# Evaluating Multicast Message Authentication Protocols for Use in Wide Area Power Grid Data Delivery Services

Carl H. Hauser, Thanigainathan Manivannan, David E. Bakken
School of Electrical Engineering and Computer Science, Washington State University
{hauser, tmanivan, bakken}@eecs.wsu.edu

## Abstract

*Modern computer communication technologies make it possible to create very flexible Data delivery service (DDS) design that can efficiently meet the quality-of-service needs of a wide variety of power system control and monitoring applications, some which only become possible with better communications. One aspect of the smart grid is that multiple applications will use data from each location. Furthermore, the use of the data for important power grid control functions requires each message's source be authenticated, leading to the need for multicast message authentication. All known multicast message authentication approaches carry trade-offs between quality-of-service (QoS) aspects such as added latency, computational cost, and the precise authentication guarantees they afford. This paper surveys several multicast authentication approaches for use in a power grid DDS based on an experimental evaluation of their latency and computational costs and an assessment of the appropriateness of their authentication guarantees for use in power grid applications.*

## 1. Introduction

Data delivery services (DDS) are now being designed to support the 21st century smart grid. Modern computer communication technologies make it possible to create very flexible DDS design that can efficiently meet the quality-of-service needs of a wide variety of power system control and monitoring applications, both those of today and those which will only become possible with better communications [1]. One aspect of the smart grid is that multiple applications will use data from each location. Furthermore, the use of the data for important power grid control functions requires that each message's source be authenticated, leading to the need for *multicast message authentication* (or *data origin authentication*). All known multicast message authentication approaches carry trade-offs between quality-of-service (QoS) aspects such as added latency, computational cost, and the precise authentication guarantees that they afford [2]. This paper surveys several multicast authentication approaches for use in a power grid DDS based on an experimental evaluation of their latency and computational costs and an assessment of the appropriateness of their authentication guarantees for use in power grid applications.

The evaluated approaches include those based on public-key cryptography, on symmetric key cryptography and hashing, on timed release of authentication keys (e.g. TESLA [3]) and those that mix the two approaches [e.g. Time-valid One Time Signatures [4]). We also consider the interactions between DDS system design decisions and authentication design decisions.

## 2. Related Work

Fuloria *et al.* address the problem of securing communications within electric utility substations [5]. This topic is also the subject of the IEC 62351-6 standard. Fuloria concludes that the draft standard's requirement to use public-key based authentication for messages inside the substation is unnecessary and impractical if not impossible. The main barrier is the computational cost associated with public key methods even with the assistance of expensive, hard-to-integrate at scale, hardware support. In a later work some of the same authors address the key management issues associated with using various message authentication techniques in substations and again conclude that public-key based methods are more costly than other available approaches for little to no practical increase of security.

Challal *et al.* provides a comprehensive taxonomy and overview of multicast data origin authentication techniques, along with their strengths and weaknesses [2]. He et al. provides a briefer survey [6]. These works are not specialized for the needs of the power grid, but provide much initial guidance.

Complementing the above work, we consider the data origin authentication needs for multicast in the *wide area* control and monitoring systems that will deliver the benefits of the smart grid.

## 3. Overview of Real-Time Multicast and Security

Increasing numbers of so-called intelligent electronic devices are being deployed throughout transmission and distribution levels of electric grids. These devices, which measure and report on the grid at high rates and with global time synchronization are a key enabler for the emerging smart vision. Using these data effectively to improve the reliability and efficiency of the grid requires that the data be made available to monitoring and control applications. So a second key smart grid enabler is creation of a data delivery service (DDS) for securely moving data from where it is produced to where it is needed. Proposed (and existing) applications for the data demonstrate that the DDS must be able to deliver data from any given source to multiple destinations. Thus logical multicast—the ability to send each measurement to multiple recipients—is a fundamental requirement [1]. While the requirement might be met by sending the measurement separately to each recipient, implementing a multicast publish-subscribe framework within the DDS offers many opportunities for efficiency and simplification. In the following discussion therefore we assume that a multicast mechanism is part of the DDS, though we emphasize again that the need for multicast message authentication arises from the application space, which inherently requires 1:many communication, and not because a multicast mechanism is implemented in the DDS.

### 3.1. Requirements for Highly Available Real Time Multicast

Electric power grids require a vast array of operational data to provide efficient power delivery and to help avoid blackouts. The diversity of the different kinds of data and the power applications using them has increased recently with the advent of synchrophasors, and this tread is expected to continue and possibly increase. Additionally, today a relatively small (but quickly growing) fraction of overall traffic is used with a computer in the loop. The vast majority of data are used with a human in the loop, namely as input for visualization tools for operators. However, this too is changing. It is widely agreed that power grids today in the industrial world are getting more stressed each year for a variety of reasons including lack of new transmission lines to keep up with the growth of load and generation, retiring operators, integrating renewable energy sources which have power characteristics and interactions that are not well understood, and other factors. These stressors are all forces driving the use of even more sensor data being used with a wider array of power applications, geographic scopes, patterns of sharing data, etc.

Some of these operational data are condition-based, for example triggered by an abnormal condition or other thresholds being crossed by sensor data. However, the majority of today's operational data, and the vast majority of newly added and emerging operational data, is rate-based, namely the periodic delivery of remote sensor data [1]. In our analysis here we focus on rate-based operational data, which is increasingly being delivered via publish-subscribe or other multicast mechanisms. We also are only describing data that is flowing outside the scope of a single substation, not within only one substation.

Rate-based operational data has very challenging data delivery requirements [1], which require great care to implement [7]. Key requirements include the following:

**Latency**: latencies required can be as low as 5 msec to several seconds or longer.

**Rate:** Rates can be as low as once per second or minute, and as high as 720 Hz. Higher rates are thinkable; for example, the FNET frequency recording device internally samples and then computes frequencies at 1440 Hz [8].

**Availability**: The most critical data require 99.9999% availability while the least critical require only 99% [9]

We also note that such delivery requirements must be met over a range of geographic areas (up to an entire power grid, which can be in the neighborhood of half a continent), and for millions of data streams, each with multiple subscribers which can have different delivery requirements (for the same sensor data stream) which must be met [1]. Additionally, these delivery requirements must be met for each update for each subscriber, not over wider periods of time and sets of subscribers or sensor streams.

### 3.2. Multicast Security Issues for Wide-area Power Grid Applications

The variety of functional and quality-of-service DDS requirements for smart grid applications, described in the previous section, leads naturally to a variety of requirements for message origin authentication. Because the DDS is intended to support a wide variety of applications, and because to be economical and effective security mechanisms must be matched to the threat, it is apparent that different applications may well require different approaches to security, and specifically to message authentication, especially when the capabilities of the different authentication schemes are considered.

As in other domains, the security considerations revolve around confidentiality, integrity (including authentication), and availability, but in the power grid the additional considerations of added latency, computational cost, and increased message size are also very important. In many cases, confidentiality is of far less concern than the other two: the ability to see the content of messages is of far less value to an attacker than the ability to insert false messages or prevent or delay delivery of legitimate messages.

Using mechanisms that rely on cryptography for any of these purposes introduces system-level issues about key management and trust [10], [11]. At best, cryptographic authentication primitives securely associate a message with an encryption key. Whether the recipient can trust that the key is correctly bound to a particular sender and still only available to that sender depends on overall systems-level characteristics that go far beyond formal cryptographic underpinnings.

In summary, data delivery for the power grid requires authentication methods that can meet a variety of security and quality-of-service requirements. Not every combination of requirements can be feasibly met. The choice of an authentication mechanism may well have to compromise on some requirements. At the same time, system-level design choices and operational considerations may well mean that the theoretically strongest mechanisms in practice provide little or no benefit over theoretically weaker but less costly approaches.

## 4. Descriptions of the Protocols

There are a variety of approaches to multicast message authentication in the literature. The ones that we survey use different cryptographic underpinnings to achieve different results in performance and security. We now describe them in turn.

### 4.1. Message Authentication Codes using Shared Keys

Block-mode ciphers, such as AES [12], and cryptographic hash functions, such as the SHA family [13] used with shared symmetric keys can each form the basis of a message authentication code (MAC). The techniques are well-researched and exist in standardized forms: HMAC [14] is based on hash functions and CMAC [15] is based on block-mode ciphers.

Because both AES encryption and SHA hashing are fast functions to compute (on modern 32-bit microprocessors a few microseconds per kilobyte of encrypted/hashed message) they are widely used for uni-cast applications where latency and computation time are matters of concern. Because the key is shared between the publisher and the subscriber these keyed MACs do not provide non-repudiation—it is obvious that the subscriber can create a message and MAC that is indistinguishable from the same message and MAC sent by the publisher).

In multicast applications, however, the keyed MAC approach suffers: if the key is shared between the publisher and all of the subscribers then any of the subscribers can send authenticated messages as if they were the publisher. If on the other hand a different key is used for each subscriber then the publisher incurs O(number-of-subscribers) computation cost and message sizes increase similarly, negating some of the benefit of using a multicast transport mechanism. (Recall that the issue is fundamentally about multiple recipients of the same message contents, regardless of whether those contents are delivered by a unicast or multicast transport mechanism.)

### 4.2. Public-key based techniques: RSA and DSA

The basic principles of public-key techniques are thoroughly covered in many texts (e.g. [16]) and the readers are assumed familiar with them. The attraction of the public-key-based techniques is that they inherently prevent subscribers from masquerading as the publisher.

These and other public-key techniques are the only ones that provide non-repudiation. It is clear that non-repudiation is at least arguably of more value in the wide-area system than in a single substation (a question addressed in [5] where the authors observe that non-repudiation is unnecessary in a substation under the control of a single entity).

On the other hand these public-key techniques come at extreme cost in added latency and computational requirements. Our own measurements as well as those of others (e.g. [5]) find that implementations of RSA and DSA on stock multiprocessors require on aggregate tens of milliseconds for signing and verification. Using these public-key techniques therefore adds substantial latency to message delivery (a message is considered to be deliverable only after authentication is complete). Moreover, the cost is in the form of active computation rather than simply waiting for sufficient time to pass (as in TESLA, below) so using these algorithms on low-powered processors or processors that are already taxed to provide their intended functionality is quite problematic. Some have proposed using special-purpose cryptographic co-processors to improve the latency associated with these protocols, but we find little reason to disagree with the conclusion reached in [5] that the co-processor approach is likely to be excessively expensive both for the hardware itself and for the systems integration engineering required to use it, difficult to manage over the long lifetime of power grid equipment, and in general not a feasible solution.

## 4.5. Timed Efficient Stream Loss-Tolerant Authentication (TESLA)

The basic idea behind the Timed Efficient Stream Loss-Tolerant Authentication (TESLA) technique [3] is that the publisher signs each message with a key that it reveals only after every subscriber should have received the message. Time synchronization between publisher and subscribers along with timestamps in the messages allow subscribers to determine whether the key was published after they received the message. If so, and the signature is verified successfully, the message could only have come from the publisher.

TESLA's signing keys are elements of a one-way hash chain generated by a publicly known function and used in reverse order of generation, as in the S/Key protocol [17]. Such chains can be generated at low computational cost. Each later-arriving key at a subscriber should be a predecessor of an earlier-arriving key in the hash chain. Therefore a subscriber can check that the later key belongs to the publisher by applying the hash function to it, possibly multiple times, to see if an earlier-received (and checked) key results. If so the key (and all of the iterated hashes of it) can be used to check the signatures of earlier-arriving messages.

Note that TESLA burdens subscribers with buffering messages from the time they arrive until the time that a key arrives that can be used to authenticate them and that for all subscribers this interval is at least the delivery latency to the subscriber most distant from the publisher. Thus, although TESLA is computationally fast, it adds significant latency to the time to authenticate every message.

TESLA prevents subscribers (even in collusion) from masquerading as the publisher provided that time synchronization between nodes remains within the designed-for tolerance. Attacks on TESLA-based authentication might be mounted through attacking the time-synchronization service (allowing masquerading) or by delaying messages for too long resulting in denial-of-service.

Power-grid sensing technologies such as PMUs are moving to using GPS-based clocks for timestamping measurements [17], [18], [19], [20], [21], [22]. Thus, it is a reasonable assumption that very good clock synchronization (within a millisecond) will be available to components of the data delivery service. Furthermore, disruption of this time-synchronization will potentially have much greater effects on the value of the measurements themselves than on the ability to authenticate the measurements.

## 4.6. Time-Valid One Time Signatures (TV-OTS)

The Time-Valid One Time Signature (TV-OTS) technique exploits research on public-key signatures with the interesting property that using the private key to sign a message reveals something about the private key [4]. If too many messages are signed with the key, or if too long elapses between signing with the key and verification of the signature, an intruder may be able to determine the key and forge messages. The authors observe that the underlying one-time signature scheme that they use is strong enough that under very conservative assumptions about the ratio between the computational resources available to the signer and an attacker it is safe to use the one-time key to sign several times over a limited period of time. (Where *several* and *limited period* depend on parameters derived from the assumptions.) As with TESLA, time synchronization between the publisher and subscribers is required, although TV-OTS is tolerant of much looser synchronization with little impact on performance.

TV-OTS also exploits hash chaining to extend the lifetime of the "one-time" keys but as with TESLA these chains are necessarily finite so periodic generation and dissemination of new public keys is required. The authors suggest that by appropriately balancing the chaining and security parameters of the system it would be reasonable to authenticate a power grid publisher's messages for about an hour using one key pair. The biggest drawback of TV-OTS is that generating key pairs is computationally expensive requiring several minutes of computation.

| Scheme | Sender Computation Cost | Receiver Computation Cost | Message Size Overhead | Packet Buffering | | Key Size | Total Latency |
|---|---|---|---|---|---|---|---|
| | | | | Sender | Receiver | | |
| AES | 1E | 1D | 1k | 1 | 1 | O(1) | 1E+1D+ND |
| RSA | 1S | 1V | 1k | 1 | 1 | O(1) | 1S+1V+ND |
| TESLA | 1H | 1H | 1k+1h | 1 | Є | O(1) | 2H+ND+KD |
| TV-OTS | 1H | 1H | 0.25*h | 1 | 1 | O(N) | 2H+ND |

**E** ⍰ Encryption **D** ⍰ Decryption **k** ⍰ #keys **H**⍰Hash **M** ⍰ Message authentication code
**Є** ⍰ #packets/key disclosure delay **N** ⍰ # hash chains used **S** ⍰ Signature
**V** ⍰ Verification **KD** ⍰ Key disclosure delay **TV-OTS** ⍰ Time Valid One time signature
**TESLA** ⍰ Timed Efficient Stream Loss-Tolerant Authentication **ND** ⍰ Network delay **h** ⍰ Hash size

**Table 1: Theoretical Performance of Authentication Protocols**

How and whether this might be securely off-loaded from the publisher is not clear at this time.

Apart from the computation needed for key pair generation and the overall implementation complexity TV-OTS is an attractive choice: both signing and verification are computationally fast and no additional latency is incurred waiting for keys to be revealed.

## 5. Assessment of Protocol Performance and Effectiveness for Power Applications

In choosing multicast authentication schemes for power grid applications the actual latency and computation times required have to be assessed in light of the applications' requirements. To help provide guidance to designers, we have conducted the assessment both in parameterized, symbolic form, both extending and condensing results of Challal et al. [2], and in absolute form to gain a sense of the magnitudes of the available performance. Table 1 expresses the computation time, message overhead, and buffering requirements associated with the various protocols in terms of more primitive, but still high-level, operations such as cryptographic hashing or symmetric encryption.

Table 2 reports our measurements of computation time associated with various authentication techniques. Recall that the computation time is important both because of its contribution to latency and because it uses computational resources that may be needed for other purposes. These measurements were made using cryptographic functions from the Java SE Runtime Environment running on the Java HotSpot 64-bit server VM from Java version 1.6.0_26, an AMD Phenom II X4 920 processor with a 2.8GHz clock, and Ubuntu Linux version 10.10. Obviously, computational performance varies greatly depending on platform but these data provide an initial, order-of-magnitude, starting point for selecting authentication techniques for different applications.

| Algorithm/ Protocol | Pub. (ms) | Sub. (ms) | Total (ms) |
|---|---|---|---|
| AES (128 bit) | 0.04 | 0.03 | 0.07 |
| SHA-256 | 0.01 | 0.01 | 0.02 |
| RSA (2048 bit) | 59.00 | 2.04 | 61.04 |
| DSA (1024 bit) | 5.10 | 9.80 | 14.90 |
| TV-OTS | 0.04 | 1.42 | 1.46 |
| TESLA | 0.03 | 0.03 | 0.06 |

**Table 2: Computational Costs for Different Algorithms at Publisher and Subscriber Nodes**

Figure 1 represents the combination of computational cost and network latency to give starting-point estimates for the overall latency of different techniques.

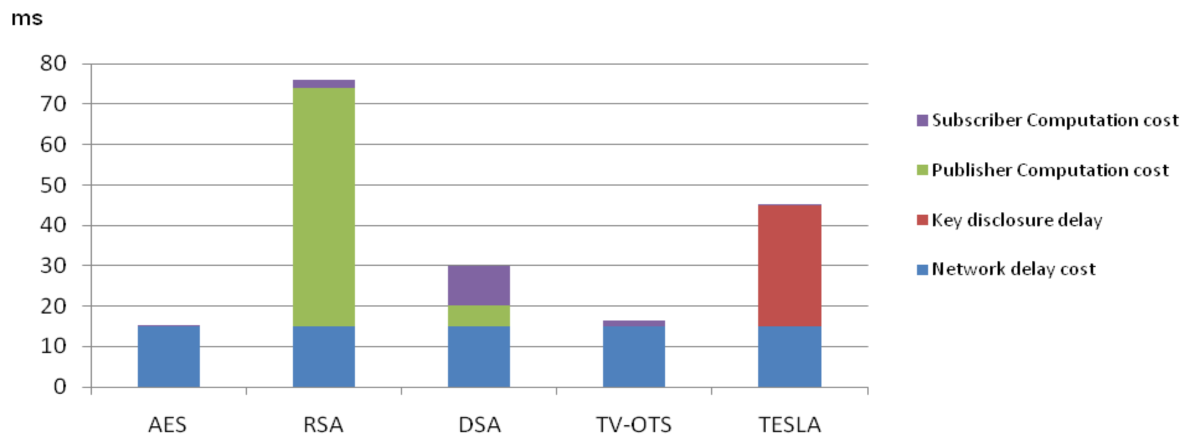

* Average network delay of the WAN network is estimated as 15ms and the maximum delay is estimated as 30ms for this illustration

**Figure 1: Performance Results of Different Multicast Authentication Algorithms**

## 6. System Context

We have shown that different message authentication techniques have different performance and security properties that must be taken into account along with an application's security requirements in choosing the technique to use. Aspects of the system context in which the application is embedded are also important influences. As an example, consider the GridStat DDS [1], [23] depicted in Figure 2. GridStat provides a rate-filtered publish-subscribe interface to power grid sensor data, as an overlay network using, for example, an IP network as the lower-level networking layer. Amongst other mechanisms, the GridStat implementation includes capabilities for multicast and for inserting security modules in the data path between publishers and subscribers, namely in their GridStat-provided proxy code [24], [25].

GridStat multicast trees are single source and change only under the control of a separate management plane. It is therefore easy, without using any cryptographic mechanism at all, to ensure that messages claiming to belong to a particular publication can only traverse specific links and in a particular direction. This means that forged messages, to have any effect, must be injected into

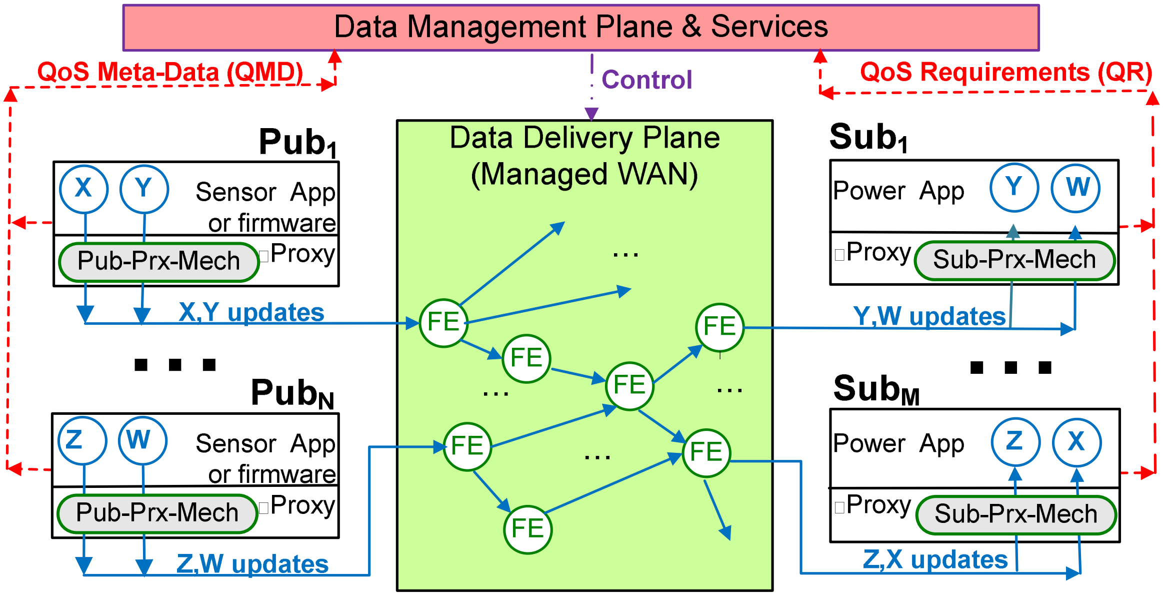

**Figure 2: Architecture and System Model of GridStat**

the tree on a valid link. Furthermore, attempts to insert messages at an inappropriate location should trigger an intrusion response, given that GridStat's forwarding engines (FE) know all traffic (including rates) that is authorized to flow through it. This raises the barrier to injection of forged messages even for techniques that don't protect against that cryptographically.

GridStat provides another example of how the encompassing system influences the choice of authentication mechanisms. Some message stream authentication mechanisms cannot cope with missing messages: missing a message prevents authentication of later messages. Two different design decisions in GridStat interact with this aspect of message authentication. First, GridStat does not provide guaranteed message delivery, preferring mechanisms such as redundant, disjoint trees to increase the probability of successful message delivery in what is fundamentally a best-effort setting. The motivation is that it is better to timely deliver the *next* measurement than to spend resources (for example to use acks/nacks) to deliver the *previous* measurement late. This proactive sending of messages in parallel, coupled with enforcement of a complete admission control perimeter, can provide extremely high availability while still supporting extremely low latencies [1]. Second, GridStat rate filtering is specifically designed *not* to deliver messages to subscribers at rates higher than are useful to them (this is a subscription parameter). Thus, some subscribers, by design, do not receive some fraction of the published messages.

## 7. Recommendations and Conclusions

Our research has shown that the variety of contexts in which multicast message authentication is required, especially the different combinations of latency and security needs, argues against any imposition through standards or regulations of a single authentication approach. In particular, we find that in the wide area, as Fuloria *et al.* found within substations, that public-key based message authentication is too costly for too little benefit for almost all use cases. Indeed, they increase end-to-end delivery latencies to rule out some of the very fast applications such as Line Current Differential Protection [26] and the faster of system integrity protection schemes [19].

Simple techniques based on cryptographic hashes or symmetric key cryptography are fast but are susceptible to key disclosure by a single recipient leading to the possibility of forged messages. System policies about where messages can be physically introduced may reduce the possibility of such attacks.

Timed key release protocols like TESLA provide better security while keeping computational cost low but also suffer from increased message delivery latency. TV-OTS has both low latency and low per-message computational cost but periodically requires a costly one-time-key generation.

Given the above, we believe that a power grid DDS needs to provide the ability to accommodate different message authentication protocols for different applications and data origins. That is, providing "one size fits all" authentication would be inappropriate for today's and tomorrow's power grids.

## 8. Acknowledgements

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097, Trustworthy Cyber Infrastructure for the Power Grid (TCIPG).

## 9. References

[1] Bakken, D.E., Bose, A., Hauser, C.H., Whitehead, D.E., and Zweigle, G.C. "Smart Generation and Transmission With Coherent, Real-Time Data", *Proceedings of the IEEE*, 99:6, June 2011, pp. 928-951.

[2] Yacine Challal, Hatem Bettahar, Abdelmadjid Bouabdallah. A taxonomy of multicast data origin authentication: Issues and Solutions, *IEEE Communications Surveys and Tutorials - COMSUR*, 6:1-4, 2004, pp. 34-57.

[3] A. Perrig, D. Song, R. Canetti, J.D. Tygar, B. Briscoe. "Timed Efficient Stream Loss-Tolerant Authentication (TESLA): Multicast Source Authentication Transform Introduction", *Internet RFC 4082*, The RFC Editor, June 2005.

[4] Qiyan Wang, Khurana, H., Ying Huang, and Nahrstedt, K. "Time Valid One-Time Signature for Time-Critical Multicast Data Authentication", *IEEE INFOCOM*, Rio de Janeiro, 2009.

[5] Shailendra Fuloria, Ross Anderson, Kevin McGrath, Kai Hansen and Fernando Alvarez. "The Protection of Substation Communications," in *Proceedings of SCADA Security Scientific Symposium*, Jan 2010.

[6] Jinxin He, Gaochao Xu, Xiaodong Fu, Xhiguo Zhuo, Jianhua Jiang, "Survey on Multicast Data Origin Authentication", *Communication Technology, 2008*. ICCT 2008. 11th IEEE International Conference on, On page(s): 749 - 752, Volume: Issue: , 10-12 Nov. 2008.

[7] H. Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications, 2ed*. Springer, 2011.

[8] Z. Zhong, C. Xu, S. Tsai, L. Zhang, V. Centeno, A, Phadke, Y. Liu, "Power System Frequency Monitoring

Network (FNET) Implementation", *IEEE Transactions on Power Systems*, 2005.Vol. 20. No.4. Nov. 2005. Pp.1914-1921.

[9] Electric Power Research Institute (EPRI), The Integrated Energy and Communication Systems Architecture, Vol. IV: Technical Analysis, 2004

[10] Himanshu Khurana, Rakesh Bobba, Tim Yardley, Pooja Agarwal, and Erich Heine. "Design Principles for Power Grid Cyber-Infrastructure Authentication", in *Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS)*, IEEE, Honolulu, Hawaii, January 5-8, 2010.

[11] Shailendra Fuloria, Ross Anderson, Fernando Alvarez and Kevin McGrath. "Key Management for Substations: Symmetric Keys, Public Keys or No Keys?" in proc. *IEEE Power Systems Conference and Exposition (PSCE)*, Phoenix, Arizona, March 20-23, 2011.

[12] National Institute of Standards and Technology. "Specification for the Advanced Encryption Standard (AES)", Federal Information Processing Standards Publication 197, U.S. Department of Commerce, November 2001 at http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf

[13] National Institute of Standards and Technology. "Secure Hash Standard (SHS)", Federal Information Processing Standards Publication 180-3, U.S. Department of Commerce, October 2008 at http://csrc.nist.gov/publications/fips/fips180-3/fips180-3_final.pdf

[14] National Institute of Standards and Technology. "The Keyed-Hash Message Authentication Code (HMAC)", *Federal Information Processing Standards Publication 198*, U.S. Department of Commerce, March 2002.

[15] Morris Dworkin, "Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication", *NIST Special Publication 800-38B*, National Institute of Standards and Technology, U.S. Department of Commerce, May 2005.

[16] Schneier, Bruce. Applied Cryptography: Protocols, Algorithms, and Source Code in C (2ed), John Wiley, October 1996.

[17] N. Haller. "The S/KEY One-Time Password System", Internet RFC 1760, February 1995.

[18] Martin, K.; Carroll, J., "Phasing in the Technology," *Power and Energy Magazine, IEEE* , vol.6, no.5, pp.24-33, September-October 2008

[19] Horowitz, S.; Novosel, D.; Madani, V.; Adamiak, M.; , "System-wide Protection," *Power and Energy Magazine, IEEE* , vol.6, no.5, pp.34-42, September-October 2008.

[20] Thorp, J.; Abur, A.; Begovic, M.; Giri, J.; Avila-Rosales, R.; , "Gaining a Wider Perspective," *Power and Energy Magazine, IEEE* , vol.6, no.5, pp.43-51, September-October 2008.

[21] Phadke, A.G.; de Moraes, R.M.; , "The Wide World of Wide-area Measurement," *Power and Energy Magazine, IEEE* , vol.6, no.5, pp.52-65, September-October 2008.

[22] Novosel, D.; Madani, V.; Bhargava, B.; Khoi Vu; Cole, J.; , "Dawn of the grid synchronization," *Power and Energy Magazine, IEEE* , vol.6, no.1, pp.49-60, January-February 2008.

[23] H. Gjermundrød, D.E. Bakken, C.H. Hauser, A. Bose "GridStat: A Flexible QoS-Managed Data Dissemination Framework for the Power Grid", *Power Delivery, IEEE Transactions on*, Vol. 24, pp. 136-143, 2009.

[24] Solum, Erik, Hauser, Carl, Chakravarthy, Rasika. "Modular over-the-wire configurable security for long-lived critical infrastructure monitoring systems", in *Proc. 3rd ACM Int'l Conf. on Distributed Event-Based Systems (DEBS 2009)*, Nashville, TN, July 2009.

[25] Chakravarthy, R., Hauser, C.H., & Bakken, D.E. (2010, December). Long-Lived Authentication Protocols for Critical Infrastructure Process Control Systems. *International Journal of Critical Infrastructure Protection*. 3(3-4), 174-181.

[26] H. Miller, J. Burger, N. Fischer, and B. Kasztenny. "Modern Line Current Differential Protection Solutions", in *Proceedings of the 2009 Western Protective Relay Conference*, Spokane, WA, October 2009. From http://www.selinc.com/WorkArea/DownloadAsset.aspx?id=6390