A Network Infrastructure for Monitoring Coastal Environments and Study Climate Changes in Marine Systems

Filippo Campagnaro, Nicola Toffolo, Alessandro Pozzebon,
Roberto Francescon, Alberto Barausse, Laura Airoldi, Michele Zorzi

***Abstract*—Climate changes have a tremendous impact on coastal and littoral areas, strongly affected by seaquakes and floods. Moreover, global warming causes a drastic change on the biodiversity of rivers, seas, lakes, including in biodiversity hotspots and protected areas, such as the Venice Lagoon in Italy. A similar impact is caused by pollutants: this called for a large-scale long-term action that aims to monitor aquatic environmental parameters in order to predict, manage and mitigate these effects. Yet, coastal systems are highly heterogeneous in space and variable over short (daily), medium and long (seasonal, inter-annual) timescales, making reliable but affordable monitoring a challenging task. This paper proposes to automate this process with the use of a low-power sustainable integrated underwater and above water Internet of Things sensor network, able to collect water measurements in a cloud database and make them available to researchers to monitor the status of a certain area and develop their predictions models. Simulation results highlight how Low-Power Wide-Area Networks can support the data collection from a dense sensor deployment.**

***Index Terms*—Underwater networks, underwater internet of things, LoRa, monitoring coastal environment.**

I. Introduction

The disastrous impact of climate changes both to coastal areas, strongly affected by seaquakes and floods, and to biodiversity, calls for large-scale long-term actions that aim to monitor, predict and limit these effects. The new European Biodiversity Strategy for 2030, for instance, is a comprehensive, ambitious, long-term plan for protecting nature and reversing the degradation of ecosystems, not only with immediate actions such as the restoration of coastal ecosystems or the creation of consortia to remove waste from coastal areas, but also with the introduction of innovative solutions to monitor water quality parameters and pollutants. Indeed, coastal systems are highly heterogeneous in space and variable over short (daily), medium and long (seasonal, inter-annual) timescales, and a hot topic related to underwater assets is the realization of low-cost, low-power and easy-to-handle equipment to support coastal monitoring and maritime operations with low-cost sensors for retrieving data related to water quality in river and coastal environments, as well as in aquaculture sites. In this scenario, where a dense sensor deployment is needed, the use of sophisticated long-range high-power acoustic modems is prohibitive due to the high cost of these devices. In fact, underwater networks are mainly used in large scale applications, such as monitoring of oil spill, remote control of underwater vehicles, and coastline protection. Their use in Internet of Things (IoT) applications, such as monitoring of aquaculture and bathing sites, is not practical due to the high cost and power consumption of most of the commercial acoustic modems, that are usually designed to achieve long range transmissions [1]–[3] rather then being used in small low-power nodes. In contrast to these legacy applications, in the last ten years both researchers [4]–[7] and modem manufactures [8], [9] have started developing low-cost low-power acoustic modems, in order to fill this technological gap that so far made the use of acoustic modems impractical in dense sensor network deployments. Moreover, the use of long range high power radio antennas to send the data collected by surface nodes to shore should be avoided to minimize the power consumption and extend the battery duration of the nodes. In this context, the use of Low-Power Wide-Area Network (LPWAN) solutions should be investigated.

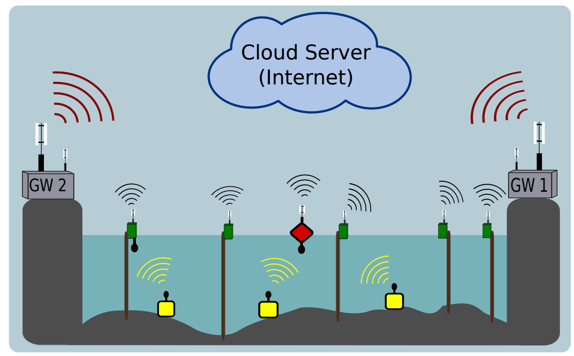

Fig. 1: Network infrastructure.

In this paper we propose a network infrastructure, depicted in Figure 1, for monitoring coastal environments and studying the effects of climate changes on biodiversity. To this aim, in Section II we first analyze the traffic requirements of the sensory systems selected with the help of the staff of the Chioggia Marine Hydrobiological Station of the University of Padova

F. Campagnaro (email: campagn1@dei.unipd.it), N. Toffolo (email: toffolonic@dei.unipd.it), A. Pozzebon (email: alessandro.pozzebon@unipd.it), and M. Zorzi (email: zorzi@dei.unipd.it) are with the Department of Information Engineering, University of Padova, Italy. R. Francescon (email: roberto.francescon@wirelessandmore.it) is with Wireless and More srl, Padova, Italy. A. Barausse (alberto.barausse@unipd.it) and L. Airoldi (laura.airoldi@unipd.it) are with the Department of Biology, University of Padova, Italy.

This work has been partially supported by the European Union - FSE REACT EU, PON Research and Innovation 2014-2020 (DM 1062/2021).

using as a use-case the Venice lagoon, and then review the most promising underwater and above water communication technologies to be applied in this context, proposing the design of three types of sensors. Then, in Section III we present the simulation scenarios and settings, and in Section IV the simulation results, in order to perform a preliminary evaluation of the integrated underwater and above water sensor network before the actual deployment. Finally, Section V concludes the paper.

II. Sensor Overview

In Figure 1 we depict the network infrastructure and the sensor nodes to collect the measurements data. Specifically, we envision the need for two types of surface nodes, a buoy node (depicted in red) equipped with sophisticated sensors to be deployed in a few strategic points of the area of interest, and a very low cost sensor node that can be installed either on the seafloor (the yellow nodes in Figure 1) or in Venetian "briccole" pillars and poles placed in the Venice lagoon (the small green nodes in Figure 1), and equipped with low power devices suitable for a long term dense deployment. While the former will focus on collecting data measurements of water properties that are usually constant in an area of 1 km^2, the latter will collect measurements of parameters that can strongly change within less than a hundred meters (due to morphological heterogeneity or habitat patchiness), for which a dense spatial granularity is required. The former can also be deployed in the lagoon salt marshes, that are a very interesting hot-spot to monitor floods and biodiversity. Both nodes are connected to a wireless gateway (gray node) that forwards the data to a cloud server accessible from the Internet by marine scientists and biologists. While buoys and nodes installed in pillars, from now onward called surface nodes, will transmit their data to the gateway through an LPWAN, the nodes deployed on the seafloor will first transmit their data to a surface node, using underwater communication devices (e.g., a low power acoustic link). The surface node will then act as a relay, forwarding the data to a gateway. Depending on the area to be covered, one or multiple gateways will be employed (for instance, two gateways are depicted in Figure 1).

The low cost sensor nodes that are part of our envisioned deployment should satisfy the system requirements in terms of types of measurements, power consumption and data transmission. The most common measurements they will collect are temperature, pressure (which can be translated to water depth), pH, turbidity, dissolved oxygen (DO) and electrical conductivity (EC). Some of these quantities are subject to rapid major variations in time (e.g., DO), while others are expected to remain more or less constant, or at least to experience very slow and smooth changes in time (e.g., temperature). This means that data measurements may be collected with a different time granularity depending on the type of sensor.

A. Low-Power Wide-Area Networks

Nodes located on the seafloor will transmit the data to surface nodes with underwater acoustic communication, which limits the maximum bitrate significantly, but allows long-range transmissions. Surface nodes, instead, could use standard protocols such as WiFi or GSM/LTE, but this would possibly cause the depletion of the scarce energy available to the nodes and constrain their deployment to locations covered by a WiFi gateway or the cellular network. In exchange, this solution would provide a very high bitrate, but this tradeoff is not convenient, as we are aiming at sending small amounts of data with a minimum waste of energy cost, as our devices run on batteries and we would like not to replace them too frequently, i.e., more than once a year. Thus, we decided to rely on robust protocols which have been developed specifically to handle networks akin to ours, making low-power communication in a wide area their key purpose. The three main competitors in this area are: NB-IoT, Sigfox and LoRaWAN. Their characteristics, as presented in [10], are summarized in Table I. NB-IoT is likely an overkill standard for our purposes;

TABLE I: Comparison of LPWAN technologies

	NB-IoT	**Sigfox**	**LoRaWAN**
Modulation	QPSK	BPSK	CSS
Frequency	Licensed LTE	Unlicensed EU: 868 MHz NA: 915 MHz AS: 433 MHz	Unlicensed EU: 868 MHz NA: 915 MHz AS: 433 MHz
Bandwidth	200 kHz	200 kHz	125 kHz, 250 kHz 500 kHz
Max.data rate	200 kbps	100 : 600 bps	50 kbps
Message/day limit	Unlimited	140(up), 4(down)	Unlimited
Max.payload length	1600 bytes	12 bytes(up) 8 bytes(down)	243 bytes
Range urban/rural	1 km/10 km	10 km/40 km	5 km/20 km
Interference robust	Low	Very good	Very good
Security	LTE encryption	Not available	AES 128B
Adaptive data rate	N/A	N/A	Available
Private networks	No	No	Yes

albeit capable of guaranteeing higher data rates and bigger payload sizes, it constrains the position of the end nodes to the locations of LTE antennas. Moreover, its maximum range is the lowest among the technologies presented in the table, its resilience to interference is low, and it relies on licensed frequency bands. This means that the channel will be less affected by external interference than the ISM radio bands (used by Sigfox and LoRaWAN), but requires either cellular connectivity or a very expensive deployment due to the high cost of licensed spectrum. When comparing Sigfox and LoRaWAN, both technologies can be an appropriate fit for our network. Nevertheless, we feel that Sigfox imposes more rigorous and strict requirements on the data transmission, which may not be high enough according to the number of sensors and amount of data we have to manage. In addition, the hard limit on the maximum number of messages that can be sent per day (both in downlink and in uplink) is tight, and, although it provides the longest communication range among the analyzed technologies, it does not appear to provide relevant benefits over LoRaWAN. On the other hand, LoRaWAN seems to be a great fit for our system. The standard, supported by the LoRa industry alliance [11], is optimal when employed to connect (constrained) nodes in a network where an extended range and low battery consumption are both prime requisites [12]: for this reason we selected LoRaWAN as the best candidate for our LPWAN. It uses a Frequency Shift Chirp modulation [13] to achieve longer communication range than Frequency Shift Keying (FSK), without increasing the power consumption [14], and provides a wide set of customizable parameters which can be different for each end device. Among these, we have the *transmission power*, the *carrier frequency*, seven *spreading factors* (SF, the higher the SF, the higher the range and the packet airtime), the *bandwidth* and the *coding rate* for the Forward Error Correction (FEC) [15]. Moreover, in [16] the authors established a LoRa communication link between a node deployed 1 meter below the water surface of a fresh water swimming pool and an in-land gateway, making it a promizing technology for shallow water sensor deployments. It is quite remarkable that LoRa devices can be separated in three different "Classes" (A, B and C) depending on their energy and transmission requirements; this would be a relevant opportunity for the system we are designing, where there are extremely constrained nodes that run on small batteries (the low cost ones, typically "Class A") coexisting with more powerful devices that still should not waste energy (the buoy nodes, typically "Class B") and nodes which are plugged in and may receive as much energy as they need (the gateways, "Class C" nodes). Thus, from this analysis, we think that LoRa is the best option for the network we are developing, among those we have considered up to now.

B. Traffic Requirements

We envision our network to be deployed in a quite peculiar and challenging environment, the Venice lagoon. The water of the lagoon is brackish and mostly shallow, with numerous salt marshes, and an intense tidal cycle. The characteristics of this unique area make it hard for researchers to perform quantitative analysis of the water parameters with respect to a more stable environment, such as the open sea, where changes are more easily predictable. Currently, the technology used to perform measurements in the lagoon is very sparse, and the data are logged a few times per day in very few locations. Our aim is to provide the researchers with data that have a much finer granularity, both spatially and temporally. This would result in more and better data to work with. On average, a commercial LoRa gateway should be reachable by sensors in a 10-15 km range if in Line of Sight (LoS), and this is compliant with our scenario, where we do not expect to have big or tall objects (buildings, trees, etc.) to prevent the sensors from communicating with the gateway. However, we certainly have sources of interference (such as boats or even ships on major access points to the lagoon), so, in order not to make the Packet Error Rate (PER) too high, we assume that the maximum distance that would allow communication between a gateway and a sensor node is 5 km. In our simulations, we focused on a single area covered by one gateway. Also, the best distance between nodes we came up with is 500 m, which consists of a good trade-off between a high enough data granularity and the economic costs. Of course, in a real scenario, this distance should be tuned with respect to local changes: on the one hand, if we expect water parameters to change frequently and unexpectedly in a small area, we may have to decrease the distance between the sensors; on the other hand, the distance could be increased in areas where parameters vary slowly and more predictably. The same reasoning could be applied to temporal variability, that is, we may want to transmit sensor data more or less frequently according to the specific areas in which sensors are deployed. Also, some data are expected to vary more often than others, so we may want to transmit these more frequently, but not the others. The average time between data transmission from the sensors has been chosen in the range 600 s-1200 s.

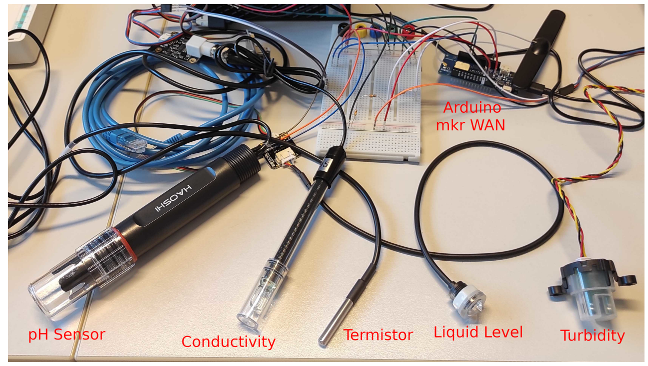

Fig. 2: Sensor in-lab prototype.

C. Sensor prototype

We decided to build a prototype of a surface sensor node according to the previous discussion; the hardware we used is shown in Figure 2. We can see how the prototype is only lab-grade up to now, since it is by no means weatherproof or waterproof, but for the sensors themselves, which are at least partially submersible. The core of the sensor is an Arduino MKR WAN 1310 [17], a microcontroller board with an integrated LoRa chip, chosen for its versatility, well-written documentation and extreme low-energy capabilities (the lowest power consumption is only $104\ \mu$A). The following sensors are attached to the Arduino:

- *pH sensor* [18]: used to measure pH values, needs calibration;
- *Electrical Conductivity sensor* [19]: used to measure EC values and to infer water salinity from these;
- *Turbidity sensor* [20]: used to measure the amount of total suspended solids (TSS) in water;
- *GPS sensor* [21]: used to determine the position of the node in GPS coordinates;
- *Termistor* [22]: used to measure temperature of the water with a ± 0.5°C accuracy;
- *Liquid Level sensor* [23]: used to determine if the sensor is in contact with water or not, possibly for tidal effects.

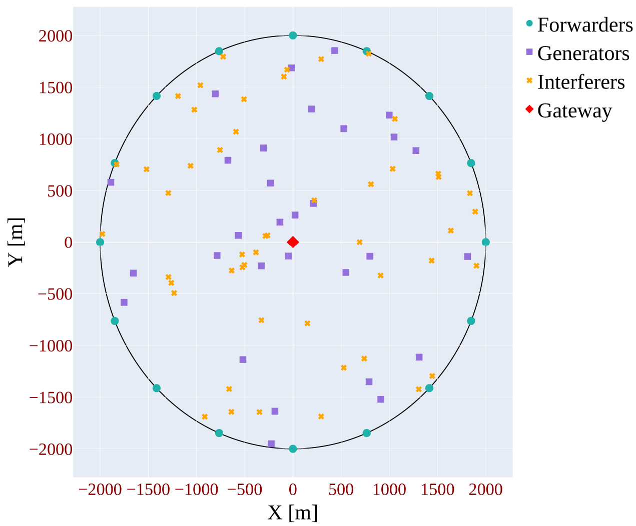

Fig. 3: An example of a challenging deployment.

III. Simulation Scenario and Settings

The simulation of this hybrid underwater and above water network is divided in two parts, using the tool developed in [24]. Specifically, the underwater network part was simulated with the DESERT Underwater framework [25], while the LoRaWAN network with the ns3 [26] module for LoRaWAN: given that the data generated by the underwater nodes are sent to surface nodes that forward it to the cloud server using the LoRaWAN network, the output of the DESERT simulations are saved in a tracefile (containing the sequence of packets received by the surface nodes) that is used as input for the ns3 simulations. Both DESERT Underwater and the LoRaWAN ns3 module are opensource tools publicly available in [27] and [28], respectively, and are developed and maintained by the University of Padova. In particular, we analyzed the performance of the envisioned network by means of a very dense deployment that can be considered quite challenging with respect to the actual scenario, considering a fraction of the Venice lagoon. We distinguished among five different types of nodes, following the requirements seen in Section II. Specifically, nodes can be either:

- *Underwater nodes*: nodes placed on the seabed collecting data from their sensors and sending them to a surface node by acoustic communication;
- *Forwarders*: surface nodes that relay the data received from a group of underwater nodes to the nearest gateway;
- *Generators*: surface nodes that generate data to send them to the nearest gateway and do not receive packet from other nodes;
- *Interferers*: surface nodes whose only role is to transmit dummy data in order to increase the load of the network, simulating the presence of LoRa nodes used for other services close to our deployment;
- *Gateways*: receiving data from all the other nodes. Note that in LoRa there is no handshake procedure between an end device and a gateway, so the former simply broadcasts relevant information, that will be received and processed by the nearest gateway.

All of these nodes communicate by means of LoRa and their Spreading Factor is set individually by the ns3 LoRa module. In Figure 3 we can observe a possible deployment, where 16 forwarders are deployed uniformly in a circle of radius 2000 m, a single gateway is located in the center of this circle and a number of 30 generators and 50 interferers is randomly placed in the area delimited by the circle. Each forwarder is connected to 3 underwater sensor nodes (not shown in the figure for the sake of simplicity), transmitting their measurements exploiting an acoustic communication protocol and deployed uniformly in a small circle whose radius decreases as the number of forwarders increases - in this case, the radius is 200 m.

TABLE II: DESERT simulations parameters

Parameter	Value
UW nodes	3 for every forwarder
Packet Size	30 Bytes
Tx Duration	100000 s
Tx Power	135 dB re 1 μPa @1 m
Frequency	25 kHz
Bandwidth	5 kHz
Bitrate	4800 bps
CBR Period	variable in [30, 960] s range
TDMA Frame	8 s
TDMA Guard Time	0.8 s

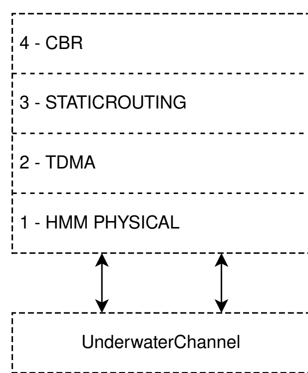

Fig. 4: Stack of the underwater nodes

The underwater nodes transmit with a Constant Bit Rate (CBR) with a Poisson mean of 30 s, 60 s, 120 s, 240 s, 480 s or 960 s. Packet size is 30 bytes, while the transmission rate is equal to 4800 bps. Transmission frequency is 25 kHz and the bandwidth is 5 kHz. The protocol stack of these nodes is shown in Figure 4: they use a CBR application layer, static routing consisting in all the nodes transmitting directly to their 1-hop nearest receiver (i.e., a forwarder node), a Time Division Multiple Access (TDMA) MAC layer and, finally, a Hidden Markov Model (HMM) physical layer, as described in [29]. In particular, according to this physical model, each underwater node is linked to its destination by a channel whose behavior is statistically determined by means of its initial condition and the set of transition probabilities, which specifies the probability that, given the current state, the channel conditions in the next slot will be better, worse, or remain the same. We have set the initial state of all the links to a good (but not excellent) condition, and the transition probabilities are randomly extracted in ranges that would characterize well the evolution of a stable link.

As far as the surface nodes are concerned, we have that the forwarders relay the packets received from the underwater nodes as soon as they are correctly delivered. The generators produce data with a CBR period equal to the one of the underwater nodes and with the same packet size, while the interferers transmit packet of 30 bytes with a CBR period of 30 s. Note that the LoRa overhead is 9 bytes for all of the packets.

TABLE III: ns3 simulations parameters

Parameter	**Value**
Packet Size	30 Bytes
Tx Duration	100250 s
CBR Period (forwarders and generators)	variable in [60, 960] s range
CBR Period (interferers)	30 s

Simulations have been run extensively using SEM [30], the ns3 simulation execution manager, for different combinations of the CBR period, number of forwarders, number of generators and number of interferers. Results are averaged over 20 runs.

IV. Results

In this section we report some relevant results from the simulations. Figure 5 compares the throughput of different numbers of forwarder nodes, deployed as previously discussed, for multiple CBR periods for 4, 15 and 20 forwarder nodes. Neither generators nor interferers are present. The offered traffic (in red) is intended as the traffic generated by the 3 underwater source nodes linked to each forwarder, the underwater throughput ("UW," in yellow) is the throughput computed after the data from the underwater nodes have been received by the forwarder by means of acoustic transmission and the final throughput (in blue) is the actual throughput computed at the final destination - that is, the gateway - achieved by means of LoRa. We can see how the major responsible for the throughput degradation is by far the underwater part, and not the above water one, and this happens in particular for low CBR periods (30 s, 60 s, 120 s) and with a higher number of forwarders. We have that, with 20 forwarders and a CBR period of 30 s, the UW throughput is just 67% of the offered traffic. Luckily, the degradation is mitigated for higher CBR periods, which is compliant with our scenario. An interesting analysis that could be carried on these data concerns the research of the best CBR period such that, in a fixed time, the probability of at least one packet reception by the gateway is maximized. Therefore, if we define p as the ratio between the throughput computed at the gateway and the offered traffic and if we set the fixed time to 960 s,

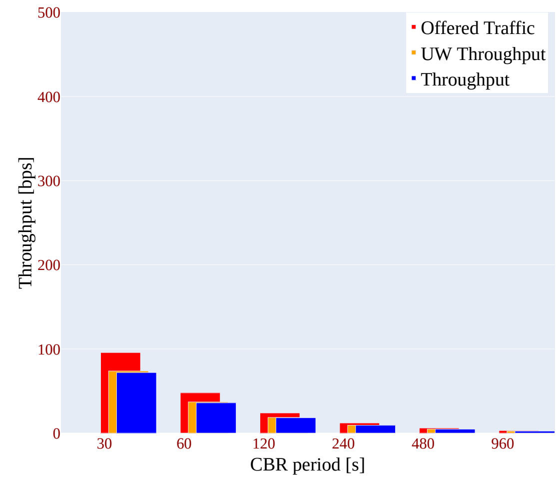

(a) 4 forwarders average throughput

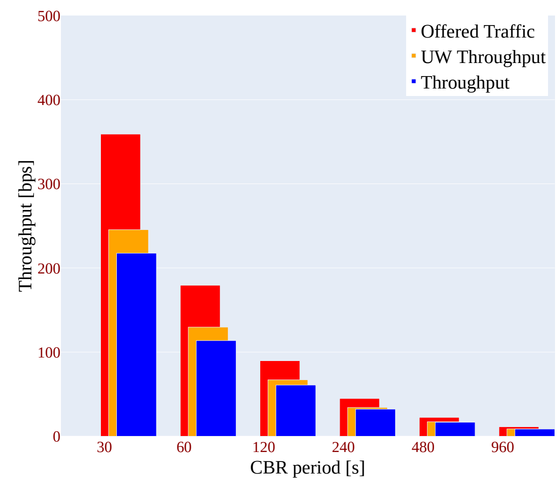

(b) 15 forwarders average throughput

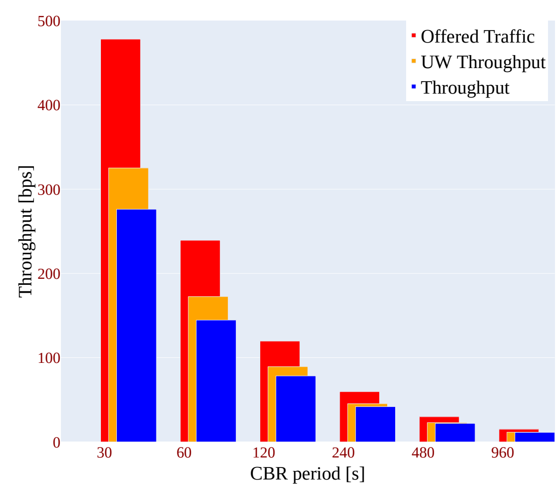

(c) 20 forwarders average throughput

Fig. 5: Forwarder nodes throughput analysis, no generators nor interferers

for a given CBR period (CBR below) we can compute the aforementioned probability as:

$$P(succ) = 1 - (1-p)^r,$$

where $r = \frac{CBR}{960s}$. Tables IV, V and VI report the probabilities computed with six different CBR periods for the deployments of Figure 5. We can see that it is not enough to send a single packet every 960 s in order to obtain an adequate packet reception probability, as it is never higher than 0.78. Conversely, a greater reception rate is achieved if we schedule four subsequent transmissions in this time lapse, that is, a transmission every 240 s. This would allow to receive at least a packet in 960 s with a probability that is never below 0.99, which is a remarkable result: a higher generation rate is then of utmost importance for the forwarders, since the throughput degradation caused by the underwater transmission is the critical aspect of the communications in the network. On the other hand, scheduling multiple transmissions will impact on the energy consumption of the devices, that is another key issue in the envisioned deployment. This means that the ideal time between transmissions will have to be chosen carefully in relation to the specific requirements. Furthermore, we may want to compare the average throughput of a forwarder with the average throughput of a generator. However, we must take into account the fact that, in the simulations, a forwarder is just a relay between the three underwater nodes connected to it and the LoRa gateway. These have the same CBR period of a generator node, and send the same amount of data, thus if we want a fair comparison between the two, a single forwarder's throughput should be compared with three different generators' throughputs. This is why in Figure 6 we report the offered traffic (in purple) and the throughput (in gray) against different CBR periods for 12, 45 and 60 generator nodes. Moreover, to understand the throughput difference between generators and forwarders (Figure 5), we have to account for the average probability of successful packet reception underwater and multiply this probability by the throughput at the gateway: the red markers in Figure 6 present what the generator throughputs would have been if they had to pass through the underwater channel, which are definitely closer to the actual forwarders throughput.

Finally, it is useful to analyze the impact of other potential networks sharing the location of our deployment. To do this, we introduced the interferer nodes, which consist of generator nodes with a fixed CBR period of 30 s whose only purpose is to create noise. Figure 7 shows how the average throughput of both forwarders and generators computed at the gateway is deteriorated by the presence of the interferers with multiple CBR periods. In particular, we consider a network with 10 forwarders and 20 generators; if we focus on the lowest CBR period, we have that a number of 400 interferers reduces the actual throughput by almost 50%, and 800 interferers reduce it to 33%. Similar considerations hold for the other generation periods. Note however, that we have set a very low period for the interferers and that they outnumber the "good" nodes in both Figure 7b and 7c. We should then consider this as a

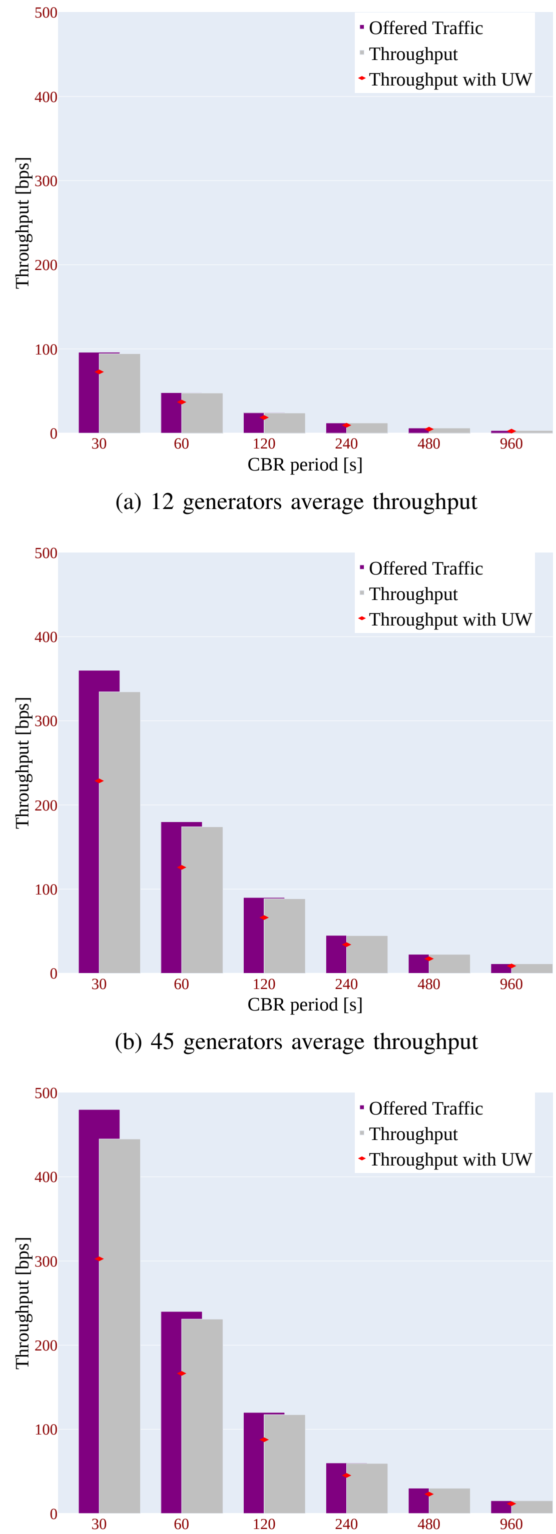

(a) 12 generators average throughput

(b) 45 generators average throughput

(c) 60 generators average throughput

Fig. 6: Generator nodes throughput analysis, 4 forwarders, no interferers

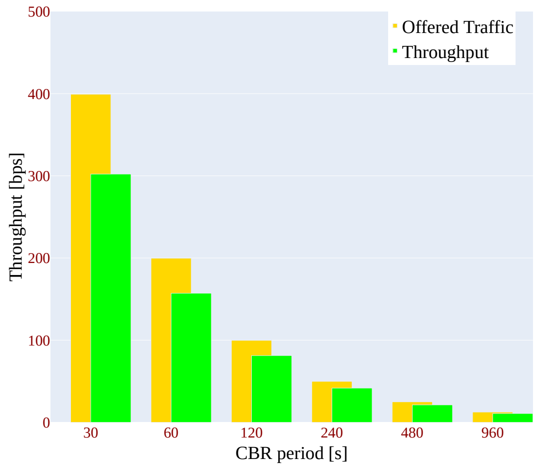

(a) 0 interferers average throughput

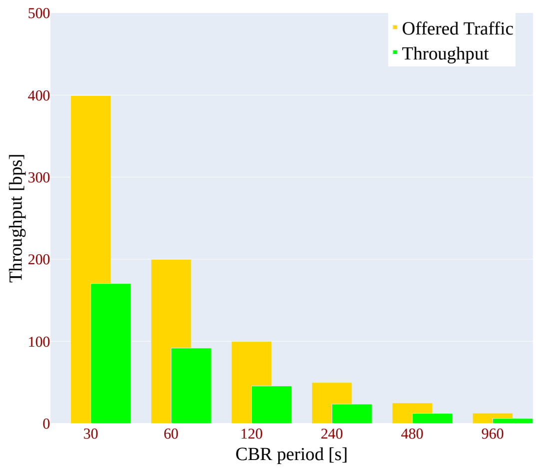

(b) 400 interferers average throughput

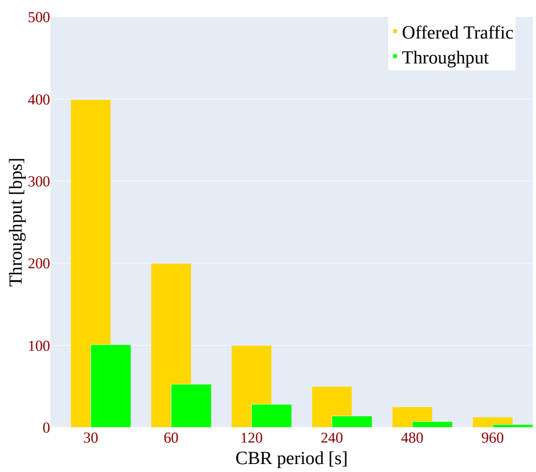

(c) 800 interferers average throughput

Fig. 7: Global throughput degradation analysis, 10 forwarders, 20 generators

worst case scenario.

In Figure 8 we plot the Packet Delivery Ratio for a network of 10 forwarders and 20 generators against the number of interferers for different generation periods. Higher periods guarantee a certain degree of robustness with respect to lower ones, but still the aggressiveness of the interferer nodes makes the PDR fall below acceptable levels quite soon when their number increases, and a number of only 100 interferers is already harmful for communication purposes. In such scenarios, we should take preventive measures, such as increasing the generation rates of the useful nodes or trying to decouple our network from the others as much as possible.

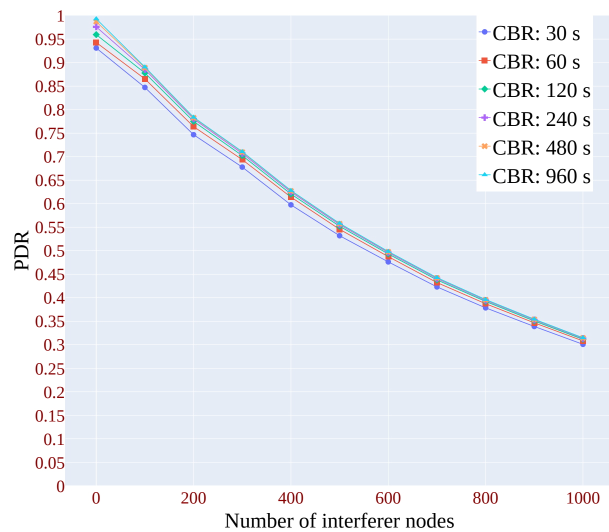

Fig. 8: PDR for different CBR periods and number of interferer nodes, 10 forwarders, 20 generators, interferers CBR Period set to 30 s

TABLE IV: Probabilities of successfully receiving a packet in 960 s, 4 forwarders

CBR	30 s	60 s	120 s	240 s	480 s	960 s
P(succ)	~ 1	0.9999	0.9998	0.9971	0.9491	0.7767

TABLE V: Probabilities of successfully receiving a packet in 960 s, 15 forwarders

CBR	30 s	60 s	120 s	240 s	480 s	960 s
P(succ)	0.9999	0.9999	0.9998	0.9937	0.9341	0.7613

TABLE VI: Probabilities of successfully receiving a packet in 960 s, 20 forwarders

CBR	30 s	60 s	120 s	240 s	480 s	960 s
P(succ)	0.9999	0.9999	0.9998	0.9921	0.9298	0.7546

V. Conclusions and Future Work

Climate changes are threatening coastal areas more and more every day. With this work, we offered an overview of a possible solution to reliably monitor the water quality parameters in these areas, and to improve the granularity of the measurements with respect to preexisting solutions. Our hybrid system consists of both underwater and above water devices, using acoustic and IoT LPWAN technologies, respectively. Traffic requirements have been carefully selected according to a specific area of interest, i.e., the Venice lagoon. Also, a working prototype of a low-cost node has been built. Simulations in DESERT and ns3 simulators have been run extensively to determine the best parameters for the devices in the network and to assess its robustness. Possible future work may include the improvement of the prototype and further testing in real world scenarios as well as simulations with a deployment fit for the Venice lagoon.

References

[1] "EvoLogics Underwater Acoustic Modems," Last time accessed: Aug. 2022. [Online]. Available: https://evologics.de/acoustic-modems

[2] "Develogic Subsea Systems," Last time accessed: Aug. 2022. [Online]. Available: http://www.develogic.de/

[3] "Subnero M25M Series Modems," last time accessed: Apr. 2022. [Online]. Available: https://subnero.com/products/modem.html

[4] B.-C. Renner, J. Heitmann, and F. Steinmetz, "AHOI: Inexpensive, Low-power Communication and Localization for Underwater Sensor Networks and μAUVs," *ACM Transactions on Sensor Networks*, vol. 16, no. 2, Jan. 2020.

[5] B. Benson, Y. Li, B. Faunce, K. Domond, D. Kimball, C. Schurgers, and R. Kastner, "Design of a Low-Cost Underwater Acoustic Modem," *IEEE Embedded Systems Letters*, vol. 2, no. 3, pp. 58–61, May 2010.

[6] B. Sherlock, J. A. Neasham, and C. C. Tsimenidis, "Implementation of a spread-spectrum acoustic modem on an android mobile device," in *Proc. MTS/IEEE OCEANS, Aberdeen, UK*, 2017.

[7] E. Coccolo, F. Campagnaro, D. Tronchin, A. Montanari, R. Francescon, L. Vangelista, and M. Zorzi, "Underwater acoustic modem for a morphing distributed autonomous underwater vehicle (MODA)," in *Proc. MTS/IEEE OCEANS, Chennai, India*, Feb. 2022.

[8] "Popoto modem," http://popotomodem.com/, Last time accessed: Aug. 2022.

[9] "Water Linked Modem M64," https://waterlinked.com/product/modem-m64/, Last time accessed: Aug. 2022.

[10] K. Mekki, E. Bajic, F. Chaxel, and F. Meyer, "Overview of Cellular LPWAN Technologies for IoT Deployment: Sigfox, LoRaWAN, and NB-IoT," in *IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, 2018, pp. 197–202.

[11] "LoRa Alliance - Homepage," https://lora-alliance.org/, Last time accessed: Apr. 2022.

[12] M. Milenkovic, *Internet of Things: Concepts and System Design*. Springer, Cham, May 2021, ch. Communications, pp. 92–94.

[13] L. Vangelista, "Frequency shift chirp modulation: The LoRa modulation," *IEEE Signal Processing Letters*, vol. 24, no. 12, pp. 1818–1821, 2017.

[14] J. M. Marais, R. Malekian, and A. M. Abu-Mahfouz, "LoRa and LoRaWAN testbeds: a review," in *IEEE AFRICON*, 2017, pp. 1496–1501.

[15] M. Bor and U. Roedig, "LoRa Transmission Parameter Selection," in *13th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, 2017, pp. 27–34.

[16] I. Cappelli, A. Fort, M. Mugnaini, S. Parrino, and A. Pozzebon, "Underwater to above water lorawan networking: Theoretical analysis and field tests," *Measurement*, vol. 196, p. 111140, June 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263224122004006

[17] "Arduino MKR WAN 1310," Last time accessed: Aug. 2022. [Online]. Available: https://docs.arduino.cc/hardware/mkr-wan-1310

[18] "Gravity: Analog pH Sensor," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-2069.html

[19] "Gravity: Analog Electrical Conductivity Sensor," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-1797.html

[20] "Gravity: Analog Turbidity Sensor," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-1394.html

[21] "GPS Module with Enclosure," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-1302.html

[22] "Gravity: Waterproof DS18B20 Sensor," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-1354.html

[23] "Gravity: Liquid Level Sensor," Last time accessed: Aug. 2022. [Online]. Available: https://www.dfrobot.com/product-1470.html

[24] D. Magrin, A. Signori, D. Tronchin, F. Campagnaro, and M. Zorzi, "Collaboration of LoRaWAN and Underwater Acoustic Communications in Sensor Data Collection Applications," in *Proc. MTS/IEEE Oceans*, Global Oceans (Virtual), Oct. 2020.

[25] F. Campagnaro, R. Francescon, F. Guerra, F. Favaro, P. Casari, R. Diamant, and M. Zorzi, "The DESERT underwater framework v2: Improved capabilities and extension tools," in *Proc. Ucomms*, Lerici, Italy, Sep. 2016.

[26] "ns3 network simulator," https://www.nsnam.org/, Last time accessed: Aug. 2022.

[27] "DESERT Underwater web site." [Online]. Available: http://nautilus.dei.unipd.it/desert-underwater

[28] "LoRaWAN ns-3 module," https://github.com/signetlabdei/lorawan, Last time accessed: Aug. 2022.

[29] F. Campagnaro, N. Toffolo, and M. Zorzi, "Modeling acoustic channel variability in underwater network simulators from real field experiment data," *Electronics*, vol. 11, no. 14, 2022. [Online]. Available: https://www.mdpi.com/2079-9292/11/14/2262

[30] D. Magrin, D. Zhou, and M. Zorzi, "A simulation execution manager for ns-3: Encouraging reproducibility and simplifying statistical analysis of ns-3 simulations," in *Proceedings of the 22nd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWIM '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 121–125. [Online]. Available: https://doi.org/10.1145/3345768.3355942