IEEE Access
Multidisciplinary : Rapid Review : Open Access Journal

Received 23 June 2022, accepted 11 August 2022, date of publication 16 August 2022, date of current version 22 August 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3198961

RESEARCH ARTICLE

Toward Widespread Remote Laboratories: Evaluating the Effectiveness of a Replication-Based Architecture for Real-World Multiinstitutional Usage

AITOR VILLAR-MARTÍNEZ[1], JAVIER GARCÍA-ZUBÍA[2], (Senior Member, IEEE), IGNACIO ANGULO[2], AND LUIS RODRÍGUEZ-GIL[1], (Senior Member, IEEE)

[1]LabsLand, 48002 Bilbao, Spain
[2]Faculty of Engineering, University of Deusto, 48007 Bilbao, Spain

Corresponding author: Aitor Villar-Martínez (aitor@labsland.com)

ABSTRACT Educational remote laboratories, when properly designed, have been demonstrated to be highly effective from a pedagogical point of view. Throughout the years their technology has evolved and they offer great capabilities, that are sometimes even superior to hands on labs. Nonetheless, despite their potential, they are not yet a widespread tool across learning institutions. One of the main reasons is that many remote laboratories have traditionally been research-oriented and have not been able to guarantee a high-enough quality of service (QoS) in a real-world educational environment. Such a QoS requires a relatively high number of students being able to access the labs concurrently, since classes may include dozens of students. Likewise, such a QoS requires reliability: the laboratory must work, be available and provide correct results. This work evaluates whether by applying an architecture oriented towards cost-effective instance replication and a model oriented towards fault-detection, it is possible to create laboratories that can provide a high QoS and that can therefore be used in a real environment, across multiple institutions. The study encompasses a period of 736 days and over 72,377 laboratory sessions, and relies on real data, from multiple institutions, of professors and students using the LabsLand Intel DE1-SoC FPGA remote laboratory, of which many instances are deployed. The results show that the QoS did indeed meet very high standards, and that such an approach can indeed lead to trusted remote laboratories, appropriate for real-world educational usage, and their eventual widespread adoption.

INDEX TERMS Remote laboratory, quality of service, availability, reliability, STEM.

I. INTRODUCTION

The effectiveness of remote and virtual laboratories has been repeatedly verified from a pedagogical point of view [1], [2], [3], [4], [5]. It has been shown that, when properly designed and applied, remote and virtual laboratories can offer similar or even superior results than conventional laboratories. Furthermore, they offer a wide range of possibilities to students, such as greater flexibility in their scheduling or the elimination of geographical barriers [6], [7]. This has led to an increase in the number of enrollments in degrees that offer distance learning methodologies [8], [9].

As shown in Figure 1, laboratories can be classified depending on their location and nature. Each type of remote laboratory has different advantages. If we consider their adoption rate, the adoption of virtual and remote laboratories at different educational levels has been very uneven. One significant advantage that remote laboratories have over virtual ones is that they offer greater realism [10]: they provide access to real hardware and equipment and allow users to have real interactions and obtain real results. Nonetheless, so far the adoption of virtual laboratories in formal courses

The associate editor coordinating the review of this manuscript and approving it for publication was Md. Abdur Razzaque.

is significantly more common than remote laboratories [11]. Despite their demonstrated effectiveness and potential, they are not currently a widespread, ubiquitous tool leveraged by instructors.

Remote experimentation has advanced significantly in recent years and has become one of the main lines of research in the field of Technology Enhanced Learning [12]. Throughout these efforts, many experimental remote laboratories have been created and deployed. The experimental nature of those remote laboratories resulting from research projects might, however, have been one of the barriers to the widespread adoption of remote experimentation [13].

Remote laboratories are often designed by research groups to meet the specific needs of their research efforts, or the needs of their research centre or university. Often, the goals of the laboratory designer are not clearly aligned with the needs of the instructor that would use the laboratory in class, whose role requires being at the center of the pedagogical process [14]. Additionally, although many remote laboratories are open for use by the educational community, they often lack funding or permanent maintenance staff. As a result, repair times in the event of a breakdown or malfunction can be extended, and not necessarily guaranteed. All of this has a negative impact on the reliability and the perception of remote laboratories, especially by third parties. It discourages adoption by third party professors, who often perceive that they do not have control over the experiment; and that they cannot dedicate the effort to adapt materials and commit to their usage if they cannot be reasonably certain that they will be able to rely on the resource. Furthermore, many remote labs lack the scalability required for consumption by large domains of students. Since, as discussed, many remote laboratories are created as one-off research prototypes and they provide control for a time to a student, they are limited at that respect. Software-based technical solutions may be implemented such as queues or calendar booking to alleviate the issue, but it can still lead to long waiting times and problems with lab reservation. These constraints have led to negative experiences in the educational community, creating an invisible barrier against the widespread adoption of remote experimentation.

Nonetheless, it has been shown that remote labs can be as educationally effective or even superior to hands-on labs [3], [4], [5], [17]. If the previously described practical and technical challenges were overcome, they could truly become a widespread and effective tool leveraged by instructors worldwide to improve technical and scientific education. In this line, previous works proposed an architecture for the creation of cost-effective scalable laboratories to support multiple-users through cost-effective replication [18]. Additional previous works focused on ensuring the reliability and Quality of Service (QoS) of remote laboratories through an automated fault-detection model [19].

The main goal of this paper is to evaluate whether by leveraging those contributions, it is possible to obtain remote laboratories that can be truly effective not only as research

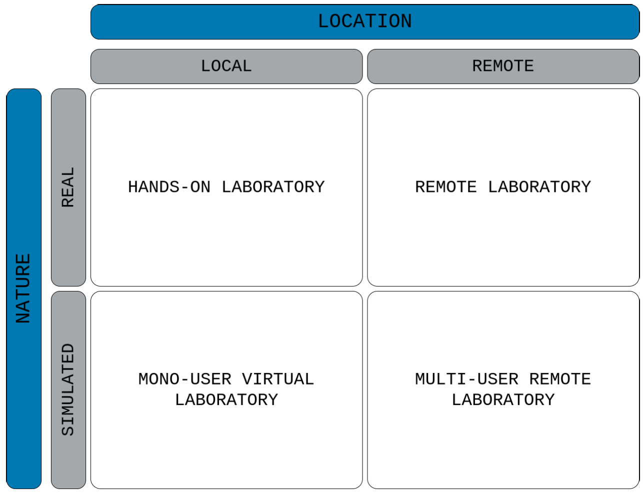

FIGURE 1. Characterization of the different types of laboratories, according to their physical location and nature. Adapted from [15], [16].

prototypes but in real-world educational operational environments, thus eventually leading to their widespread adoption. To evaluate this, the aforementioned replication-oriented architecture and the fault-detection model were applied in the creation and deployment of various remote laboratories, that will be described. All those laboratories are replicated, have multiple copies, and are deployed in a minimum of two institutions around the world, in different countries, relying on a *federated load balancing* architecture as defined in [20]. Since the goal is to evaluate whether widespread adoption for real-world multi-institutional usage is possible, the study will have the following characteristics:

- It will focus on a remote laboratory (the LabsLand FPGA lab) that is deployed in several institutions around the world and with dozens of deployed instances, and rely on data spanning more than two years (736 days).
- It will involve students from multiple universities in multiple different countries.
- The professors using the laboratories with their students will not be directly associated with the authors of this study or with the designers of the remote laboratories.

The remaining of this paper is arranged as follows: Section II describes the state of the art regarding modern remote labs, their current limitations in regards to widespread adoption, and provides detail regarding WebLab-Deusto and LabsLand. It also details the motivation and main contributions of this work. Section III summarizes the two main works and contributions upon which the remote laboratories involved in this paper are based. First is the replication-oriented architecture, designed to provide cost-effective reliability and scalability. Second is the fault-detection model, designed to guarantee Quality of Service. In addition, this section describes in some detail the remote laboratories that apply those two contributions, and upon which the analysis in this work is based. Section IV describes the criteria and methods used to analyze the solutions for developing scalable and reliable remote laboratories. Section V collects and discusses the results of the analysis, which are focused mainly on the quantitative characterization of LabsLand's DE1-SoC FPGA remote lab. Section VI summarizes the conclusions of this paper and Section VII describes future lines of work that could be pursued.

II. MOTIVATION

A. REMOTE LABORATORIES

In recent years, the number of remote laboratories accessible from the Internet has been growing [17]. From a general point of view, remote laboratories can bring many advantages such as cost savings, democratization of technology or the elimination of physical and temporal barriers [1], [2], [10], [21], [22], [23]. Multiple institutions have developed these labs in order to provide students and teachers with the ability to teach, learn and interact with these devices remotely.

An example is RemLabNet,[1] who has developed its own Remote Laboratory Management System [24], [25], and hosts multiple remote laboratories, allowing users to experiment in fields such as physics, chemistry or environmental sciences.

Similarly, iSES[2] has developed its own online teaching system [26], which hosts a number of remote laboratories for very specific purposes in physics. Among them, it is possible to verify the Heisenberg uncertainty principle, or to perform the Franck-Hertz experiment, and also, perform experimentation with microcontrollers [27].

The Grid of Online Laboratory Devices Ilmenau, or GOLDi,[3] gives access to a remote laboratory with different configuration options. Here it is possible to select different control units (a CPLD, an FPGA, a microcontroller, a FSM, etc.) and different physical systems, such as an elevator, a production cell, or a storage warehouse. This allows the user to have a great deal of flexibility in the way users perform their experiments.

The Instituto Superior Técnico (IST) of the University of Lisbon has created its own remote laboratory project [28], known as *e-lab*.[4] It is a web portal that gives access to different remote physics and chemistry [29] laboratories.

The Universidad Federal de Santa Catarina in Brazil also has its own remote laboratory project [30], known as RELLE.[5] In it we can find various labs, such as labs for experimenting with alternating current and direct current electrical panels, or engineering-oriented labs, such as programming microcontrollers based on Arduino boards [31], [32].

Another significant project is LaboREM, developed at the Technological University Institute (IUT) in Bayonne, France. This is a remote laboratory [33], [34] that allows remote experimentation with electronic components. Through a camera and a robotic arm, the user can insert different PCBs into different slots to create a range of electronic circuits, which can later be powered and characterized.

The remote laboratories described above are not meant to be an exhaustive list, but they are some of the research works that have advanced the remote laboratories state of the art very significantly. All of them present real-time remote labs, and thus provide real-time remote control of the target hardware. As real-time laboratories however, they are likely relatively expensive to develop and maintain [19], [35]. Also, most of the described laboratories support a single user at the same time. Though this barrier is overcome in certain cases through replication [18], [36] or laboratory-specific means [37].

To overcome some of those limitations and in response to the problem of user concurrency, LabsLand[6] has developed its own subset of ultra-concurrent virtual laboratories [38], [39], [40], [41], [42] in collaboration with other universities and organizations. Ultra-concurrent remote laboratories are based on a large pre-recorded dataset formed by videos and data of a real experience. Through that dataset, an ultra-concurrent laboratory can provide a laboratory experience that can closely resemble the one that the student would have with the traditional hands-on version of the lab, or with the real-time remote laboratory one. Ultra-concurrent laboratories are not simply videos; they are interactive: they allow the student to take an active role and make choices. To allow this, the dataset includes experiences that can encompass hundreds or thousands of different experimental combinations. They are not simulations either, since they are based on real data recorded with real equipment and do not rely on a simulation engine. In these labs, the user can experiment with a set of data, results, images and/or videos previously recorded in a real lab. This solution allows the user to experiment with real results, while allowing a large number of users to access the lab concurrently, since no finite equipment is handled in real time, but requests are made to a database that hosts a dataset of the lab with the set of pre-recorded samples. The advantage of these labs is that they can be accessed concurrently by hundreds or thousands of students simultaneously without needing to replicate the physical equipment. Also, that they require only software maintenance, and are thus more cost-effective. The potential limitation of these labs is that they are only suitable for experiments in which the range of variables to control is constrained. For certain chemistry practices, for example, in which the experiment is performed according to a laboratory script and there are few variations, this limitation is not significant. However, for other types of laboratories, such as those for programming, those that give students access to various development boards, or in general, those with many interaction options, make this approach impossible or ineffective. For those cases, a real-time lab and thus a more complex replication-based approach such as the one proposed in this work is required.

[1] http://www.remlabnet.eu
[2] https://www.ises.info/index.php/en/laboratory
[3] https://www.goldi-labs.net/
[4] https://www.e-lab.ist.utl.pt/
[5] https://rexlab.ufsc.br/
[6] https://labsland.com

A similar approach is the one proposed by Stanford University[7] and used for some of the online laboratories of its remote laboratory system [43], [44]. For each remote lab in the platform, they have created a digital twin. The experimental results obtained in these experiments are real, though the experimentation itself does not occur in real time but is instead recorded. This does involve a longer development time than the real-time remote lab on its own would require, since it is necessary to digitize a battery of experiments and their results. Once the process is done, the laboratory is highly scalable, without needing to replicate the hardware. Similarly to LabsLand's ultraconcurrent laboratories, this approach is only effective for those laboratories in which the potential variables and user interaction are limited, thus requiring a reasonably-sized dataset and battery of experiments.

Another similar approach is VirtualRemoteLab,[8] developed at the Faculty of Physics of the Ludwig Maximilian University of Munich, Germany. This laboratory [45] allows experiments in optical spectrometry to be performed remotely. Since the laboratory equipment is expensive and precise calibration is necessary, they have chosen to provide the user with two types of experimental setups: a real one, and a virtual one (similar to the ultra-concurrent approach). The real one allows a real experimentation, controlling the equipment in real-time, while the virtual setup allows experimentation with pre-recorded contents of the laboratory itself.

Apart from the aforementioned concurrency challenge, significant research efforts have been dedicated towards developing architectures that can be leveraged to create different remote laboratories more effectively [46], [47], and some of those architectures are in fact the basis of previously described remote laboratories [48], [49], [50], [51].

However, despite the existence of all the tools described above, the usage of remote laboratories as a substitute or complementary tool in classrooms is not yet widespread. This might be at least partially due to practical limitations that still remain [13], [14], [52], [53], and which should be resolved or attenuated to obtain greater use of remote laboratories.

B. WEBLAB-DEUSTO REMOTE LAB MANAGEMENT SYSTEM

WebLab-Deusto [54] is a Remote Laboratory Management System (RLMS) created by the research group of the same name[9] to facilitate the creation of remote laboratories.

If a remote laboratory were to be created from scratch, it would not only be necessary to develop the specific components for each copy of the laboratory itself (such as the hardware control, the programming system of the microcontroller or electronic device, the real-time camera system or the web client), but it would also be necessary to develop a series of components that are common to many different remote laboratories, regardless of their nature, such as:

- User management and authentication
- User group management
- User authorization
- Queue and user flow management
- Integration of labs into Learning Management Systems
- Lab sharing between different entities through federation
- Data extraction and learning analytics

Remote Laboratory Management Systems provide most of those common components, along with other tools to facilitate the development of laboratories. That way, remote laboratory developers only need to focus on the components that are specific to their remote lab, and can save thousands of development hours. As a consequence, the resulting laboratory is typically better and more reliable, since that way developers can more efficiently use their time, and most of the components are already well-tested in previous laboratories.

Remote laboratories created under the scalability-focused architecture described in subsequent sections base one of their layers on the WebLab-Deusto RLMS [54] as the laboratory control and management system. They also rely on WeblabLib [55]. That is an Open Source Python-based library, part of WebLab-Deusto, which is specifically intended to facilitate the development of remote labs and their integration with WebLab-Deusto.

Various remote laboratory initiatives around the world have used WebLab-Deusto for their projects. This FPGA lab [56], [57], [58], created by the University and District Library Bonn-Rhein-Sieg or this physics/electronics lab [59] created by the Faculty of Exact Sciences, Engineering and Surveying National University of Rosario, are some examples. There are several other alternative remote laboratory management systems, which share some of the features, such as [24], [25], [60], [61], [62], [63], [64], [65], [66].

C. LABSLAND

LabsLand[10] is an EdTech company that provides remote laboratories and it is also the name of its remote labs platform [67], [68]. LabsLand provides products and services such as subscriptions to its large global network of remote laboratories, sale of remote laboratory equipment that clients can deploy in their own institutions to be integrated into the platform, and development of new remote laboratories. Those laboratories are often developed in collaboration with universities. LabsLand has developed various technologies oriented towards facilitating the development of remote laboratories and particularly towards ensuring that they are both replicable and reliable. The goals of this contribution are aligned with this perspective. The remote laboratories that are analyzed in this study are part of the LabsLand network. LabsLand has provided the usage data for this study.

D. CHALLENGE AND CONTRIBUTIONS

As described above, remote labs can be as educationally effective as hands-on labs [3], [4], [17], as long as they maintain a certain QoS.

[7]https://stanford.ilabsolutions.com/
[8]http://virtualremotelab.net
[9]https://weblab.deusto.es
[10]https://labsland.com

When remote labs are down, have connection problems, are not well maintained or are unable to support a large number of users, it is not possible to maintain adequate QoS, which can lead users and instructors into giving up remote laboratories for other alternatives such as virtual labs or traditional hands-on labs, even if they have disadvantages for their use cases. Studies with instructors [52] remarked that reliability was the fundamental characteristic expected in a remote laboratory. Also, Salzmann and Gillet [69] state that

> *Robustness toward hardware faults or unavailability is also a key issue for the acceptability of the remote experimentation paradigm by the students. If at connection time they are not able to access the chosen experiment, they may lose motivation and interest.*

and that

> *Remote laboratories maintenance is a difficult and time-consuming task when a 24/7 availability is targeted. The first step in providing a wide availability is to detect problems; this implies that the physical equipment and its associated software are capable of self-diagnoses.*

Traditionally, remote laboratories have been derived from research prototypes. As Figure 2 shows, the fundamental goal of these prototypes is for a device to be remotely controllable, and nothing more. Issues such as QoS or scalability are not considered during their design process.

Some of these proof-of-concepts evolve into higher-quality labs, which may or may not have multiple instances, and achieve greater stability. The goal of these labs is not only to enable remote control, but also to make the lab useful to potential users. These labs are often used by professors and students of the institution (university, community college, or other) that has developed and maintains them. However, in these cases both QoS and scalability are limited, and make them unsuitable for widespread use. Saenz *et al.* [53] have evaluated these issues in greater detail in their study, and state that

> *Due to the fact that there are no clear road maps or recommendations concerning the best selection of technologies and that different standardization attempts coexist, the development of online labs has become very challenging.*

In previous works, we proposed two complementary solutions that address these issues. On the one hand, [18] reports the creation of an architecture for the creation of remote laboratories (for experimentation with microcontrollers, FPGAs and other electronic components) focused on high scalability through replication, designed to lower costs and to share hardware components among different laboratory instances. On the other hand, [19] reports the creation of a model based on fault-detection to promote remote laboratory reliability. It uses a multi-layer solution to evaluate the performance of a laboratory, detect failures, solve them if feasible disconnect the failed instance otherwise, and enable it again when it is repaired and works again. All of this is done automatically and is oriented towards ensuring QoS even if it means disconnecting instances. This is important since a) It is better to disconnect an instance and prevent access, than for the student to waste time and get frustrated getting invalid results b) The fault-detection model is intended to complement a replication-based architecture, in which disconnecting a faulty instance implies simply that the load is redirected to a working one.

These solutions, as shown in Figure 2, allow the creation of remote laboratories with the capacity to support widespread use. Full-scale labs have higher QoS and QoE, are easily scalable and can be used in a production environment by hundreds or thousands of users concurrently. They employ techniques such as multiple instances, hardware resource sharing, automatic failure detection and interactive instance control based on number of active and failed instances.

To this end, the contributions of this work are:

1) Description of a set of remote laboratories created under the set of technological solutions previously introduced.
2) Quantitative analysis performed on an FPGA remote lab, which is based on the previously developed high scalability architecture and high reliability model.
3) Conclusions, based on the previous quantitative analysis, taking into account mainly the availability and quality of service provided by the laboratory from the point of view of the potential user.

It is noteworthy that though as mentioned this work leverages a cost-effective replication-oriented architecture and a fault-detection model, and that both are summarized in more detail in Section III, they are not contributions of this study and are previously published works.

III. COMPONENTS OF THE FUNDAMENTAL ARCHITECTURE

As explained in previous sections, the objective of this paper is to promote the use of remote laboratories in real educational environments to effectively meet the challenges of practical STEM education.

To this end, the main contribution of this paper will be to confirm that by employing a certain model, developed in previous works, and which will be explained below, it is possible to meet the practical educational requirements of real-world institutions, as discussed earlier.

Since the objective is precisely to confirm it in real educational environments, the analysis will involve several remote laboratories implementing the model in question, as well as several international institutions, using them in a real way, by instructors not directly linked to the authors of this paper.

The model on which the laboratories used for the analysis are based is developed, technically evaluated and detailed in previous works [18], [19]. However, we summarize the model here as well, especially those aspects that are most relevant for understanding and conveying the analysis. It should be noted

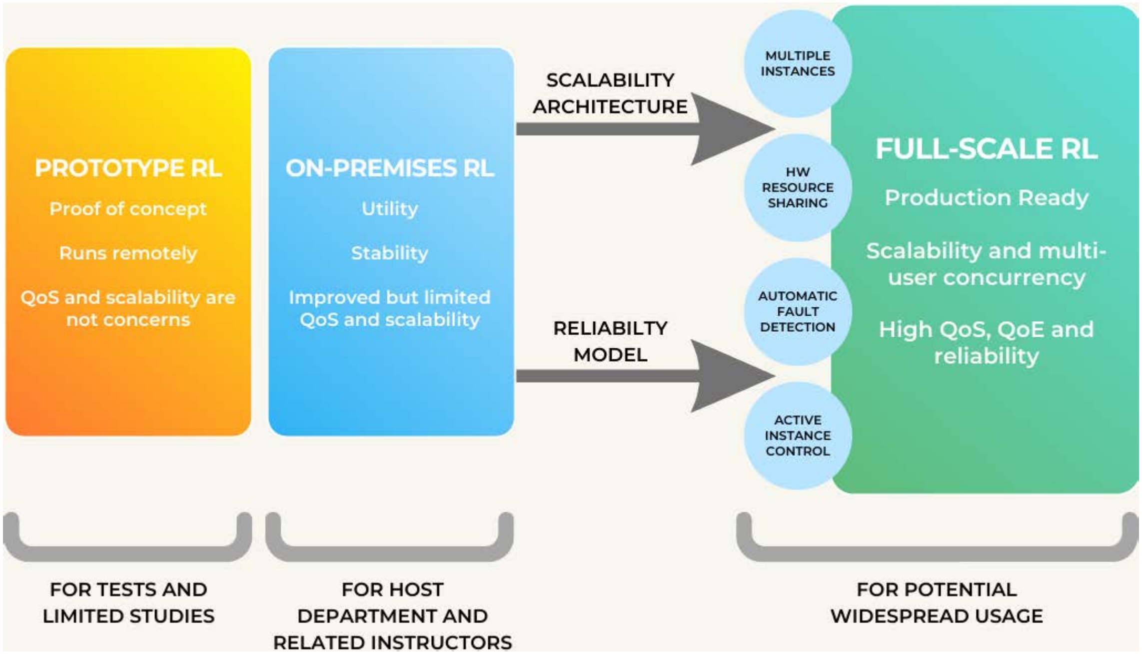

FIGURE 2. Categorization and evolution of remote laboratories, from prototype versions to laboratories capable of widespread use. The scalability architecture and reliability model are key components that enable the creation of remote labs with the potential for widespread use.

that this is not the main contribution of this paper, and that describing them technically in detail is beyond its scope.

The following subsections summarize, first, the two fundamental aspects of the model used:

- Replication: Replicable, and for this purpose distributed, scalable, and cost-effective laboratories.
- Reliability: Reliable laboratories, with the ability to automatically detect faults without exposing them to the user, repairable, and able to guarantee a good Quality of Service (QoS).

Finally, in the last subsection, we will briefly describe the various remote laboratories that implement this architecture and these principles, which will serve as the basis for the multi-institutional analysis, which is the main contribution of this work.

A. ARCHITECTURE FOR REMOTE LABS REPLICATION

One of the works previously conducted by the authors includes the creation of an architecture for the development of remote laboratories [18]. The architecture allows the creation of highly scalable, multi-instance remote laboratories, that enables remote experimentation for multiple users concurrently.

The architecture also contemplates the ease of adaptation to different remotizable objects and allows software and hardware sharing, which maintains high cost efficiency while multiplying the capacity of concurrent users.

The architecture is divided into four layers, each of which has a specific role. Each layer contains both physical and logical components, which can often be supported by low-cost single-board computers (such as Raspberry Pi, for example). Communication between layers is done via common and well established connection standards, such as Ethernet, Internet, USB, SPI and others, allowing easy synchronization and control of the components involved, as well as fast and modular lab development and deployment.

Figure 3 shows an overview of the layers of the architecture, with all the physical and logical components involved. Thanks to its modularity, it is adaptable to a large number of remotizable objects or devices under test, and allows the creation of different deployment topologies, depending on the location and number of components in each layer.

For those laboratories where visual feedback is necessary, an interactive live-streaming platform [70] is deployed in a parallel way to the other four layers. This component, which is also formed by various layers to ensure scalability and reliability, is in charge of providing the user of the remote laboratory with a real time stream, which is captured on-site by a camera.

This architecture is applied for the creation of several different remote laboratories. They use different lower-level topologies and they provide control over various different target devices. Those include FPGAs, microcontroller development boards and robots, among others.

The architecture and the lower-level topologies are designed with replicability in mind. For example, in many of these laboratories a single hardware setup includes several target devices to be controlled as separated instances; and certain components are nonetheless shared. This way it is possible:

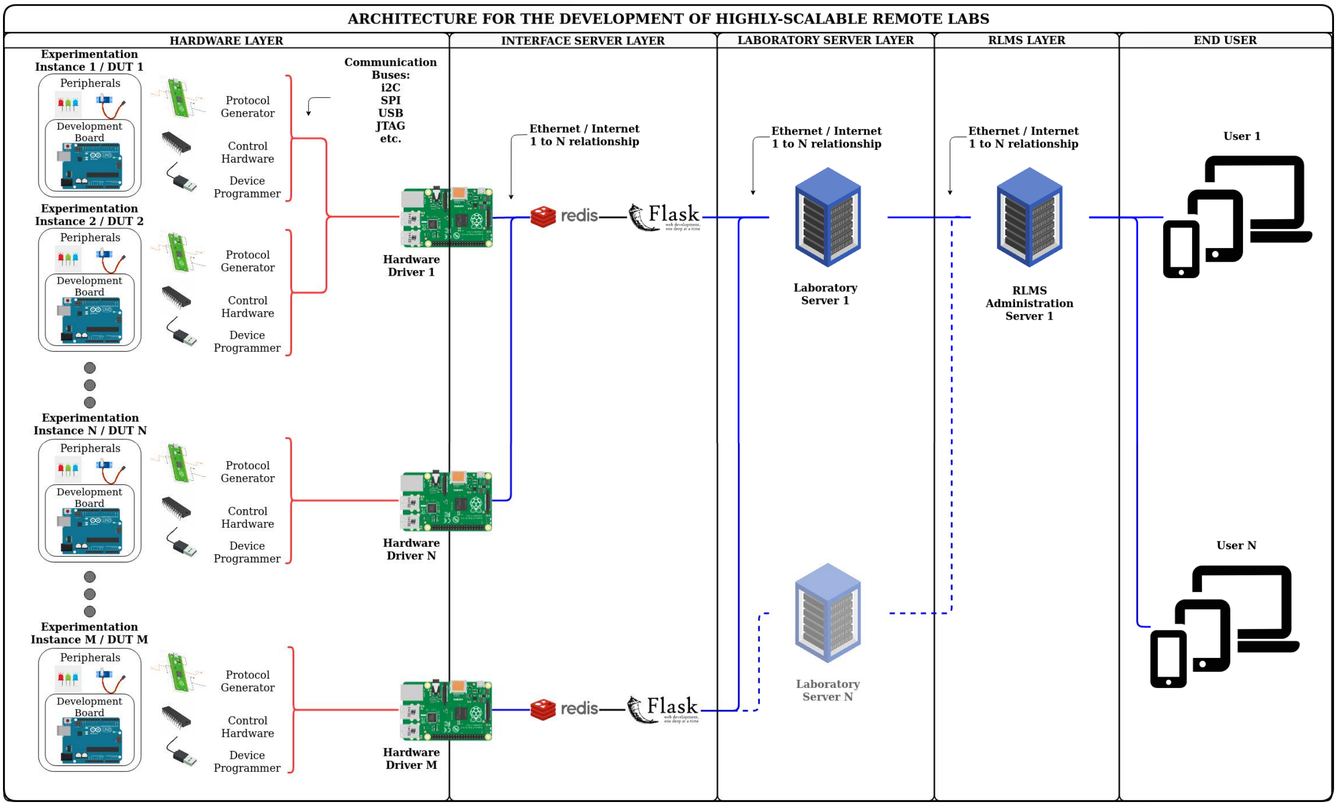

FIGURE 3. Overview of the architecture for the development of highly-scalable, multi-instance remote labs, split into the several layers that compose it. This architecture is detailed in [18] and is not a contribution of this article. It describes the architecture within a single institution.

- To facilitate and promote the presence of more than a single experimentation instance.
- To lower the cost per experimentation instance.

In conclusion, the ultimate goal of the architecture is to transmit information in a bidirectional way between the end user and the experimental setup. The user's actions are captured, transmitted and generated in the real experiment, while the laboratory behavior is captured, measured and/or detected and transmitted and displayed in real time to the end user. The advantage of this particular architecture is to be able to do these tasks in a scalable way, in order to support multiple users at the same time.

Naturally, such a remote laboratory has a capacity for as many concurrent users as there are instances. However, typically students use the laboratory in a program-evaluate cycle. That is, first they design their program in a different environment, and only when they believe it is ready, they test it in the real hardware. Then they go back to the design stage, and repeat the process. Since for many practical sessions students spend the most time programming rather than controlling the hardware itself, by creating a queue of users and limiting the hardware-accessing sessions in time (e.g. to 2 to 10 minutes), the number of students in a class that can be working at the same time is actually significantly higher than the number of physical instances. This fact makes scalability significantly easier and less costly than otherwise would be, for many types of laboratory.

B. MODEL FOR REMOTE LABS RELIABILITY

In order to have a proper experience with a remote laboratory, professors and their students certainly need to be able to have remote hardware available, which in group settings as the one described in the previous section, can be achieved through replication. However, it is also critical to be able to provide a good enough Quality of Service (QoS). The hardware that they access not only needs to be available, but it also needs to work as intended. It needs to be controllable, the results need to be accurate, and users should not be exposed to failing instances. Otherwise, the QoS can degreade significantly, and users lose confidence in the remote laboratory.

Hardware issues are relatively common in remote laboratories. To counteract this and maintain a high level of perceived laboratory quality, it is necessary to detect failed instances, and act accordingly, i.e., provide information to the laboratory maintenance staff team, remove the failed instances from the pool of instances available to the user, and redirect all users to the new pool of active instances. In this way, although waiting times may increase, the laboratory continues to provide service to the user, and the user perceives it as 100% functional and therefore reliable.

This is precisely what the model [19] mentioned focuses on. It allows for the detection of failed instances and the automatic distribution of users among the active instances. For this purpose, the model has four layers of failure detection. These are to be partially or fully applied depending on the laboratory, and applied with different latency. Once faults are detected, in any of the layers, the service acts automatically. First, it is evaluated whether the failure is persistent or not. If it is not, an attempt is made to restore or repair the experimentation instance. If the fault condition persists, the experimentation instance is removed from the set of active experimentation instances until the fault is repaired. That way, the Quality of Service from a user perspective, is safeguarded.

C. REMOTE LABORATORIES DEVELOPED APPLYING THESE MODELS

The aforementioned model is applied in various remote laboratories, which are part of the LabsLand network. LabsLand develops these laboratories, builds them and often deploys them; but the equipment is owned and deployed in different institutions around the world. As per LabsLand's policy, every real-time remote laboratory available in their network has multiple copies and is deployed in at least two institutions, to guarantee a minimum level of federated load balancing [20]. This is, in line with that model, to guarantee both scalability and reliability, in case an institution were to, for example, lose their power or internet connection.

The laboratories that will be described here are mostly for providing access to different electronic devices such as microcontrollers or FPGAs. These electronic devices are usually embedded in development boards that contain the device and some other electronic components necessary for its operation. An example of a development board is the Arduino UNO, which is perhaps the most well-known development board today. It allows users to easily control an Atmel ATmega328P microcontroller to practice on embedded systems.

These development boards are commonly used in hands-on labs because they are easy to use to connect to other devices and peripherals for lab exercises. They generally require users to write their own code on the device using a computer and then test its behavior, sometimes by generating and consuming electrical input and output signals and sometimes using electronic peripherals such as LEDs, displays, motors, servo motors, actuators, potentiometers, switches, push buttons, etc. In the remote laboratory environment, development boards are also used, although the situation is more complex. In this case, the electrical signals and behavior of the peripherals that may be encountered in a hands-on situation must be transmitted in real time to and from the user so that the user can perform similar laboratory practices.

LabsLand now has several different remote labs built on the above technological solutions that allow users to experiment with various embedded system development boards and FPGAs. Some of these labs have been developed in collaboration with other institutions, such as universities or colleges around the world. This allows the creation of a series of federated remote labs.

1) ARDUINO ROBOT LABORATORY

One of the first remote labs to be developed was the Arduino Robot lab, shown in Figure 4. In it, the user can program an Arduino-based Pololu Zumo32u4 tracked robot. Thanks to the many sensors it is equipped with (line detection, obstacle detection, accelerometer, gyroscope, motor rotation detection), the user can develop complex driving algorithms. In this lab, the user can program the robot with his own code. Once programmed, users can control it through the lab interface (shown in Figure 5) while visually evaluating the robot's behavior, which is recorded by a camera in real time.

This lab has multiple instances around the world. In total, there are instances distributed among facilities in the following countries: Spain, South Africa [71], Colombia [72], Costa Rica, and the United States of America. It is expected that the number of operational instances will continue to increase during 2022.

2) ARDUINO BOARD LABORATORY

Another lab that applies the described model is the Arduino Board lab from LabsLand. This lab is similar in concept to the previous one, except that in this case users can program and interact with an Arduino UNO development board connected to a set of peripherals, including a servo motor, several LED diodes, a RGB LED diode, an OLED display, a serial communication console, and remote buttons, switches, and potentiometers. The peripherals that can be visually evaluated are captured in real time by a camera, while the other peripherals are displayed on the lab interface, shown in Figure 6. With this set of devices, users can perform all the exercises that can be done with an Arduino Starter Kit, but remotely. As can be seen in the following sections, this helps users realize Arduino-based designs, focusing on programming first and understanding the operation of the electronic components at a lower level later. Both this lab and the robotics lab are Arduino-based and can be programmed directly from the LabsLand web platform, via their own IDE. The latter allows users to write Arduino-compatible code, either in written form or through block-based programming, code that can later be uploaded to both labs to evaluate their behavior. Figure 7 shows the differences between the block-based IDE and the standard written-code IDE.

The Arduino board lab also has several globally distributed instances. This instances are grouped in structures that contain 4 or 8 experimentation instances each, as those shown in Figure 8. The instances are currently distributed across the following countries: Spain, South Africa, Costa Rica, and the United States of America. It is expected that the number of operational instances will continue to increase during 2022.

3) LABSLAND's INTEL FPGA DE1-SOC LABORATORY

The described replication-oriented architecture can be adapted to support different types of devices-under-test (DUTs). The pool of remote labs is thus expanded to other

FIGURE 4. Supporting structure for a couple of LabsLand's Arduino Robot experimentation instances. The tracked robot, which is controlled by the user's code, can be moved over the entire surface of each rectangular tray.

development boards. A popular example is the Intel DE1-SoC FPGA lab. This lab is based on the Intel DE1-SoC development board, which is equipped with an Altera System-on-Chip Cyclone V FPGA. The development board itself, by Terasic, already has several visual peripherals, such as LEDs or 7-segment displays captured by a camera in real time. LabsLand adds additional peripherals, such as audio capture and injection, video capture and injection, protocol injection (PS2 keyboard, Nintendo N8 controller, etc.), and virtualized buttons and pushbuttons. As with the Arduino labs, users of this lab can also develop their VHDL or Verilog code in LabsLand's own IDE without having to download, install, or use any additional software. However, they can still synthesize their own binaries through Intel's Quartus, if they wish. Currently, there are 62 instances of this lab around the world. They can be found in various institutions around the world, in the following countries: Spain (26) and the United States of America (36). In a similar way to the Arduino Board lab, the hardware for the lab is fixed in a 3D-printed structure, shown in Figure 9.

4) LABSLAND's INTEL FPGA DE2-115 LABORATORY

Once the DE1 SoC development board lab was developed, it was not too complicated to develop a new remote lab with a similar family of development boards and the same remote lab support architecture. The LabsLand Intel FPGA DE2-115 lab allows users to program and practice with an Altera Cyclone IV on an Intel DE2-115 development board. The external peripherals are similar to those of the DE1 SoC lab, while the development board itself has more LEDs, more 7-segment displays, and a two-line 16-character display from Hitachi LCD. As with the other FPGA lab, code can be developed in the web-based LabsLand IDE. In this case, there are instances distributed worldwide in the following countries: Spain, Malaysia, Brazil, and the United States of America.

In order to homogenize and keep the development costs of all the above labs (except the Arduino Robot lab), a modular assembly system was developed that uses 3D printed parts to support multiple lab instances in a single physical structure.

This structure, which can be seen below in Figures 8 and 9, can hold 4 or 8 development boards, depending on size, along with power supplies, Ethernet switches, control hardware, peripherals, lighting, and camera. These components are typically shared across the 4 or 8 development boards, minimizing development costs by not duplicating hardware for each instance.

IV. METHODOLOGY

As described in Section I, the overarching goal of this paper is to evaluate whether remote labs, when properly designed and applying the aforementioned model (described in Section III), can indeed be reliable and practical teaching tools that are ready for widespread real-world usage.

For this, we will focus on two questions: First, whether the labs created under the proposed architecture are capable of scaling the number of instances without compromising their functionality and keeping costs under control. Second, whether the laboratories created under the proposed architecture are indeed reliable.

Additionally, to ensure the results are indeed representative and can confirm whether wide-spread real-world usage is possible and effective, the study will apply the characteristics mentioned in Section I:

- It will focus on a remote laboratory (the LabsLand FPGA lab) that is deployed in several institutions around the world and with dozens of deployed instances, and rely on data spanning more than two years (736 days).
- It will involve students from multiple universities in multiple different countries.
- The professors using the laboratories with their students will not be directly associated with the authors of this study or with the designers of the remote laboratories.

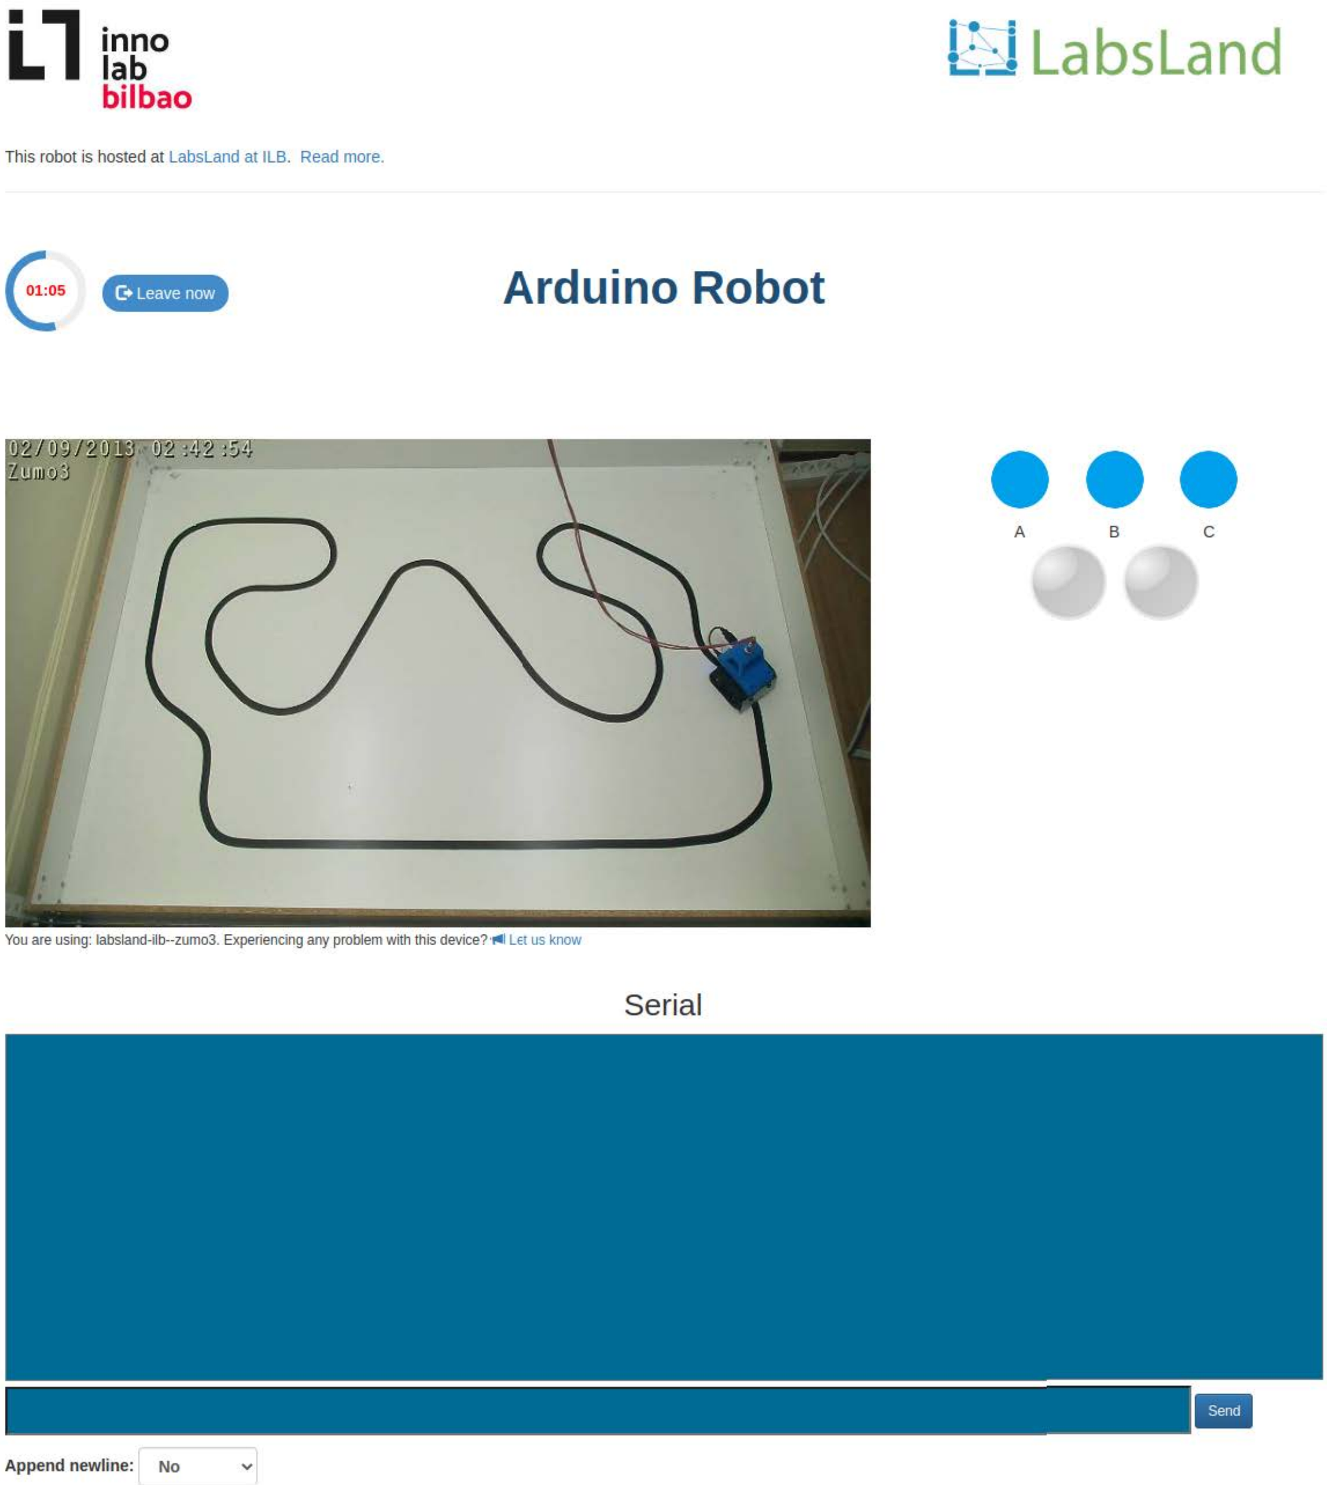

FIGURE 5. **Control panel for the LabsLand's arduino robot remote lab.**

The quantitative analysis examines whether the remote labs created as part of the architecture described earlier meet the goals of scalability and reliability. For this, it is necessary to analyze whether and how the laboratory is able to support a high number of experimental sessions within a given time period.

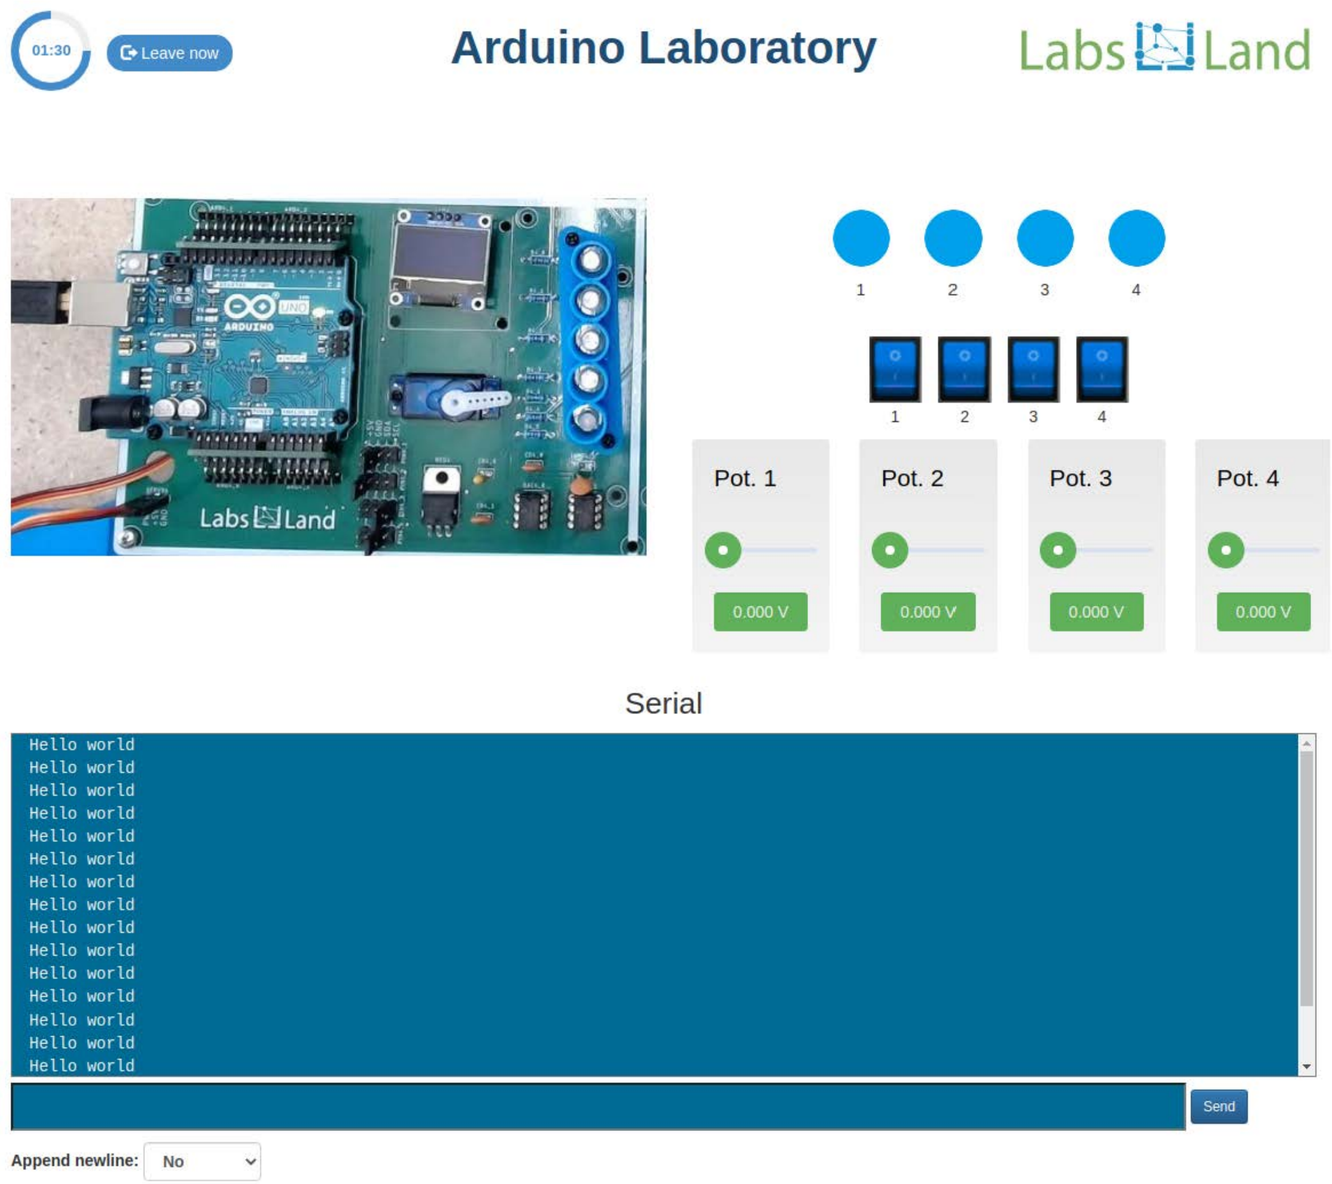

FIGURE 6. Control panel for the LabsLand's Arduino Board remote lab.

To this end, various parameters related to each experimental session were recorded and stored daily over a 24-month period. These parameters are:

- Date
- Number of available experimentation instances
- Number of experimentation sessions carried out
- Number of users served
- Maximum position reached in the queue
- Maximum time reached in the waiting queue, in seconds

It is important to evaluate both the capacity and reliability of the laboratory. A high capacity laboratory is one that is capable of performing a large number of experiments in a given time frame. Furthermore, a reliable lab is not one that does not fail, but one that, even if multiple experiment instances fail, is able to automatically adapt to the new number of available instances without the user noticing the failure.

An excessive number of faulty instances or an undersized lab will have the same result with this architecture: a potentially frustrated user due to excessive queue wait times and/or high queue positions. Low wait times or a low number of maximum queue positions reached result in the user perceiving the remote lab as reliable, knowing that they can check the performance of their code quasi-instantaneously each time they access the lab, just as they would in a hands-on lab with their own equipment.

In cases where the waiting time or the maximum reached position in the queue are high, it is necessary to compare the volume of users and experiment sessions with the number of

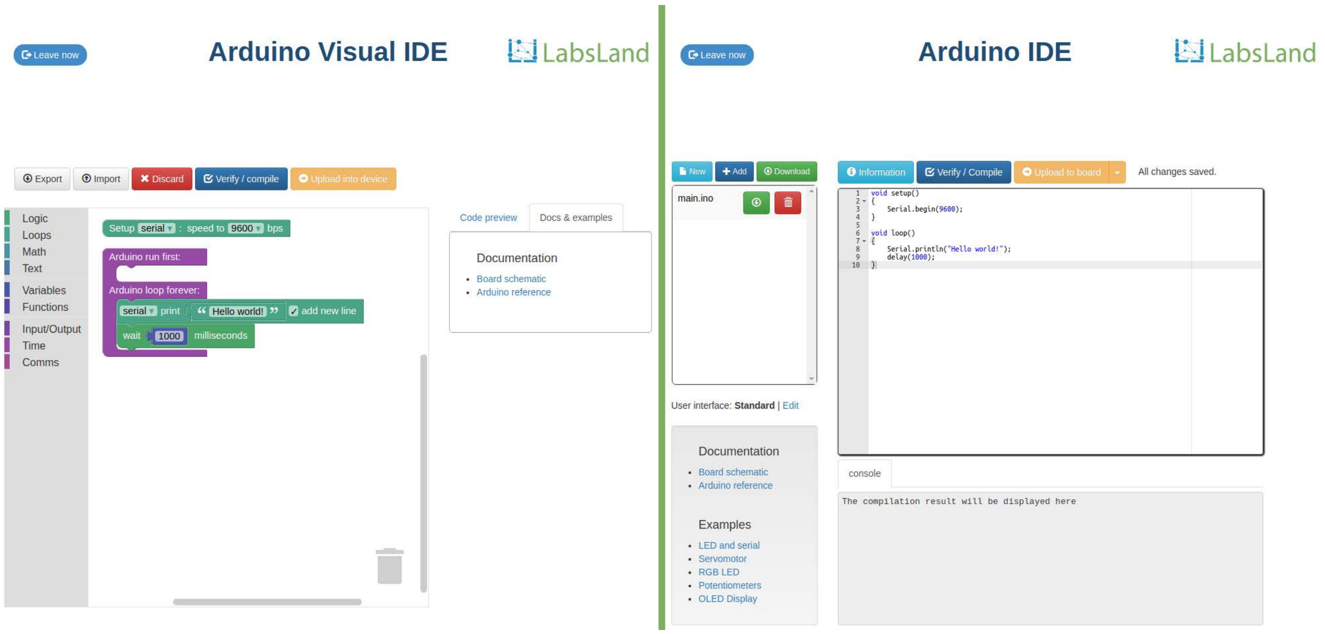

FIGURE 7. Graphical comparison between LabsLand's Arduino Visual IDE and Arduino (standard) IDE. Both IDEs generate the same microcontroller code.

available experiment instances. In this way, it is possible to determine whether it is an unusual situation (peak load at a certain time, e.g., during a class assignment, or when part of the available experiment instances fail, e.g., due to a power outage) or whether it is a usual situation that would indicate undersizing of the laboratory.

The parameters described above result in the following metrics:

- Number of available experimentation instances.
- Number of days of the subperiod.
- Number of days of the subperiod in which there was at least one use.
- Total number of experimentation sessions carried out in the subperiod.
- Maximum number of experimentation sessions carried out in one day.
- Average number of experimentation sessions carried out in one day.
- 90th percentile of number of experimentation sessions carried out in a day.
- Maximum number of users served in one day.
- Average number of users served in one day.
- 90th percentile of number of users served in one day.
- Maximum position reached in the queue in one day.
- Average position reached in the queue in one day.
- 90th percentile of position reached in the queue in one day.
- 99th percentile of position reached in the queue in one day.
- Maximum time reached in the waiting queue, in seconds, in one day.
- Average time reached in the waiting queue, in seconds, in one day.
- 90th percentile of time reached in the waiting queue, in seconds, in one day.

The maximum values help identify the lab's worst situations, while the average and 90th percentile of the data help understand how the lab performs in general and 90 percent of the time.

The results of this analysis are presented and discussed in the next section.

V. RESULTS AND DISCUSSION

This section presents the results of the analysis conducted through the aforementioned methodology (see previous Section IV).

As described, the quantitative analysis focuses on two key aspects: scalability and reliability. To this end, we analyzed a set of data on laboratory use collected over a 736-day time span across multiple institutions. The usage data is from real-world production usage. As such, the professors and students using the laboratory are not directly involved in this study or associated with the authors, and are using the laboratory freely, without following any specific guideline. This is as intended, since the goal is precisely to evaluate the suitability for real-world multi-institutional usage.

The period consists of four subperiods throughout which the total number of available experiment instances increased. In subperiod 1, there were 10 experimentation instances online in the laboratory, while in subsequent subperiods there were 18, 34, and 62 experimentation instances, respectively. This variation can be useful since it allows us to evaluate how the model's performance varies with different numbers of replicated instances and different loads. It is noteworthy, nonetheless, that it is not a result of the experimental design itself, but rather a consequence of the fact that we are using

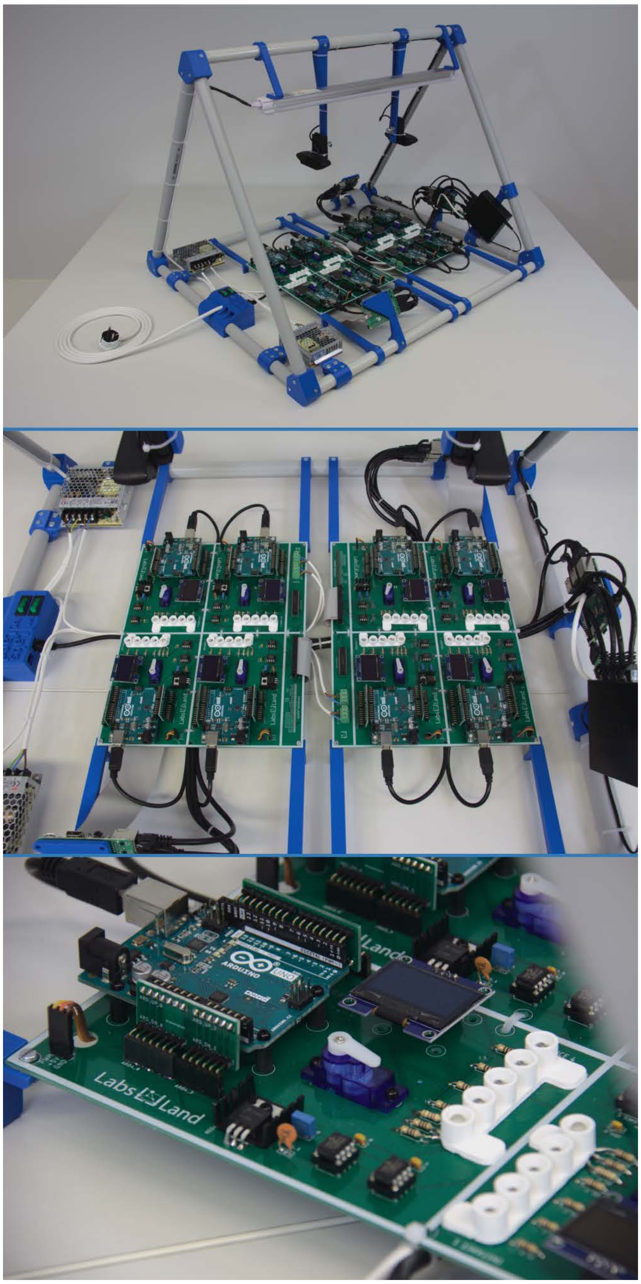

FIGURE 8. Multiple views of the physical structure for LabsLand's Arduino Board remote lab. The structure supports eight Arduino UNO development boards, the peripherals associated with each experimentation instance, and the remotization hardware.

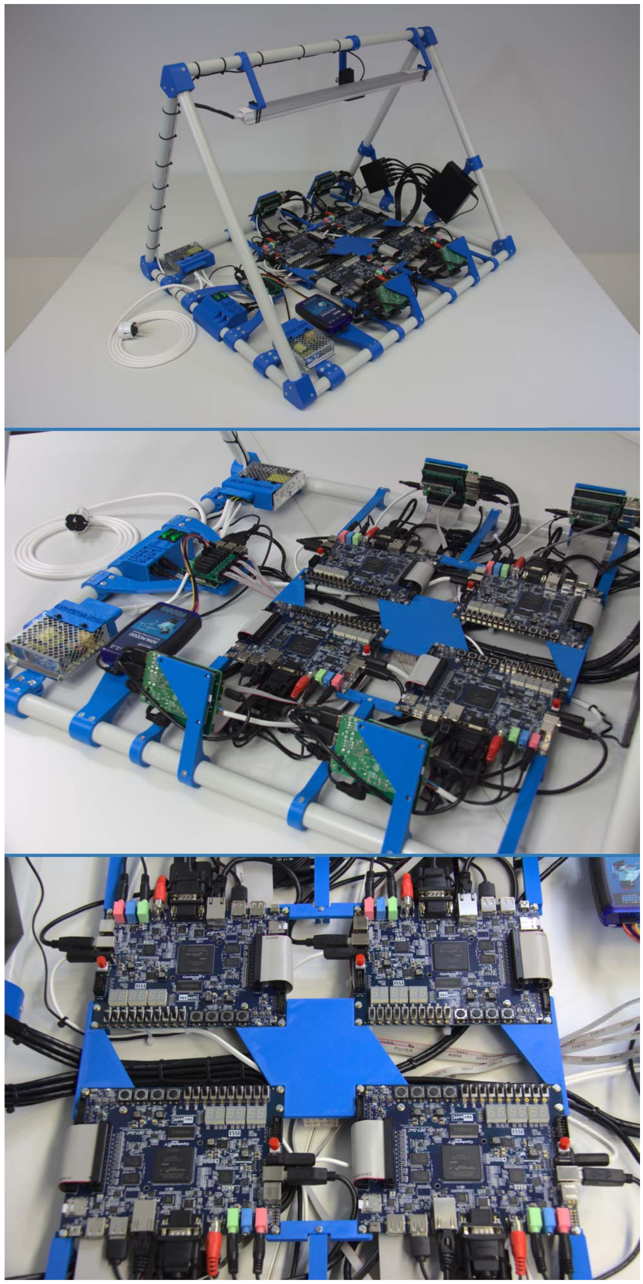

FIGURE 9. Multiple views of the physical structure for LabsLand's Intel DE1-SoC FPGA remote lab. The structure supports four Intel DE1-SoC FPGA development boards and the remotization hardware.

real-world usage data, and throughout this relatively long period (nearly two years) new institutions purchased new copies of the laboratory, deployed them, and they were integrated into the LabsLand's network cluster.

Table 1 summarizes the data obtained during the analysis. Rows 1, 2, 3, and 4 provide context, characterizing the subperiod. Rows 5, 6, and 7 collect results related to the total number of experimentation sessions performed in the laboratory for each day of the subperiod. Rows 8, 9, and 10 show the results in terms of the total number of users for each day of the subperiod. Rows 11, 12, 13 and 14 show the results in terms of the maximum position in the queue that allows access to the lab for each day of the subperiod. The position 0 in the queue indicates that they have no one in front of them and that they will gain access to the laboratory as soon as one of the existing experiment instances are available. Rows 15, 16, and 17 show the results in terms of the maximum queue waiting time reached each day (in seconds). This waiting time can be increased by various causes, such as the time needed to restore lab equipment between sessions, session time for previous users, or session allocation time by the lab management system.

Figure 10 graphically shows the maximum queue position (in red) and the maximum waiting time (in blue) for each day within the period in which at least one use of the lab occurred. The left axis represents the time in seconds, with a range between 0 and 240 seconds, while the right axis represents the highest reached position in the queue, with values between 0 and 4 (included). The light green line is a linear trend line, calculated from the maximum queue waiting time data. It is a practically horizontal line, with a slight upward trend.

A quick visual analysis reveals two things: one, that there are some days when the values are far from the typical values, and two, that on most days, the values remain below the 60 seconds threshold that we have set for the maximum queue waiting time and practically 0 waiting positions reached in the queue. However, Figure 10 lacks the precision needed to see those common values in detail due to the peak values. Therefore, Figure 11 has been included, which shows a zoomed-in version in which the peak cases have been cropped out. The horizontal line in Figure 10, set in the horizontal axis at the point of 60 seconds maximum waiting time, is an arbitrary threshold and also the new boundary that is used for cropping the data for Figure 11.

Peak cases are cases in which values deviate significantly from the average. To identify these cases, 45 seconds of waiting time in the queue and a position in the queue greater than 1 were chosen as cutoff values. The cases that have been identified have been collected in a new figure, Figure 12. These cases have not been frequent, but are nonetheless to be analyzed in detail. Additional data associated to these days have been collected in Table 2.

In general, the results of the quantitative analysis are positive and indicate that the original goal is met. The remote laboratory is indeed scalable and reliable, and has been able to provide the intended production-level service for multiple institutions and thousands of students and experimentation sessions, maintaining a high QoS the great majority of the time across the 736-days period.

More than 70,000 experiments were performed throughout the complete testing period. In the first subperiod, there were days when the lab was not used, while in the other subperiods, usage was significantly more regular, occurring almost every day and including weekends.

Looking at the number of sessions performed in the laboratory, we can see that the maximum number of sessions per day has remained close to around 1000 daily sessions. Looking at the mean and 90th percentile of this metric, we find that the use of the laboratory tripled in the last subperiod, while it remained nearly constant throughout the first three subperiods. In the first and second subperiods, 90% of the days, up to 232 experimentation sessions were served, while in the fourth subperiod, during 90% of the days, up to 653 daily sessions were served.

Looking at the number of users served in the lab, the results show a different trend. Both the total number of users served and the 90th percentile increased by only 50%. If the number of sessions served increased by 300%, but the number of users served only increased by 50%, it is reasonable to assume that users were using the lab much more intensively.

As for the position reached in the queue, the results show that in 90% of the cases the users did not find anyone ahead of them in the queue.

The maximum values are particularly low and are in fact 0 during subperiod 2. Since the 90th percentile is 0.00 in all subperiods (90% of users were able to access without waiting for anyone in the queue), the 99th percentile was calculated. In the first two sub-periods, no user was found second in the queue 99% of the time, while in sub-periods 3 and 4, only 1% of users found someone ahead of them in the queue.

It is important to remember that despite not having anyone in front in the queue, users may still have to wait a short time. In this case, it is not possible to discern whether the user has directly accessed the laboratory, or is waiting for the first instance to be freed. These waiting times can be increased by factors such as connection delays or recovery processes between experimentation sessions, so they cannot always be directly controlled. For this reason, the results are not completely regular in terms of waiting time.

There are some rather high maximum values, but these are few of the total values. The mean value is close to 9 seconds in the first three subperiods, when the laboratory was used less intensively. The fourth subperiod has an average value of around 15 seconds. Considering the 90th percentile, it can be ensured that in 90% of the cases the users had to wait at most 19.60 seconds in the subperiod in which the use of the laboratory was more extensive, which is a positive fact that strengthens the feeling of reliability of the laboratory.

Using the graphical data shown in Figure 10, it is possible to observe the behavior of the laboratory throughout the test period. In general, both the waiting times and the queue positions reached are kept low, except for a few peak cases.

Table 2 shows in detail the data related to the days when peaks were detected. We can observe peaks in 7 different days. 3 of them have relatively high queue wait times (a user spent 98, 57 and 226 seconds waiting, respectively, shown in blue) while nonetheless having a maximum queue position of 0 (shown in red). Those are rare circumstances, since normally a long waiting time would be accompanied by a higher-than-zero queue position, and also the number of sessions and users are not particularly far from the average. It is likely that those rare instances (in 3 days out of the 736) were due to temporary issues in the cloud-based cluster or temporary network problems.

In the days of the 5th and 6th peak (in red), the maximum reached positions in the queue were higher than 1 (2 for both days). This is higher than usual and is caused by high concurrent user loads. It implies that during those days, there were times in which every single online instance available for the laboratory was serving a student at a specific point in time. At this point in time there were 34 instances of the laboratory, so that is a significant load. The maximum

TABLE 1. LabsLand DE1-SoC FPGA lab quantitative analysis results. Four subperiods are distinguished, in which the number of available experimental instances varied.

		Subperiod 1	Subperiod 2	Subperiod 3	Subperiod 4
1	Number of available experimentation instances	**10**	**18**	**34**	**62**
2	Number of days of the subperiod	361	115	205	55
3	Number of days of the subperiod in which there was at least one use	252	107	196	55
4	Total number of experimentation sessions carried out in the subperiod	22635	10033	24835	14874
5	Maximum number of experimentation sessions carried out in one day	807	1393	1242	1104
6	Average number of experimentation sessions carried out in one day	89.82	93.77	126.71	270.43
7	90th percentile of number of experimentation sessions carried out in one day	232.60	202.00	287.50	653.20
8	Maximum number of users served in one day	75	84	98	118
9	Average number of users served in one day	15.88	16.36	19.19	34.72
10	90th percentile of number of users served in one day	38.00	40.80	44.50	58.60
11	Maximum position reached in the queue in one day	4	0	2	1
12	Average position reached in the queue in one day	0.030	0.000	0.025	0.018
13	90th percentile of position reached in the queue in one day	0.00	0.00	0.00	0.00
14	99th percentile of position reached in the queue in one day	0.00	0.00	1.05	0.46
15	Maximum waiting time reached in the queue, in seconds, in one day	98	31	35	226
16	Average waiting time reached in the queue, in seconds, in one day	8.73	8.29	9.34	15.018
17	90th percentile waiting time reached in the queue, in seconds, in one day	12.00	13.00	16.50	19.60

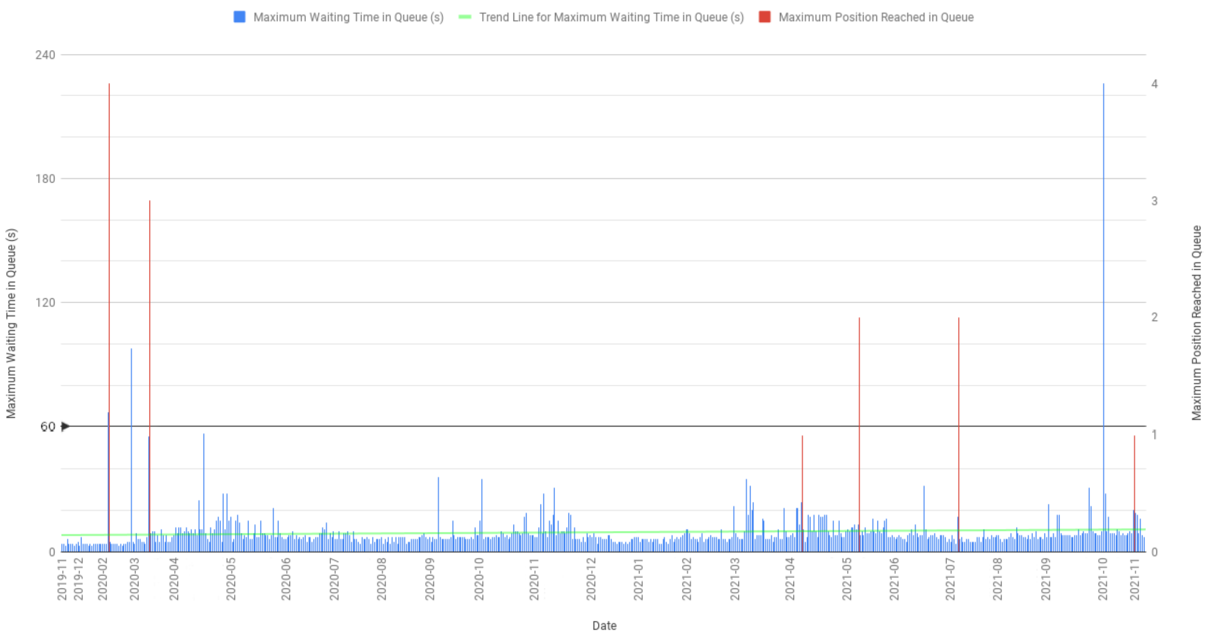

FIGURE 10. Graph showing the maximum queue position (in red) and the maximum recorded waiting time (in blue) for each day in the 24-month analysis period.

waiting times, nonetheless, were only 13 and 17 seconds, respectively. As such, the line in the zoomed-in Figure 12 is red only, since those times are below the threshold. This indicates that even in those specific days under such a high load, they were able to successfully access the laboratory after waiting a very reasonable time.

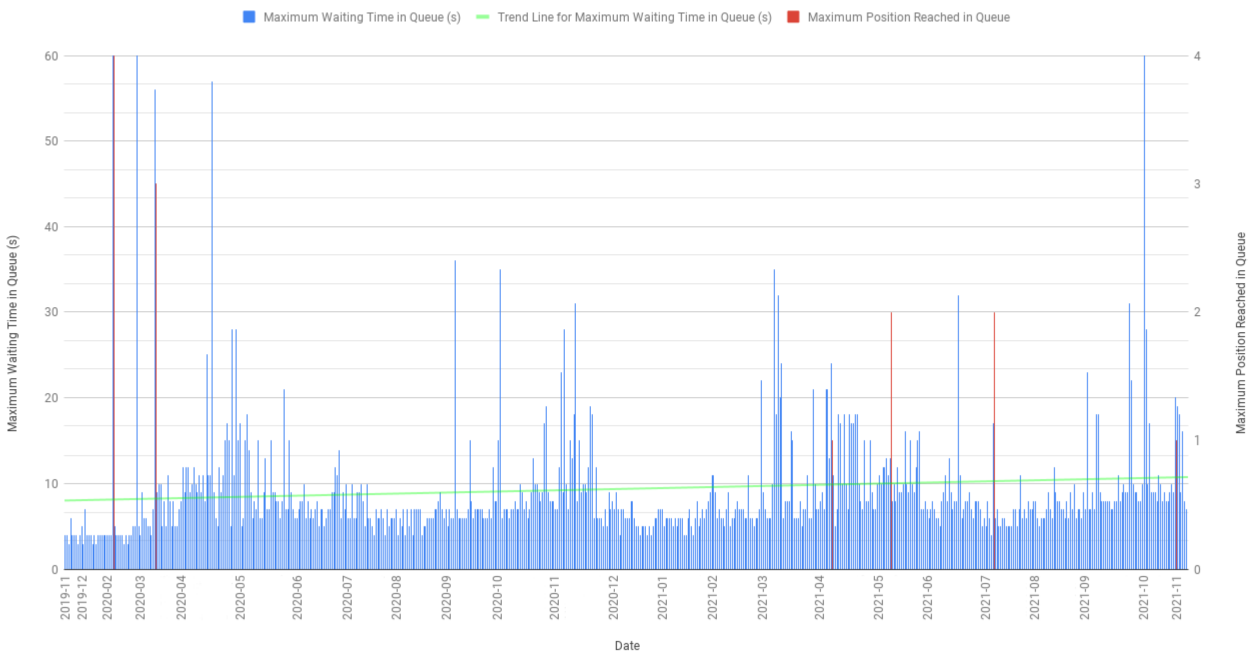

FIGURE 11. Zoomed-in version of the Figure 10 graph. The vertical axis has been re-scaled leaving out peak times higher than 60 seconds so that the more frequent cases can be observed better.

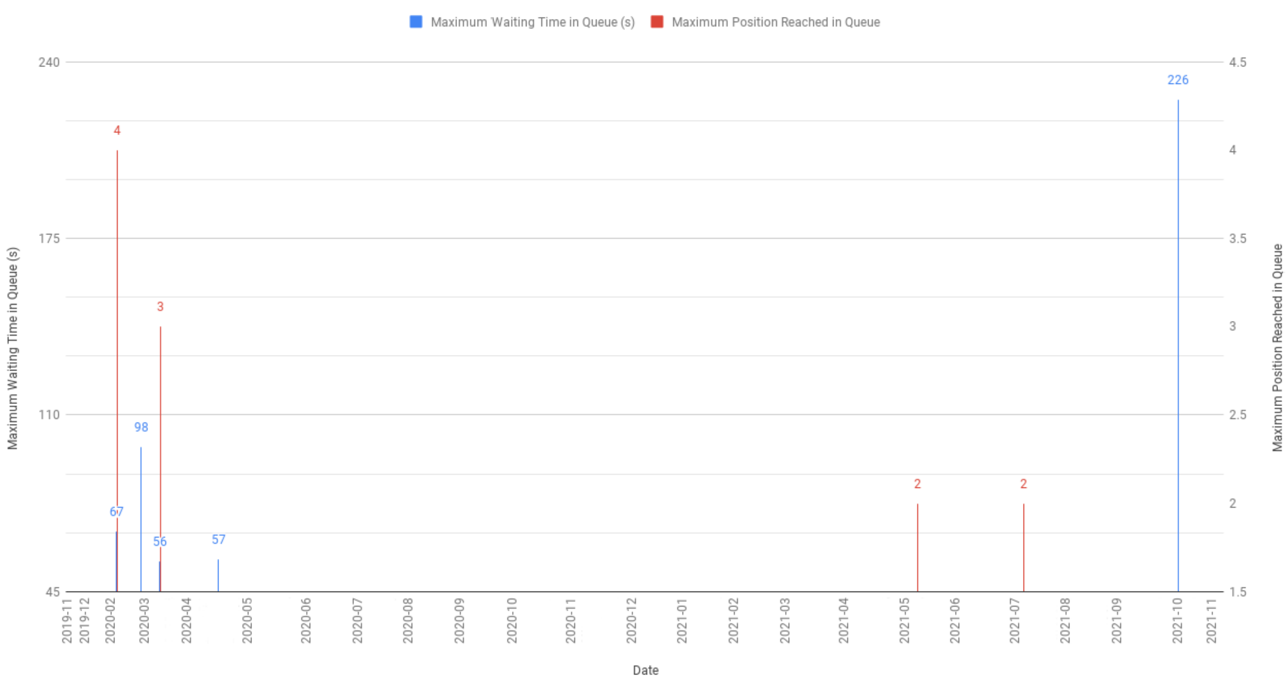

FIGURE 12. Zoomed-in version of the Figure 10 graph. This version highlights the peaks, starting at 45 seconds (vertical axis). In blue, the maximum waiting time. In red, the highest position in the queue.

Finally, days with the first and third peaks represent a situation in which laboratory capacity was reached and the Quality of Service was partially compromised. As such, the maximum position in the queue was high, and the actual maximum waiting times were also relatively high. In the day of the first peak, a user was 4th in the queue and a user

TABLE 2. Peak cases detected during the quantitative analysis of LabsLand's DE1-SoC FPGA lab.

Day	Date	Available Experimentation Instances	Experimentation Sessions Carried Out	Users Served	Maximum Position Reached in Queue	Maximum Waiting Time Reached in Queue
1	**2020-02-13**	10	168	55	4	67
2	**2020-03-18**	10	25	12	0	98
3	**2020-03-30**	10	53	30	3	56
4	**2020-04-30**	10	305	47	0	57
5	**2021-05-24**	34	363	58	2	13
6	**2021-07-19**	34	80	25	2	17
7	**2021-10-18**	62	643	105	0	226

(probably but not necessarily the same one) had to wait up to 67 seconds to gain access to the lab. On the day of the third peak, something similar occurred: a user was at the 3rd position in the queue and a user had to wait up to 56 seconds to access his experiment session.

In the period of those two peaks the number of available experiment instances was 10. These 10 instances served 55 users on the first day and 30 users on the second day. Since they were not spread out over the day but were instead concentrated in a smaller window of time (probably because they were using the laboratory concurrently in-class) this situation was more likely to occur. This is less likely with higher numbers of instances (as in the next periods) since it is less likely for classes from different institutions to be using the laboratory at the same time. It is also noteworthy that this situation occurred on only 2 of the 736 days of the period if study, did not recur in subperiods in which there were a larger number of instances available, and did affect the QoS, but probably not very significantly (they had to wait at most 67 seconds to access the laboratory).

VI. CONCLUSION

The main goal of this work was to evaluate whether by applying the model that has been described in the first sections it is indeed possible to obtain remote laboratories that can be effective for real-world educational operational environments encompassing multiple institutions. This may eventually lead to their widespread adoption.

As discussed, for a remote lab to become a widespread and standard tool in an educational institution setting, its potential users must trust it and perceive it as a reliable tool. In a hands-on lab this is also important: the equipment must be working. Though since users maintain a physical connection to the equipment, it is often possible for them to fix issues that would be much harder to fix remotely. Also, even if issues occur, users will be less likely to feel that they are not in control of the experiment, which is an extremely harmful perception for the QoS. To be trusted and perceived as reliable, the standard is therefore even higher for a remote laboratory.

The study that was conducted, encompassing a period of slightly more than 2 years (736 days), suggests that the main goal was met and that it is indeed possible to create production-level laboratories for real-world multi-institutional usage by applying the described model. Professors and their classes, from multiple different institutions, successfully used the FPGA remote laboratory for their courses. As the results in Section V show, they were used extensively, throughout thousands of sessions (72,377 times across the four described periods). At all times throughout that 736 days period the laboratory remained available and provided quick and effective access and a high QoS. Only in 2 days out of the 736 ones a user experienced minor degradations in the QoS, by having to wait up to 67 seconds to access the laboratory. Thus we conclude that in general, the availability and QoS goals are indeed clearly met.

Although it is not part of this study, it is noteworthy that the individual reports of the professors that used the laboratory throughout this period and that reported to us their experience were indeed positive. In a separate independent study that one of the professors conducted using the LabsLand FPGA laboratory, for example, the learning results were in fact superior to the traditional hands-on course of their previous year [5].

In summary, we can conclude that by applying a model such as the proposed one it is possible to create remote laboratories that are indeed useful for real-world multi-institutional usage and that can meet the QoS standards. Meeting such standards and being trusted and reliable is an important step for remote labs to become a truly ubiquitous tool for scientific and technical education across universities and other learning institutions, and therefore to reach their potential and become truly widespread.

VII. FUTURE WORK

This contribution and related ones have shown that it is indeed possible to create remote laboratories that are both scalable to multiple users and reliable; and that they can thus be useful for real-world educational usage, such as engineering courses in both hands-on and distance universities. This has been demonstrated with laboratories for various fields (e.g. basic electronics [73], [74] or digital electronics [75]). In the future, research and development efforts will be dedicated to extend the reach of such remote laboratories towards additional fields, while still focusing on the basic goals of promoting reliability and scalability.

In this same line, we intend to keep working towards improving the reliability and scalability of existing and future remote labs. This will mainly be approached in various ways:

- **Reducing error rates.** If the individual instances of remote laboratories can become less prone to failure, the likelihood of failures going undetected tends to be smaller, and at the same time, scalability to a given number of simultaneous users can be reliably achieved with a lower number of laboratory instances (copies of the laboratory).
- **Improving error detection.** The ability to automatically realise when a laboratory is failing, and to do so as fast as possible, and before a student gets exposed to the failure is important for ensuring proper Quality of Service and making it easier for the failing copy of the laboratory to be repaired swiftly.
- **Reducing replication costs.** As shown, replication of the laboratory instances is critical for being able to provide a reliable service that is scalable to multiple users. However, replicating a laboratory is significantly costly, due to the equipment cost itself, and also the deployment and maintenance effort required. Efforts will be dedicated towards lowering that cost, by exploring ways to streamline the process, to share certain components among different instances, or to use low-cost and easy-to-assembly components.

ACKNOWLEDGMENT

The authors would like to thank LabsLand for supporting this research and providing the usage data. They also would like to thank the LabsLand community in general. In particular, they would like to thank Pablo Orduña (LabsLand), who assisted during the writing of this article and the provision of data; and also both Candido Aramburu (Professor at the Public University of Navarre, UPNA[11]) and Rania Hussein (Professor at the University of Washington and the Founder and the Director of the Remote Hub Lab[12]), who are leaders of the remote laboratory research groups in their universities, and who own and host the DE1-SoC FPGA remote laboratory instances that were involved in this study, and that students from around the world accessed.

REFERENCES

[1] J. V. Nickerson, J. E. Corter, S. K. Esche, and C. Chassapis, "A model for evaluating the effectiveness of remote engineering laboratories and simulations in education," *Comput. Educ.*, vol. 49, no. 3, pp. 708–725, Nov. 2007.

[2] J. E. Corter, J. V. Nickerson, S. K. Esche, C. Chassapis, S. Im, and J. Ma, "Constructing reality: A study of remote, hands-on, and simulated laboratories," *ACM Trans. Comput.-Hum. Interact.*, vol. 14, no. 2, p. 7, Aug. 2007, doi: 10.1145/1275511.1275513.

[11]https://unavarra.es

[12]https://rhlab.ece.uw.edu

[3] T. De Jong, M. C. Linn, and Z. C. Zacharia, "Physical and virtual laboratories in science and engineering education," *Science*, vol. 340, no. 6130, pp. 305–308, 2013.

[4] J. R. Brinson, "Learning outcome achievement in non-traditional (virtual and remote) versus traditional (hands-on) laboratories: A review of the empirical research," *Comput. Educ.*, vol. 87, pp. 218–237, Sep. 2015.

[5] R. Hussein and D. Wilson, "Remote versus in-hand hardware laboratory in digital circuits courses," in *Proc. ASEE Virtual Annu. Conf. Content Access*, 2021, pp. 1–19.

[6] C. Batanero, L. De-Marcos, J. Holvikivi, J. R. Hilera, and S. Oton, "Effects of new supportive technologies for blind and deaf engineering students in online learning," *IEEE Trans. Educ.*, vol. 62, no. 4, pp. 270–277, Nov. 2019.

[7] J. Garcia-zubia. (2021). *Remote Laboratories: Empowering STEM Education With Technology*. World Scientific Publishing Company. [Online]. Available: https://books.google.es/books?id=eo8bzgEACAAJ

[8] K. P. Nuci, R. Tahir, A. I. Wang, and A. S. Imran, "Game-based digital quiz as a tool for improving students' engagement and learning in online lectures," *IEEE Access*, vol. 9, pp. 91220–91234, 2021.

[9] A. W. Muzaffar, M. Tahir, M. W. Anwar, Q. Chaudry, S. R. Mir, and Y. Rasheed, "A systematic review of online exams solutions in E-learning: Techniques, tools, and global adoption," *IEEE Access*, vol. 9, pp. 32689–32712, 2021.

[10] Z. Nedic, J. Machotka, and A. Nafalski, "Remote laboratories versus virtual and real laboratories," in *Proc. 33rd Annu. Frontiers Educ. (FIE)*, 2003, p. T3E.

[11] V. Potkonjak, M. Gardner, V. Callaghan, P. Mattila, C. Guetl, V. M. Petrovic, and K. Jovanovic, "Virtual laboratories for education in science, technology, and engineering: A review," *Comput. Educ.*, vol. 95, pp. 309–327, Apr. 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360131516300227

[12] J. E. Froyd, P. C. Wankat, and K. A. Smith, "Five major shifts in 100 years of engineering education," *Proc. IEEE*, vol. 100, pp. 1344–1360, May 2012, doi: 10.1109/JPROC.2012.2190167.

[13] K. Jona, A. D. Walter, and S. N. Pradhan, "Designing remote labs for broader adoption," in *Proc. 12th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, Feb. 2015, pp. 74–80.

[14] N. Lima, C. Viegas, and F. Garcia-Peñalvo, "Didactical use of a remote lab: A qualitative reflection of a teacher," in *Proc. 7th Int. Conf. Technol. Ecosyst. Enhancing Multiculturality*. New York, NY, USA: Association for Computing Machinery, Oct. 2019, pp. 99–108, doi: 10.1145/3362789.3362891.

[15] S. D. Bencomo, "Control learning: Present and future," *IFAC Proc. Volumes*, vol. 35, no. 1, pp. 71–93, 2002.

[16] L. Gomes and S. Bogosyan, "Current trends in remote laboratories," *IEEE Trans. Ind. Electron.*, vol. 56, no. 12, pp. 4744–4756, Dec. 2009.

[17] J. Ma and J. V. Nickerson, "Hands-on, simulated, and remote laboratories: A comparative literature review," *ACM Comput. Surveys*, vol. 38, no. 3, p. 7, Sep. 2006.

[18] A. Villar-Martínez, L. Rodriguez-Gil, I. Angulo, P. Orduña, J. García- Zubía, and D. López-De-Ipiña, "Improving the scalability and replicability of embedded systems remote laboratories through a cost-effective architecture," *IEEE Access*, vol. 7, pp. 164164–164185, 2019.

[19] A. Villar-Martínez, J. García-Zubía, I. Angulo, and L. Rodríguez-Gil, "Towards reliable remote laboratory experiences: A model for maximizing availability through fault-detection and replication," *IEEE Access*, vol. 9, pp. 45032–45054, 2021, doi: 10.1109/ACCESS.2021.3065742.

[20] P. Orduña, "Transitive and scalable federation model for remote laboratories," Ph.D. dissertation, Fac. Eng., Univ. Deusto, May 2013. [Online]. Available: https://weblab.deusto.es/pub/dissertation_pablo.pdf

[21] E. Lindsay and B. Stumpers, "Remote laboratories: Enhancing accredited engineering degree programs," in *Proc. 22nd Annu. Conf. Australas. Assoc. Eng. Educ., (AAEE)*. Barton, ACT, Australia, Engineers Australia, 2011, pp. 588–593.

[22] A. Cordeiro, V. Fernão Pires, and D. Foito, "Combining local and remote laboratories for the interactive learning of industrial automation," *Comput. Appl. Eng. Educ.*, vol. 26, no. 3, pp. 675–687, May 2018.

[23] F. Lustig, P. Brom, P. Kuriscak, and J. Dvorak, "'Hands-on-remote' laboratories," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2018, pp. 118–127.

[24] F. Schauer, M. Krbecek, P. Beno, M. Gerza, L. Palka, and P. Spilakova, "REMLABNET—Open remote laboratory management system for e-experiments," in *Proc. 11th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, Feb. 2014, pp. 268–273.

[25] F. Schauer, M. Krbecek, P. Beno, M. Gerza, L. Palka, P. Spilakova, and L. Tkac, "REMLABNET III—Federated remote laboratory management system for university and secondary schools," in *Proc. 13th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, Feb. 2016, pp. 238–241.
[26] F. Schauer *et al.*, "Creative laboratory experiments for basic physics using computer data collection and evaluation exemplified on the intelligent school experimental system (ISES)," in *INNOVATIONS 2006: World Innovations in Engineering Education and Research*, vol. 1, no. 1, W. Aung, C. Crosthwaite, R. V. Espinosa, J. Moscinski, S.-H. Ou, and L. M. S. Ruiz, Eds., Jul. 2006, pp. 305–312.
[27] P. Brom, F. Lustig, P. Kuriscák, and J. Dvorák, "DIY 'hands-on-remote' experiment in physics with Arduino," in *Proc. 15th Int. Conf. Remote Eng. Virtual Instrum. (REV)*. Düsseldorf, Germany: Univ. Applied Science Düsseldorf, Mar. 2018, pp. 617–624.
[28] H. Fernandes, S. Leal, and J. P. Leal, "E-lab: O laboratório online," *Gazeta Da Física*, vol. 33, no. 3, pp. 37–40, 2010.
[29] S. Leal and J. P. Leal, "A new chemistry e-lab experiment chemical equilibirum reaction," in *Proc. 2nd Exp.@ Int. Conf.*, Sep. 2013, pp. 154–155.
[30] J. B. Da Silva, S. M. S. Bilessimo, and K. C. N. da Silva, "Mobile remote experimentation workgroup (GTMRE) strategy for integrating technology into education [a estratégia de integração de tecnologia na educação do Grupo de Trabalho em Experimentação Remota Móvel (GTMRE)]," *Revista Tecnologias na Educação*, vol. 17, no. 8, pp. 1–14, Dec. 2016.
[31] J. P. C. De Lima, L. M. Carlos, J. P. S. Simão, J. Pereira, P. M. Mafra, and J. B. Da Silva, "Design and implementation of a remote lab for teaching programming and robotics," *IFAC-PapersOnLine*, vol. 49, no. 30, pp. 86–91, Jan. 2016.
[32] *RELLE RExLab Arduino Remote Laboratory*. Accessed: Jul. 23, 2019. [Online]. Available: http://relle.ufsc.br/labs/4
[33] B. Letowski, C. Lavayssière, B. Larroque, and F. Luthon, "An open source remote laboratory network based on a ready to use solutioN: Laborem," in *Proc. ICERI*, Seville, Spain, Nov. 2019, pp. 1–7. [Online]. Available: https://hal.archives-ouvertes.fr/hal-02405954
[34] C. Lavayssière, B. Larroque, and F. Luthon, "Laborem Box: A scalable and open source platform to design remote lab experiments in electronics," *HardwareX*, vol. 11, Apr. 2022, Art. no. e00301. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2468067222000463
[35] D. A. H. Samuelsen and O. H. Graven, "Design of a general purpose platform for easy setup of low-cost remote laboratories in electronics," in *Proc. 10th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, Feb. 2013, pp. 1–6.
[36] S. K. Esche, C. Chassapis, J. W. Nazalewicz, and D. J. Hromin, "An architecture for multi-user remote laboratories," *Dyn. Typical Class Size 20 Students*, vol. 5, no. 6, pp. 1–6, 2003.
[37] I. Gustavsson, J. Zackrisson, L. Håkansson, I. Claesson, and T. L. Lagö, "The VISIR project—An open source software initiative for distributed online laboratories," in *Proc. REV*, 2007, pp. 1–6.
[38] G. A. Pokoo-Aikins, N. Hunsu, and D. May, "Development of a remote laboratory diffusion experiment module for an enhanced laboratory experience," in *Proc. IEEE Frontiers Educ. Conf. (FIE)*, Oct. 2019, pp. 1–5.
[39] I. Idoyaga, C. Moya, R. Sanchéz-Brenes, J. Maeyoshimoto, C. A. Matarrita, and E. Montero-Miranda, "Acid-base validation remote laboratory in a university chemistry course [El laboratorio remoto de validación Ácido-Base en un curso de química en la universidad]," in *Proc. 11th Int. Congr. Appl. Adv. Inf. Commun. Technol. (ATICA)*, Dec. 2020, pp. 792–799.
[40] K. C. Narasimhamurthy, P. Orduña, L. Rodríguez-Gil, G. C. Bharath, C. N. S. Srivatsa, and K. Mulamuttal, "Analog electronic experiments in ultra-concurrent laboratory," in *Cross Reality and Data Science in Engineering* (Advances in Intelligent Systems and Computing), M. E. Auer and D. May, Eds. Cham, Switzerland: Springer, 2021, pp. 37–45.
[41] M. S. Apel, F. J. C. Odebrett, C. Paz, and N. Perozo, "Multi-phase flowloop remote laboratory," in *Cross Reality and Data Science in Engineering*, M. E. Auer and D. May, Eds. Cham, Switzerland: Springer, 2021, pp. 184–192.
[42] C. Arguedas-Matarrita, E. Montero-Miranda, L. Vargas-Badilla, R. Sánchez-Brenes, E. Rios-Badilla, P. Orduña, and L. Rodríguez-Gil, "Design and development of an ultra-concurrent laboratory for the study of an acid-base titration (ABT) at the Universidad Estatal a Distancia (UNED), Costa Rica," in *Online Engineering and Society 4.0*, M. E. Auer, K. R. Bhimavaram, and X.-G. Yue, Eds. Cham, Switzerland: Springer, 2022, pp. 122–130.
[43] M. A. Zaman, L. T. Neustock, and L. Hesselink, "iLabs as an online laboratory platform: A case study at Stanford university during the COVID-19 pandemic," in *Proc. IEEE Global Eng. Educ. Conf. (EDUCON)*, Apr. 2021, pp. 1615–1623.
[44] L. T. Neustock, M. A. Zaman, G. K. Herring, F. J. Garcia Clemente, and L. Hesselink, "Scalable laboratory experimentation using iLabs—The digital twins for experiments," in *Proc. Educ. Training Opt. Photon. Conf.* Washington, DC, USA: Optica, 2021, p. F2A.7. [Online]. Available: http://opg.optica.org/abstract.cfm?URI=ETOP-2021-F2A.7
[45] L.-J. Thoms and R. Girwidz, "Experiments on optical spectrometry in the VirtualRemoteLab," in *Proc. 4th Exp.@Int. Conf.*, Jun. 2017, pp. 147–148.
[46] J. Bermúdez-Ortega, E. Besada-Portas, J. López-Orozco, J. Bonache-Seco, and J. De La Cruz, "Remote web-based control laboratory for mobile devices based on EJsS, Raspberry Pi and Node.js," *IFAC-PapersOnLine*, vol. 48, no. 29, pp. 158–163, 2015, [Online]. Available: https://www.sciencedirect.com/science/article/pii/S240589631502488X
[47] N. Wang, X. Chen, Q. Lan, G. Song, H. R. Parsaei, and S.-C. Ho, "A novel Wiki-based remote laboratory platform for engineering education," *IEEE Trans. Learn. Technol.*, vol. 10, no. 3, p. 1, Jan. 2016.
[48] R. B. Henriques, J. Pereira, F. Dias, H. Fernandes, A. S. Duarte, T. Pereira, and J. Fortunato, "A generic communication protocol for remote laboratories: An implementation on e-lab," in *Proc. 4th Int. Conf. Advancements Nucl. Instrum. Meas. Methods Appl. (ANIMMA)*, Apr. 2015, pp. 1–6.
[49] K. Henke, T. Vietzke, R. Hutschenreuter, and H.-D. Wuttke, "The remote lab cloud 'GOLDi-labs.net,'" in *Proc. 13th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, 2016, pp. 37–42.
[50] K. Henke, T. Vietzke, H.-D. Wuttke, and S. Ostendorff, "GOLDi-Grid of online lab devices Ilmenau," *Int. J. Online Eng.*, vol. 12, no. 4, pp. 11–13, 2016.
[51] T. Fáth, K. Henke, R. Hutschenreuter, F. Seidel, and H.-D. Wuttke, "On effective maintenance of distributed remote laboratories," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2018, pp. 80–89.
[52] H. A. Lahoud and J. P. Krichen, "Networking labs in the online environment: Indicators for success," *J. Technol. Stud.*, vol. 36, no. 2, pp. 31–40, May 2010.
[53] J. Sáenz, L. De La Torre, J. Chacon, and S. Dormido, "A study of strategies for developing online laboratories," *IEEE Trans. Learn. Technol.*, vol. 14, no. 6, pp. 777–787, Dec. 2021.
[54] P. Orduña, J. Garcia-Zubia, L. Rodriguez-Gil, I. Angulo, U. Hernandez-Jayo, O. Dziabenko, and D. López-de Ipiña, "The WebLab-Deusto remote laboratory management system architecture: Achieving scalability, interoperability, and federation of remote experimentation," in *Cyber-Physical Laboratories in Engineering and Science Education*. Cham, Switzerland: Springer, 2018, pp. 17–42, doi: 10.1007/978-3-319-76935-6_2.
[55] P. Orduña, L. Rodriguez-Gil, I. Angulo, U. Hernandez, A. Villar, and J. Garcia-Zubia, "WebLabLib: New approach for creating remote laboratories," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2019, pp. 477–488.
[56] M. Winzker, R. Kiessling, A. Schwandt, C. S. Paez, and S. A. Shanab, "Teaching across the ocean with video lectures and remote-lab," in *Proc. IEEE World Eng. Educ. Conf. (EDUNINE)*, Mar. 2018, pp. 1–4.
[57] M. Winzker and A. Schwandt, "Open education teaching unit for low-power design and FPGA image processing," in *Proc. IEEE Frontiers Educ. Conf. (FIE)*, Oct. 2019, pp. 1–9.
[58] A. Schwandt and M. Winzker, "Make it open—Improving usability and availability of an FPGA remote lab," in *Proc. IEEE Global Eng. Educ. Conf. (EDUCON)*, Apr. 2019, pp. 232–236.
[59] S. T. Marchisio and F. G. Lerro, "The FCEIA-UNR remote laboratory: Integration of resources and collaborative networking for Engineering Education [El laboratorio remoto FCEIA-UNR: Integración de recursos y trabajo en redes colaborativas para la enseñanza de la ingeniería]," in *Proc. Quinta Conferencia de Directores de Tecnología de Información TICAL 2015, Gestión de las TICs Para la Investigación y la Colaboración*, Oct. 2015. [Online]. Available: http://hdl.handle.net/10786/1001
[60] H. Benmohamed, A. Leleve, and P. Prevot, "Generic framework for remote laboratory integration," in *Proc. 6th Int. Conf. Inf. Technol. Based Higher Educ. Training*, 2005, pp. T2B/11–T2B/16.
[61] D. Lowe, S. Murray, L. Weber, M. De La Villefromoy, A. Johnston, E. Lindsay, W. Nageswaran, and A. Nafalski, "LabShare: Towards a national approach to laboratory sharing," in *Proc. 20th Annu. Conf. Australas. Assoc. Eng. Educ.* Adelaide, SA, Australia: Univ. Adelaide, 2009, pp. 458–463.
[62] H. Yeung, D. Lowe, and S. Murray, "Interoperability of remote laboratories systems," *Int. J. Online Biomed. Eng. (iJOE)*, vol. 6, no. 5, p. 71, Aug. 2010.
[63] A. Maiti, A. D. Maxwell, and A. A. Kist, "Features, trends and characteristics of remote access laboratory management systems," *Int. J. Online Eng.*, vol. 10, no. 2, pp. 30–37, 2014.

[64] P. Spilakova and F. Schauer, "Remote laboratory management system RemLabNet and its booking system," in *Proc. 4th Int. Conf. e-Technol. Netw. Develop. (ICeND)*, Sep. 2015, pp. 1–5.
[65] F. Schauer, M. Krbecek, P. Beno, M. Gerza, L. Palka, and P. Spilaková, "REMLABNET II—Open remote laboratory management system for university and secondary schools research based teaching," in *Proc. 12th Int. Conf. Remote Eng. Virtual Instrum. (REV)*, Feb. 2015, pp. 109–112.
[66] F. Schauer, M. Ozvoldova, M. Gerza, M. Krbecek, P. Beno, T. Komenda, S. Das, and M. I. Archibong, "REMLABNET IV—LTI federated remote laboratory management system with embedded multiparameter simulations," *Int. J. Online Eng. (iJOE)*, vol. 13, no. 10, p. 103, Nov. 2017.
[67] P. Orduña, L. Rodriguez-Gil, J. Garcia-Zubia, I. Angulo, U. Hernandez, and E. Azcuenaga, "LabsLand: A sharing economy platform to promote educational remote laboratories maintainability, sustainability and adoption," in *Proc. IEEE Frontiers Educ. Conf. (FIE)*, Oct. 2016, pp. 1–6.
[68] P. Orduña, L. Rodriguez-Gil, J. Garcia-Zubia, I. Angulo, U. Hernandez, and E. Azcuenaga, "Increasing the value of remote laboratory federations through an open sharing platform: LabsLand," in *Online Eng. Internet Things*. Springer, 2018, pp. 859–873.
[69] C. Salzmann and D. Gillet, "Challenges in remote laboratory sustainability," in *Proc. Int. Conf. Eng. Educ. (ICEE)*, 2007, pp. 1–6.
[70] L. Rodríguez-Gil, J. García-Zubia, P. Orduña, and D. López-de-Ipiña, "An open and scalable web-based interactive live-streaming architecture: The WILSP platform," *IEEE Access*, vol. 5, pp. 9842–9856, 2017.
[71] P. M. Kwinana, P. Nomnga, M. Rani, and M. L. Lekala, "Real laboratories available online: Establishment of ReVEL as a conceptual framework for implementing remote experimentation in South African higher education institutions and rural-based schools—A case study at the University of Fort Hare," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2020, pp. 128–142.
[72] P. A. Buitrago, R. Camacho, H. E. Pérez, O. Jaramillo, A. Villar-Martinez, L. Rodriíguez-Gil, and P. Orduña, "Mobile Arduino robot programming using a remote laboratory in UNAD: Pedagogic and technical aspects," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2020, pp. 171–183.
[73] C. Arguedas-Matarrita, M. Conejo-Villalobos, F. Ureña Elizondo, O. Barahona-Aguilar, P. Orduña, L. Rodriguez-Gil, U. Hernandez-Jayo, and J. García-Zubia, "Experience with the VISIR remote laboratory at the Universidad Estatal a Distancia (UNED)," in *Proc. Int. Conf. Remote Eng. Virtual Instrum.* Cham, Switzerland: Springer, 2020, pp. 162–170.
[74] R. Li, J. R. Morelock, and D. May, "A comparative study of an online lab using Labsland and zoom during COVID-19," *Adv. Eng. Educ.*, vol. 8, pp. 1–10, Dec. 2020.
[75] C. A. Mayoz, A. L. Da Silva Beraldo, A. Villar-Martinez, L. Rodriguez-Gil, W. F. M. De Souza Seron, and P. Orduña, "FPGA remote laboratory: Experience of a shared laboratory between UPNA and UNIFESP," in *Proc. 14th Technol. Appl. Electron. Teaching Conf. (TAEE)*, Jul. 2020, pp. 1–8.

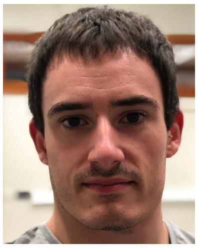

AITOR VILLAR-MARTÍNEZ received the Bachelor of Science degree in telecommunication engineering and the Master of Science degree in telecommunication engineering from the University of Deusto, in 2016 and 2018, respectively, where he is currently pursuing the Ph.D. degree, which involve work and research in relation to remote laboratories and the technologies employed in the latter. Since 2014, he has collaborated in various projects at the University. Since 2017, he has been working as a Hardware/Software Developer with LabsLand (spin-off of the WebLab-Deusto Project). Since 2018, he has been collaborating with the MORElab Research Group, DeustoTech—Deusto Institute of Technology.

JAVIER GARCÍA-ZUBÍA (Senior Member, IEEE) received the Ph.D. degree in computer science from the University of Deusto, Spain. He is currently a Full Professor with the Faculty of Engineering, University of Deusto. He is the Leader of the WebLab-Deusto Research Group. His research interest includes remote laboratory design, implementation, and evaluation.

IGNACIO ANGULO has been with the Department of Information Technology, Electronics and Communication, University of Deusto, since 2002. He has participated in over 25 research projects. He has collaborated in the writing of 32 scientific articles published in magazines of international impact. He presented his Ph.D. thesis on Open Architecture for the Deployment of Remote Laboratories on Embedded Systems, in 2015.

LUIS RODRÍGUEZ-GIL (Senior Member, IEEE) received the dual degree in computer engineering and industrial organization engineering, the M.Sc. degree in information security, and the Ph.D. degree in computer science from the University of Deusto, in 2013, 2014, and 2017, respectively. During his Ph.D., he co-founded the LabsLand remote labs company, where he currently works full-time as its CTO. Since 2009, he has been a part of the WebLab-Deusto Research Group, collaborating in the development of the WebLab-Deusto RLMS. Throughout this time, he has authored various peer-reviewed publications and contributed to several open source projects.