Real-Time Compression of IEC 61869-9 Sampled Value Data

Steven M. Blair, Andrew J. Roscoe, and James Irvine

University of Strathclyde

Glasgow, UK

steven.m.blair@strath.ac.uk

Abstract—Fast-acting, yet cost-effective, communications is critical for smarter grid monitoring, protection, and control. This paper demonstrates a new approach for the real-time compression of Sampled Value (SV) data based on the IEC 61869-9 recommendations. This approach applies simple compression rules, yet yields excellent compression performance—typically compressing data to less than half of the original size. This leads to a significant and beneficial reduction in encoding time (in the merging unit producing the SV data) and decoding time (at the end application), as well as the main benefit of reduced Ethernet transmission times resulting from the reduced frame size. As well as reducing the absolute bandwidth requirements in typical applications, this has system-wide benefits due to reducing Ethernet queuing delays and the consequent network jitter. The approach has been validated on a real-time platform to accurately measure all contributions to the end-to-end delay. This work will help enable low-latency and bandwidth-sensitive applications involving the SV protocol, such as phasor measurement units and wide-area protection.

Index Terms—Communications, IEC 61850-9-2, IEC 61869-9, merging units, phasor measurement units, power system protection, time synchronization.

I. INTRODUCTION

The use of communications and the analysis of increasingly larger quantities of real-time data are key facilitators of the Smart Grid [1], [2]. However, there are two conflicting requirements to be addressed: a desire for a vast number of power system measurements at high-fidelity and low-latency, versus minimising communications bandwidth requirements and the associated costs for electrical utilities. Therefore, minimising communications bandwidth and latency requirements, where possible, is advantageous.

The Sampled Value (SV) protocol is intended to transfer high sample rate sensor data within substations, and potentially over wide-area networks. This paper analyses the potential for losslessly compressing SV data in real-time to provide two key benefits: reduced bandwidth requirements and reduced end-to-end delay from the measurement point to the end application. As illustrated in Fig. 1, intelligent electronic devices (IEDs) commonly use 100 Mbps (100BASE-TX or 100BASE-FX) network interfaces due to the additional cost and complexity of Gigabit Ethernet (1000BASE-LX) interfaces [3]; higher bandwidth links are more likely between Ethernet switches forming the backbone of the local area network (LAN) or, in some cases, in wide-area networks. The communications
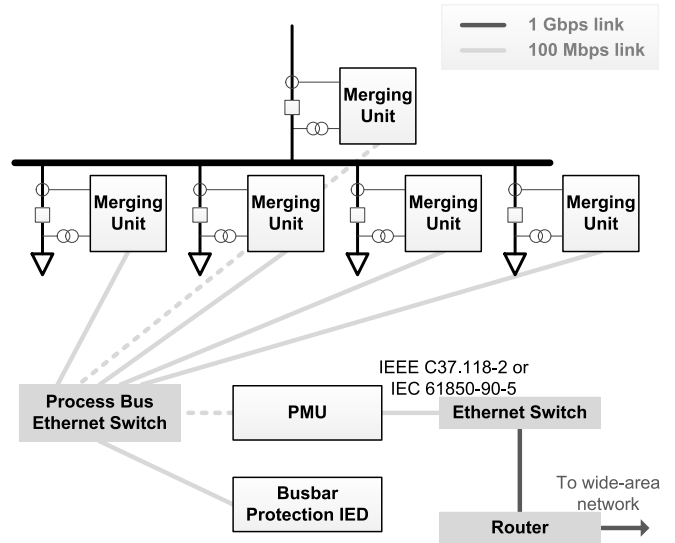

Fig. 1. Simplified substation example with SV communications for protection and PMU applications. In this example, only the dashed SV data stream is required by the PMU. Redundancy provisions, VLANs, time synchronisation, and GOOSE tripping functionality are not shown for simplicity.

transmission time, for each communications link, is directly proportional to the link speed and the size of data to be transmitted—and can be the most significant component of the overall end-to-end delay. These potential benefits arising from compressing SV data are attractive for a number of critical smart grid monitoring, protection, and control applications, including:

- Power quality monitoring, fault recorders, and the use of merging units to provide data to phasor measurement units (PMUs), which all require high sampling rates.
- Wide-area applications using SV data, such as for current differential protection, where the bandwidth requirements many be prohibitive. This is especially important for networks which involve multiple links or low bandwidth links [4], [5].
- Reducing PMU reporting latency, particularly for P-class PMUs, is highly beneficial to real-time applications which require PMU data, such as for island-mode detection [6], wide-area protection [7], and disturbance identification and location [8]. There is also growing interest in the use of PMUs in distribution systems which are especially

cost-sensitive [9].

- Real-time encryption of SV traffic is possible, but the method demonstrated in [10] can increase bandwidth requirements by approximately 30%. Therefore compression mitigates this drawback.
- Compression will mitigate the penalty on bandwidth use due to communications redundancy mechanisms which require duplicating data over two or more paths [5].
- High-voltage direct current (HVDC) measurement applications using the SV protocol require a maximum delay time, as defined in IEC 61869-9, of less than 25 μs [11]. Delivering fast-responding, yet robust, measurements is also increasingly important for virtual inertia and frequency response functions.
- Busbar protection schemes can involve many measurement locations, with each location transmitting an SV data stream. This requires relatively high bandwidth and careful system engineering to avoid affecting other process bus functions in the same substation.

In this paper, a novel compression algorithm, which is tailored to the requirements and data of the SV protocol, is proposed and analysed. A real-time platform is used to accurately measure the end-to-end performance of the proposed method.

II. PROPOSED SV COMPRESSION APPROACH

A. Analysis of SV Protocol

The SV protocol, as defined in IEC 61850-9-2 [12], provides a generic method for transferring arbitrary datasets using multicast Ethernet, with the intention of streaming raw sensor data in real-time. It was originally designed to transfer data within substations, but can also be delivered over wide-area networks using multiprotocol label switching (MPLS) or IEC 61850-90-5 [5], [10]. An implementation guideline of IEC 61850-9-2, known as "LE" [13], has been defined; this includes, among other specifications, recommended sampling rates and a dataset which is likely to be useful in many three-phase power system applications. It has become the de facto implementation reference for so-called merging units, as illustrated in Fig. 1. The essence of IEC 61850-9-2 "LE" has been recently standardised in IEC 61869-9, which provides further technical and pragmatic guidance, including harmonisation of sampling rates regardless of the nominal power system frequency. The recommended sampling rates for future applications are summarised in Table I. One of the key characteristics of these recommendations is to group two or six consecutive measurements (known as Application Service Data Units (ASDUs)) into a single Ethernet frame to improve bandwidth efficiency at the expense of delay.

Fig. 2a illustrates part of the dataset encoding format used for an SV frame with multiple ASDUs per frame. This highlights the redundancy of the repeated data (such as "svID") and the fact that the raw sensor values (32-bit signed integers in two's complement format) and quality attributes (a 32-bit field with multiple indicators) are encoded in full in each ASDU.

B. Compression Method

Using knowledge of the SV data format, a custom compression algorithm has been designed to eliminate some of the redundancy in the data. There are two main stages to the compression algorithm:

1) For SV Ethernet frames containing more than one ASDU, only the differences between the previous ASDU values need to be recorded, using differential encoding [14]. Therefore, the differences can be encoded using fewer bits than the original values, using a variable-length integer representation as described in the Appendix. This is especially effective for the "quality" values which will typically have a differential value (or the "delta") of zero, and therefore can be encoded in one byte instead of the conventional four bytes. Even if the quality value changes between ASDUs, the most significant bit (MSB) which can be used is the 14th bit to indicate that the neutral value is derived rather than measured [13] (where the 1st bit is the least significant bit (LSB)), and this difference can be encoded using a maximum of two bytes. It is important that the first ASDU in the frame is encoded in full (i.e. not with differential encoding) to ensure that the data from each frame can always be fully recovered, and such that a delayed or "dropped" Ethernet frame does not jeopardize the decoding of future frames.

2) Remove other redundant information from all ASDUs (except for the first ASDU), unless there is a difference in a value compared to the previous ASDU. This includes the mandatory "svID", "smpCnt", and "smpSynch" fields and several optional fields. The smpCnt values are implicit from the order of the ASDUs in the frame, except during a time adjustment process following resynchronisation [11]; in such a case, the smpCnt values should be given in full. This stage eliminates a further significant number of bytes, but ensures that the data can be unambiguously decoded.

Fig. 2 compares the compression method with conventional SV encoding. Note that some parts of the encoding structure have been omitted from Fig. 2 for clarity. It is clear that the proposed compression method can reduce the encoded size of each ASDU (except the first ASDU) by approximately 60 bytes. This accumulates to a saving of $5 \times 60 = 300$ bytes for the recommended 14.4 kHz sample rate—eliminating more than half of the frame size compared with conventional SV encoding. Therefore, reducing a 589 byte SV frame to approximately 289 bytes equates to a transmission time reduction of 24 μs for a single 100 Mbps link. Assuming a minimum of two spans of 100 Mbps in a practical network such as the example in Fig. 1, and ignoring the delays associated with switching fabric and queuing delays, a total reduction of at least 48 μs would be achieved. The smaller frame sizes will also benefit other traffic on the same network—even higher priority traffic—due to reducing the jitter associated with head-of-line blocking in Ethernet switches and routers [15].

TABLE I
SUMMARY OF PREFERRED SAMPLE RATES GIVEN IN IEC 61869-9

Digital output sample rates (Hz)	Number of ASDUs per frame	Digital output publishing rate (frames/s)	Typical Ethernet frame size (bytes)	Bandwidth requirement (Mbps)	Ethernet transmission time at 100 Mbps (μs)	Notes
4,800	2	2,400	~219	4.2	17.5	For general measuring and protection applications classes
14,400	6	2,400	~589	11.3	47.1	For power quality and metering accuracy class applications

(a) Conventional SV frame encoding
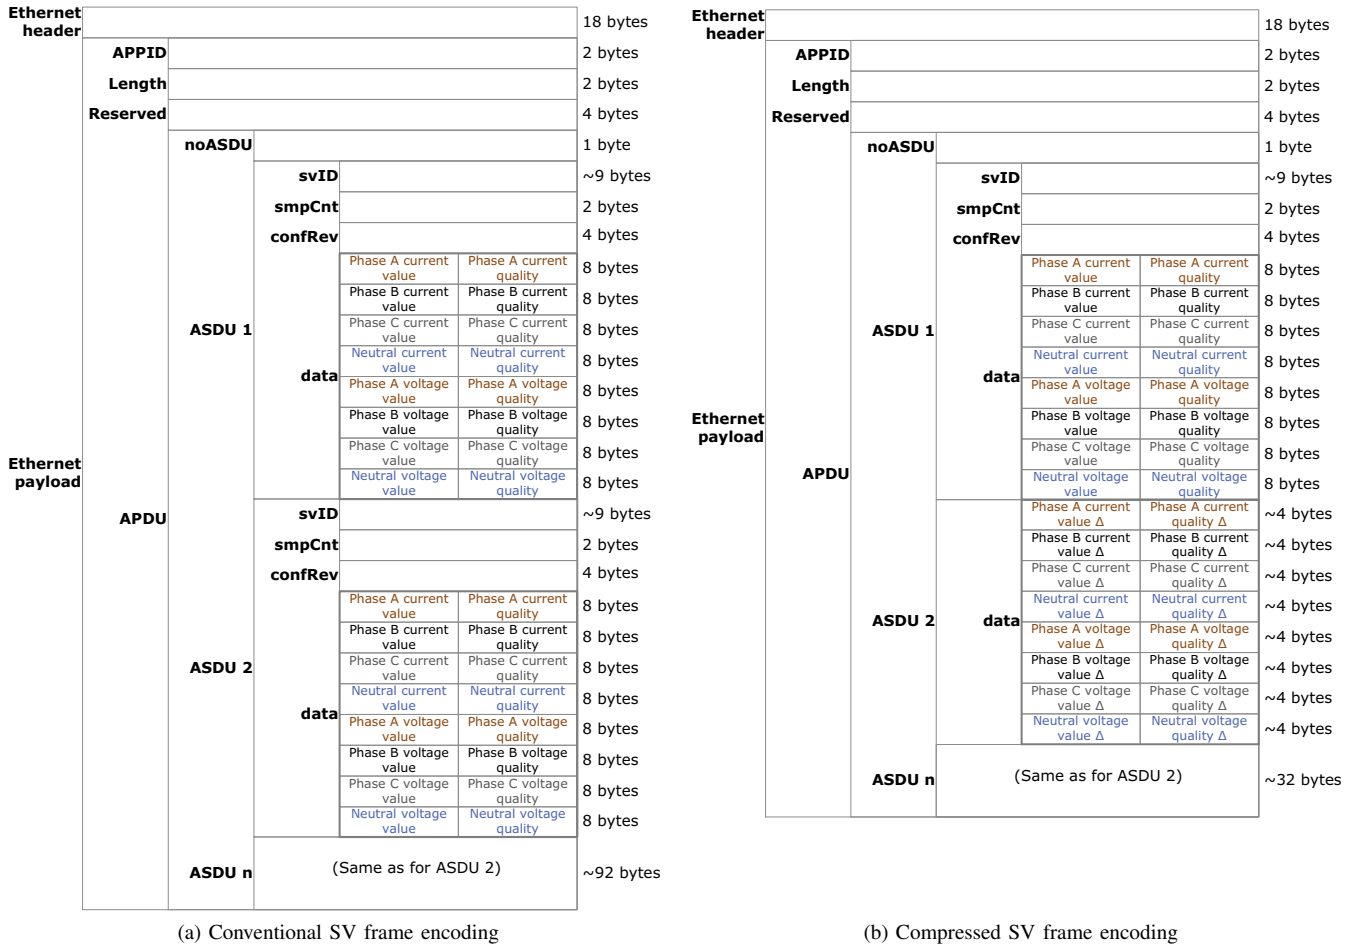

(b) Compressed SV frame encoding

Fig. 2. Simplified comparison of SV frame encoding

C. Backwards Compatibility Approach

It should be noted that the proposed compression method is non-standard, but could be standardised in the future. Due to the nature of the multicast publisher-subscriber communications model and the IEC 61850-6 configuration approach [16], it is essential the SV publishers which support compression are explicitly identified in the appropriate substation configuration language (SCL) files (e.g. using an additional SV Control Block attribute). This allows SV subscribers to be configured appropriately.

Furthermore, it is assumed that one of the reserved fields within the SV protocol can be used to indicate the use of real-time compression at run-time. At present, three reserved bits are unallocated [12], and only one bit is required to indicate that an SV publisher is compressing data. This method would allow SV subscribers to distinguish between conventional and compressed SV frames to ensure correct decoding. The proposed approach has the benefit that only SV subscribers need to be updated to ensure backwards compatibility.

III. MEASUREMENT AND VALIDATION

A. End-to-End Timing Measurement

Section II-B indicates the theoretical communications bandwidth and latency benefits to be expected from SV compres-
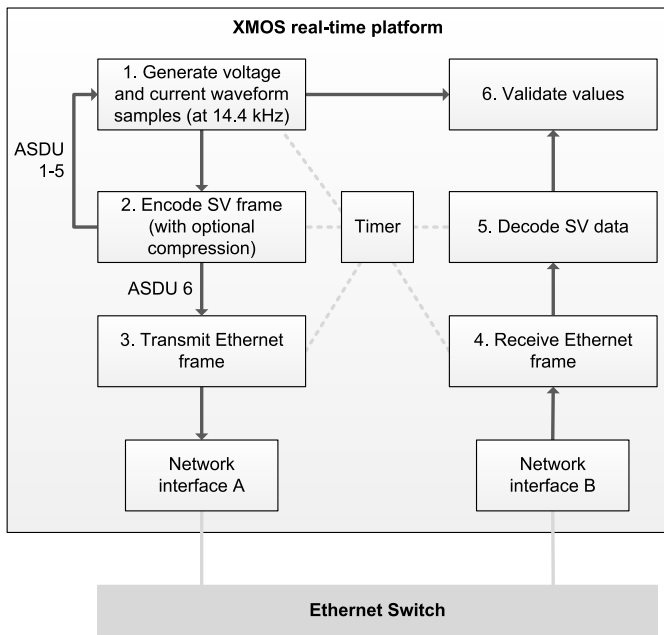
Fig. 3. Real-time validation of compression process for a 14.4 kHz sampling rate with 6 ASDUs per frame

Fig. 4. XMOS hardware platform

sion. However it is important to test the compression process on real-time hardware to establish the full behaviour, including the computational impact of SV compression. The XMOS platform [17] is designed for real-time applications that require deterministic behaviour. This is achieved by the provision of dedicated CPU registers for each thread of computation, such that there is no significant penalty for a context switch between threads. Furthermore, multiple Ethernet interfaces can be connected to the same device, so that the device can function as both a merging unit and an SV subscriber in

parallel. This allows each processing stage to be timed using a common on-board timer (a 100 MHz clock providing 10 ns accuracy); the use of a common time reference simplifies making accurate, consistent, and convenient measurements compared to using separate devices. The Ethernet interfaces also provide hardware timestamping required for the IEEE 1588 protocol, which provides a simple method for evaluating the "maximum delay time" of the merging unit as defined by IEC 61869-9. The process used for measuring and validating the timing is illustrated in Fig. 3. The hardware configuration is shown in Fig. 4.

B. Validation Results

Fig. 5 compares the performance of the proposed compression method to conventional SV encoding, with reference to the steps defined in Fig. 3. There are two clear benefits of the proposed compression method (in addition to reducing communications bandwidth requirements):

1) The proposed compression method does not increase computational requirements; counter-intuitively the compressed SV frames are faster to encode and decode than the highly-efficient SV implementation described in [1]. This benefit is due to the fact that fewer bytes are required to be encoded using the proposed method.

2) The time taken to transmit and receive Ethernet frames is substantially reduced due to the significantly smaller SV frame sizes. For the relatively simple LAN example with two 100 Mbps links, and combined with reduced encoding and decoding times, a total of at least 133 μs is saved from the end-to-end delay due to the use of compression.

It should be noted that both the conventional encoding and compressed approaches are suitable for implementation on a field-programmable gate array (FPGA) to further reduce encoding and decoding times [18].

IV. COMPARISON WITH GENERAL PURPOSE LOSSLESS COMPRESSION ALGORITHMS

Table II compares the performance of several open source general purpose lossless compression algorithms for SV data. These algorithms are designed for relatively large datasets and therefore the entire APDU section of the SV frame (see Fig. 2a) is used, rather than each ASDU individually. A 14.4 kHz sampling rate with 6 ASDUs per frame has been used. The tests have been performed over one cycle of data with 50 Hz nominal frequency (i.e. for 48 SV frames, each containing 6 ASDUs), using a 4 GHz Intel i5 processor. The dictionary-based Lempel-Ziv (LZ) algorithms achieve reasonable levels of compression when applied to SV data, but are far too slow to execute—even on a relatively fast PC platform—to offer any overall benefit. Even the fastest algorithm, run-length encoding (RLE), executes significantly slower on an embedded platform (by a factor approximately 100 on a 96 MHz ARM Cortex-M3 microcontroller [1]) to be useful in practice.

Fig. 5. Timing diagrams for minimum SV encoding and decoding processes (14.4 kHz sampling, with 6 ASDUs per SV frame)

Therefore, none of the general purpose algorithms are feasible for real-time applications. The tailored compression algorithm proposed in this paper has the unique ability to achieve extremely good compression performance whilst also reducing computational requirements.

V. CONCLUSIONS

This paper has proposed a new approach for the real-time compression of SV data based on the IEC 61869-9 recommendations. This approach applies simple compression rules, yet yields excellent compression performance—typically compressing data to less than half of the original size. Counter-intuitively, this leads to a significant and beneficial reduction in encoding time (in the merging unit) and decoding time (at the end application), as well as the main benefit of reduced Ethernet transmission times resulting from the reduced frame size. As well as reducing the absolute bandwidth requirements in typical applications, this has system-wide benefits due to reducing Ethernet queuing delays and the consequent network jitter. The compression is performed losslessly and therefore does not affect measurement uncertainty.

One of the advantages of the SV protocol is its simplicity in order to facilitate low-latency encoding and decoding; this work proves that a slightly more complex encoding method can achieve superior performance. The approach has been validated on a real-time platform to accurately measure all contributions to the end-to-end delay in a representative application. Latency can be reduced by at least 130 μs, or significantly greater for wide-area applications.

Based on the results in this paper, the following recommendations can be made:

- The proposed approach is only applicable if using multiple ASDUs per SV Ethernet frame.
- The approach has greatest benefit in networks with 100 Mbps (or lower) links.
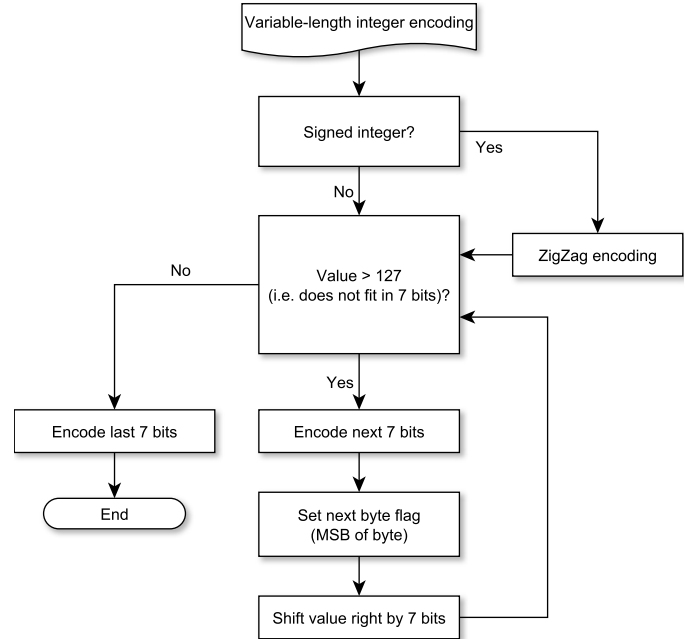

Fig. 6. Variable-length integer encoding

- To maintain backwards compatibility, the "reserved" field in the SV protocol should be used to indicate the use of SV compression, and SV subscribers must be able to decode both conventional and compressed data streams.
- The approach is most useful for applications where minimising measurement latency or communications bandwidth is critical.

APPENDIX

Fig. 6 defines the process for encoding signed and unsigned integer values in a variable-length format. The reverse of this process is used to decode data into the original format. To improve the efficiency of encoding signed integers which use

TABLE II
PERFORMANCE COMPARISON OF GENERAL PURPOSE COMPRESSION ALGORITHMS AND THE PROPOSED APPROACH

Algorithm	Original data size (bytes)	Mean compression ratio (lower is better)	Std. dev. of compression ratio	Total processing time (μs)	Encode time, decode time (μs)	Score (ratio $\times$ total time)	Code origin
RLE	561	0.85	0.01	2.9	2.2, 0.7	2.5	[19]
Huffman coding	561	0.8	0.01	53.1	40.3, 12.7	42.5	[19]
LZ77	561	0.59	0.01	59.9	58.8, 1.2	35.6	[19]
LZ77 (fast)	561	0.59	0.01	34.4	33.6, 0.8	20.4	[19]
FastLZ (level 1)	561	0.72	0.01	10.2	8.2, 2.0	7.4	[20]
FastLZ (level 2)	561	0.71	0.02	5.9	4.5, 1.4	4.2	[20]
lzfx	561	0.69	0.01	45.3	43.9, 1.4	31.2	[21]
TurboRLE	561	0.8	0.01	18.3	16.1, 2.3	14.6	[22]
Mespotine RLE	561	0.82	0.01	7.0	5.6, 1.3	5.7	[23]
Proposed method	561	0.54	0.00	0.0	0.0, 0.0	0.0	[24]

two's complement format, this process uses ZigZag encoding [25] to map negative values to positive integers.

REFERENCES

[1] S. M. Blair, F. Coffele, C. D. Booth, and G. M. Burt, "An Open Platform for Rapid-Prototyping Protection and Control Schemes with IEC 61850," *IEEE Trans. Power Deliv.*, vol. 28, no. 2, pp. 1103–1110, 2013.

[2] C. Gellings, "The Future of the Power Grid," 2014. [Online]. Available: http://spectrum.ieee.org/podcast/energy/the-smarter-grid/clark-gellings-the-future-of-the-power-grid

[3] IEC TC 57, "IEC/TR 61850-90-4: Network engineering guidelines," Tech. Rep., 2013.

[4] S. M. Blair and C. D. Booth, "Real-time teleprotection testing using IP/MPLS over xDSL," Glasgow, 2013. [Online]. Available: https://pure.strath.ac.uk/portal/files/26184600/001{_}DSL{_}Testing.pdf

[5] S. M. Blair, F. Coffele, C. Booth, B. De Valck, and D. Verhulst, "Demonstration and analysis of IP/MPLS communications for delivering power system protection solutions using IEEE C37.94, IEC 61850 Sampled Values, and IEC 61850 GOOSE protocols," in *CIGRE Paris Sess. B5*, aug 2014.

[6] M. S. Almas and L. Vanfretti, "RT-HIL Implementation of the Hybrid Synchrophasor and GOOSE-Based Passive Islanding Schemes," *IEEE Trans. Power Deliv.*, vol. 31, no. 3, pp. 1299–1309, jun 2016.

[7] E. Udren, "Principles for Practical Wide-Area Backup Protection with Synchrophasor Communications," in *CIGRE Paris Sess. B5*, 2014.

[8] W. Li, J. Tang, J. Ma, and Y. Liu, "Online Detection of Start Time and Location for Hypocenter in North America Power Grid," *IEEE Trans. Smart Grid*, vol. 1, no. 3, pp. 253–260, dec 2010.

[9] M. P. Andersen, S. Kumar, C. Brooks, A. von Meier, and D. E. Culler, "DISTIL: Design and implementation of a scalable synchrophasor data processing system," in *2015 IEEE Int. Conf. Smart Grid Commun.* IEEE, nov 2015, pp. 271–277.

[10] S. M. Blair, C. D. Booth, B. D. Valck, D. Verhulst, C. Kirasack, K. Y. Wong, and S. Lakshminarayanan, "Validating Secure and Reliable IP/MPLS Communications for Current Differential Protection," in *Dev. Power Syst. Prot.*, 2016.

[11] IEC, "IEC 61869-9: Digital interface for instrument transformers," Tech. Rep., 2016.

[12] IEC TC 57, "Communication networks and systems in substations Part 9-2: Specific Communication Service Mapping (SCSM) - Sampled values over ISO/IEC 8802-3 (IEC 61850-9-2:2011)," 2011.

[13] UCA International Users Group, "Implementation Guideline for Digital Interface to Instrument Transformers Using IEC 61850-9-2," Tech. Rep., 2004. [Online]. Available: http://iec61850.ucaiug.org/ImplementationGuidelines/DigIF{_}spec{_}9-2LE{_}R2-1{_}040707-CB.pdf

[14] J. Irvine and D. D. Harle, *Data communications and networks: an engineering approach.* Wiley, 2002.

[15] Nokia, "Mission-critical communications networks for power utilities," Tech. Rep., 2016. [Online]. Available: http://resources.alcatel-lucent.com/?cid=180690

[16] IEC TC 57, "Communication networks and systems for power utility automation - Part 6: Configuration description language for communication in electrical substations related to IEDs (IEC 61850-6:2009)," 2009.

[17] G. Martins, D. Lacey, A. Moses, M. J. Rutherford, and K. P. Valavanis, "A case for I/O response benchmarking of microprocessors," in *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.* IEEE, oct 2012, pp. 3018–3023.

[18] W. Wei, H.-b. Li, and H.-m. Cheng, "A multi-function IEC 61850 packet generator based on FPGA," *Meas. Sci. Technol.*, vol. 27, no. 7, p. 075901, jul 2016.

[19] M. Geelnard, "Basic Compression Library," 2016. [Online]. Available: http://bcl.comli.eu/

[20] A. Hidayat, "FastLZ," 2016. [Online]. Available: http://fastlz.org/

[21] K. Hamasaki, "lzfx," 2016. [Online]. Available: https://github.com/hmsk/lzfx

[22] Powturbo, "TurboRLE," 2016. [Online]. Available: https://github.com/powturbo/TurboRLE

[23] M. Mespotine, "Mespotine-RLE-basic v0.9 - An overhead-reduced and improved Run-Length-Encoding Method," p. 16, jan 2015. [Online]. Available: http://arxiv.org/abs/1501.05542

[24] S. M. Blair, "rapid61850," 2012. [Online]. Available: https://github.com/stevenblair/rapid61850

[25] Google, "Protocol Buffers - Encoding," 2016. [Online]. Available: https://developers.google.com/protocol-buffers/docs/encoding{#}varints