

# Experimental Validation of the Usability of Wi-Fi over Redundant Paths for Streaming Phasor Data

Maaz Mohiuddin, Miroslav Popovic, Athanasios Giannakopoulos, Jean-Yves Le Boudec  
School of Computer Science and Communication Systems  
École Polytechnique Fédérale de Lausanne, Switzerland  
{firstname.lastname}@epfl.ch

**Abstract**—Applications performing streaming of phasor-measurement data require low latency and losses from the communication network. Traditionally, such requirements are realized through wired infrastructure. Recently, wireless infrastructure has gained attention due to its low-cost and ease of deployment, but its poor quality-of-service is a strong deterrent for use in mission-critical applications. Recent studies have used measurements to explore the use of packet replication over redundant Wi-Fi paths, for obtaining the desired loss performance without hampering the end-to-end latency. However, these studies are done in a controlled, laboratory environment and do not reflect the real, in-field performance. In this paper, we perform extensive measurements using two co-located directional Wi-Fi links in a real-life setting, to experimentally validate the use of packet replication over Wi-Fi for streaming phasor data. In the setting that we evaluated, we find that the two channels are not fail-independent but the performance achieved with replication is very close to what it would be if they were to be independent. From the loss and latency statistics after replication, we conclude that replicating the phasor data over redundant Wi-Fi paths is a viable option for achieving the desired quality-of-service.

## I. INTRODUCTION

Real-time streaming of phasor measurements is often mission-critical and demands extremely low packet-loss rates (PLRs) and latency from the communication network [1]. Although applications such as energy metering can tolerate a PLR of  $10^{-3}$  and latencies up to 1 s [2], applications such as phasor measurement unit (PMU) based state-estimation and wide-area protection require a PLR of  $10^{-5}$  and latencies  $< 4$  ms [3]. Also, such mission-critical PMU-streaming applications require an availability greater than five-9's ( $> 0.99999$ ).

Traditionally, this is realized using wired infrastructure. However, wired networks are often laden with slow deployment and high-installation costs. Consequently, wireless infrastructure, Wi-Fi (IEEE 802.11) in particular, has recently gained traction [4], [5]. The main reserve in the use of wireless infrastructure is the low quality-of-service (QoS) in terms of losses and latency, due to low reliability of wireless links.

In the literature, the desired QoS target is achieved by replication of packets over two or more fail-independent paths using PRP at the MAC-layer [6], [7] or iPRP at the IP-layer [8], [9]. Needless to say, these protocols rely on the existence of fail-independent paths to reap maximum benefits from the replication. In wired networks, fail-independent paths are obtained by using physically separated networks or by appropriate routing rules in a larger, parent network. Hence, by ensuring that the replicas of the packets share no common

links or devices, it is easy to guarantee that the loss of one packet does not affect the reception of its replicas.

Obtaining fail-independent paths, however, is no longer trivial in wireless networks, as the Wi-Fi links share a common medium. Consequently, replication of packets over redundant Wi-Fi does not necessarily guarantee that the QoS requirements of mission-critical PMU-streaming applications will be satisfied. Hence, to be able to use Wi-Fi for PMU-streaming applications, the performance of parallel redundancy protocols over Wi-Fi needs to be experimentally characterized and compared against the desired QoS requirements.

Recently, simulation- [10]–[12] and measurement-based [13], [14] studies were employed to quantify the loss and latency performance of packet replication over redundant Wi-Fi paths. The simulation-based studies fail to capture the real performance of redundant Wi-Fi paths and the measurement-based studies have two shortcomings. First, the measurements were conducted in a laboratory environment and do not represent a real-life setting. Second, the measurement studies focus on generic real-time applications. Thus, the traffic profiles used are quite different from those of PMU-streaming applications, thereby rendering these studies non-representative.

We perform measurements on a test bed that is designed to closely imitate a real-life deployment of a campus-wide active distribution-network that uses PMU-based state estimation as described in [15]. Concretely, using commodity hardware, we designed a test bed that uses Wi-Fi technology (IEEE 802.11b standard). The test bed consists of nodes communicating with each other by using two spatially co-located Wi-Fi links that use directional antennas. The traffic profile used in the test bed is as same as that of PMUs in the distribution-network. Furthermore, the sending and receiving nodes are placed at the same locations as the PMUs and phasor data concentrator (PDC) in the distribution-network.

Using the measurements, we quantify the PLR, end-to-end latency, jitter in latency and availability of the effective channel, as perceived by the receiving application, after replication. Using statistical inference techniques, we verify if the two links are truly fail-independent. From the setting we evaluated, we conclude that the losses on the two links are in fact not independent. However, the effective PLR is similar to what it would be if they were to be independent. Consequently, we conclude from the setting we evaluated, that using replication over redundant Wi-Fi paths is a viable option for streaming

mission-critical PMU data.

The rest of the paper is organized as follows. In Section II, we review the related work. In Section III, we describe the experimental setup and methodology. We present the results of the measurements, in Section IV, and the statistical analysis of fail-independence of the links in Section V. We make our concluding remarks in Section VI.

## II. RELATED WORK

In the literature, the loss and latency performance of individual Wi-Fi links has been extensively studied through measurements [16], [17]. However, there exist very few measurement-based studies [13], [14] that characterize the performance of replicated Wi-Fi links in the context of real-time applications. Both these measurement-based studies [13], [14] share similar shortcomings that render the results non-representative of streaming applications for mission-critical PMU data.

First, the measurements in [13], [14] were conducted under a controlled, laboratory environment that is radically different from an in-field deployment. In contrast, our measurements were conducted on campus roof-tops using directional antennas. The locations of the measurement sites are the same as that of the PMUs in an existing campus-wide distribution-network that relies on the mission-critical PMU data [15].

Second, as the existing studies do not focus on validating the feasibility of using Wi-Fi for PMU-based applications, the traffic profiles used are quite different from those of PMU-streaming applications. Incidentally, the main finding of these papers is that the PLR is strongly dependent on the traffic profile. In light of these results, we employ the same traffic profile as used by the PMU-based state-estimation in [15]. Recently, traffic from streaming applications was used for measurements in a similar study [18]. However, we perform a more formal statistical analysis on the data to test fail-independence of the two links.

Lastly, in contrast to the short duration (of a few days) of these measurement campaigns, our measurements were conducted over period of 45 days, thereby increasing the amount of data at hand. Longer duration also captures a wider spectrum of fading effects, electro-magnetic disturbances, cross-talk, etc.; they are likely to surface in real deployments.

## III. EXPERIMENTAL SETUP

### A. Description of the test bed

Figure 1 shows the map of the EPFL campus with the three roof-top measurement sites, namely A, B and C. The distances between sending and receiving antennas are 180 m (site B to site A) and 230 m (site C to site A) and there is a line of sight between the sending and receiving sites. We use directional antennas (shown in Figure 2) for transmission and reception, a common practice when Wi-Fi is used as a replacement for a cable in mission-critical networks. The antennas have an 8 dBi gain. The test bed is based on 2.4 GHz Wi-Fi-technology (802.11b standard) and we use different channels and polarizations for the two links to minimize the mutual interference between them.

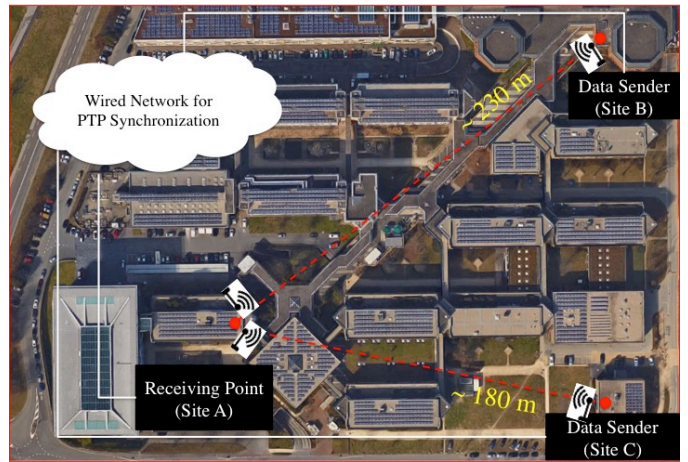


Fig. 1: Map of the campus with the antenna locations.




Fig. 2: Hardware used for the test bed.

Sites B and C are used for traffic generation, and site A is used for reception and logging. The locations of sending nodes at sites B and C is same as that of the PMUs at locations ELL and CM in the campus smart-grid described in [15]. Also, site A is present at the same location as the PDC in [15].

For reception, at site A, we use a ruggedized PC<sup>1</sup>, equipped with two Wi-Fi cards to support the two desired wireless links. For sending, at sites B and C, we have two Alix2d2 system boards<sup>2</sup>. The ruggedized PC runs 64-bit Ubuntu operating system and Alix2d2's run OpenWrt 10.3<sup>3</sup> operating system. The machines used in the test bed are shown in Figure 2.

For logging all packets exchanged over both the Wi-Fi links, we use the packet capture tool `tcpdump`<sup>4</sup> on sites A, B and C. To measure the end-to-end latency in transmission of the packets, we have time synchronization among the machines. To this end, we connect the machines through the wired, campus network that is used for network time-synchronization with the precision time protocol (PTP). In this way, we can compute the end-to-end latency with a 0.1 ms accuracy.

### B. Experimental Methodology

In order to study the effect of different Wi-Fi parameters (such as raw data-rate, beacon interval, choice of the channel

<sup>1</sup><http://www.logicsupply.com/da-1000/>

<sup>2</sup><http://www.pcengines.ch/alix2d2.htm>

<sup>3</sup><https://openwrt.org/>

<sup>4</sup><http://www.tcpdump.org/>

number) on the recorded PLR, we carried out measurements by varying different parameters. Through preliminary measurements, we observed that the choice of the channel number and the beacon interval have practically no effect on the observed PLR and latency. Hence, we omit these factors from the analysis that follows.

The two factors taken into account in our analysis are (1) MAC-layer retransmission enabled/disabled, and (2) the raw data-rate. MAC-layer retransmissions are expected to reduce the PLR because a lost packet might be repaired by the MAC-layer in subsequent transmissions. However, it is worth noting that each retransmission adds to the end-to-end latency. Hence, receiving a packet as a result of MAC-layer retransmissions is a case of trading-off latency for PLR. This trade-off is further elaborated on in Section IV.

Last, we analyze three different raw data-rates supported by the standard: 1 Mbps, 5.5 Mbps and 11 Mbps. In theory, lower raw data-rates are more robust, due to a higher degree of redundancy from channel coding. Hence, the effect of channel fading is expected to be lower for lower data rates thereby resulting in a lower overall packet-loss probability.

These factors translate to conducting experiments in six different scenarios: presence or absence of MAC-layer retransmissions, and raw data-rates of 1 Mbps, 5.5 Mbps or 11 Mbps. We randomize these parameters in the measurement scenarios so as to normalize the bias due to individual parameters. Each scenario lasts for nearly 30 minutes, with both the senders sending approximately 90,000 UDP packets of size 300 bytes each, with a packet every 20 ms. This traffic profile is the same as that of the PMUs that stream mission-critical measurements used for state estimation in [15].

The measurement campaign lasted for 45 days resulting in approximately 1500 scenarios, which amounts to 300 million packets being transmitted over the two replicated paths. Furthermore, there were 250 instances of each scenario amounting to nearly 50 million packets for each scenario.

The packets from both links AB and AC are independently labeled with an increasing sequence number (SN) and a timestamp of the instant at which each packet was generated ( $t_g$ ). The time at which the packet was actually sent ( $t_s$ ) is recorded by `tcpdump` at the sender. Note that, due to possible non-negligible delays in processing of the packet by the operating system,  $t_s$  is not necessarily equal to  $t_g$ . As the goal of this measurement campaign is to characterize the latency due to Wi-Fi links, using  $t_s$  is more apt than using  $t_g$ . Hence, the use of `tcpdump` in our measurement test-bed.

At the receiver, for each received packet, the time of reception ( $t_r$ ) is also logged by `tcpdump`. In the post-processing, the logs of links AB and AC are both analyzed for losses and latency, as described in Section IV.

#### IV. MEASUREMENT RESULTS

From the measurements of each scenario  $i$ , we evaluate the PLR for link AB ( $P_{ab}^i$ ) and link AC ( $P_{ac}^i$ ). We compute the availability (A) as  $A = MTF/(MTTR + MTF) = 1 - \text{PLR}$ , where MTF and MTTR are the mean-times to

failure and recovery, respectively. As a result, we obtain availability for link AB ( $A_{ab}^i$ ) and link AC ( $A_{ac}^i$ ).

Consider two packets with SNs  $x$  and  $y$  sent over links AB and AC, respectively. If  $x = y$ , then the two packets can be regarded as belonging to the same “generation”, i.e., they mimic replicas of each other. Consequently, packet  $x$  is said to be lost after replication, if and only the packet is lost on both links AB and AC. As a result, for each scenario  $i$ , we obtain the effective PLR as seen by the receiver after replication ( $P_{rep}^i$ ) and the corresponding availability ( $A_{rep}^i$ ).

Table I shows mean, 95<sup>th</sup> percentile (95%-ile) and 99<sup>th</sup> percentile (99%-ile) values of  $P_{ab}$ ,  $P_{ac}$ ,  $P_{rep}$ . We find that the mean PLRs for individual links AB and AC are  $9.69 \times 10^{-4}$  and  $2.4 \times 10^{-3}$  respectively. As expected, these values are much higher than  $10^{-5}$ , the PLR required for mission-critical PMU-streaming applications. However, the mean  $P_{rep}$  is  $3.58 \times 10^{-6}$ , is well within the acceptable value for streaming of mission-critical PMU data.

	Mean	95%-ile	99%-ile
$P_{ab}$	$9.69 \times 10^{-4}$	0.0013	0.0293
$P_{ac}$	$2.4 \times 10^{-3}$	0.0191	0.0466
$P_{rep}$	$3.58 \times 10^{-6}$	$1.10 \times 10^{-5}$	$2.22 \times 10^{-5}$

TABLE I: PLR statistics obtained from measurements

The performance improvement achieved by replication over redundant Wi-Fi paths is more prominent when either of the individual links experiences high losses. We observe, in Table I, that the 99%-ile values of  $P_{ab}$ ,  $P_{ac}$  and  $P_{rep}$  are 0.0293, 0.0466 and  $2.22 \times 10^{-5}$ , respectively. We see that even the 99%-ile value of  $P_{rep}$  is comparable with the desired PLR of the PMU-streaming applications.

Furthermore, we find the mean availability  $A_{rep} = 0.999996$  and its 99%-ile value is 0.99997, which is comparable to the five-9’s requirement of mission-critical PMU-streaming applications. Next, we study the variation of PLR as a function of different parameters and check if the availability can be improved by tuning the parameters of the Wi-Fi links.

##### A. Effect of MAC-Layer Retransmission

We studied the effect of MAC-layer retransmissions on loss and latency. Figure 3 shows the box plot of  $P_{ab}$ ,  $P_{ac}$  and  $P_{rep}$ , with and without MAC-layer retransmissions. Table II shows the 99%-ile PLR values with and without MAC-layer retransmissions. Although the use of MAC-layer retransmissions provides an order of magnitude improvement in PLRs over individual links, we find absolutely no effective losses over the redundant Wi-Fi paths among the 50 million transmitted. This strongly asserts the use of MAC-layer retransmissions for streaming mission-critical PMU applications.

Next, we quantify the end-to-end latency and verify if the latency performance with MAC-layer retransmissions conforms to the requirements of the mission-critical PMU-streaming applications (about 4 ms). For this purpose, let the delay due to packet with SN  $x$  on link AB is given by  $d_{ab}^x = t_r^{x,ab} - t_s^{x,ab}$ . Similarly,  $d_{ac}^x = t_r^{x,ac} - t_s^{x,ac}$ . Then, the effective delay for packet  $x$  after replication is given by  $d_{rep}^x = \min(d_{ab}^x, d_{ac}^x)$ .

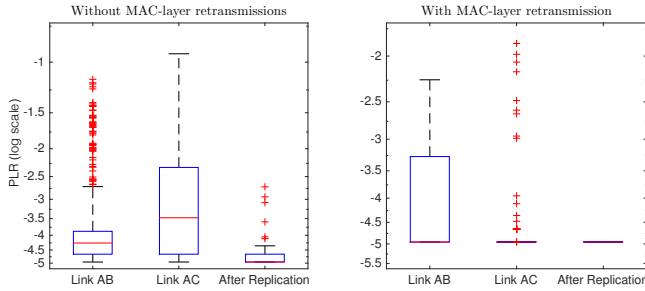


Fig. 3: Box plot of the PLRs for link AB, link AC and after replication for scenarios with and without MAC-layer retransmissions, shown in log-scale.

	Without Retransmissions	With Retransmissions
$P_{ab,99\%}$	0.0402	0.0010
$P_{ac,99\%}$	0.0610	0.0018
$P_{rep,99\%}$	$4.43 \times 10^{-5}$	0

TABLE II: 99%-ile value of PLRs with and without MAC-layer retransmissions

Table III shows the mean and quantiles of  $d_{ab}$ ,  $d_{ac}$  and  $d_{rep}$ . For most packets, the latencies of both the individual links and the latency after replication is within the admissible latency of 4 ms, required by mission-critical PMU-streaming applications. However, the tail latencies at 99.9%-ile and 99.99%-ile for individual links exceed the 4 ms mark, indicating that although very rare, there are cases where the real-time requirements are not respected by individual links. Additionally, replication over redundant paths brings down the delay for most of the packets to within the admissible range, barring a minuscule fraction of  $10^{-4}$ .

We find that the mean jitter in  $d_{ab}$ ,  $d_{ac}$  is 0.455 ms and 0.055 ms, whereas that in  $d_{rep}$  is 0.025 ms. Hence, although individual Wi-Fi links have a high jitter in latency, which is undesirable for PMU-streaming applications, replication over redundant Wi-Fi paths significantly reduces the jitter in latency.

Hence, from the setting we studied, we find that using MAC-layer retransmissions satisfies the loss- and latency-requirements of mission-critical PMU-streaming applications. Thus, we conclude that replicating packets over directional Wi-Fi links is a viable option for streaming PMU measurements.

### B. Effect of Raw Data-Rate

To study the effect of raw data-rate on the PLR over directional Wi-Fi links, we performed experiments with three data-rates: 1 Mbps, 5.5 Mbps and 11 Mbps. As lower data-rates have a higher degree of redundancy due to channel-coding, the PLR for a lower data-rate is expected to be lower than that of higher data-rates.

We found that the mean  $P_{rep}$  for scenarios with raw data-rate of 1 Mbps, 5.5 Mbps and 11 Mbps are  $2.21 \times 10^{-6}$ ,  $1.77 \times 10^{-6}$  and  $6.52 \times 10^{-6}$  respectively. The 99%-ile values of the same were  $2.04 \times 10^{-5}$ ,  $1.11 \times 10^{-5}$  and  $3.60 \times 10^{-5}$

	Mean	95%-ile	99%-ile	99.9%-ile	99.99%-ile
$d_{ab}$ (in ms)	0.87	2.58	3.82	4.94	6.32
$d_{ac}$ (in ms)	0.61	0.73	1.26	4.2	6.67
$d_{rep}$ (in ms)	0.58	0.59	0.92	2.69	4.3

TABLE III: Latency statistics with MAC-layer retransmissions

respectively. We find that the PLR for 1 Mbps and 5.5 Mbps are quite close to each other, whereas the ones for 11 Mbps are slightly worse. As expected, we suggest the use of a lower raw data-rate.

### V. FAIL-INDEPENDENCE TEST

In this section, we characterize the correlation between losses on the two wireless links AB and AC in order to verify whether the directional Wi-Fi links are truly fail-independent.

As a first step towards verifying whether the losses on both the links are statistically correlated, we evaluate the cross-correlation co-efficient of losses on both links. We find the cross-correlation co-efficient to be -0.0005 with a standard deviation of 0.0504. The coefficient of variation is  $1.117 \times 10^3$ , indicating that this first-order statistic cannot be used for a conclusive answer and a formal statistical test, such as the likelihood-ratio test [19] needs to be applied.

The prerequisite for the application of the likelihood-ratio test is the knowledge of the distribution of losses. Furthermore, to be able to apply the test, the distributions of losses on the links AB and AC need to be stationary. The estimation of these distributions for link AB and link AC, and their stationarity is discussed below.

#### A. Estimating the Distribution of Losses

We know from the literature, that packet losses over wireless links are bursty, i.e., losses are correlated in time. We consider the two-state Gilbert model [20], for representing the observed losses over individual wireless links. Figure 4 shows the two-state Markov chain representing the Gilbert model. It consists of the Good state (1) where no packets are lost and the Bad state (0) where the link drops all packets. The transition probabilities from Good state to Bad state, and from Bad state to Good state, are  $q$  and  $p$ , respectively. Then, the average PLR is given by  $q/(p+q)$ .




Fig. 4: Two-state Gilbert model for bursty losses

From the measurement, for each scenario  $i$ , we obtain the parameters of the two-state Gilbert model for link AB and AC:  $p_{ab}^i, q_{ab}^i, p_{ac}^i, q_{ac}^i$ . For link AB, we find that the mean  $p_{ab}$  ( $\hat{p}_{ab}$ ) is 0.4217 and mean  $q_{ab}$  ( $\hat{q}_{ab}$ ) is  $6.99 \times 10^{-5}$ . For link AC, we find that the mean  $p_{ac}$  ( $\hat{p}_{ac}$ ) is 0.2962 and mean  $q_{ac}$  ( $\hat{q}_{ac}$ ) is  $3.27 \times 10^{-5}$ . Figure 5 shows the boxplots of parameters for links AB and AC respectively. As seen from the plots, the variance in the parameters of the model is quite high.






Fig. 5: Box plot of parameters of the two-state Gilbert model for links AB and AC,  $p_{ab}^i$ ,  $q_{ab}^i$ ,  $p_{ac}^i$ ,  $q_{ac}^i$ , shown in log scale




Fig. 6: QQ-plot of the number of losses on each link showing the non-stationary of distribution of losses across scenarios

Furthermore, from the QQ-plots in Figure 6, we see that the distribution of the observed number of losses in the  $n$  scenarios differs significantly from the distribution of number of losses if the losses on the links AB and AC were to follow Gilbert( $\hat{p}_{ab}, \hat{q}_{ab}$ ) and Gilbert( $\hat{p}_{ac}, \hat{q}_{ac}$ ), respectively. We conclude that the distribution of losses across scenarios is non-stationary and a standard statistical test such as the likelihood ratio test cannot be applied. Therefore, we develop an information-theoretic test, based on normalized mutual-information between the losses on the link AB and AC; it uses the bootstrap technique for statistical inference [21]. This information-theoretic test is described below.

### B. Information-Theoretic Test of Independence

If the losses on the two channels were to be independent, then the Wi-Fi links AB and AC can be represented as two-state Markov chains with transition matrices given by  $\alpha = \text{Gilbert}(p_{ab}, q_{ab})$  and  $\beta = \text{Gilbert}(p_{ac}, q_{ac})$ , respectively. Then, the joint channel is obtained by the product of the Markov chains given by  $\Theta_o = \alpha \otimes \beta$ , where  $\otimes$  represents the tensor product. Alternatively, if the losses on the two links are dependent, then the joint channel of the two wireless links AB and AC is represented by a 4-state Markov chain with states: 1) (Good, Good) 2) (Good, Bad) 3) (Bad, Good) and 4) (Bad, Bad). Let its transition matrix be  $\Theta$ .

Then, the test of independence is

$H_0$  : “ $\alpha \sim \mathcal{F}_{ab}$ ,  $\beta \sim \mathcal{F}_{ac}$  and the losses follow  $\Theta_o = \alpha \otimes \beta$ ” over  $H_1$  : “The losses follow  $\Theta \neq \Theta_o$ ,  $\Theta \sim \mathcal{G}$ ”, where  $\mathcal{F}_{ab}, \mathcal{F}_{ac}, \mathcal{G}$  are arbitrary distributions.

The test statistic is mean normalized mutual-information  $\mathbb{J} = \frac{\sum_{i=0}^n J^i}{n}$ , where ( $J^i$ ) is the normalized mutual-information

### Algorithm 1: Bootstrapping the generation of losses

---

```

1 currentRun  $\leftarrow$  0;
2  $\hat{\mathcal{F}}_{ab} = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$ ;
3  $\hat{\mathcal{F}}_{ac} = \{\beta_1, \beta_2, \dots, \beta_n\}$ ;
4 while currentRun < numRuns do
5    $i, j \leftarrow \text{random\_integer}(1, n)$ ;
6    $\Theta_{sim} = \alpha_i \otimes \beta_j$ ;
7    $\text{Pkts}_{ab}, \text{Pkts}_{ac} \leftarrow \text{generate\_packets}(\Theta_{sim})$ ;
8   currentRun++;
9 end

```

---

between the losses on both the links in the  $i^{\text{th}}$  scenario and is given as follows. Recall that the PLRs for link AB and AC in the  $i^{\text{th}}$  scenario are  $P_{ab}^i$  and  $P_{ac}^i$ , respectively. Then, the entropies of the links in the  $i^{\text{th}}$  scenario are given by

$$H^i(AB) = P_{ab}^i \log\left(\frac{1}{P_{ab}^i}\right) + (1 - P_{ab}^i) \log\left(\frac{1}{1 - P_{ab}^i}\right)$$

$$H^i(AC) = P_{ac}^i \log\left(\frac{1}{P_{ac}^i}\right) + (1 - P_{ac}^i) \log\left(\frac{1}{1 - P_{ac}^i}\right)$$

Let  $P_{xy}^i$ , where  $x = \{0, 1\}$  and  $y = \{0, 1\}$ , represent the empirical probability of packet loss (0) and reception (1) on the both the links AB and AC in the  $i^{\text{th}}$  scenario. Then, the mutual information and the normalized mutual-information in the  $i^{\text{th}}$  scenario is given by

$$I^i(AB; AC) = \sum_{x=\{0,1\}} \sum_{y=\{0,1\}} P_{xy}^i \log\left(\frac{P_{xy}^i}{P_x^i \times P_y^i}\right)$$

$$J^i = \frac{I^i(AB; AC)}{\sqrt{H^i(AB) \times H^i(AC)}}$$

It is worth noting that normalized mutual-information, consequently  $\mathbb{J}$ , is the information-theoretic analogue to covariance. In the ideal case, i.e., in the absence of measurement noise, if the losses were independent, then  $\mathbb{J} = 0$ . Hence, if  $H_0$  is true, we expect  $\mathbb{J}$  to be small. Furthermore,  $\mathbb{J} \geq 0$ . Therefore, the test’s rejection region is  $\mathbb{J} > \eta$ , where  $\eta$  is needed to be computed as a function of the confidence level.

For a confidence level of 99%, we have to evaluate  $\eta_{0.01}$  such that, under  $H_0$ ,  $\mathbb{P}_{H_0}(\mathbb{J} > \eta_{0.01}) < 0.01$ . However,  $\mathbb{P}_{H_0}(\mathbb{J} > \eta_{0.01}) = \mathbb{P}(\mathbb{J} > \eta_{0.01} | \mathcal{F}_{ab}, \mathcal{F}_{ac})$ . From the theory of bootstrapping [21], this can be approximated as  $\mathbb{P}(\mathbb{J} > \eta_{0.01} | \hat{\mathcal{F}}_{ab}, \hat{\mathcal{F}}_{ac})$ , where  $\hat{\mathcal{F}}_{ab}$  and  $\hat{\mathcal{F}}_{ac}$  are the empirical distributions of the parameters of the Gilbert model for link AB and link AC, respectively.

To evaluate  $\eta_{0.01}$ , we perform generation of losses on the two links under  $H_0$ , according to the bootstrap method shown in Algorithm 1. In each run, we randomly select (with replacement) one set of parameters of the Gilbert model for link AB and link AC, from among the parameters observed from real measurements. From these, under  $H_0$ , we obtain  $\Theta_{sim}$ , the distribution of the losses on the two channels in the current simulation run.  $\Theta_{sim}$  is used to generate a large number of packets ( $10^6$ ) on the two links ( $\text{Pkts}_{ab}, \text{Pkts}_{ac}$ ).

From the generated packets,  $\mathbb{J}_{sim}$  is evaluated. This process is repeated `numRuns` times to obtain several independent values of  $\mathbb{J}_{sim}$ . Then,  $\eta_{0.01}$  is the 99%-ile value of  $\mathbb{J}_{sim}$ .

Using `numRuns` = 10<sup>6</sup>, we obtain  $\eta_{0.01} = 3.82 \times 10^{-5} \pm 1.4 \times 10^{-5}$  at a 99% confidence level. From the measurements, we have  $\mathbb{J} = 8.91 \times 10^{-4} > \eta_{0.01}$ . Hence, we reject  $H_0$  with 99% confidence level and conclude that the two links are not fail-independent.

### C. Impact of Dependent Losses on $P_{rep}$

In Section V-B, we found that losses on the two Wi-Fi links AB and AC are not independent. The mean normalized mutual-information calculated from the measurements ( $\mathbb{J} = 8.91 \times 10^{-4}$ ) is a measure of mutual information between the two channels. Its low value indicates that, although the losses are not independent, the dependence is low.

The low degree of dependence is also noticeable from the PLRs. From the measurements, we have  $P_{ab} = 9.69 \times 10^{-4}$ ,  $P_{ac} = 2.4 \times 10^{-3}$ . Hence, if the losses were independent, the effective PLR would be  $P_{indep} = P_{ab} \times P_{ac} = 2.32 \times 10^{-6}$ . Additionally, from the measurements we have  $P_{rep} = 3.58 \times 10^{-6}$ , which is close to  $P_{indep}$ . Therefore, we conclude that although the two Wi-Fi links are not fail-independent, the effective loss performance as observed after replication is close to what it would have been, if they were to be fail-independent.

## VI. CONCLUSION

In this paper, we have experimentally validated the feasibility of using Wi-Fi measurements over redundant paths for streaming mission-critical PMU data. In the measurement test-bed we used, the location of the measurement sites and the traffic profile are the same as those used in the PMU-based state estimation in the campus-wide active distribution-network described in [15]. Such a setting is also commonly encountered when the last-hop of a wide-area network is realized using Wi-Fi, where directional antennas are used to boost the reliability of individual links.

In the setting we evaluated, we have concluded that although the PLRs of individual Wi-Fi links were far from admissible for streaming mission-critical PMU data, packet replication over the two Wi-Fi paths provided the desired level of reliability (PLR  $\sim 10^{-5}$ ). The effective PLR and availability can be further improved by enabling MAC-layer retransmissions on each Wi-Fi link. The end-to-end latency was observed to be within 4 ms, the required latency by mission-critical PMU-streaming applications. Moreover, the mean jitter in the latency was measured to be very low (0.02 ms), further bolstering the usability of Wi-Fi over redundant paths.

From the measurements, we observed that although the two links were found to not be fail-independent, the effective PLR as observed after replication was very close to the product of the two PLRs. Thus, for all practical purposes, the losses on the two links can be thought of as being independent.

We conclude, based on these results, that replication over redundant Wi-Fi links, formed from off-the-shelf components, is a viable option for achieving the loss and latency performance required for streaming mission-critical PMU data.

## REFERENCES

- [1] J. Gao, Y. Xiao, J. Liu, W. Liang, and C. P. Chen, "A Survey of Communication/Networking in Smart Grids," *Future Gener. Comput. Syst.*, vol. 28, no. 2, pp. 391–404, Feb. 2012.
- [2] M. Kuzlu, M. Pipattanasomporn, and S. Rahman, "Communication Network Requirements for Major Smart Grid Applications in HAN, NAN and WAN," *Computer Networks*, vol. 67, pp. 74–88, 2014.
- [3] K. C. Budka, J. G. Deshpande, and M. Thottan, "Communication Networks for Smart Grids," in *Computer Communications and Networks*. Springer, 2014.
- [4] T. Kropp, "Assessment of Wireless Technologies in Substation Functions- Part II: Substation Monitoring and Management Technologies," Technical Report, Electrical Power Research Institute, 2006.
- [5] F. Cleveland, "Use of Wireless Data Communications in Power System Operations," in *Power Systems Conference and Exposition, 2006. PSCE'06. 2006 IEEE PES*. IEEE, 2006, pp. 631–640.
- [6] M. Rentschler and H. Heine, "The Parallel Redundancy Protocol for Industrial IP Networks," in *Industrial Technology (ICIT), 2013 IEEE International Conference on*, Feb 2013, pp. 1404–1409.
- [7] H. Kirmann, M. Hansson, and P. Muri, "IEC 62439 PRP: Bumpless Recovery for Highly Available, Hard Real-Time Industrial Networks," in *Emerging Technologies and Factory Automation, 2007. ETFA. IEEE Conference on*, Sept 2007, pp. 1396–1399.
- [8] M. Popovic, M. Mohiuddin, D. C. Tomozei, and J. Y. L. Boudec, "iPRP: Parallel Redundancy Protocol for IP Networks," in *Factory Communication Systems (WFCS), 2015 IEEE World Conference on*, May 2015, pp. 1–4.
- [9] M. Popovic, M. Mohiuddin, D. C. Tomozei, and J. Y. L. Boudec, "iPRP - the Parallel Redundancy Protocol for IP Networks: Protocol Design and Operation," *IEEE Transactions on Industrial Informatics*, vol. PP, no. 99, pp. 1–1, 2016.
- [10] M. Rentschler, O. A. Mady, M. T. Kassis, H. H. Halawa, T. K. Refaat, R. M. Daoud, H. H. Amer, *et al.*, "Simulation of Parallel Redundant WLAN with OPNET," in *Emerging Technologies & Factory Automation (ETFA), 2013 IEEE 18th Conference on*. IEEE, 2013, pp. 1–8.
- [11] M. T. Kassis, O. A. Mady, H. H. Halawa, M. Rentschler, R. M. Daoud, H. H. Amer, and H. M. ElSayed, "Analysis of Parallel Redundant WLAN with Timing Diversity," in *Computer and Information Technology (WCIT), 2013 World Congress on*. IEEE, 2013, pp. 1–6.
- [12] M. Hendawy, M. ElMansoury, K. N. Tawfik, M. M. ElShenawy, A. H. Nagui, A. T. Elsayed, H. H. Halawa, R. M. Daoud, H. H. Amer, *et al.*, "Application of Parallel Redundancy in a Wi-Fi-based WNCs using OPNET," in *Electrical and Computer Engineering (CCECE), 2014 IEEE 27th Canadian Conference on*. IEEE, 2014, pp. 1–6.
- [13] M. Rentschler and P. Laukemann, "Towards a Reliable Parallel Redundant WLAN Black Channel," in *Factory Communication Systems (WFCS), 2012 9th IEEE International Workshop on*. IEEE, 2012, pp. 255–264.
- [14] M. Rentschler and P. Laukemann, "Performance Analysis of Parallel Redundant Wlan," in *Emerging Technologies & Factory Automation (ETFA), 2012 IEEE 17th Conference on*. IEEE, 2012, pp. 1–8.
- [15] M. Pignati, M. Popovic, S. Barreto, R. Cherkaoui, G. Dario Flores, J.-Y. Le Boudec, M. Mohiuddin, M. Paolone, P. Romano, *et al.*, "Real-Time State Estimation of the EPFL-Campus Medium-Voltage Grid by using PMUs," in *Innovative Smart Grid Technologies Conference (ISGT), 2015 IEEE Power & Energy Society*. IEEE, 2015, pp. 1–5.
- [16] J. Korhonen and Y. Wang, "Effect of Packet Size on Loss Rate and Delay in Wireless Links," in *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 3. IEEE, 2005, pp. 1608–1613.
- [17] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz, "Measurements of a Wireless Link in an Industrial Environment using an IEEE 802.11-Compliant Physical Layer," *Industrial Electronics, IEEE Transactions on*, vol. 49, no. 6, pp. 1265–1282, 2002.
- [18] G. Cena, S. Scanzio, and A. Valenzano, "Experimental Characterization of Redundant Channels in Industrial Wi-Fi Networks," in *2016 IEEE World Conference on Factory Communication Systems (WFCS)*. IEEE, 2016, pp. 1–4.
- [19] J.-Y. Le Boudec, *Performance Evaluation of Computer and Communication Systems*. Epfl Press, 2010.
- [20] E. N. Gilbert, "Capacity of a Burst-Noise Channel," *Bell system technical journal*, vol. 39, no. 5, pp. 1253–1265, 1960.
- [21] C. Z. Mooney, R. D. Duval, and R. Duval, *Bootstrapping: A Nonparametric Approach to Statistical Inference*. Sage, 1993, no. 94-95.