

Enabling the Maritime Internet of Things: CoAP and 6LoWPAN performance over VHF links

David Palma

Abstract—The dissemination of digital devices across the seas and oceans of the world is becoming a reality, laying the foundations for sophisticated operations such as autonomous shipping or exposed fish-farming. This reflects the importance of maritime regions as well as the growing demand for more data and knowledge in these remote areas. Currently, the presence of Information and Communication Technologies (ICTs) in offshore settings is limited by the lack of infrastructures, with communications being typically decoupled from the standard Internet or relying on expensive satellite solutions. This paper proposes integrating a commonly available maritime radio-technology, namely Very High Frequency (VHF) communications, with protocols used in the Internet of Things (IoT). By using the Internet Protocol version 6 (IPv6) over a Low-Power Wireless Personal Area Networks (6LoWPAN) and the Constrained Application Protocol (CoAP), interoperability is ensured between already existing maritime systems and the Internet. An experimental setup is evaluated under different settings and using different configurations. In addition, an analytical assessment of the solutions' reliability is presented. The obtained results prove the feasibility of using IoT protocols over a VHF link and demonstrate that such an approach outperforms IPv4-based solutions. The conducted assessment reveals that, in constrained settings, configurations at the application layer (e.g. CoAP's Block size in Block-wise transfers) strongly impact the overall performance and reliability of the system. This motivates the development of new maritime solutions based on IPv6 for ensuring a sustainable and ubiquitous development of ICTs even in the most remote locations.

Index Terms—Internet of Things, VHF devices, Computer Network Reliability, IP Networks, Maritime Communication, 6LoWPAN, IP Fragmentation.

I. INTRODUCTION

THE digitalisation of maritime operations is becoming increasingly more important. It enables innovative, safer and more sustainable ways of navigating the seas and oceans of the world, which account for 70% of its surface and have a deep impact on its biosphere [1]. This innovation includes autonomous shipping [2], offshore and exposed precision fish-farming [3], sensor networks [4], among other systems which strongly rely on the development of Information and Communication Technologies (ICTs) in maritime settings. However,

This work was partially funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 699924, SINet. The author would also like to thank Maritime Robotics, the Coastal and Arctic Maritime Operations and Surveillance lighthouse project (CAMOS) and the valuable comments from the reviewers.

(Corresponding author: David Palma.)

David Palma is with the Department of Information Security and Communication Technology at Norwegian University of Science and Technology (NTNU), Trondheim, Norway (email: david.palma@ntnu.no).

Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

due to the lack of infrastructures and harsh conditions in remote maritime areas communications are limited [5].

Very High Frequency (VHF) radios are commonly used by vessels [6] but are generally confined to voice communications or to simple pre-defined digital messages. This limits communications to few purposes such as collision prevention by sending information regarding positioning, speed and heading. In this paper, a proposal to extend the use of VHF communications is presented, towards a maritime Internet of Things (IoT).

Due to the limitations of VHF links, such as low bandwidth and small packet size, the Internet Protocol version 6 (IPv6) [7] over Low-Power Wireless Personal Area Networks (6LoWPAN) was considered to be an appropriate solution, as it has been in the past for other constrained link-layers [8]. The hypothesis is that this allows for an efficient application of IPv6 in Low-Rate Wireless Personal Area Networks (LR-WPANs) as specified by IEEE 802.15.4 [9], benefiting from the existing adaptation layer [10] and compression mechanisms [11]. The validation of this approach was achieved by using a real setup with the VHF-based OWL radio¹, which by default only supports IPv4.

In addition to the use of IPv6, the Constrained Application Protocol (CoAP) [12] was used due to its suitability for IoT constrained nodes, since it was designed to avoid the negative impacts of IP fragmentation [13], likely to occur in LR-WPANs due to their small Physical-layer Service Data Unit (PSDU). This is achieved by limiting the protocol's message overhead [12], as well as by introducing Block-wise transfers [14], because fragmentation may also result from the size of payload being transmitted, either by source or destination nodes. In this paper, we verify that the choice of Block size can strongly affect the performance of a link, where smaller sizes will increase overhead while larger ones increase unreliability, especially in lossy links.

By taking advantage of widely adopted IoT technologies and protocols, the proposed approach enables maritime communications from a feasible and realistic perspective. This allows taking advantage of many existing features of IP-based communications, including security options, which were also tested by combining CoAP with the Datagram Transport Layer Security (DTLS) protocol [15]. In addition, an assessment of the proposed solution is presented, considering its performance and the trade-off between overhead and reliability for different values of the CoAP Block size, from 32 to 1024 bytes

The source code developed to enable a LR-WPAN over the VHF radio-link was based on the already existing 6LoWPAN

¹<http://skagmoelectronics.com>

support of the Linux Kernel and is freely available online². Similarly, the used CoAP implementation, which is based on FreeCoAP, can also be found online³.

Related literature addressing maritime communications and the use of IoT in such scenarios is discussed in Section II, including an overview of the limitations of constrained networks. The proposed approach and methodology used to enable and evaluate 6LoWPAN over a VHF radio is detailed in Section III, followed by a theoretical performance analysis in section IV. This analysis is confronted and complemented against results obtained in a real testbed, presented in Section V. Finally, concluding remarks are provided in Section VI.

II. RELATED WORK

The importance of communications and networking in maritime environments has been widely addressed by researchers over the years. Existing works study, for instance, several aspects regarding the implementation of Wireless Sensor Networks (WSNs) in order to instrument the oceans [4], [16]. They consider the challenges and opportunities resulting from the unexplored and harsh conditions typically found in these settings, resorting to different communication technologies, including VHF radios [17], as well as distinct networking protocols. Several works focus on the principles of Delay Tolerant Networking (DTN) [18], [19], Mobile Ad-hoc Networking (MANET) [20], or even on the combination of different protocols [21], in order to address the lack of available infrastructures.

In addition to typical communication infrastructures, the increasingly popular use of unmanned vehicles (e.g. aerial, surface, underwater) is also being acknowledged as a potential solution for improving data acquisition in remote areas [5], [22]. These vehicles can be used not only to acquire data with higher-levels of precision and accuracy [23], [24], being able to support operations such as autonomous shipping [25], but also to interconnect different communication systems [26]. Still oriented towards maritime environments, small satellites or SmallSats, are also perceived as excellent candidates due to their reduced cost [27], increased coverage of remote areas and to the suitability of communication technologies such as VHF [28].

While VHF is widely used in maritime communication systems, several authors often consider other options as a consequence of the myriad of possible ICT applications. These can be selected for several reasons, such as bandwidth, range, energy efficiency, commercial-of-the-shelf availability or even due to features such as builtin privacy and security [29]. The diversity of available technologies, in addition or as an alternative to VHF, include Wireless Fidelity (Wi-Fi), LR-WPANs or even cellular networks such as the General Packet Radio Service (GPRS), Long-Term Evolution (LTE) and Worldwide Interoperability for Microwave Access (WiMAX).

The co-existence of multiple radio access technologies (RATs) in maritime settings has the potential of efficiently utilising resources according to existing needs, however it may

also lead to increased complexity and fragmentation [30]. This issue is being addressed by the ongoing activities towards the 5th Generation of Wireless/Mobile Systems (5G) in order to enable communication transparency in distinct use cases [31], while also supporting smart remote operations. 5G has already been studied for the coverage of maritime scenarios [32]–[34] with IPv6 as an enabler for seamless and interoperable operations. This is aligned with the existing availability of IPv6 across the spectrum of suitable maritime networking options, from LR-WPANs to satellite-based communications [35].

Focusing specifically on the definition of LR-WPAN within the VHF spectrum, the IEEE 802.15.4m amendment [36] includes VHF and UHF frequencies from the so called Television White Space (TVWS), suitable for smart utility networks, machine-to-machine communications, among other use cases [37]. LR-WPAN and 6LoWPAN are usually associated to devices with limited battery, memory and processing power, composing networks referred to as Low-Power and Lossy Networks (LLNs) [38]. In these networks, an efficient use of resources can be critical and the fragmentation of packets in lossy links is known to be a cause of wasted resources due to the forwarding fragments regardless of the ability to reconstruct a packet (e.g. when a fragment is dropped) [39].

The CoAP protocol was designed following a client/server representational state transfer-style (RESTful), focusing on constrained devices and simplifying its messages with a header of only 4 B [12]. CoAP uses the User Datagram Protocol (UDP) and provides additional mechanisms such as confirmable messages or Block-wise transfers, where the latter was defined “...in order to minimize the number of fragments needed for each message (to maximize the probability of delivery of the message)...” [14]. From the application layer’s perspective, Block-wise transfers in CoAP are impacted by the chosen Block size, delays and lost packets [40], resulting in longer completion times and additional overhead. However, in addition to this, the use of other features such as DTLS or even specific characteristics from lower layers (e.g. 6LoWPAN header compression) may also affect the amount of transmitted data and the need for fragmentation.

III. PROPOSAL AND METHODOLOGY

Driven by the myriad of applications and communication technologies considered for offshore operations, this work proposes extending the use of already existing IoT protocols to such environments. Specifically, the use of 6LoWPAN and CoAP as enablers of the maritime IoT was studied over VHF links and exploiting their support for LLNs. A real VHF radio was used with a small PSDU of 127 B, as specified by the IEEE 802.15.4 Medium Access Control (MAC), and the feasibility of using IPv6 was compared against the already existing IPv4 alternative.

The use of CoAP with 6LoWPAN allows a lightweight interface for Hypertext Transfer Protocol (HTTP) requests, ensuring interoperability with the Internet while simultaneously meeting the requirements of resource-constrained networks. This matches the goal of connecting maritime scenarios to the Internet of Things, where normal HTTP requests may

²<https://github.com/PalmaITEM/6lochar>

³<https://github.com/PalmaITEM/FreeCoAP> (maritime_robotics branch)

be forwarded to maritime nodes through several links, each with distinct characteristics. The interface provided by CoAP is derived from its proxy functionality, maintaining a request/response interaction model between endpoints with low-overhead messages and using UDP instead of the Transmission Control Protocol (TCP).

By configuring a testbed consisting of two nodes running a Linux distribution, each one equipped with an OWL VHF radio, the proposed solution was assessed in a controlled environment (i.e. with near-perfect link conditions to guarantee repeatability). This eliminates variability in interferences, latency or congestion, which could be handled by more sophisticated CoAP mechanisms such as defined by the Simple Congestion Control/Advanced (CoCoA) [41].

The IP addresses of the used nodes were automatically configured based on their interface identifier (IID) – created using modified IEEE EUI-64 format interface identifiers [42] – so that 6LoWPAN features could be fully exploited, ensuring IP header compression. In our setup, this results in IP headers being reduced from 20 B in IPv4 and 40 B in IPv6 to only 2 B in 6LoWPAN since it uses IEEE 802.15.4 MAC-layer data, corresponding to 9 B. Additionally, CoAP’s source and destination ports were defined to be within a specific range (between 61 616 and 61 631), taking advantage of 6LoWPAN’s Next Header Compression (NHC) mechanisms and saving 3 bytes per request as a consequence of compressed UDP headers.

In order to evaluate the proposed setup, a *GET* request was issued by an HTTP client for retrieving a 13 kB image located in a maritime node. This request was forwarded through a CoAP proxy towards a CoAP server using Block-wise transfers over a real VHF link. While the initial HTTP request was always sent using IPv4, the exchange of CoAP messages was assessed using different configurations. In particular, the setup was tested while using both IPv4 and 6LoWPAN/IPv6 with and without encryption (i.e. DTLS), demonstrating the co-existence of multiple technologies while also comparing their performance. The used CoAP proxy was configured to use Confirmable requests (CON), following the recommendations for Block-wise transfers [14] and as a way of providing better guarantees of reliable communication over lossy links.

The adequate configuration of protocols and their parameters is tightly connected to the radio link properties and limitations such as the maximum PSDU. In addition, other mechanisms such as channel-access also have a direct impact on the performance of networks (e.g. latency). The used radio allows using both Time-division multiple access (TDMA) and Carrier-sense multiple access (CSMA), which were both evaluated in the performed experiments for the sake of completeness. TDMA was configured in two different settings, one where frames were 300 ms long and another with 1000 ms of length. In both configurations each direction of the link was allocated one TDMA-slot, with a total of 2 slots per frame (1 per node), where one of the nodes was configured as TDMA master and was responsible for transmitting a *time of hour* packet for establishing a common time reference.

The chosen TDMA configuration guaranteed symmetric links, fitting a generic scenario where data may be equally

downloaded or uploaded over a constrained link. For example, since Block-wise CON connections rely on a set of requests and responses/acknowledgements per block, link requirements should be symmetric. Indeed, considering that each request must include the content’s Uniform Resource Identifier (URI), a CoAP request may be as long as a response for smaller Block sizes. Nonetheless, if URI shortening techniques are used [43], or if the defined Block size leads mostly to the fragmentation of responses, then the radio downlink should be given more slots within the available TDMA-frame length, optimising the overall usage of “air time”.

IV. PERFORMANCE ANALYSIS

As discussed by previous works available in the literature, the performance of CoAP Block-wise transfers is affected by the chosen Block size. Additionally, the underlying physical layer used for communication may influence the overall performance if fragmentation is to occur. In particular, since we consider an IEEE 802.15.4 MAC layer, any Block larger than 64 B will be fragmented and affect performance differently depending on whether or not link-losses exist. This section analyses the total overhead of a complete Block-wise transfer according to the probability of a packet being lost or dropped by the network, as well as the reliability of a transfer for a given number of CoAP retransmissions of a confirmable request. Finally, a trade-off analysis between overhead and reliability is given.

A. Overhead Analysis

In a Block-wise transfer, content to be transferred with a total size of T_{reply} bytes is split into q queries, depending on the chosen Block size (T_{block}) as defined in (1).

$$q = \lfloor T_{reply}/T_{block} \rfloor, \\ T_{block} \in \{32, 64, 128, 256, 512, 1024\} \text{ bytes} \quad (1)$$

For each confirmable query, a *GET* and an *ACK* will be issued by CoAP, adding to the total amount of necessary bytes needed for one transfer. The best, or lowest, theoretical amount of bytes to be transmitted per transfer (T_{best}) is defined by (2), where the overhead from the *GET* request ($GET_{overhead}$) will be constant per query and the *ACK* overhead ($ACK_{overhead}$) will be higher for the first fragment (*Headers*) and lower for subsequent ones (*Frag*).

$$T_{best} = T_{reply} + q \times (GET_{overhead} + ACK_{overhead}) \text{ (bytes)} \quad (2)$$

In addition to the Block with the desired content and necessary CoAP headers, networking headers are also added. All these headers contribute to the final size of each frame, which may have to be fragmented in order to meet the PSDU defined by the IEEE 802.15.4 MAC layer ($F_{len} = 127$ B). The number of fragments, f , is defined by (3), which depends on the number of available bytes for the first fragment (F_1) and subsequent fragments ($F_{>1}$). These values were respectively 91 B and 105 B in the tested CoAP and 6LoWPAN setup.

$$f = \left\lceil \frac{T_{block} - F_1}{F_{>1}} \right\rceil + 1 \text{ (bytes)}, \quad F_1 = F_{len} - Headers, \\ F_{>1} = F_{len} - Frag \quad (3)$$

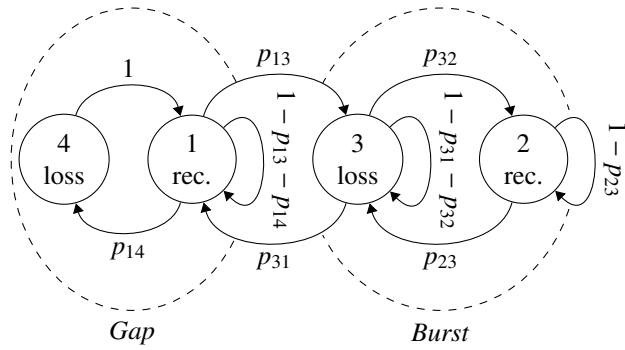


Fig. 1. Loss represented by a 4-state Markov Model (inspired by [44])

Packet loss in wireless links can often occur in bursts, which means that the probability of dropping a given packet or fragment is correlated to previous ones. A well-known approach for modelling this behaviour considers the use of a 4-state Markov chain extending the Gilbert-Elliot model [45]. This Markov chain is depicted by Fig. 1, with two 2-state Markov sub-models representing a *gap* or good period with isolated losses, as well as a *burst*, or bad, period with bursts of losses [44]. Specifically, in State 1 packets are received successfully, while in State 4 an isolated packet is lost, always returning to State 1 (i.e. $p_{41} = 1$). In State 3 packets are lost within a burst, while in State 2 packets are received also in a burst.

In the used loss model, state-transition probabilities p_{ij} can be determined by analysing real link statistics to correctly characterise state probabilities ($\pi_i, i = [1, 4]$). In particular, $p_{13}, p_{31}, p_{23}, p_{32}$ and p_{14} will be independent transition probabilities, while p_{11}, p_{22} and p_{33} are defined as represented in Fig. 1 and $p_{44} = 0$. These will be reflected in the overall mean loss probability $P(D)$, as well as in the average burst length, loss density within a burst, isolated loss probability and mean good burst length [44]. In particular, $P(D)$ corresponds to the combined state probability of State 3 and State 4 (i.e. $P(D) = \pi_3 + \pi_4$), respectively defined by 4 and 5.

$$\pi_3 = \frac{p_{13}p_{23}}{p_{13}p_{23} + p_{23}p_{31} + p_{14}p_{23}p_{31} + p_{13}p_{32}} \quad (4)$$

$$\pi_4 = \frac{p_{14}p_{23}p_{31}}{p_{13}p_{23} + p_{23}p_{31} + p_{14}p_{23}p_{31} + p_{13}p_{32}} \quad (5)$$

$$(6)$$

Knowing the average probability of dropping any given packet in bursty links, the probability of not dropping a packet is given by $P(\bar{D}) = 1 - P(D)$. In addition, considering how CoAP queries are handled, the probability of sending a *GET* and receiving all the fragments f of the respective *ACK*, without dropping any of them, is defined by (7).

$$c(f) = P(\bar{D})^{f+1} \quad (7)$$

Conversely, the probability of dropping at least the *GET* or one of the fragments from the *ACK* in a given query, is defined by (8).

$$\bar{c}(f) = 1 - c(f) \quad (8)$$

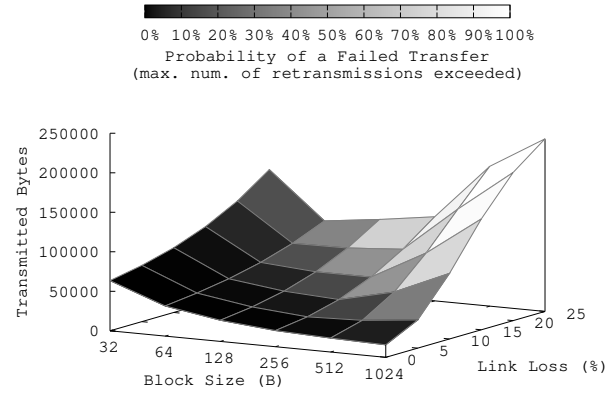


Fig. 2. Overhead for different CoAP Block sizes and link quality (probability of a failed connection represented by the shaded area)

Since confirmable CoAP requests support a maximum number of r retransmissions limited by the maximum allowed latency [12], the probability of a query reaching this limit is $\bar{c}(f)^r$. With this, and taking into account the lowest number of transmitted bytes per request, the worst case scenario regarding the total number of transmitted bytes is given by the sum of the expected overhead for each possible number of retransmission attempts, considering the probability of their occurrence, as defined by (9).

$$w(f, r, T_{best}) = T_{best} + \sum_{i=1}^r \bar{c}(f)^i \times i \times T_{best} \text{ (bytes)} \quad (9)$$

B. Reliability Analysis

A broken connection, or an unsuccessful query completion, implies that at least the *GET* or one of the fragments from the *ACK* are dropped $r + 1$ consecutive times, resulting in a probability of $\bar{c}(f)^{r+1}$ per query. Therefore, the probability of successfully completing a query in less than r retransmissions is defined by (10) and the probability of successfully completing q queries, in less than $r + 1$ consecutive attempts per query, is defined by (11).

$$s(r, f) = 1 - \bar{c}(f)^{r+1} \quad (10)$$

$$s(r, f, q) = s(r, f)^q \quad (11)$$

Finally, (12) corresponds to the probability of a CoAP transfer completely failing, for a given number of q queries, with f fragments and a maximum number of r retransmissions.

$$\bar{s}(r, f, q) = 1 - s(r, f, q) \quad (12)$$

C. Overhead vs. Reliability

By combining the presented equations with different Block sizes and different link-loss probability, namely (9) and (12), it is possible to more clearly analyse the impact of fragmentation regarding both overhead and reliability. Fig. 2 combines these variables into a 4-dimensional graph where the shaded area




Fig. 3. Total transmitted length for different Block sizes in CoAP(s)




Fig. 4. Overhead per CoAP(s) packet

represents the probability of a broken connection in CoAP, for the default maximum number of CoAP retransmissions ($r = 4$). The figure shows that, while a Block size of 512 B or 1024 B performs well for links with less than 10% losses, for higher loss values the overhead becomes larger than for smaller Block sizes. In addition to the increased overhead, represented by the lighter areas, the probability of a broken CoAP transfer rapidly increases for larger Block sizes, resulting in extremely unreliable connections.

V. RESULTS

All the results presented in this section were obtained using a real testbed under controlled conditions, with virtually non-existing losses. This means that for the performed repetitions no significant changes occurred. In particular, data related to fixed parameters, such as the overhead of one request, was verified to be constant in all repetitions. However, when analysing the absolute performance of transferring a file with artificially added bursty losses, variations do occur. Therefore, in Section V-B, such results are presented as an average of three repetitions with error bars corresponding to the standard deviation.

The conducted assessment takes into account two main aspects. The first considers the total amount of transmitted data and the corresponding overhead per transmitted CoAP block, while the second pertains to the total amount of time taken when retrieving a 13 kB image, analysing the effective goodput and comparing alternative radio-access configurations. All the performed experiments used six different Block sizes, which must be a power of 2, showing the impact of different levels of fragmentation in a realistic setting.

A. 6LoWPAN vs. IPv4

The performance comparison between using 6LoWPAN and IPv4 focuses on the overhead of each alternative while also reflecting on their limitations. It is based only on a radio configuration with TDMA using a 1000 ms frame, since the actual overhead of these solutions is independent of the radio access mode.

The presented assessment was done by retrieving HTTP data over VHF while using CoAP and included sending both unencrypted and encrypted data. Encrypted connections, from now on referred to as *CoAPs*, used DTLS with a modification of the default timeout value (i.e. increased limit) in order to support our constrained link. However, when performing the experiments with the IPv4 protocol it failed to establish an encrypted connection. This was due to the complexity of the initial handshake process, which transmits larger packets that require fragmentation in order to meet the radio's internal maximum transmission unit (MTU) of 220 bytes. Conversely, 6LoWPAN and its compliance with the IEEE 802.15.4 127 B PSDU limit handled this seamlessly.

The experimental results of using both the 6LoWPAN and IPv4 approaches are presented in Fig. 3 and 4 respectively named as *6lo* and *IPv4*. These figures depict the total amount of bytes and overhead per CoAP packet with different Block sizes, which were used to determine F_1 and $F_{>1}$. Additionally, the connection failure of *IPv4+CoAPs* is illustrated in these figures by a column with a thicker and dashed border, with its height corresponding to the expected value. This value was determined by calculating the equivalent overhead of *6lo+CoAPs*, compared to *6lo+CoAP*, and by adding it to the value obtained for *IPv4+CoAP*.

When analysing the total amount of transmitted bytes for different Block sizes it becomes clear that 6LoWPAN offers a competitive solution to the standard IPv4. The performance of each solution is similar, with 6LoWPAN being slightly outperformed for Block sizes of 128, 512 and 1024 bytes. Nonetheless, for smaller Block sizes, and for 256 B, 6LoWPAN outperforms IPv4 by larger percent margins. Moreover, 6LoWPAN's adaptation layer was more robust than IPv4 when having to handle more complex transactions, such as the DTLS handshake.

Fig. 4 shows the percentage of overhead per transmitted Block. Only without fragmenting, for a Block size up to 64 B, does the 6LoWPAN version present a lower overhead when compared against IPv4 for 128 B. However, as expected, fragmentation allows a lower overall percentage of overhead for both solutions and 6LoWPAN requires only marginally more

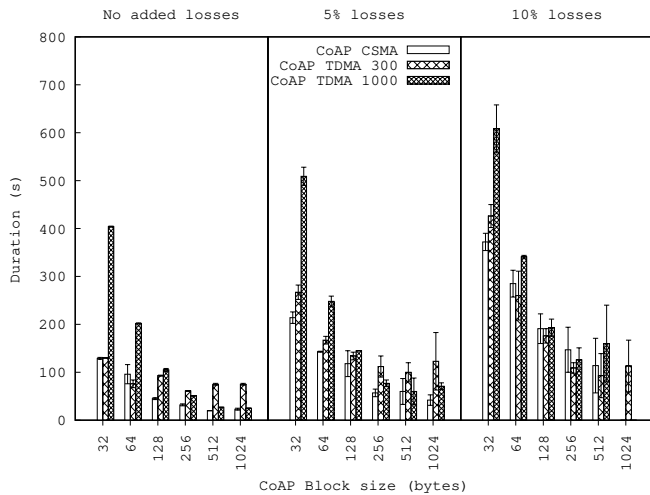


Fig. 5. Total time for different link qualities (13 kB image)




Fig. 6. Goodput for different link qualities

data than IPv4. This suggests that, even though fragmenting should in general be avoided, the option of using larger Block sizes when links exhibit good quality may be worth exploring.

Another noteworthy observation of the obtained results, in particular for the 6LoWPAN solution, is that none of the available Block sizes (power of 2) fully utilises the 127 B from the available PSDU. For a Block size of 64 B, 6LoWPAN frames had a total length of 91 B, corresponding to the payload size plus 9 B from the IEEE 802.15.4 frame header, 6 B from the IPHC and NHC and 12 B from CoAP headers. This means that there were 36 B left available, even though this number could be reduced to as low as 15 B if link-layer security was added⁴, which would take at most 21 B regardless of the chosen Block size. The impact of these unused bytes is only noticeable when there is a fixed amount of “air time”, as with TDMA, where resources end up being wasted.

B. Overall Performance

Considering that 6LoWPAN has an overhead comparable to that of IPv4, and driven by advantages of IPv6-based solutions for IoT, this subsection focuses on the performance of 6LoWPAN. Moreover, bearing in mind that losses due to reduced link-quality are likely to occur in the harsh conditions found offshore, this evaluation must consider networking in such conditions (i.e. LLNs). For this purpose, in addition to the real performance of the VHF link, losses were added into the system by using the *netem qdisc*, which was configured to drop packets by following the 4-state Markov model previously presented (c.f. Section IV-A). Two versions of this model were created with the five necessary independent state-transition probabilities, as presented in Table I. These can represent two different link configurations with a mean burstiness length of 3 packets, one with 5% and another with 10% losses. This approach allowed maintaining a controlled environment and the repeatability of each performed test.

Fig. 5 shows the average duration for retrieving the previously used image, including the registered standard deviation.

⁴This feature is not currently supported by the Linux Kernel and therefore was not tested (c.f. http://wpan.cakelab.org/#_open_tasks)

TABLE I
STATE-TRANSITION PROBABILITIES

	5% Losses	10% Losses
p13	2.1959 %	4.8591 %
p31	66.6(6) %	66.6(6) %
p32	33.3(3) %	33.3(3) %
p23	33.3(3) %	33.3(3) %
p14	1.0101 %	2.0408 %

It shows the total transfer time when using a link in near-perfect conditions, as well as when using links with 5% and 10% added artificial losses. These 5% and 10% losses respectively increased the duration up to 355% and 470%, for a Block size of 512 B with CSMA. Moreover, the variation between each experiment increases not only with the number of losses but also with the Block size, which resulted in failed transfers for larger sizes (i.e. 512 B and 1024 B) in both 5% and 10% loss links. This confirms the drawbacks of using fragmentation and the importance of features such as Block-wise transfers.

The achieved goodput is directly correlated to the total duration of a transfer and was also significantly affected by introducing a small number of losses in the link, as seen in Fig. 6. Even without artificial losses, the negative impact of fragmentation is clearly depicted by the standard deviation of goodput for Block sizes larger than 256 B. On one hand, the overall improvement from using larger Block sizes is undeniable with links in almost perfect conditions, with a speedup of up to 14.8 for TDMA using a frame length of 1000 ms. On the other, in links with around 5% losses, the performance and stability of 1024 B Block sizes is surpassed by smaller ones (e.g. with 128 B). In addition, while all experiments were successful for Block sizes below 256 B, in links with approximately 10% losses, few successful transmissions were completed for Block sizes of 512 B and 1024 B. This performance degradation results from lost fragments to which CoAP is unaware, requiring the repetition of the entire process (i.e. *GET* and *ACKs*).




Fig. 7. Total theoretical and measured overhead

Concerning the correlation between the radio’s access mode and the chosen network parameters, it is also possible to see how different Block sizes can impact performance. Indeed, for longer TDMA frames, smaller Block sizes represent a significant loss of performance, with up to three times longer transfers when compared against CSMA. This results from unused “air-time” between each request and response, while CSMA only uses the strictly necessary time to transmit each message. However, as the Block size increases, longer TDMA frames surpass shorter ones and perform similarly to CSMA, improving the overall performance for near-perfect links. In spite of the apparent advantage of increasing Block sizes for improving the goodput, in links with 10 % losses the variation in performance for all access modes is also increased and CSMA is, on average, outperformed by TDMA.

C. Overhead vs. Reliability

The presented performance analysis reveals a significant degradation of the connection for larger Block sizes whenever losses exist. As discussed in the presented reliability analysis (c.f. Section IV), this degradation is a consequence of unnecessarily retransmitted fragments whenever a loss occurs. This subsection compares the registered overhead with the theoretical number of transmitted bytes required for retrieving the 13 kB image.

Fig. 7 compares the theoretical values (T) for the number of transmitted bytes against the measured results in the performed experiments (M). The theoretical values were calculated from (9), knowing that F_1 and $F_{>1}$ were respectively 91 B and 105 B, measured in the used setup. These values include a link in the 3 previously presented conditions, one where no losses are added, one with 5 % losses and another with 10 % losses.

Even though an increasing Block size is intuitively expected to reduce the number of transmitted bytes since it requires a lower number of exchanged messages, it is possible to see that this is not the case when losses exist. In fact, the instability caused by fragmentation and lost fragments with larger Block sizes, even in a near-perfect link, shows that the measured overhead is higher than the theoretical value,

while for smaller block sizes the results are very similar. The impact of fragmentation is also seen in the stability of the link, leading to unexpected results. For example, in links with 10 % losses, 512 B and 1024 B Block sizes registered a similar overhead to links with 5 % losses. This is a consequence of outliers, since in links with 10 % losses only a few experiments were successfully completed, while in links with 5 % nearly all the experiments were successful, with some reaching the maximum number of CoAP retransmissions, thus having a large overhead. As an example, all the performed experiments with CSMA and TDMA1000 failed when using Block size values of 1024 B with 10 % link losses since the maximum pre-defined number of CoAP retransmissions was always exceeded.

The obtained results confirm that, despite the possible overhead reduction from fragmenting packets, this is not true for non-perfect links, since larger Block sizes increase the overhead as well as the unreliability of communications. This is clearly seen when analysing the total measured overhead for retrieving a small image using Block sizes of 512 B or 1024 B, which even without added losses account for a higher overhead than if using Block sizes of 128 B or 256 B. In addition to increased overhead, the use of larger Block sizes also resulted in unreliable transmissions, with a failure rate of approximately 70 % for the performed experiments.

VI. CONCLUSION

The performance analysis of 6LoWPAN over a VHF link revealed that HTTP-like requests are feasible in a commonly available maritime communication technology, outperforming an existing IPv4 alternative. This motivates the adoption of an IPv6-based approach for the development of the Maritime Internet of Things. Moreover, the proposed solution benefits from the interoperability between technologies and features from already existing protocols, such as CoAP and DTLS.

A theoretical evaluation was provided, demonstrating that a careful selection of parameters, namely CoAP’s Block sizes in Block-wise transfers, strongly affects the overall performance of similar solutions. This assessment concerns not only the total overhead but also the likelihood of successfully completing a transfer.

Through a testbed evaluation, the theoretical results were validated and confirmed that, while fragmentation is generally not desirable, it can substantially improve the performance of good quality links. However, the obtained results also verify that in links subject to losses the benefits of fragmentation can be quickly surpassed by the added unreliability, creating more overhead and instability of connections. This motivates a closer relationship between different networking layers, suggesting that cross-layer optimisations may be fundamental for networking solutions in constrained settings.

REFERENCES

- [1] N. Campbell, *Biology: Concepts & Connections*. Pearson/Benjamin Cummings, 2006.
- [2] L. Elkins, D. Sellers, and W. R. Monach, “The autonomous maritime navigation (amn) project: Field tests, autonomous and cooperative behaviors, data fusion, sensors, and vehicles,” *Journal of Field Robotics*, vol. 27, no. 6, 2010.

- [3] M. Føre, K. Frank, T. Norton, E. Svendsen, J. A. Alfredsen, T. Dempster, H. Eguiraun, W. Watson, A. Stahl, L. M. Sunde, C. Schellewald, K. R. Skjøien, M. O. Alver, and D. Berckmans, "Precision fish farming: A new framework to improve production in aquaculture," *Biosystems Engineering*, 2017.
- [4] G. Xu, W. Shen, and X. Wang, "Marine environment monitoring using wireless sensor networks: A systematic review," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Oct 2014.
- [5] T. A. Johansen, A. Zolich, T. Hansen, and A. J. Sørensen, "Unmanned aerial vehicle as communication relay for autonomous underwater vehicle – field tests," in *IEEE Globecom Workshops*, Dec 2014.
- [6] Y. Kim, J. Kim, Y. Wang, K. Chang, J. W. Park, and Y. Lim, "Application scenarios of nautical ad-hoc network for maritime communications," in *OCEANS 2009*, 2009.
- [7] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 2460, Dec. 1998.
- [8] Y.-G. Hong, C. Gomez, A. R. Sangi, T. Aanstoot, and S. Chakrabarti, "IPv6 over Constrained Node Networks (6lo) Applicability & Use cases," Internet Engineering Task Force, Internet-Draft draft-ietf-6lo-use-cases-03, Oct. 2017, work in Progress.
- [9] "IEEE standard for local and metropolitan area networks—part 15.4: Low-rate wireless personal area networks (lr-wpans)," *IEEE Std 802.15.4-2011 (Revision of IEEE Std 802.15.4-2006)*, Sept 2011.
- [10] G. Montenegro, J. Hui, D. Culler, and N. Kushalnagar, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," RFC 4944, Oct. 2015.
- [11] P. Thubert and J. Hui, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks," RFC 6282, Sep. 2011.
- [12] C. Bormann, K. Hartke, and Z. Shelby, "The Constrained Application Protocol (CoAP)," RFC 7252, Jun. 2014.
- [13] C. A. Kent and J. C. Mogul, "Fragmentation considered harmful," *SIGCOMM Comput. Commun. Rev.*, vol. 25, no. 1, Jan. 1995.
- [14] C. Bormann and Z. Shelby, "Block-Wise Transfers in the Constrained Application Protocol (CoAP)," RFC 7959, Aug. 2016.
- [15] E. Rescorla and N. Modadugu, "Datagram Transport Layer Security Version 1.2," RFC 6347, Jan. 2012.
- [16] C. Albaladejo, P. Sánchez, A. Iborra, F. Soto, J. A. López, and R. Torres, "Wireless sensor networks for oceanographic monitoring: A systematic review," *Sensors*, vol. 10, no. 7, 2010.
- [17] R. Al-Zaidi, J. C. Woods, M. Al-Khalidi, and H. Hu, "Building novel vhf-based wireless sensor networks for the internet of marine things," *IEEE Sensors Journal*, vol. 18, no. 5, March 2018.
- [18] L. Lambrinos, C. Djouvas, and C. Chrysostomou, "Applying delay tolerant networking routing algorithms in maritime communications," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), IEEE 14th International Symposium and Workshops on a*, June 2013.
- [19] L. Torgerson, S. C. Burleigh, H. Weiss, A. J. Hooke, K. Fall, D. V. G. Cerf, K. Scott, and R. C. Durst, "Delay-Tolerant Networking Architecture," RFC 4838, Oct. 2015.
- [20] R. J. Mohsin, J. Woods, and M. Q. Shawkat, "Density and mobility impact on manet routing protocols in a maritime environment," in *Science and Information Conference (SAI)*, 2015, July 2015.
- [21] M. Auzias, Y. Mahéo, and F. Raimbault, "Coap over bp for a delay-tolerant internet of things," in *2015 3rd International Conference on Future Internet of Things and Cloud*, Aug 2015.
- [22] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke, "Data collection, storage, and retrieval with an underwater sensor network," in *In Proceedings of the International Conference on Embedded Networked Sensor Systems (ACM SenSys)*, 2005.
- [23] A. S. Woodget, P. E. Carbonneau, F. Visser, and I. P. Maddock, "Quantifying submerged fluvial topography using hyperspatial resolution uas imagery and structure from motion photogrammetry," *Earth Surface Processes and Landforms*, vol. 40, no. 1, 2015.
- [24] A. Sivertsen, S. Solbø, R. Storvold, A. Tøllefsen, and K.-S. Johansen, "Automatic mapping of sea ice using unmanned aircrafts," in *ReCAMP Flagship Workshop Book of Abstracts*, ReCAMP Flagship Workshop. ReCAMP Flagship Workshop, 2016.
- [25] T. A. Johansen and T. Perez, "Unmanned aerial surveillance system for hazard collision avoidance in autonomous shipping," in *2016 International Conference on Unmanned Aircraft Systems (ICUAS)*, June 2016.
- [26] A. Zolich, J. A. Alfredsen, T. A. Johansen, and K. R. Skjøien, "A communication bridge between underwater sensors and unmanned vehicles using a surface wireless sensor network - design and validation," in *OCEANS 2016 – Shanghai*, April 2016.
- [27] R. Birkeland, "Freely drifting cubesat constellations for improving coverage for arctic sensor networks," in *2017 IEEE International Conference on Communications (ICC)*, May 2017.
- [28] A. G. Guerra, F. Francisco, J. Villate, F. A. Agelet, O. Bertolami, and K. Rajan, "On small satellites for oceanography: A survey," *Acta Astronautica*, vol. 127, pp. 404 – 423, 2016.
- [29] A. A. Allal, K. Mansouri, M. Youssfi, and M. Qbadou, "Toward a new maritime communication system in detroit of gibraltar where conventional and autonomous ships will co-exist," in *International Conference on Wireless Networks and Mobile Communications*, 2017.
- [30] D. Jiang and G. Liu, *An Overview of 5G Requirements*. Cham: Springer International Publishing, 2017, pp. 3–26.
- [31] A. NGMN, "Ngmn 5g white paper," Tech. Rep., March 2015. [Online]. Available: https://www.ngmn.org/fileadmin/ngmn/content/downloads/Technical/2015/NGMN_5G_White_Paper_V1_0.pdf
- [32] A. García-Domínguez, "Mobile applications, cloud and bigdata on ships and shore stations for increased safety on marine traffic; a smart ship project," in *2015 IEEE International Conference on Industrial Technology (ICIT)*, Mar. 2015.
- [33] Y. Xu, "Quality of Service Provisions for Maritime Communications Based on Cellular Networks," *IEEE Access*, vol. 5, 2017.
- [34] H. K. Son and Y. J. Chong, "Analysis of the interference effects from maritime earth station in motion to 5g mobile service," in *2017 International Conference on Information and Communication Technology Convergence (ICTC)*, Oct. 2017.
- [35] M. D. Sanctis, E. Cianca, G. Araniti, I. Bisio, and R. Prasad, "Satellite communications supporting internet of remote things," *IEEE Internet of Things Journal*, vol. 3, no. 1, Feb 2016.
- [36] "IEEE standard for local and metropolitan area networks - part 15.4: Low-rate wireless personal area networks (lr-wpans) - amendment 6: Tv white space between 54 mhz and 862 mhz physical layer," *IEEE Std 802.15.4m-2014 (Amendment to IEEE Std 802.15.4-2011 as amended by IEEE Std 802.15.4e-2012, IEEE Std 802.15.4f-2012, IEEE Std 802.15.4g-2012, IEEE Std 802.15.4j-2013, and IEEE Std 802.15.4k-2013)*, April 2014.
- [37] C. S. Sum, M. T. Zhou, L. Lu, R. Funada, F. Kojima, and H. Harada, "IEEE 802.15.4m: The first low rate wireless personal area networks operating in tv white space," in *2012 18th IEEE International Conference on Networks (ICON)*, Dec 2012.
- [38] J. Vasseur, "Terms Used in Routing for Low-Power and Lossy Networks," RFC 7102, Jan. 2014.
- [39] A. Ludovici, A. Calveras, and J. Casademont, "Forwarding techniques for ip fragmented packets in a real 6lowpan network," *Sensors (Basel, Switzerland)*, vol. 11, no. 1, 2011.
- [40] B. Schütz and N. Aschenbruck, "Adding a network coding extension to CoAP for large resource transfer," in *2016 IEEE 41st Conference on Local Computer Networks (LCN)*, Nov 2016.
- [41] C. Bormann, A. Betzler, C. Gomez, and I. Demirkol, "CoAP Simple Congestion Control/Advanced," Internet Engineering Task Force, Internet-Draft draft-ietf-core-cocoa-03, Feb. 2018, work in Progress.
- [42] D. S. E. Deering and R. M. Hinden, "IP Version 6 Addressing Architecture," RFC 4291, Feb. 2006.
- [43] Wikipedia, "Url shortening — wikipedia, the free encyclopedia," 2017, [Online; accessed 9-November-2017]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=URL_shortening&oldid=807713318
- [44] S. Salsano, F. Ludovici, A. Ordine, and D. Giannuzzi, "Definition of a general and intuitive loss model for packet networks and its implementation in the Netem module in the Linux kernel," University of Rome – "Tor Vergata", Technical Report, Aug. 2012.
- [45] E. N. Gilbert, "Capacity of a burst-noise channel," *The Bell System Technical Journal*, vol. 39, no. 5, Sept 1960.



David Palma is an Associate Professor at the Department of Information Security and Communication Technology, from the Norwegian University of Science and Technology (NTNU). He was an H2020 Marie Skłodowska-Curie Postdoctoral fellow at NTNU and has worked in the past as a Researcher and Project Manager at OneSource, as well as an invited Assistant Professor at the University of Coimbra. He holds a PhD in Information Science and Technology received from the University of Coimbra. His current research interests are on Cognitive IoT, Networking in Remote Areas, Routing, Cloud-Computing and Software-Defined Networks, subjects on which he has authored and co-authored multiple papers in refereed conferences and journals.