

A Fast Prekeying Based Integrity Protection for Smart Grid Communications

Amitangshu Pal*, Alireza Jolfaei†, Krishna Kant*

* Temple University, Philadelphia, PA † Macquarie University, Australia

Abstract—In this paper, we propose a prekeying based integrity protection (PreKIP) mechanism for critical smart grid communications that are often left unprotected due to tight timing constraints. Our mechanism computes the key for the next message in advance followed by a simple exclusive-or operation with the message when it is generated. This provides both integrity and confidentiality at a very low latency cost. The rigorous security analysis shows that the proposed method is secure against CRC and message replay attacks. The extensive evaluation shows that the method is up to 21 times faster than standard integrity protection algorithms, and can do the message encryption in under 1 ms even on a very low-end microcontroller.

Index Terms—Cyber security, GOOSE, IEEE C37.118, Integrity protection, Smart grid communications, SV protocol.

I. INTRODUCTION

The emerging smart grid architecture uses real-time monitoring and control of the power grid in order to provide high efficiency, stability and robustness in the power supply. A series of communications protocols have been defined for this purpose over the years. In particular, there are 3 protocols in existence currently. The earliest one is the IEEE C37.118, which is now split into two parts, with C37.118.2 being aligned with more widely deployed IEC 61850-9-5 for synchrophaser communications [1]. It is currently the most widely deployed protocol for synchrophasors but lacks security. On the other end is the DoE developed STTP protocol [2] which has been recently picked up by IEEE for standardization, but currently not deployed except on an experimental basis. Therefore, we will largely focus on IEC 61850-9-5, although a wide-spread deployment of even this protocol is likely to take many years.

Although IEC 61850-9-5 includes integrity protection mechanisms for the communications; they are optional for critical protection messages requiring very low latency. Without any integrity protection, if a protection message is falsely set to indicate abnormal voltage or current value, it could trigger protective relays and/or the generation control equipment to react, potentially leading to blackouts.

In this paper, we propose a lightweight prekeying based approach, named PreKIP, that generates the keys between successive sample generations and then simply XoRs it with the sample. We apply this method both to the regular PMU data (where the available key generation time is strictly determined by the sampling rate) and for event based data with certain minimum inter-event time. The rigorous security analysis shows

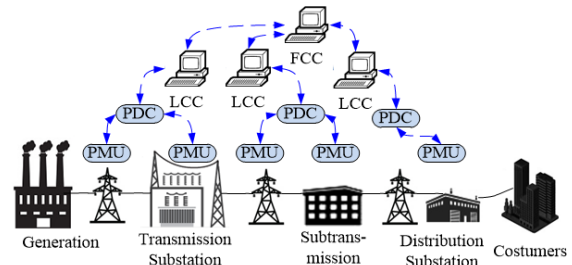


Fig. 1: Typical model of a power grid architecture.

that the method can successfully thwart both ciphertext-only attacks and known/chosen plaintext attacks. Our extensive evaluation shows that the proposed method is upto 21 times faster than the HMAC scheme (and even faster than others), and has acceptable latency for typical power protection applications even with a low-end microcontroller.

The remainder of this paper is organized as follows. Section II discusses the gaps in integrity protection mechanisms in smart grid protocols. Section III discusses the proposed method, Section IV analyzes its security, and Section V evaluates its performance. Section VI summarizes the related work, and Section VII concludes the paper.

II. SMART GRID ARCHITECTURE AND COMMUNICATIONS

A. Smart-Grid Architecture

The emerging smart grid architecture uses PMUs to continuously monitor line data (for example, voltage, phase, frequency, and GPS location) and communicates them to the supervisory control and data acquisition (SCADA) systems to ensure that any issues related to grid health are handled promptly. Fig. 1 shows the overall architecture where the data from PMUs is “concentrated” through Phasor data concentrators (PDCs) installed in key substations. PMUs collect and send samples 30 or 60 times a second through a publish-subscribe mechanism, where the PMUs work as publishers to which the PDCs subscribe. A PDC receives data from many (typically 3 to 32) PMUs, and then sorts and aggregates the received data based on the time-tag. The aggregated data is then relayed using a two-way communication system to a number of local control centers (LCCs), which coordinate their actions interacting with a federated control center (FCC). Subsequently, LCCs draw the best overall snapshot solution using all PMU measurements [3]. Control centers use the IEC 61850-90-5 standard to communicate with smart measurement units [1].

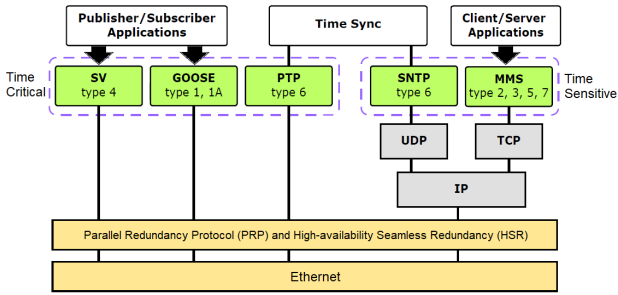


Fig. 2: Communications protocols used in smart grid

The IEC 61850-9-5 standard is itself a collection of protocols, each defined for a different set of smart grid applications. This is shown in Fig. 2. The MMS (Manufacturing Message Specification) protocol connects communications centers and gateways through a client-server connection with the IEDs (Intelligent Electronic Devices) inside the substation and TimeSync provides time synchronization. The two important protocols of interest here are SV or SMV (sampled value) and GOOSE (generic object oriented station event), and both were originally defined to operate directly on top of Ethernet and thus are not routable. The SV protocol is used to exchange messages containing samples of electrical quantities such as the very fast rate sampling of voltage/current at a relay. While SV was originally designed for intra-substation use, it is also useful for sending the PMU data stream to PDCs and LCCs. The GOOSE protocol is used to exchange information between IEDs (Intelligent Electronic Devices) based on the important events that occur in the system. Such messages include power measurements going between protection relays, status updates, sending command requests, and critical messages that demand an immediate action, e.g., a relay trip. The use of GOOSE for transporting PMU data is also reported in several publications [1], but in this case, repetition of the message would not be useful. Both GOOSE/SV messages require very low transmission latency, which is the main reason for running them directly on layer-2. However, more recently the routable version of these protocols known, respectively, as R-SV and R-GOOSE have been defined, so that it is possible to transport data from several substation networks to a concentrator or a control center over UDP or TCP. Tunneling can also be used to carry layer-2 GOOSE/SV messages to points outside a substation.

Assuming that the PDCs are located on the local substation network, SV/GOOSE can be used for transmission from PMU to PDC. However, since LCC is must necessarily be located remotely and connect to multiple substations, all exchanges between PDCs and LCC must be using either a routable protocol such as R-SV/R-GOOSE or by tunneling the SV/GOOSE protocols. In the following, we largely focus on non-routable versions of SV/GOOSE which can be carried over Ethernet frames riding a SONET or similar lower-level network. These are appropriate for low-latency, and low-packet-loss communications. R-SV/R-GOOSE instead usually ride TCP and are more appropriate for less critical communications. Table I gives the size of payload data in transmission and distribution substations under various PMU/PDC configurations.

TABLE I: Payload data size under various configurations

#PMU voltage channels	PMU data size	PDC data size with n PMUs			Applicability
		$n = 3$	$n = 6$	$n = 10$	
1	40	104	200	328	1-phase power distribution
3	56	120	216	344	3-phase power distribution
6	80	144	240	368	Power transmission lines
8	96	160	256	384	Power transmission lines
10	112	176	272	400	Power transmission lines

B. GOOSE Protocol

The GOOSE protocol operates on a publish-subscribe principle, that is, each IED publishes its data stream, and other interested entities (IEDs, PDCs, LCCs) subscribe to it. Thus, the communication is largely one-way; the subscriber does not send any ACK/NACK to the publisher and thus the publisher cannot tell if the data was received correctly. The publisher can, however, retransmit the data.

When GOOSE is used for communicating events, it will typically generate bursts of messages (e.g., on an over-voltage detection), with significant quiet periods in between. It continues to retransmit these events (with successively increasing gap between retransmissions) until a new event occurs. In such a situation, the transmission interval exponentially increases to the normal periodic interval. Assume that T_0 is the time interval between GOOSE messages in periodic mode, and $T_1, T_2, \dots, T_n = T_0$ are the time intervals in the burst mode, with $T_1 < T_2 < \dots < T_n$. In IEC 61850, T_0 is typically in the range of 5–100 ms, whereas T_1 is in the range of 0.5–5 ms [4]. After the first retransmission after T_1 , each successive interval doubles until it reaches T_0 . The requirements for the delivery of GOOSE messages are pretty stringent; the messages should be delivered within 4 ms from the time an event occurs to the time the message is received for protection and control applications; this requirement is revised to 3 ms in IEC 61850-5 [4].

GOOSE tries to deliver messages in sequence by using two fields called `stNum` (state number) and `sqNum` (sequence number). Every time a new event occurs, the transmitter increments `stNum` and transmits a new message with this `stNum`. If no event occurs and a time `timeAllowedtoLive` elapses, the transmitted simply repeats the last message. Note that if the next event happens rather quickly, it is considered to override the previous one and hence no repetition of previous message is done. The counter `sqNum` is incremented each time the previous message is repeated.

C. Sample Value Protocol

As stated earlier, SV protocol is primarily intended for fast transmission of samples of analog measurements from various substation devices on to the local bus for which it may use very high sampling rates (e.g., 80 to 256 samples per cycle). However, for use in the PMU context, the rates will usually be far lower, perhaps less than once per cycle. Like GOOSE, it also uses the publisher/subscriber model, it does not retransmit data, since retransmission is not meaningful for data stream delivery. Each SV PDU can include several measurements, one per ASDU (application service data unit). Each ASDU contains the measurement data (e.g., voltages and currents for each phase

in a 3-phase circuit) plus some predefined fields, only one of which is of interest here. This is `smpCnt`, which is incremented each time a new data sample in ASDU is taken. `smpCnt` is only 2-bytes and it is reset every second through the TimeSync based synchronization protocol.

The sampling rate of SV involves two factors: measured signal frequency and Samples Per Period (SPP). IEC 61850-9-2LE defines two SPP values of 80 and 256. Thus, if the measured signal frequency is 50 Hz and SPP is 80, then the sending time interval is $1/50/80$, or $250 \mu\text{s}$ [5].

D. IEEE C37.118.2 Protocol

Even though this protocol is being replaced by IEC 81650-9-5, it remains the most widely deployed protocol for PMU communications and also lacks security. Each synchrophasor measurement is tagged with a UTC timestamp consisting of three components: 1. Second-Of-Century (SOC), 2. FRAction-of-SECond (FRACSEC), and 3. Message time quality flag. The SOC count is a 4B integer count in seconds from UTC midnight (00:00:00) on January 1, 1970. Each second is divided into an integer number of subdivisions by the TIME BASE parameter that is defined in configuration frame. The FRACSEC count is an integer representing the numerator of the FRACSEC with TIME BASE as the denominator. There is also a time quality flag provides the accuracy of the time measurement.

E. Protecting Messages

In general, the adversaries may be *passive* or *active*. Passive attackers eavesdrop on the communications between devices by wiretapping the links between end devices and the substation switches to obtain the message contents and traffic characteristics of the substation; they neither modify the messages in the channel nor communicate with the end devices. The objective of active attacks includes learning more about substation's operations which can be helpful in later disruptions. In fact, with the dramatic rise in analytics capabilities, a long-term silent monitoring could derive important information about the types and sources of messages. For example, an attacker may be able to determine what type of perturbation will cause the greatest disruption in the power flows through active attacks. Thus, encryption may be desirable for privacy and integrity purposes, although ensuring communication availability is regarded as the primary goal in smart grids.

However, the current smart grid deployments generally do not use any encryption or integrity protection to reduce communication latency in time-critical applications; for example, the standard for formatting and delivery of PMU data (IEEE Standard C37.118 [6]) includes no end-to-end security mechanisms. Although a standard IEC 62351 introduces several message authentication code (MAC) algorithms to protect GOOSE integrity, but it still allows an escape route for latency critical messages by specifying an identifier in the message header to zero. Thus, for time critical messages requiring a latency of 3 ms, no encryption is likely to be used. This includes GOOSE




Fig. 3: Timing diagram of PreKIP

and SV messages that are used to transmit both the streaming PMU data and the event based data concerning load shed and synchrophasor-assisted transfer trip [7]. The latter involves sending a trip signal from one substation to another that could be more than 100 miles away.

We make no assumption on the attacker's ability to capture, examine, or alter transmitted data or inject new data into the system; however, we do not consider a persistent attack resulting in denial of service (DoS). While our mechanism can detect DoS attacks, it cannot prevent them; other mechanisms such as redundant transmission or isolation may be necessary to deal with DoS attacks. Our mechanism aims to harden the communication protocol to ensure end-to-end confidentiality and integrity under the assumption of trusted and uncompromised IEDs. Attacks that exploit the hardware or firmware vulnerabilities to bypass the integrity protection are beyond the scope of this paper as they require hardware and firmware attestation techniques.

Because of the one-way communication and a quiet overriding of previous message with a new one in GOOSE/SV, it is simply not possible to synchronize the two sides precisely or account for every message. Note that the PDC does have the capability to send a message to PMU for special purposes (e.g., key exchange) which are outside the scope of the proposed mechanism. The goal of our mechanism is to have the keys for encrypting the next few messages ready before those messages are generated, so that the encryption operation is a simple XoR that can be done very fast. Such a mechanism can work well in most situations, the only exception being cases where an avalanche of messages is generated in short-time, possibly due to many simultaneous faults.

III. PROPOSED INTEGRITY PROTECTION SCHEME

The purpose of the proposed mechanism is to ensure integrity via the use of CRC and confidentiality via a fast encryption mechanism that generates a new key for almost every message between transmissions so that the latency can be largely hidden. This is summarized in Fig. 3. As depicted in Fig. 3, the key K_i corresponding to P_i is calculated before it arrives; thus, after the packet arrives, the payload is XOR-ed with the key to produce the ciphertext C_i before transmitting. The XOR operation is very fast and does not have any significant contribution to the encryption latency. Unfortunately, the use of XoR along with CRC introduces some vulnerabilities that we address by obfuscating the CRC, as discussed below.

A. Key Stream Generation

Assume that the lengths of the original message and the CRC checksum be L_m and L_c bits respectively. Thus, we have a plaintext of length $L = L_m + L_c$ bits. Now, both the transmitter and receiver need to generate the same key-stream KS of length at least L bits for a message without any explicit handshake. In the following, we discuss the generation of KS , which will be ultimately XOR'ed with the plaintext to generate the ciphertext.

To generate KS for the j th new message, we start with a one-time key sk_j which is defined as

$$sk_j = \mathcal{K} || salt_j, \quad (1)$$

where $||$ refers to the concatenation of a pre-distributed key \mathcal{K} and a per-transmission ‘‘salt’’ to make the key unique to each transmission. We assume that \mathcal{K} is either manually configured on both transmit and receive sides, or communicated in an out-of-band manner. Note that once we establish an encrypted communication channel, \mathcal{K} can be changed easily through a handshake, and the best practice demands that this be done periodically (e.g., once a month). The length of the preconfigured key \mathcal{K} should be adequate to avoid brute-force attacks, e.g., 128 bits.

However, $salt_j$ needs to be generated correctly for each message by both transmit and receive sides without any further message exchanges, as we shall discuss shortly in Subsections III-B to III-D for different communication protocols. The resultant one-time key sk_j is not used directly for encryption; instead, we compute a one-way, secure hash H over it so that the KS does not reveal the key. The hash H could be SHA1, SHA256 or another suitable one-way function. Notice that the resulting message digest, say d_j , will be a fixed length value L_d (e.g., 160 bits for SHA1 or 256 bits for SHA256).

$$d_j = H(\mathcal{K} || salt_j) \quad (2)$$

Thus, to generate a long enough key-stream, we utilize the *avalanche effect* of the hash function H to output significantly changing digests by adding an additional counter into the salting. By concatenating the counter's value cnt from 0 to $N - 1$ to $salt_j$, we obtain N salts for each message, which are used to generate N hash digests $d_j^{(i)}$, for $i = 1, \dots, N$. The final key-stream is then a concatenation of the hash digests.

$$KS = d_j^{(1)} || d_j^{(2)} || \dots || d_j^{(N)}, \quad (3)$$

where N is chosen as $N = \lceil \frac{L}{L_d} \rceil$ so that the KS is at least as long as L . The KS is then XOR'ed with the plaintext of length L . Our salting scheme differs based on the communication protocols which we discuss in the following subsections.

B. Salting Scheme for GOOSE

Here we consider general GOOSE messages that are initiated by specific events, as well as the retransmitted ones. The *intent* of the salting scheme is to generate a unique KS for every message, including both new and retransmitted messages; unfortunately, given the lack of acknowledgements (ACK or NACK), it is very difficult to realize this intent. As stated

earlier, stNum is incremented by 1 when the message transmits a new event and sqNum is incremented by 1 if the message retransmits an old event. Thus, to realize the intent, we need to set $salt = \text{stNum} || \text{sqNum}$ for every message, which means that stNum and sqNum must be transmitted in the clear to the receiver, so that it too can construct KS . However, that will expose the keys to the attacker, thus, the receiver must be able to predict the pair $(\text{stNum}, \text{sqNum})$ *before* it decodes the message, which we discuss in Subsection III-C.

One other situation to consider is the rollover of stNum . First, the rollover of 32 bit sqNum is a nonissue, since that can only happen if the same message continues to be retransmitted without success for a very long time. Note that sqNum will be reset to 0 each time a transmission succeeds or the next overriding event occurs. Furthermore, the rollover of stNum is easily handled by the modulo arithmetic on both sender and receiver sides, since there is no chance that a rollover would collide with a previous message with $\text{stNum} = 0$.

C. Key Synchronization

To decrypt and verify the received message, the subscribers need to use the same key-stream KS . To this end, the subscribers need to use the same salts as those on the publisher's side. Recall that each salt is computed from 3 parameters, i.e. stNum , sqNum and cnt . The cnt is easy to synchronize since it takes the same values from 0 to $N - 1$ for all messages.

The challenge for synchronizing stNum and sqNum in key-stream generation is that both a publisher and its subscribers cannot predict the next message is to transmit a new event or to re-transmit an old event (recall that key-streams are generated ahead of time and that the publisher keeps retransmitting messages for an event until a new event comes). From the publisher's side, if a new event happens then it needs to send a packet *instantly* with $\text{stNum} = \text{stNum} + 1$ and $\text{sqNum} = 0$. In an extreme scenario, multiple events can occur one after other (or within a interval of very short time), the publisher needs to prepare key-streams for the next several event driven messages with increasing stNum values. The publisher can do so by proactively maintaining a certain number (say \mathcal{N}) of key-streams and storing it in a queue. Whenever a new event occurs, the publisher can dequeue a new key-stream, execute the XOR operation and transmit. This is required only for critical event reporting and not for PMUs. Furthermore, generation of too many critical events at the same time is very unlikely, therefore a rather small value of \mathcal{N} should be enough. The publishers can generate the key-streams for the retransmitted messages in between the message transmission intervals, with $\text{sqNum} = \text{sqNum} + 1$. As the minimum time interval for fast retransmission is around 0.5–5 ms, such key-stream generation process needs to be fast enough. Algorithm 1 summarizes the key-stream generation for the publisher.

On the other hand, the subscriber side also needs to generate the same key-stream to decrypt the incoming messages. To ensure this, the subscriber follows the following set of steps after receiving the message. We assume that the subscriber

Algorithm 1: Key stream generation at the publisher

```
Input :  $L \leftarrow$  the required key-stream length;  
           $msg \leftarrow$  last sent GOOSE message;  
           $\mathcal{K} \leftarrow$  secret key;  
Output: Key stream  $KS$ ;  
 $cnt \leftarrow 0$ ;  
 $stNum \leftarrow (msg.stNum == 2^{32} - 1)?1 : msg.stNum$ ;  
if  $msg.isNewEvent()$  then  
   $sqNum \leftarrow 0$ ;  
else  
   $sqNum \leftarrow (sqNum + 1 \leq 2^{32} - 1)?(sqNum + 1) : 1$ ;  
while  $cnt < \lceil L/L_d \rceil$  do  
   $salt_{cnt} \leftarrow (stNum || sqNum || cnt)$ ;  
   $KS \leftarrow H(sk_i || salt_{cnt})$ ;  
   $cnt = cnt + 1$ ;
```

is equipped with a fast computing unit, so it can generate multiple key-streams and try them out for decryption. Assume a normal operating condition where the receiver knows $stNum$ for the last received message and $sqNum = 0$ (i.e., no message loss/retransmission). Now we have 4 possible situations:

(1) Next message also makes it to the receiver without corruption. In this case, the receiver can predict $(stNum+1, 0)$, decrypt the message, and do the integrity check. This is the normal case and likely to occur in all but a small number of cases.

(2) Next k messages are lost (for some $k \geq 1$), but the $k+1$ st message, which is correctly received, is still generated before the retransmission time of first lost message. In this case, the receiver can estimate the number of lost messages and predicts $(stNum + k + 1, 0)$ for a suitable k . The probability of this case should vanish rapidly as k increases.

(3) No new messages are generated until it is time for retransmission of the last message. This is the normal case of retransmission, so the receiver can predict the key-stream with $sqNum = sqNum + 1$.

(4) Several messages are lost, including some new messages and some retransmitted ones. In such a scenario, the receiver tries all combinations of $(stNum + k_1, sqNum + k_2)$, where (k_1, k_2) vary from 0 to a certain threshold \mathbb{K} . A very small \mathbb{K} (e.g., 2) should be adequate in most cases.

In all cases, after decoding, the receiver will know the correct $stNum$ and $sqNum$, which means that the operation can continue. With PMU data, as long as the retransmission time is larger than the time of next event, there is no situation where a retransmission will be received. In other cases, the receiver can infer from the timings how to predict $stNum$ and $sqNum$. The only situation where a prediction could be incorrect is if the transmission delays vary a lot, so that the elapsed time does not provide a reliable basis for prediction. This is unlikely to be the case in SCADA networks.

The most time-consuming operation in the key generation is the secure hash; therefore, to ensure that the key is ready for each transmission, we require that the key-stream generation time T_{KS} should satisfy the inequality:

$$T_{KS} = N \cdot T_H = \left\lceil \frac{L}{L_d} \right\rceil \cdot T_H \leq T_1, \quad (4)$$

where T_1 is the minimum retransmission interval and T_H is the time to execute the hashing once.

Note that in case of message loss, the receiver may need to try decoding the message with multiple potential keys. This is acceptable for the following reasons: (a) the receiver, being a PDC or LCC, has substantially more computing power than the sender (a PMU), for example, a desktop/server level machine as opposed to a micro-controller, (b) the message loss probability is expected to be quite low for non-routable critical communications, and thus trying with multiple keys is needed only occasionally, and (c) when a message loss does occur, the 3 ms latency objective is unlikely to be met already, and a small additional delay should not be significant.

D. Salting Scheme for SV and IEEE C37.118.2

Since SV is specifically intended for streaming data without any retransmissions, the proposed mechanism is ideally suited for it. The salting scheme for SV is a little more complex since SV does not have an explicit sequence number. We can try to generate it by using $smpCnt$ of the first ASDU as the salt; however, $smpCnt$ is only 2-bytes, and it is reset every second. Thus, $smpCnt$ by itself is unable to provide a unique sequence number for a ASDU. However, this issue can be addressed by defining a virtual sequence number $VsqNum$, which is incremented each time the $smpCnt$ of first ASDU is reset to zero. Note that for this to work, we need to make the implicit assumption that the first ASDU sent by a PMU always concerns the same entity (e.g., the same bus). A 32-bit counter is quite adequate as it overflows in about 138 years. Thus, the pair $(smpCnt, VsqNum)$ can be a unique sequence number.

IEEE C37.118.2 messages are transmitted at relatively lower frequencies, around 30–60 frames/seconds [1], which gives a gap of 16.67–33.33 ms. Thus the key generation key of around 10 ms should be sufficient for IEEE C37.118.2. As for IEEE C37.118.2, we can use a combination of SOC and $FRACSEC$ as salt. The “TIME BASE” that defines the range of $FRACSEC$ is a configuration parameter and should not be changed during operation (it may be changed by taking the IED offline and making configuration changes). As the IEEE C37.118.2 messages are generated and transmitted periodically, the receiver can predict the next salt (or set of salts in case of packet loss) for generating the key-stream. Effectively, the pair $(SOC, FRACSEC)$ can act like a unique sequence number.

E. CRC Obfuscation

The common practice of appending CRC code to the message introduces a vulnerability in XOR based encryption due to the linearity property of CRC [8]. It is easy to verify that $CRC(X \oplus Y) = CRC(X) \oplus CRC(Y)$. Thus, an attacker could XOR the ciphertext X with an arbitrary message Y through a Man-in-the-Middle (MitM) attack and XOR CRC bits with $CRC(Y)$. This is easy to do if the CRC bits appear in a known position in the message (e.g., at the end). The attacker could also choose Y such that $CRC(Y) = 0$. To address this, we employ a keyed obfuscation algorithm to shield the CRC bits, disabling CRC

TABLE II: Keyed obfuscation and de-obfuscation

Key Bits	Obfuscation Operation	De-obfuscation Operation
00	Do nothing	Do nothing
01	Flip the current CRC bit	Flip the current CRC bit
10	Flip all CRC bits except the current bit	Flip all CRC bits except the current bit
11	Rotate the current CRC byte	Reverse-rotate the current CRC byte

cracking attacks. The obfuscation function should alter the CRC significantly and yet should be easily recoverable using the key.

In the obfuscation approach used here, we apply a transformation involving the n -th bit of the CRC by using the $2n$ and $(2n + 1)$ -th bits of the keystream KS as shown in Table II. As a result, the bits in the CRC segment become obscure and “non-linear” to attackers (unless they have access to the key). It also becomes impractical to perform brute-force attacks since repeated CRC verification failures is unusual and can be used to trigger alarm.

F. Message Embedding and Verification

After obfuscating the CRC, the publisher uses the first L bits of KS to encrypt the concatenation of the message and the obfuscated CRC by performing an XOR operation. Upon receiving the message, the subscribers reverse the above process by: (1) decrypting the message with the first L bits of KS ; (2) de-obfuscation the CRC based on the operations in Table II; (3) verifying the derived CRC. The received message is untampered if the CRC checking is passed.

Notice that this mechanism achieves the confidentiality using the one-time key encryption mechanism, whereas the message integrity is provided in two ways. First, the receiver predicts the salt for the next message to decrypt it, and after the decryption it matches the salt with the corresponding portion of the message. For example, in case of GOOSE, the salt consists of $stNum$ and $sqNum$, which can be matched with the original message after decryption. Similarly, in case of SV, the $smcCnt$ is used in the salt, which can also be checked with that of the decrypted message. In addition to matching the salt fields, CRC can further check for integrity.

G. Secret Key Management

Recall that the initial secret key \mathcal{K} is shared among the publisher and its subscribers. In XOR-based stream ciphers, attackers who have access to the original plaintext afterwards (from published PMU datasets) can gain access to the key-stream by XOR-ing the plaintext and the sniffed ciphertext, namely the N hash digests d_j . However, because of the one-way hash function, it is very difficult to derive the original string from there (which is required to get at the underlying fixed key). Nevertheless, there is some risk in using the same fixed secret key for extended periods. Therefore, the transmitter can exploit our encryption mechanism to occasionally provide a new key to the subscribers.

IV. SECURITY ANALYSIS

The proposed scheme provides confidentiality and integrity in the presence of passive or active attacks. A powerful attacker (for example, an inside attacker) can acquire both the plaintext and ciphertext of messages transmitted in the substation system. Thus, it can obtain the historical key-streams easily by performing an XOR calculation over the plaintext and ciphertext of the same messages. Note that the key-streams are the digests output by hash functions with the shared secret key and the synchronized salts. Through the historical key-streams, the attacker can collect hash samples, mappings between hash salts, and hash digests. Then, the attacker can make attempts to find out the secret keys sk_i with these samples.

With a 128 bit preconfigured key, a brute-force attack needs 2^{127} attempts to discover it. The SHA1 secure hash used as the $H(\cdot)$ function is known to have a collision attack length of 63 bits. In practice, this is more than adequate since the collision attacks are unstructured. However, we did use of SHA-256 as well, which not only raises the collision attack length to 128 bits [9], but turns out to be more efficient overall for reasons mentioned later. We have also tried SHA-512 which may or may not be desirable depending on the availability of long-word arithmetic. In the following, we analyze the security of the proposed method with respect to CRC and replay attacks.

A. CRC Attacks

We assume an active, strong adversary who has access to the encryption machine without the knowledge of the secret key. The adversary can control the parameters $stNum$, $sqNum$, and cnt , and can input arbitrary measurements to the encryption machine and generate corresponding ciphertexts. Such an adversary is able to encrypt different measurement payloads using the same key-stream. In particular, to mount a successful attack and perturb a message y without disturbing the integrity, the adversary only needs to generate a message x such that $CRC(x) = 0$. If $CRC(x) = 0$, then for all y , $CRC(x \oplus y) = CRC(x) \oplus CRC(y) = CRC(y)$. The adversary can compute $(KS \oplus (x||CRC(x))) \oplus (KS \oplus (y||CRC(y))) = (x||CRC(x)) \oplus (y||CRC(y))$, where KS is the key-stream. If $CRC(x) = 0$, then for all y , $CRC(x \oplus y) = CRC(x) \oplus CRC(y) = CRC(y)$.

The CRC obfuscation avoids such an attack by altering the bit value or the position of every CRC bit to the key KS . The result is an obfuscated CRC with length more than 32 bits. Without the key, attackers cannot determine which bits belong to the original CRC and therefore cannot leverage the linear properties of CRC to perform active attacks. Moreover, the intercepted messages are encrypted with a stream cipher, which further randomizes the bit value distribution of the obfuscated CRC. When intercepting a ciphertext, the attackers cannot determine the correct positions and thus cannot modify the original message at will.

B. Replay Attack

The adversary can launch a *replay attack* by overhearing a legitimate message and replaying it at some later time. In PreKIP, as the hashing is salted by incorporating the state number, sequence number, and a counter, a replayed message can easily be identified and discarded at the control center. The adversary may only succeed if he replays a message within a short time window of its origin, however, such an attack will not have a detrimental impact on knowing the current state of the grid and it cannot perturb the state as it would be easily detected and discarded.

V. PERFORMANCE EVALUATION

A key issue in the performance of an integrity protection algorithm is the resource-constrained microprocessors used by PMUs. Although the power of these processors is increasing, it will always be limited because of the small form factors. In this section, we show that PreKIP is substantially faster as compared to other well-known cryptographic algorithms and can meet the timing requirements of IEEE C37.118 even when implemented on a rather low-end microprocessor such as LPC2148 [10].

Comparison between PreKIP and other schemes: We first compare between PreKIP and other well-known security and integrity techniques are summarized them in Fig. 4a. The results are obtained from an Intel Core i5-6500T @ 2.50GHz processor with 16 GB RAM, with compiler optimizations set to minimize execution time. The motivation to use a desktop-class processor is to check how well the algorithm will perform for intra-substation use of the GOOSE/SMV protocols, which are likely to have such processors. We will later study performance on a microprocessor as well.

All the simulations are averaged over 100 runs. In PreKIP, we define the post-processing stage as the stage after the event (or packet) arrival, which includes CRC computation, obfuscation, and encryption. The post-processing stage on the subscriber side is symmetric.

In Fig. 4a, we have compared the post-processing stage of our proposed PreKIP scheme along with (a) AES-CBC-128 (128 bit AES with block chaining) and (b) HMAC-256 (256 bit hash based MAC). As compared to AES encryption, HMAC signature generation is $\sim 3x$ faster with a payload size of 400 Bytes. However, AES decryption performs poorly which results in $\sim 16x$ slower decryption than HMAC verification.

While comparing between PreKIP and HMAC-256, we observe that PreKIP is roughly 21x faster with a payload length of 10 Bytes, and $\sim 4x$ faster with 400 Bytes payload, than its nearest competitor HMAC-256, and even more as compared to others. In fact, PreKIP only takes 3–4 μs in the post-processing stage with is much lesser than the inter-packet transmission time of SV (i.e. 173–416 μs) and the minimum retransmission interval of GOOSE (i.e. 0.5–5 ms). Notice that in Fig. 4a HMAC embedding and verification takes up to 12 μs and 14 μs respectively, which also fulfills the timing requirement of both SV and GOOSE. However, PreKIP is several times faster, and


thus can be implemented even older/cheaper processors than HMAC. At the same time, PreKIP also provides confidentiality of the information in addition to integrity. This experiment clearly shows the lightweight nature of PreKIP that makes it suitable for such low-latency crypto applications.

Key Generation Latency of PreKIP: Fig. 4b shows the key generation latency of PreKIP. We have compared SHA-1, SHA-256 and SHA-512 for generating the digests. From Fig. 4b, we can observe that *the key generation using SHA-256 is faster as compared to SHA-1, whereas SHA-512 is the fastest.* This is because of the fact that for a message with length L , and a digest length of length L_d , the hash function is called $\lceil \frac{L}{L_d} \rceil$ times. Even if SHA-1 is faster than SHA-256, the $\lceil \frac{L}{L_d} \rceil$ for SHA-1 is larger due to its smaller digest size (20 bytes as compared to 32 bytes in case of SHA-256). Hence, SHA-256 performs much faster than SHA-1 in key-generation phase, especially for larger message length. For the same reason, the key generation time of SHA-512 is also faster than the other two. Also notice that the key generation of PreKIP is actually slower than HMAC embedding. This is because of the fact that in PreKIP the hash function is called $\lceil \frac{L}{L_d} \rceil$ times, as opposed to just once in HMAC. However, this does not matter for smart grid message forwarding so long as the key-stream is generated in between the samples (or retransmissions).


Performance of PreKIP on LPC2148: Fig. 4c shows the performance of PreKIP on LPC2148 microprocessor. LPC2148 ARM7TDMI microcontroller has a limited memory with 40 kB of on-chip static RAM and 512 kB of on-chip flash memory. Because of this limitation, it is not possible to implement arbitrary algorithms on it. Therefore, we have shown our evaluations only for PreKIP. Fig. 4c also shows faster key generation with SHA-256 as compared to SHA-1. Even with the largest message size of 400 bytes, the key generation is only 11 ms with SHA-256. With typical PMU reporting rate of 60/second, the total inter-sampling time is 16.67 ms, which is adequate for key generation, CRC computation, obfuscation and encryption using the LPC2148 microprocessor. Furthermore, the post-processing stage is extremely fast, and just takes less than a millisecond. This shows the feasibility of implementing PreKIP on a low-end microprocessor for real-life applications.

From Fig. 4c, we can also observe that LPC2148 microprocessor will not be suitable for SV or GOOSE message integrity checking when the sampling rate is 173–416 μs or the minimum retransmission interval is 0.5–5 ms. However, such high rates are unnecessary for transmitting PMU data (30-60 samples/sec or 1/2 sample/cycle). In such applications, PreKIP can provide a secured communication even with an inexpensive microprocessor like LPC2148.


It is seen that the key generation with SHA-512 is slower than that of SHA-256 in Fig. 4c. This is due to the use of 16/32-bit ARM7TDMI-S micro-controller [10] in LPC2148, which does not do well on long-word arithmetic as compared to a 64-bit Intel processor used in Fig. 4b.



(a) PreKIP vs. AES and HMAC.



(b) SHA-1 vs. SHA-256 for key generation



(c) PreKIP performance on LPC2148

Fig. 4: Performance of PreKIP

VI. RELATED WORK

Cyber security in smart-grid communication system is a well-mined area, and multiple research studies have been conducted in this area [11], [12]. As already indicated, IEEE C37.118 does not provide any security features [6]. Reference [13] provides a very recent review of cybersecurity challenges in IEC 61850-substation network. To alleviate the security issues, IEC 62351-6 [14] recommends RSA digital signature algorithm for signing and verifying the substation messages. However, the timing related performances are studied in [15]–[18], which raised concerns over the applicability of RSA-based signature scheme for GOOSE messages.

References [15], [18] have further studied that HMAC based schemes provide faster integrity check as compared to RSA-based schemes and can satisfy the timing requirements for the GOOSE messages with today’s commodity processors. In [19], the authors have implemented and compared a framework for RSA and MAC-based digital signature schemes. Reference [20] has presented a sequence hopping algorithm for securing the GOOSE messages. However, this requires a separate sequence synchronization and monitoring server, thus needs a separate infrastructure to be installed. As opposed to these contributions, we proposed a novel lightweight solution for the confidentiality and integrity problem mechanism, where the key is generated within subsequent message transfer, and thus the post-processing stage is about 4–21 times faster than an HMAC scheme. The scheme also provides extra confidentiality and is secure from well-known attacks.

VII. CONCLUSIONS

The paper presents PreKIP, a lightweight and secure integrity protection algorithm for critical smart-grid communications. It achieves low latency by generating a new key between samples that can be simply XORed with the generated sample. We show that our mechanism incurs negligible latency during the post-message arrival stage and is secure against powerful adversarial attacks. The proposed scheme is 4–21 times faster than competitive HMAC schemes, and even faster than other crypto algorithms. The proposed algorithm can also be used in other applications so long as there is a minimum gap between successive messages to allow for key generation without increasing the latency (e.g., vehicle-to-vehicle communications in intelligent transportation systems).

REFERENCES

- [1] S. R. Firouzi *et al.*, “Design, Implementation and Validation of an IEC 61850-90-5 Gateway for IEEE C37. 118.2 Synchrophasor Data Transfer,” Ph.D. dissertation, Escola Tècnica Superior d’Enginyeria, 2015.
- [2] J. R. Carroll *et al.*, “A comparison of phasor communications protocols,” Pacific Northwest National Lab.(PNNL), Richland, WA (United States), Tech. Rep., 2019.
- [3] Y. Weng *et al.*, “Robust data-driven state estimation for smart grid,” *IEEE Transactions on Smart Grid*, pp. 1–12, 2016.
- [4] Stefan Nohe, Oliver Hartmann, Farel Becker and Cedric Harisporu, “Designing non-deterministic pac systems to meet deterministic requirements,” http://rtpis.org/psc13/files/PSC2013_final_1358893882.pdf.
- [5] “IEC 61850 Sampled Values protocol,” https://www.typhoon-hil.com/documentation/typhoon-hil-schematic-editor-library/References/iec_61850_sampled_values_protocol.html.
- [6] I. Ali *et al.*, “Performance comparison of IEC 61850-90-5 and IEEE C37.118.2 based wide area PMU communication networks,” *Journal of Modern Power Systems and Clean Energy*, vol. 4, pp. 487–495, 2016.
- [7] P. Kundu *et al.*, “Synchrophasor-assisted zone 3 operation,” *IEEE Transactions on Power Delivery*, vol. 29, no. 2, pp. 660–667, 2014.
- [8] L. K. *et al.*, “Active attacks on stream ciphers with cyclic redundancy checks (crcs),” <http://www.cix.co.uk/~klockstone/crcheck.htm>, 2000.
- [9] G. Leurent *et al.*, “From collisions to chosen-prefix collisions application to full sha-1,” in *EUROCRYPT*, 2019, pp. 527–555.
- [10] “Arm7 lpc2148 development board,” <http://www.nex-robotics.com/products/development-tools/arm7-lpc2148-development-board.html>.
- [11] A. Abuadbba *et al.*, “Resilient to shared spectrum noise scheme for protecting cognitive radio smart grid readings- BCH based steganographic approach,” *Ad Hoc Networks*, vol. 41, pp. 30–46, 2016.
- [12] S. Garg *et al.*, “Secure and lightweight authentication scheme for smart metering infrastructure in smart grid,” *IEEE Trans on Industrial Informatics (IEEE-TII)*, vol. 16, no. 5, pp. 3548–3557, 2020.
- [13] T. S. Ustun *et al.*, “A review of cybersecurity issues in smartgrid communication networks,” in *ICPECA*, 2019, pp. 1–6.
- [14] IEC TS 62351-6, “Power systems management and associated information exchange - data and communications security - part 6: Security for IEC 61850,” 2007.
- [15] D. Ishchenko *et al.*, “Secure communication of intelligent electronic devices in digital substations,” in *IEEE/PES Transmission and Distribution Conference and Exposition*, 2018, pp. 1–5.
- [16] S. M. Farooq *et al.*, “Performance Evaluation and Analysis of IEC 62351-6 Probabilistic Signature Scheme for Securing GOOSE Messages,” *IEEE Access*, vol. 7, pp. 32 343–32 351, 2019.
- [17] S. S. Hussain *et al.*, “Analysis and implementation of message authentication code (mac) algorithms for goose message security,” *IEEE Access*, vol. 7, pp. 80 980–80 984, 2019.
- [18] G. Elbez *et al.*, “Authentication of goose messages under timing constraints in iec 61850 substations,” in *ICS-CSR*, 2019.
- [19] S. M. Farooq *et al.*, “S-GoSV: Framework for Generating Secure IEC 61850 GOOSE and Sample Value Messages,” *Energies*, vol. 12, no. 13, pp. 1–13, 2019.
- [20] Osama Mohammed, “Security aware microgrids:securing goose messages against cyberattacks,” in *APPA and FMPAA Southeast Regional Municipal Utility Cybersecurity Summit*, 2019.