ALSET Lab: Designing Precise Timing and Communications for a Digital Power Grid Laboratory

Hossein Hooshyar
Power Delivery and Utilization Sector
Electric Power Research Institute
White Plains, NY, US
hhooshyar@epri.com

Luigi Vanfretti, Joe H. Chow
Randall Fitsik, Prottay M. Adhikari
Rensselaer Polytechnic Institute
Troy, NY, US
{vanfrl, chowj, fitsir, mondap2}@rpi.edu

***Abstract*— This paper describes the design and implementation of the ALSET Lab, a simulator-based laboratory that can host real-time hardware-in-the-loop experiments on time-critical applications for cyber-physical power networks. The paper makes emphasis on precise timing and communication networks, which can be used as a guide to build similar facilities.**

***Index Terms*—ALSET Lab, cyber-physical power system, real-time hardware-in-the-loop simulation.**

I. INTRODUCTION

Digitalization and renewable energy sources bring tremendous possibilities to revolutionize the way we produce, transport, distribute and consume electricity. However, they also bring tremendous challenges to maintain adequate resiliency and cyber-physical security of electrical power networks. Hence, new concepts, models, methods and tools are required to facilitate their operational design, and testing/verification/validation of their performance. In order to address these challenges, the Analysis Laboratory for Synchrophasor and Electrical energy Technology (ALSET Lab) was established in 2018. ALSET Lab is a simulator-based laboratory that can host real-time hardware-in-the-loop experiments on time-critical applications for cyber-physical power networks. This paper elaborates the design and implementation of precise timing and communications used in the ALSET Lab.

II. DESIGN OF THE ARCHITECTURE

Inspired by the Smart Grid Architecture Model (SGAM) [1], ALSET is designed in three main layers, as explained below. Each of the layers describes how the components in ALSET are interconnected from a particular point of view.

A. *Precise Time Layer*

This layer is all about supply and distribution of precise timing in the ALSET Lab. Through this layer, all equipment in ALSET has access to precise time with maximum inaccuracy of less than a microsecond. The overall view of this layer is shown in Figure 1. As the figure shows, the layer is partitioned in horizontal dimension into two domains of Simulator and Actual Hardware. In the vertical dimension, it spans the hierarchical levels of power system management, partitioned into six zones of Process, Field, Station, Operation (not shown), Enterprise (not shown) and Market (not shown). The layer displays: **1)** All equipment, engaged in the realization of this layer of the architecture, with their names and labels. Note that the labels can be used to search for more information in the ALSET inventory file on specific equipment, **2)** Information about the type and the length of the coax cables and the signals they carry, **3)** Information about the type of connectors that connect the coax cables to the equipment.

As shown in Figure 1, four antennas from Schweitzer Engineering Laboratories (SEL) provide the lab with Global Navigation Satellite System (GNSS) signals that are needed to extract the precise time [2]. While all four antennas receive signals from the GPS satellites, two of them are capable of receiving GLONASS signals as well to increase reliability. GLONASS is a Russian space-based satellite navigation system and provides an alternative to the American GPS system. The GNSS signals provided by the antennas are first passed through the surge protectors and are then fed to two satellite-synchronized network clocks from SEL and to one 1-to-16 GPS splitter from GPSSource [2,3]. The clocks extract the precise time out of the GNSS signals and make it available in different formats of PPS, modulated and demodulated IRIG(Inter-Range Instrumentation Group)-B signals. Unlike the clocks, the GPS splitter does not process the GPS signal, but it makes multiple copies of it. The precise time in the format of IRIG-B are then provided to all end users such as SEL relays and RTAC, lab server, OPAL-RT simulators, etc. [2,4]. In addition, there are some end users who need the raw GPS signal instead of IRIG-B, such as the CompactRIO Systems from National Instruments and the PMU from Arbiter Systems [5,6]. This is because they have their own GPS receivers and need to be fed from the GPS splitter.

B. *Communication Layer*

This layer is all about the Local Area Network (LAN) of the ALSET Lab. The LAN is utilized by the equipment to exchange different types of data. Since most data types are time-sensitive (i.e. they should be exchanged with minimal latency), VLANs (Virtual LANs) and QoS (Quality of Service) rules have been implemented to avoid any data congestion.

Having a VLAN for each data type means that the path, each data type travels within the LAN, has been virtually separated from that of the other data types. In ALSET, five VLANs, each assigned to one of the following data types, have been implemented: GOOSE, PTP (Precision Time Protocol), SV (Sampled-Value), PMU (Synchrophasor data), and Station. The VLAN *Station* is used for any data type that is not listed among the first four. For example, if one sends a

---

This work was supported in whole or in part by the ERC Program of the National Science Foundation and the Department of Energy under Award EEC-1041877, in part by GE Global Research on behalf of Nicholas Miller through the 2017 Edison Pioneer Award, in part by Dominion Energy, and in part by the ECSE Department at RPI.

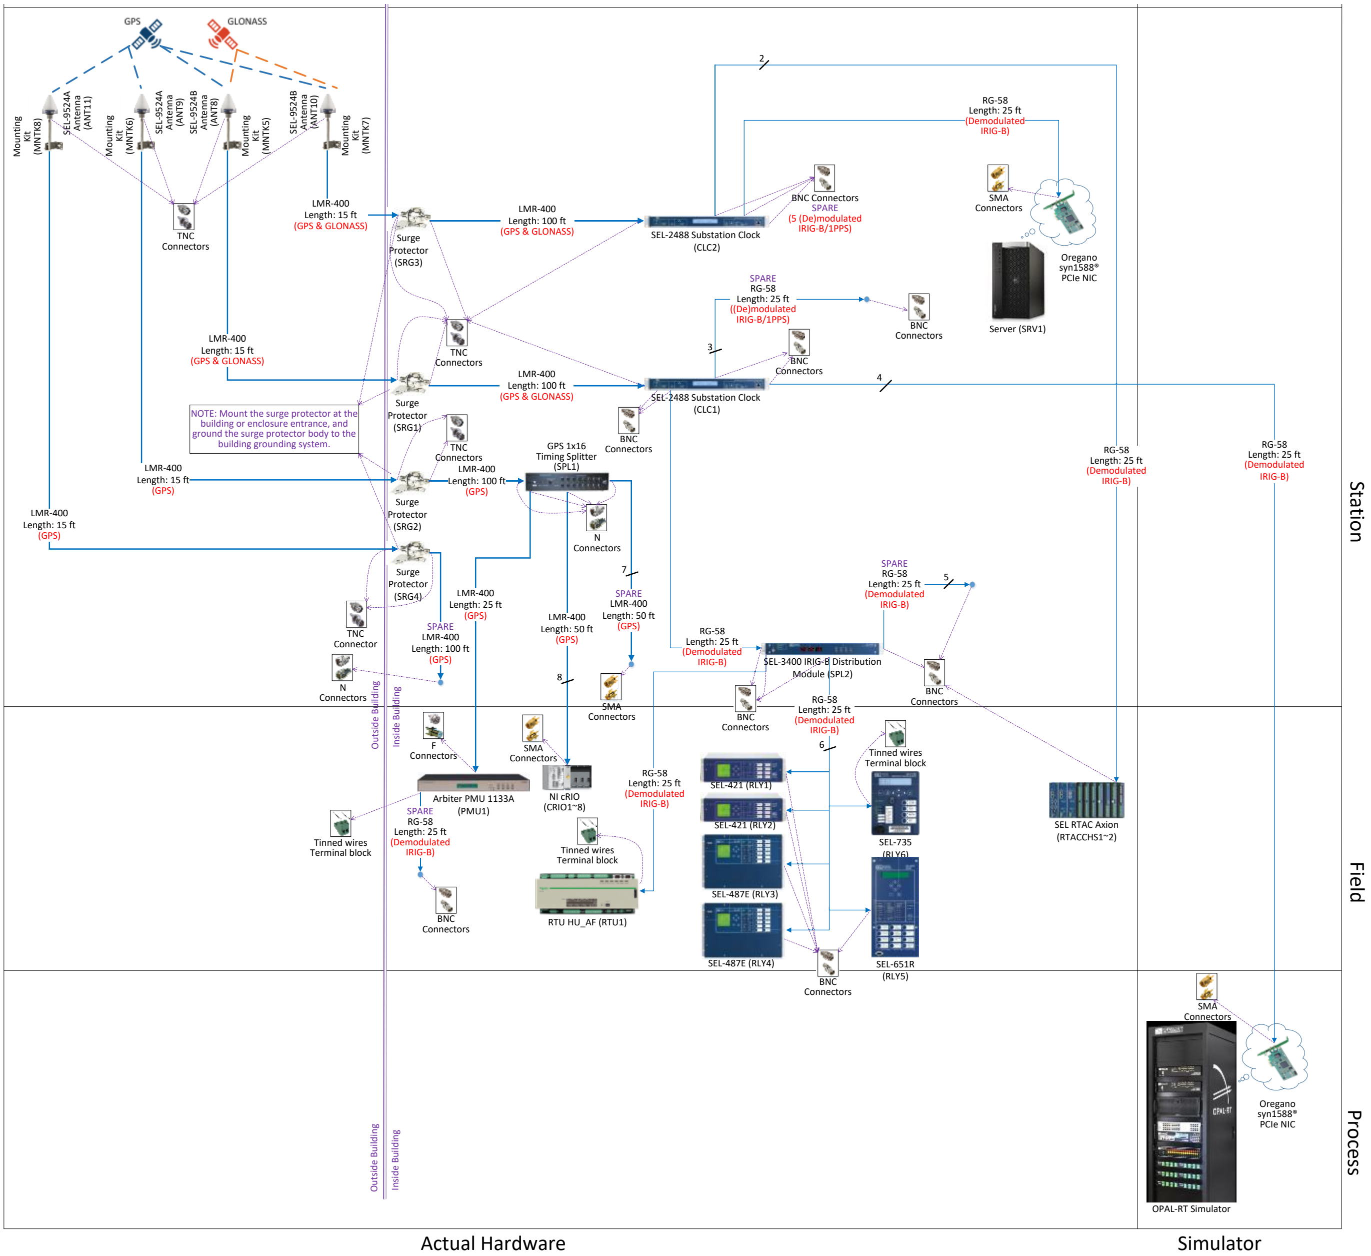

Figure 1. Precise time layer of the ALSET Lab architecture.

print to the lab printer or if one loads a model from the RT-Lab software running on his/her laptop to the OPAL-RT simulator, the data flows through the VLAN *Station*. In other words, while the data flowing through the other VLANs are tagged, the data flowing through the VLAN Station remains untagged.

The QoS rules, implemented in the ALSET LAN, basically consist of two set of constraints on the VLANs. The first one determines the priority of the VLANs, as shown in Table I. For example, *GOOSE* has the highest priority (i.e. the data flowing through this VLAN will be handled first by the equipment) and *Station* has the lowest one (i.e. the data flowing through this VLAN will be handled only if other data types are being handled and there is still processing power left). This implies that all equipment in ALSET are configured to respect this set of rules. The second set of rules is about the minimum bandwidth that should be guaranteed for each VLAN, as shown in the table. As seen in the table, while *PTP* occupies a small portion of the available bandwidth of the LAN (5.2kbps out of 1Gbps), *SV* demands nearly half of the available bandwidth (480Mbps out of 1Gbps). Note that, as can be inferred from the table, the sum of the allocated bandwidths is less than the 1 Gbps available bandwidth of the Cat5e Ethernet cables. The unallocated bandwidth remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority. The following subsections elaborate how the minimum bandwidth, required for each data type, is calculated.

1. *GOOSE*

GOOSE messages are typically 200 to 300 bytes long, so they can be serialized quickly [7]. As mentioned in the IEC 61850-8-1 Standard, a GOOSE message is retransmitted with gradually increasing retransmission time to achieve additional reliability [8]. According to the standard, the minimum “timeAllowedtoLive” of a GOOSE message, i.e. the time that the subscriber needs to wait for the next GOOSE message, is 1 ms. Note that, in practice, the “timeAllowedtoLive” is typically set to 4-5 ms or longer. Taking the "timeAllowedtoLive" set to 1 ms to obtain the maximum bandwidth required, one can compute the maximum number of retransmissions within a second, as shown in Table II. As the table shows, a GOOSE message would be retransmitted at most 11 times within a second. Targeting the worst-case scenario, let's assume three simultaneous events occurring in the system, each event triggering four IEDs, and each IED publishing four GOOSE messages. The bandwidth required for the worst-case scenario is computed as 3 (the number of simultaneous events) times 4 (the number of triggered IEDs per event) times 3 (the number of published GOOSE messages per IED) times 11 (the number of retransmissions per GOOSE message) times 300 (the number of bytes) times 8 (the number of bits per byte), that is equal to 950400 bps ≈ 1 Mbps.

2. *PTP*

PTP is an industry-standard protocol that enables the precise transfer of time to synchronize clocks over packet-based Ethernet networks [9]. There are two mechanisms used in PTP to measure the propagation delay between the clocks, the Delay Request_Response mechanism and the Peer Delay mechanism, each consisting of certain messages to keep the clocks synchronized. Table III lists the messages used in the mechanisms [10]. Note that the two mechanisms do not interwork on the same communication path; therefore, the total number of bytes exchanged in the Peer Delay mechanism is used to compute the required bandwidth. The bandwidth is computed as 646 (the number of bytes exchanged in the Peer Delay mechanism) times 8 (the number of bits per byte), that is equal to 5168 bps ≈ 5.2 kbps. Note that some messages are transmitted less frequent than once per second. For example, the Sync message is transmitted once every two seconds and the Delay_Req message once every minute; however, the bandwidth is computed assuming that all messages are transmitted once every second to account for some margin.

3. *SV*

Commercial products available on the market support SV messaging, as defined in the IEC 61850-9-2 Standard, through its "Light Edition", IEC 61850-9-2LE [11,12]. A SV message frame as specified in IEC 61850-9-2LE with a sampling rate of 80 samples per cycle for protection applications have an approximate size of 140 bytes, consuming a bandwidth of approximately 5 Mbps for systems operating at 50 Hz or 6 Mbps for system operating at 60 Hz, per source IED. For monitoring applications, the sampling rate is 256 samples per cycle, but each eight points are grouped and sent in a single packet in a lower rate, resulting in a bandwidth of up to 10Mbps for systems operating at 50 Hz and 12 Mbps for systems operating at 60 Hz [13]. Targeting the worst-case scenario, let's assume a communication path transmitting 40 simultaneous monitoring-class SV streams for a 60 Hz system. This assumption is in accordance with the capability of the OPAL-RT simulator as it can publish or subscribe to 40 SV streams per simulation core, provided a proper optimization and an adequate number of Ethernet interfaces. The bandwidth required for the worst-case scenario is computed as 12 (the bandwidth consumed by a monitoring-class SV stream) times 40 (the number of simultaneous SV streams), that is equal to 480 Mbps.

TABLE I. QOS RULES IMPLEMENTED IN THE COMMUNICATION LAYER

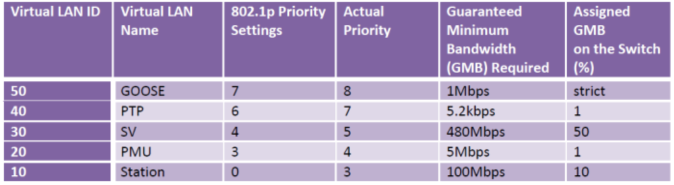

Virtual LAN ID	Virtual LAN Name	802.1p Priority Settings	Actual Priority	Guaranteed Minimum Bandwidth (GMB) Required	Assigned GMB on the Switch (%)
50	GOOSE	7	8	1Mbps	strict
40	PTP	6	7	5.2kbps	1
30	SV	4	5	480Mbps	50
20	PMU	3	4	5Mbps	1
10	Station	0	3	100Mbps	10

TABLE II. MAXIMUM NUMBER OF GOOSE MESSAGE RETRANSMISSION IN ONE SECOND

No. of re/transmissions	Instant of time (ms)
1	0
2	0×2+1 = 1
3	1×2+1 = 3
4	2×2+3 = 7
5	4×2+7 = 15
6	8×2+15 = 31
7	16×2+31 = 63
8	32×2+63 = 127
9	64×2+127 = 255
10	511
11	1023

TABLE III. MESSAGES EXCHANGED IN PTP MECHANISMS

Delay Request_Response Mechanism			
Type of Message	**Size of Message (bytes)**	**UDP/IP Overhead (bytes)**	**Total Size (bytes)**
Announce message	64	28	92
Sync message	44	28	72
Delay_Req message	44	28	72
Follow_Up message	44	28	72
Delay_Resp message	54	28	82
Signalling message	54	28	82
Management message	54	28	82
Sum			554
Peer Delay Mechanism			
Type of Message	**Size of Message (bytes)**	**UDP/IP Overhead (bytes)**	**Total Size (bytes)**
Announce message	64	28	92
Sync message	44	28	72
Follow_Up message	44	28	72
Pdelay_Req message	54	28	82
Pdelay_Resp message	54	28	82
Pdelay_Resp_Follow_Up message	54	28	82
Signalling message	54	28	82
Management message	54	28	82
Sum			646

TABLE IV. PMU DATA FRAME

Field	**Description**	**No. of bytes**
SYNC	Sync byte followed by frame type and version number.	2
FRAMESIZE	Number of bytes in frame.	2
IDCODE	Stream source ID number.	2
SOC	SOC time stamp.	4
FRACSEC	Fraction of second and time quality.	4
STAT	Bit-mapped flags.	2
PHASORS	Phasor of phase current in floating-point format.	48
FREQ	Frequency in floating-point format.	4
DFREQ	ROCOF in floating-point format.	4
CHK	CRC-CCITT	2
Total Number of Bytes		**74**
UDP/IP Overhead (bytes)		**28**
Total Size (bytes)		**102**

4. *PMU*

A PMU data frame, containing three current phasors and three voltage phasors, would be 74 bytes long, as defined in the IEEE C37.118.2 Standard [14,15]. The structure of the data frame is detailed in Table IV. Furthermore, the communication media, i.e. UDP/IP, adds 28 bytes to each PMU data frame. Targeting the worst-case scenario, let's assume a communication path loaded with 100 streams with reporting rate of 60 frames per second. The bandwidth required for the worst-case scenario is computed as 102 (the total number of bytes of each PMU data frame) times 60 (the reporting rate) times 100 (the number of streams) times 8 (the number of bits per byte), that is equal to 4896000 bps ≈ 5 Mbps.

5. Station

The 100 Mbps bandwidth is an arbitrary value reserved for the VLAN *Station* to support fast exchange of untagged data with high volumes. In addition, this is to prevent applications, generating lowest-priority traffic in *Station*, to be "starved" by high volumes of higher-priority traffics in other VLANs.

The overall view of the communication layer is shown in Figure 2. The layer displays: **1)** All equipment engaged in the realization of this layer of the architecture. Like the precise time layer, one can see both names and labels of the equipment, **2)** The VLANs each equipment has access to, **3)** The IP address(es) assigned to each equipment, **4)** The ports of the main switch to which the equipment is connected. As the figure shows, each equipment has one or multiple IP addresses. The number of IP addresses assigned to each equipment depends on the number of VLANs it has access to and on how VLAN processing is implemented in that particular equipment. For example, while protection relays from SEL process multiple VLANs by only one network card (requiring one IP address), the server with Redhat operating system requires multiple virtual network cards each processing one of the VLANs (requiring multiple IP addresses).

C. Component Layer

This layer describes the electrical connections in the ALSET Lab (not shown here due to space limit). Most of the equipment including the main network switch, the network clocks, the server, the relays, and the OPAL-RT simulators are fed by the APC Smart-SMT3000C UPS through three 120V outputs [16]. Note that most of the equipment are electrically connected to the OPAL-RT simulator via their low-level interfaces or communication interfaces. This is to remove the need for having voltage/current amplifiers in the loop, as the high current/voltage is risky to handle.

III. IMPLEMENTATION OF THE ARCHITECTURE

This section describes the common issues faced to implement the architecture, described in Section II.

A. Precise Time Layer

The implementation of this layer is quite straightforward; however, it's worth noting the following:

- It is important to use appropriate type of coax cable, ideally LMR-400, for transmission of GNSS signals. Other types of coax cable such as RG-58 are not suitable as their attenuation is severe at the frequency of the GNSS signals, which, in turn, leads to timing errors (intermittent loss of GPS/GLONASS). Note that RG-58 is still good enough for transmitting signals with lower frequencies such as IRIG-B.
- OPAL-RT simulators are synchronized through Oregano Systems syn1588® PCIe network cards. If the cards are fed with IRIG-B signals from the network clock, the 'Time Reference' in the network clock should be set as 'UTC' and not 'Local'. This is because the cards, as they are configured, cannot distinguish the offset between the local time and the UTC time. This is not an issue with other end users such as the protection relays, and they can be safely fed with the local time.

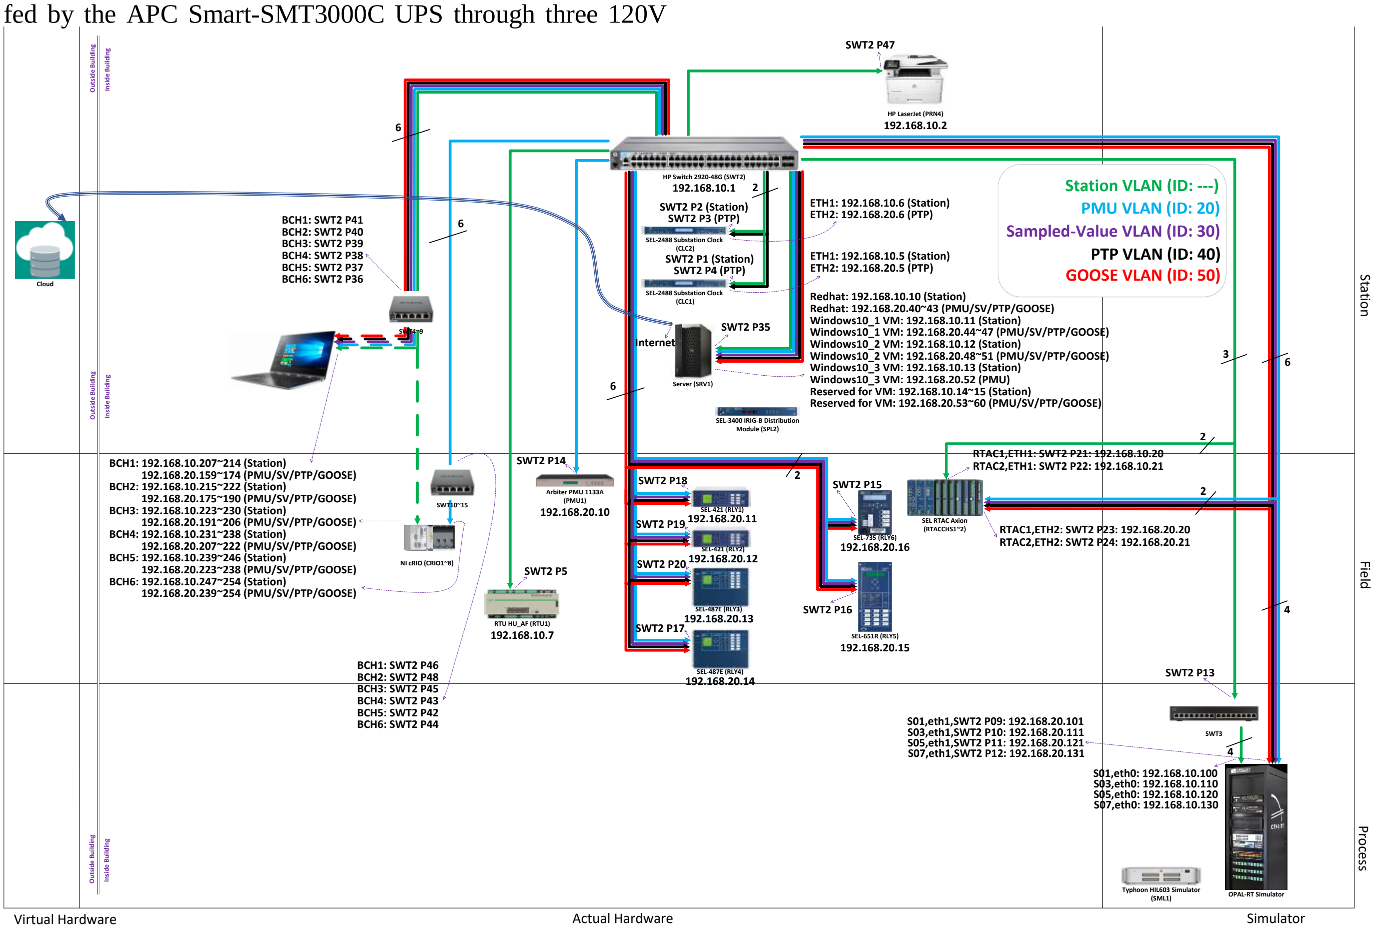

Figure 2. Communication layer of the ALSET Lab architecture.

B. *Communication Layer*

The implementation of this layer is often the most challenging one. This is because while the QoS rules should be implemented on the main network switch, the definition of VLANs should be implemented on both the main network switch and the equipment. However,

- Not all equipment support the VLAN technology. For example, while Windows-based computers fully support VLAN-tagging, the protection relays from SEL do it for SV, PTP, and GOOSE messages, but not for PMU data. In addition, other equipment such as the PMU from Arbiter Systems has no VLAN-tagging capability.
- Different equipment support the VLAN technology differently. For example, while Windows-based computers process each VLAN through a separate virtual network card, the protection relays from SEL do it through one network card that handles multiple VLANs.

Due to the above-mentioned complexities, the authors suggest the following for the implementation of this layer:

1. It is highly advisable to configure the (virtual) network cards, processing VLAN-tagged data, to be on a different subnet than the network cards handling untagged data. For example, in ALSET, while the IPs assigned to the network cards processing VLANs *GOOSE*, *PTP*, *SV*, and *PMU* are all 192.168.**20**.XXX (subnet mask: 255.255.255.0), the IPs assigned to those handling *Station* are 192.168.**10**.XXX (subnet mask: 255.255.255.0).
2. If an equipment does not fully support VLAN-tagging (i.e. it does not tag one or more of the following data types: synchrophasor, SV, PTP, GOOSE), additional programming is needed on the Ethernet port of the main network switch, to which the equipment is connected, to fully implement the VLANs, as described below:
   a) If the equipment ***partially*** supports VLAN-tagging (e.g. the SEL relays tag GOOSE messages but do not tag synchrophasors), close the Ethernet port to VLAN Station and make the VLAN corresponding to the unsupported data type (e.g. PMU) accept untagged data. This will configure the switch to VLAN-tag the untagged data. Note that with this configuration, any untagged data coming out of the equipment will be tagged by the switch; so different media should be used for other untagged communication (such as equipment configuration messages). For example, in ALSET, the SEL relays are configured through the serial ports.
   b) If the equipment does not support VLAN-tagging but it is supposed to exchange ***both*** untagged and tagged communication messages, the same approach as explained in the scenario ‘a’ must be taken. However, if the untagged communication is very limited, it is still okay to let the untagged communication get tagged (e.g. the Arbiter PMU does not support VLAN-tagging at all. Both the PMU data and configuration messages are exchanged through its Ethernet port, but the configuration is done only once and is not supposed to change frequently).
   c) If the equipment does not support VLAN-tagging and it is supposed to exchange ***only*** untagged communication messages (e.g. the printer), close the Ethernet port to all VLANs except the VLAN *Station*.

C. *Component Layer*

The implementation of this layer is quite straightforward and does not need any special consideration. However, the authors suggest implementing a patch panel to electrically interface the real-time simulators with the other equipment, as the patch panel considerably saves time in preparation of the hardware-in-the-loop simulation setups.

IV. CONCLUSIONS

The design and implementation process of a simulator-based facility that can host real-time hardware-in-the-loop experiments for time-critical applications was presented in this paper. The paper can act as a guide for other researchers who would like to build similar facilities, in particular, when considering precise timing and communication needs.

ACKNOWLEDGMENT

The authors would also like to acknowledge the generous support of Dr. Greg Zweigle from Schweitzer Engineering Laboratories; and Drs. Matthew Gardner, Kevin Jones and Mr. Kyle Thomas from Dominion Energy; for facilitating generous donations in the form of equipment for our laboratory. The kind assistance of Mr. Jerry Dziuba during the lab implementation is also gratefully acknowledged.

REFERENCES

[1] Smart Grid Coordination Group, *Smart Grid Reference Architecture*, CEN-CENELEC-ETSI, Tech. Rep., 2012.

[2] SEL Products, Available online at https://selinc.com/

[3] GPS Rack Mount Splitter 1 x 16 outputs - RMS116, Available online at https://www.gpssource.com/products/rms116-1x16-rack-mount-gps-splitter

[4] OPAL-RT Simulator, Available online at https://www.opal-rt.com/simulation-systems-overview/

[5] NI CompactRIO Systems, Available online at http://www.ni.com/en-us/shop/compactrio.html

[6] Arbiter Systems 1133A Power Sentinel™, Available online at https://www.arbiter.com/catalog/product/model-1133a-power-sentinel.php

[7] J. Wen, C. Hammond and E. A. Udren, "Wide-area Ethernet network configuration for system protection messaging," *65th Annual Conference for Protective Relay Engineers*, College Station, TX, 2012.

[8] *Communication networks and systems for power utility automation - Part 8-1: Specific communication service mapping (SCSM) - Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3*, IEC 61850-8-1, 2011.

[9] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE 1588, 2008.

[10] “Implementing IEEE 1588v2 for use in the mobile backhaul,” *Calnex Solutions Ltd Technical Brief*, 2009.

[11] *Communication networks and systems for power utility automation - Part 9-2: Specific communication service mapping (SCSM) - Sampled values over ISO/IEC 8802-3*, IEC 61850-9-2, 2011.

[12] *Implementation guideline for digital interface to instrument transformers using IEC 61850-9-2*, UCA Interntnl Users Group, 2004.

[13] L. B. Oliviera, M. Zapella, A. Sarda, W. Zanatta, “Managed Ethernet switches performance over IEC 61850 networks: Applications with high traffic flow,” *13th International Conference on Development in Power System Protection (DPSP)*, Edinburgh, UK, 2016, pp. 1-6.

[14] *IEEE Standard for Synchrophasor Data Transfer for Power Systems*, IEEE Std C37.118.2, 2011.

[15] H. Hooshyar, M. E. Baran, S. R. Firouzi, L. Vanfretti, “PMU-assisted overcurrent protection for distribution feeders employing solid state transformers,” *Elsevier Journal of Sustainable Energy, Grids and Networks (SEGAN)*, vol. 10, pp. 26-34, June 2017.

[16] APC SMT3000C UPS, Available online at https://www.apc.com/shop/us/en/products/APC-Smart-UPS-3000VA-LCD-120V-with-SmartConnect/P-SMT3000C