle-GCS azimuth angle for link breaks differed between the circular and racetrack patterns.

6 Conclusions and future work

In this report, we have evaluated the link and network performance of the LINE EW-UAS sensor network. We have derived the results from passive logging of the sensor application data.

6.1 Radio link and network performance

While a lot of focus in this report has been on the problems with UAV1, the radio system on-board UAV2 performed very well. We saw good connectivity at all operating ranges, extending outwards to 6 km, at an altitude of 460 m, and in all vehicle attitudes. We also saw good effect of the dynamic ad hoc network mechanism in the radios, forwarding a lot of the packets from UAV1 via UAV2 which would otherwise have been lost.

The KDC TacLAN radios produced a lot of duplicate packets. The most plausible reason for this is a problem in the implementation of the MAC-layer in the SR600 and WM600 radios. While the application can handle these duplicates in a proper manner, a more important reason to fix this is the potential waste of resources if these duplicates are transmitted over the air unnecessarily.

6.2 Antenna radiation challenges

The results show that there was a large difference in transmission loss between UAV1-GCS and UAV2-GCS links. The plausible explanation for this is differences between the communication antennas, either in production or with regards to mounting. The antennas were custom built at FFI and were mounted exposed to the weather and potential damage while handling the UAVs before the flight. Although this impacted the system performance negatively, it provided us with an opportunity to investigate the challenges of repetitive movement patterns for fixed-wing UAVs. The link breaks were clearly confined to certain positions during the circular movement. This points towards the antenna radiation pattern impacting the link performance.

6.3 Other lessons learned

Ad hoc networking with dynamic forwarding appeared in this experiment to be crucial in achieving a robust use of UAV-mounted distributed sensors. If the communication network had not supported forwarding, UAV1 would have had a 29.7 % packet loss instead of 2.3 % in Phase H.

During the experiment we observed linear clock drift between the two radios and the Wireshark logging PC at the GCS from the first packets received to the last. For UAV1, the clock drift was 593 ms (for a duration of 7818 s), and for UAV2 it was 464 ms (for a duration of 6097 s), where the clocks in the SR600 radios were faster than the clock in the Wireshark PC. Both results correspond to a drift of 76 µs per second for the SR600 radios, compared to the Wireshark PC.

The choice of using only user application packets to measure the link and network performance has been both advantageous and disadvantageous. The advantage has been that we have not affected the overall goal of the experiment negatively. If we were to send competing traffic while the experiment was ongoing, to make sure that the link situation was tested as much as we would have wanted, this could have caused congestion and increased the packet loss of the application data traffic. If we had allotted time during the experiment purely for communication tests, this would have limited the time available, and also limited the understanding of the interplay between user application and the use of the underlying communications infrastructure.

The disadvantage of passively measuring the link and network performance is that we may not have continuous measurements throughout the experiment, since the application had its own policies for transmitting traffic. A surprise in this respect was the difference in traffic volume produced by the sensors in the two UAVs.

6.4 Future work

The results in this report represent initial lessons learned from working experimentally with UAV-UAV and UAV-Ground communications at FFI from a communications infrastructure perspective.

Further work will focus both on distributed sensor experiments, such as the LINE EW-UAS system, but also on experiments aimed at improving ground communication systems. Autonomy is expected to allow extensive use of UAVs for communication purposes, and understanding the challenges of UAV-communication and translating this into improved control algorithms regarding communications limitations and possibilities for autonomous UAVs will be a priority.

Moving forward, range, coverage and LoS challenges in various terrain and environments should be tested, in addition to link rate experiments.

7 Abbreviations

ACK	Acknowledgment
AIS	Automatic Identification System
BER	Bit Error Rate
CBR	Constant Bit Rate
CESMO	Cooperative ESM Operations
CoNSIS	Coalition Networks for Secure Information Sharing
CSMA/CA	Carrier Sense Multiple Access/Collision Avoidance
DSSS	Direct Sequence Spread Spectrum
ESM	Electronic Support Measures
EW	Electronic Warfare
FANET	Flying Ad Hoc Network
FFI	Norwegian Defence Resesarch Establishment
GCS	Ground Control Station
ICMP	Internet Control Message Protocol
KDC	Kongsberg Defence and Communications
LAN	Local Area Network
LINE	Light Navigation-radar ESM
LINE EW-UAS	Light Navigation-radar ESM Electronic Warfare Unmanned Aircraft System
LLN	Link Layer Notification
LoS	Line-of-Sight
MAC	Medium Access Control
MANET	Mobile Ad Hoc Network
MoD	Ministry of Defence
MTU	Maximum Transmission Unit
NAF	Norwegian Armed Forces
NorCCIS	Norwegian Command, Control & Information System

OS	Operating System
OSPFv2	Open Shortest Path First version 2
pps	Packets per second
R&D	Research and Development
TacLAN	Tactical Local Area Network
TTL	Time-to-Live
UAV	Unmanned Aerial Vehicle
UHF	Ultra High Frequency (300-3000 MHz)
VANET	Vehicle Ad Hoc Network
VHF	Very High Frequency (30-300 MHz)

Bibliography

[1] I. Bekmezci, O. K. Sahingoz, and S. Temel, "Flying Ad-Hoc Networks (FANETs): A survey," *Ad Hoc Networks*, vol. 11, no. 3, pp. 1254–1270, 2013.

[2] E. S. Grimstvedt, M. Aronsen, E. Bergh-Nilssen, F. Gulbrandsen, Ø. Hoelsæter, B. Jahnsen, M. Kloster, Ø. Kure, E. Larsen, K. Lund, R. MacDonald, J. Moen, J.-R. Nilssen, J. Sander, T. Smestad, and T. Thoresen, "LINE EW-UAS. An experimental unmanned system for coastal surveillance using ESM technology," FFI, FFI-RAPPORT 2015/02442, 2016.

[3] Kongsberg. (2016) Kongsberg TacLAN UHF radio. [Online]. Available: http://www.kongsberg.com/en/kds/products/defencecommunications/taclan/

[4] *Cooperative Electronic Eupport Measure Operations*, North Atlantic Treaty Organization (NATO) Std. STANAG 4658, Edition 1 Version 1.1 DRAFT, 2014.

[5] J. Moy, "RFC 2328: OSPF Version 2," pp. 1–244, 1998. [Online]. Available: http://ietf.org/rfc/rfc2328.txt

[6] IEEE, "IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," pp. 1–2793, 2012. [Online]. Available: http://ieeexplore.ieee.org/servlet/opac?punumber=6178209

[7] Wireshark. [Online]. Available: https://www.wireshark.org/

[8] Z.-S. Su, "RFC 781: A Specification of the Internet Protocol (IP) Timestamp Option," 1981. [Online]. Available: http://ietf.org/rfc/rfc781.txt

[9] The Click Modular Router Project. [Online]. Available: http://www.read.cs.ucla.edu/click/

[10] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click Modular Router," *ACM Trans. Comput. Syst.*, vol. 18, no. 3, pp. 263–297, aug 2000. [Online]. Available: http://doi.acm.org/10.1145/354871.354874

[11] B. Hugsted, J. Sander, and M. Hauge, "Taktisk radio i felteksperiment og metode for rekkeviddemåling for trådløs datakommunikasjon," Norwegian Defence Research Establishment, FFI-RAPPORT 2009/01560, 2010.

About FFI

The Norwegian Defence Research Establishment (FFI) was founded 11th of April 1946. It is organised as an administrative agency subordinate to the Ministry of Defence.

FFI's MISSION

FFI is the prime institution responsible for defence related research in Norway. Its principal mission is to carry out research and development to meet the requirements of the Armed Forces. FFI has the role of chief adviser to the political and military leadership. In particular, the institute shall focus on aspects of the development in science and technology that can influence our security policy or defence planning.

FFI's VISION

FFI turns knowledge and ideas into an efficient defence.

FFI's CHARACTERISTICS

Creative, daring, broad-minded and responsible.

Om FFI

Forsvarets forskningsinstitutt ble etablert 11. april 1946. Instituttet er organisert som et forvaltningsorgan med særskilte fullmakter underlagt Forsvarsdepartementet.

FFIs FORMÅL

Forsvarets forskningsinstitutt er Forsvarets sentrale forskningsinstitusjon og har som formål å drive forskning og utvikling for Forsvarets behov. Videre er FFI rådgiver overfor Forsvarets strategiske ledelse. Spesielt skal instituttet følge opp trekk ved vitenskapelig og militærteknisk utvikling som kan påvirke forutsetningene for sikkerhetspolitikken eller forsvarsplanleggingen.

FFIs VISJON

FFI gjør kunnskap og ideer til et effektivt forsvar.

FFIs VERDIER

Skapende, drivende, vidsynt og ansvarlig.

FFI's organisation

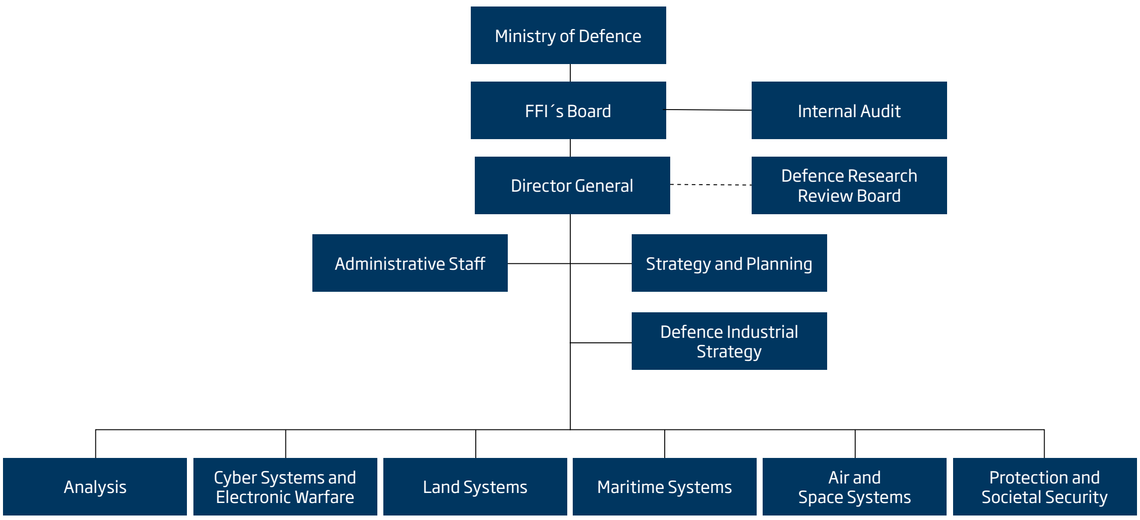

Forsvarets forskningsinstitutt
Postboks 25
2027 Kjeller

Besøksadresse:
Instituttveien 20
2007 Kjeller

Telefon: 63 80 70 00
Telefaks: 63 80 71 15
Epost: ffi@ffi.no

Norwegian Defence Research Establishment (FFI)
P.O. Box 25
NO-2027 Kjeller

Office address:
Instituttveien 20
N-2007 Kjeller

Telephone: +47 63 80 70 00
Telefax: +47 63 80 71 15
Email: ffi@ffi.no

FFI Forsvarets forskningsinstitutt
Norwegian Defence Research Establishment