This project has received funding from "HORIZON 2020" the European Union's Framework Programme for research, technological development and demonstration under grant agreement no 645220

Road-, Air- and Water-based Future Internet Experimentation

Project Acronym:	RAWFIE		
Contract Number:	**645220**		
Starting date:	**Jan 1st 2015**	**Ending date:**	**Dec 31^{2t}, 2018**

Deliverable Number and Title	UNSURPASSED D1.4: 2nd Progress Report		
Confidentiality	RAWFIE internal	**Deliverable type**[1]	R
Deliverable File	D14_UNSURPASSED.pdf	**Date**	31/5/2018
Approval Status[2]		**Version**	V1.0
Contact Person	Stavros Toumpis	**Organization**	AUEB
Phone	+30 210 8203551 +30 6974 740063	**E-Mail**	toumpis@aueb.gr

1 Deliverable type: P(Prototype), R (Report), O (Other)
2 Approval Status: WP leader, 1st Reviewer, 2nd Reviewer, Advisory Board

AUTHORS TABLE

Name	Company	E-Mail
Esmerald Aliaj	AUEB	aliai16@aueb.gr
Georgia Dimaki	AUEB	ginadimaki135@gmail.com
Nikos Fotiou	AUEB	fotiou@aueb.gr
Petros Getsopoulos	AUEB	pgetsopoulos@hotmail.com
Iordanis Koutsopoulos	AUEB	jordan@aueb.gr
Vasilios Siris	AUEB	vsiris@aueb.gr
Stavros Toumpis	AUEB	toumpis@aueb.gr

REVIEWERS TABLE

Name	Company	E-Mail

DISTRIBUTION

Name / Role	Company	Level of confidentiality[3]	Type of deliverable

CHANGE HISTORY

Version	Date	Reason for Change	Pages/Sections Affected
1.0	31/5/2018	Original Submission	

[3] Deliverable Distribution: PU (Public, can be distributed to everyone), CO (Confidential, for use by consortium members only), RE (Restricted, available to a group specified by the Project Advisory Board).

Abstract: In this deliverable we describe the progress in the project in the months since the previous deliverable, D1.3 (31/3/2017). Notably, we describe experiments conducted in three visits to the testbed, on 17/5, 24/5, and 31/5.

Keywords: **EDL, Experiments, UNSURPASSED**

Table of Contents

List of Tables

List of Figures

1. Description of Testbed Experiments

As outlined in our previous deliverable, D1.3, we visited the Skaramagas testbed on three occasions, on 17/5, 24/5, and 31/5. On all occasions, our team comprised E. Aliaj, G. Dimaki, and S. Toumpis. The experiences were very useful, as, on one hand, we verified that parts of the software we have developed worked as intended and we collected preliminary measurements, and, on the other hand, we spotted problems with the integration of our software to RAWFIE, which will have to be resolved in the coming few months. In the following, we organize the exposition around the dates of the experiments.

Visit 1 (17/5/2018)

The aims of the first visit were firstly to verify that the wireless environment at the test site did not interfere with our hardware and secondly to verify that our Raspberry Pis are able to communicate with our broker using the Apache Kafka software. Following the lead of the Glasgow team, we decided to proceed modestly, and for the experiments of that day, we used a Kafka server installed on a laptop that communicated with the Raspberry Pis through a wireless router, both laptop and router brought by us.

Regarding the RAWFIE hardware we used, we used three Flexus USVs, on which we attached zip-lock bags with our nodes. We used a total of five nodes, one on each of the three USVs, as well as two more, which were placed on the shore. However, the majority of our measurements involved only two boats, due to the fact that the primary aim was to measure the quality of the wireless communication

Overall, the experiments of that date were very successful. We collected numerous measurements, which helped us understand the channel, and did discover a number of hurdles that must be resolved.

The measurements we collected were of two types:

- Link costs, as estimated by the BABEL routing protocol, between existing links. These measurements were continuous and were stored in log files in the Raspberry Pis as well as in the Dedalus client that was monitoring the topology of the network.
- Throughputs, as measured by inserting traffic in the network using the iperf tool.

In more detail, in Table 1 we have collected measurements on the throughput of single links (no multihop experiments were involved). Throughput was estimated by collecting measurements on received data packets for one second. The measurements were divided in three groups.

- The first group involved measurements on land, where two Pis were placed on a table next to each other, and various power levels were explored. Therefore, there was no source of mobility.
- The second group involved measurements with the Pis placed on the boats while these were not travelling to any given destination, but placed in the water, and controlled either by the EDL or by the remote control unit. The only source of mobility was the rocking of the boats.
- The third group involved measurements with the Pis placed on the USVs while these were traveling towards some destination, either through EDL instructions or the remote control. Therefore, there were two sources of mobility, i.e., the rocking of the boat and the traveling of the USVs.

Two major, rather unexpected, findings of these experiments are:

- Firstly, mobility reduces the throughput between the transmitter and the receiver in such a manner that increasing the transmitter power does not provide any meaningful compensation. We attribute this to the fact that there are no obstacles to provide a multihop environment, and as a result when the line of sight is gone, communication is no longer possible, irrespective of the transmitter power. It would be interesting to compare the results with similar results obtained indoors.
- Secondly, even a gently rocking motion of the USVs can reduce the achieved throughput significantly.

In Figure 1 we have plotted the snippet of our measurements of the link cost between a transmitter and a receiver, when these are placed on two boats of which the first is stationary and the second is moving in circles around the first, manually, using the remote control. The link cost is reported by the BABEL protocol. Note that the link costs are highly asymmetric in the rate with which they change. This was another unexpected result, that should probably be attributed to the fact that one of the boats was moving and the other was not. This puzzling issue will be investigated further. We note that our logs contain these measurements for the whole duration of the experiments (which all-in-all lasted for around two hours).

Regarding the problems we faced, one issue was the fact that the GPS receivers of the boats had an accuracy somewhat worse than what we anticipated, on the order of 5 meters. This poses a challenge in deriving plots of KPIs of a link versus the distance, and also places some limits in specifying some experiments on the EDL. (For example, it is not possible to instruct three nodes to stay on a line so that the middle one obscures the line of sight between the other two.) As a suggestion to the EDL developers, it might make sense to provide the experimenters with some feedback about whether their experiment demands too much location precision. (For example, it might be possible to display, together with a node, a circle showing the area where the node might actually be found during the experiment.)

A second problem we faced was that due to the variability of the channel and the fact that the boats are moving, it is very hard to specify beforehand the topology of the network in terms of which USV pairs can support traffic. In particular, it was hard to place three boats on a line so that the one in the middle can communicate with both the other boats, and these two cannot communicate directly. What was happening was that for some time all boats could communicate directly with each other, and for the rest of the time one of the boats was disconnected from the other two. In fact, we managed to succeed in only one multihop traffic exchange (although this was by no means the focus of the experiments), where a modest 251kbps was supported.

On the less technical side, a final lesson learned was that it would have useful to be able to tell the boats apart, when they are at a distance. Due to the distances involved, this would require a flag or other sign of significant size, that might affect the performance (e.g., speed and controllability) of the USVs, especially in windy conditions.

Apart from these problems, the experiments were very successful. Firstly, our GUI-based client was very useful in providing information on the online topology of the network, which accelerated progress, and we are sure other experimenters will find it very useful as well. Secondly, the Raspberry Pis were able to maintain high-throughput links even under mobile conditions, and even using the internal WiFi cards.

Visit 2 (24/5/2018)

The second visit was far less successful than the first. Its first aim was to verify that we were able to perform experiments using the Kafka server of the testbed. To this effect, an important first step was to connect our nodes to the RAWFIE WiFi, which we were not able to achieve. Following an extensive debugging process, we reached the conclusion that there was interference from the antennas of the base station communicating with the Flexus boats. Indeed, turning off that base station solved the problem, but, in the meantime, the day was lost. We have been investigating numerous workarounds to this problem, each with its own advantages and disadvantages. Specifically:

- One option is to change the 2.4 GHz channel used by the RAWFIE WiFi infrastructure. This requires action from the side of the RAWFIE operators and introduces overhead; we are yet to ascertain if this option is viable, as other parts of the network might be affected.
- Another option is to use, for the communication of our boats with the Kafka server, the access points provided by the Flexus boats themselves. This option is viable only when we use the Flexus boats. Again, this requires action on the side of the RAWFIE operators, and we have requested the pertinent information.
- A third option is to switch our hardware platform from the Raspberry Pi 3 Model B to the recently introduced Raspberry Pi 3 Model B+, which is equipped with dual band (2.4 GHz/5 GHz) NICs, hoping that the interference of the Flexus Base Station is limited to the 2.4GHz ISM band. We are currently exploring this possibility.

Visit 3 (31/5/2018)

The third visit was more successful. One complication was that the sea was rougher than during the first two visits, and for security purposes it was possible to only have one boat on the water. (On the plus side, we verified that the zip-lock bags provided protection, despite being repeatedly drenched in salt-water). We compensated by having another two nodes close to the beach, but distant from each other, so that the boat in the middle could provide connectivity to them. A number of experiments were conducted with this arrangement.

- Firstly, we conducted data muling experiments using the IBR-DTN protocol, and using the boat in the middle as a data mule between the two boats.
- Secondly, we used the BMX ad hoc protocol to carry traffic from one node on the beach to the other one, using the boat as a relay (on our first visit, we used the BABEL routing protocol).
- Thirdly, we took measurements of the link quality, as reported by the BMX routing protocol, versus the link distance, by placing the two nodes on the shore next to each other, and moving the USV from a location close to them to a location far out to sea, where the connectivity was lost.
- Finally, we experimented with obstructing the line of sight of nodes by using improvised aluminum deflectors. No measurements were collected for these experiments. It was found that, even in this relatively scatter-free environment, obstructing the communication of WiFi devices is not easy, and the results are not predictable.

2. Other Progress

Other experiments

We also note that we have conducted numerous other experiments, related to all four tasks of our project, in our laboratory. In most of these, we have used the Kafka server of RAWFIE, in order to emulate, as much as possible, the conditions on the testbed.

Future steps

In the coming months, and as we move closer to the end of the project, we anticipate further visits to the testbed.

We plan to use the coming month of June to digest the lessons learned and prepare for another round of experiments in the month of July.

The immediate steps envisioned for experiments in the month of July are:

- Experiments involving more boats, i.e., around ten, possibly also of different types, to explore the limits of our various components, as also described in past deliverables.
- Experiments involving larger distances, in order to ascertain the limits of the wireless infrastructure.
- Experiments involving traffic specified by the EDL, as described in past deliverables.
- A comparison of our experimental results with those obtained indoors under similar settings.
- A streamlining of the parsing and processing of experimental results.
- Other measurements, specifically link qualities as measured by the NIC cards; such measurements would be of great immediate practical use to other experimenters, even those not able to use the other RAWFIE infrastructures.

Publications

During the period since the previous deliverable, D1.3, our paper titled “Analysis of Hybrid Geographic /Delay-Tolerant Routing Protocols for Wireless Mobile Networks”, co-authored by R. Cavallari, S. Toumpis, and R. Verdone, and reported in D1.1, was presented in Infocom 2018 (Honolulu, Hawaii, USA).

Also, our paper titled “Asymptotics of the Packet Speed and Cost in a Mobile Wireless Network Model,” co-authored by I. Kontoyiannis, S. Toumpis, R. Cavallari, and R. Verdone, and reported in Deliverable D1.2, was accepted to IEEE ISIT 2018 (Vail, Colorado, USA).

Finally, another paper, titled “Analysis of a One-Dimensional Continuous Delay-Tolerant Network Model,” co-authored by D. Cheliotis, I. Kontoyiannis, M. Loulakis, and S. Toumpis, was submitted and accepted in IEEE SPAWC 2018 (Kalamata, Greece), to be presented as a poster and included in the proceedings. That paper is appended to this deliverable.

3. Acknowledgment

We would like to gratefully acknowledge the crucial help of Kostas Kolomvatsos, Kakia Panagidi, Kostis Gerakos, Spyros Bolis, and the HN personnel for helping us with our

debugging, providing information about the RAWFIE testbed, and helping us with conducting our experiments.

Table 1 Throughput measurements

MEASURE MENT	Source ID	Dest. ID	Throughput (Kbps)	Topology
PREPARATORY MEASUREMENTS ON LAND				
1	10	12	3076	2 nodes placed next to each other on a table, transmitting with 0 dBm
2	10	12	2855	As above
3	12	10	4771	As above
4	12	10	2524	As above
5	11	12	43712	2 nodes placed next to each other on a table, transmitting with 30 dBm
6	12	11	44315	As above
7	10	12	33715	As above
8	12	10	33327	As above
9	11	10	42990	As above
10	10	11	42489	As above
11	10	12	30790	As above
12	12	10	31678	As above
MEASUREMENTS WITH NO UAVs TRAVELING				
13	10	11	2454	5 meters apart, 30dBm
14	10	12	25666	3 meters apart, 30dBm
15	12	10	25116	As above
16	12	10	21779	As above
17	10	12	22528	As above
18	10	12	24486	5-8 meters, 10dBm
19	12	10	24983	As above
MEASUREMENTS WITH UAVs TRAVELLING				
20	12	10	5314	4-8 m apart moving in parallel, 10dBm
21	10	12	3195	As above
22	12	10	19758	5m apart, one boat moving in circles around the other one in manual mode, 10dBm
23	10	12	20466	As above
24	12	10	21306	As above
25	10	12	24971	As above
26	12	10	25516	As above
27	10	12	24496	As above
28	10	12	24112	As above
29	12	10	24438	As above

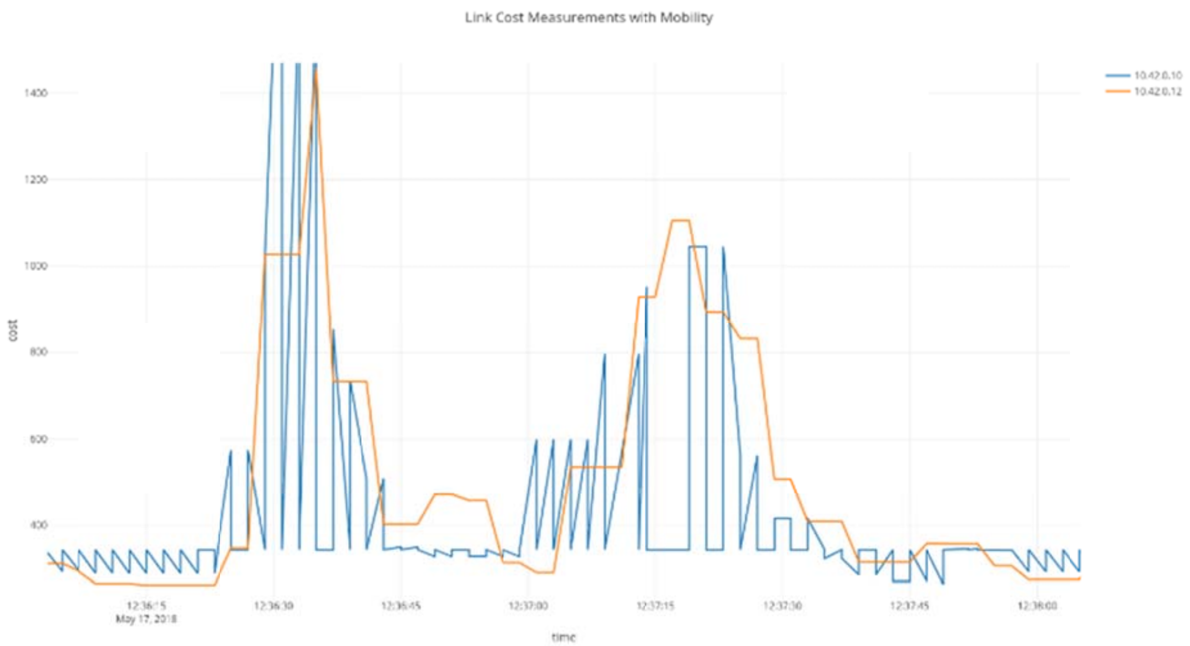

Figure 1: Snippet of link cost measurements