

Parallel Vectorized Algebraic AES in MATLAB for Rapid Prototyping of Encrypted Sensor Processing Algorithms and Database Analytics

Jeremy Kepner, Vijay Gadepally, Braden Hancock, Peter Michaleas, Elizabeth Michel, Mayank Varia
MIT Lincoln Laboratory, Lexington, MA, U.S.A.

Abstract— The increasing use of networked sensor systems and networked databases has led to an increased interest in incorporating encryption directly into sensor algorithms and database analytics. MATLAB is the dominant tool for rapid prototyping of sensor algorithms and has extensive database analytics capabilities. The advent of high level and high performance Galois Field mathematical environments allows encryption algorithms to be expressed succinctly and efficiently. This work leverages the Galois Field primitives found in the MATLAB Communication Toolbox to implement a mode of the Advanced Encrypted Standard (AES) based on first principals mathematics. The resulting implementation requires 100x less code than standard AES implementations and delivers speed that is effective for many design purposes. The parallel version achieves speed comparable to native OpenSSL on a single node and is sufficient for real-time prototyping of many sensor processing algorithms and database analytics.

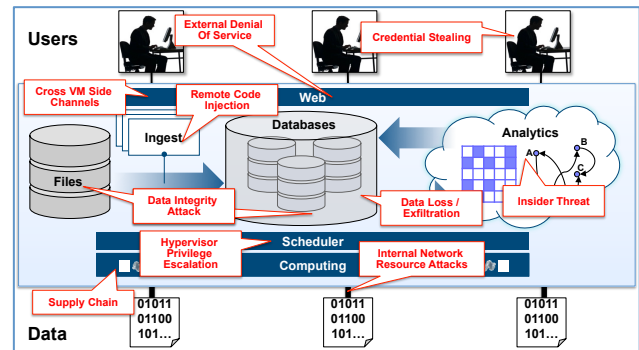
Keywords—Signal Processing; Databases; Encryption; Security; Algorithms; MATLAB

I. INTRODUCTION

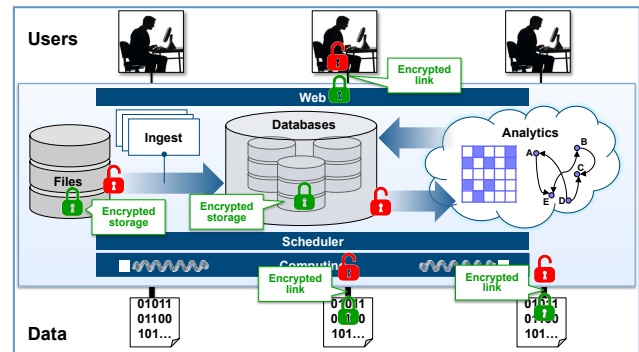
Modern buildings [Hubbell 2012] and vehicles [Gadepally 2014] contain many networked sensors to monitor and control their environmental, security, and safety systems. The decreasing size and increasing capability of these networked sensors has led to their increased usage on people [Starnar 2014]. These wearable technologies provide a wide variety of data that will extend and improve human lives. This trend towards increased use of networked sensors is often referred to as the Internet-of-Things (IoT) [Atzori 2014].

IoT data is stored and analyzed in networked databases accessible to a wide range of stakeholders: users, family members, care providers, operators, manufacturers, and regulators. Sensor algorithms and database analytics are the primary mechanisms for transforming raw IoT data into useful information. Incorporating encryption directly into sensor algorithms and database analytics is one method for adding security protections to IoT systems.

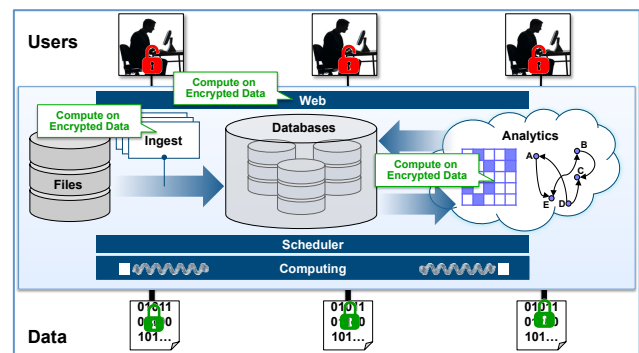
The critical role that IoT data plays in controlling key aspects of society necessitates that a variety of protection measures be included in the design of IoT systems. Figure 1 depicts a common architecture for connecting IoT data sources with users via a range of technologies.



(a) IoT Security challenges



(b) IoT Security current approaches



(c) IoT Security vision

Figure 1. Standard architecture for connecting diverse data and users. Shown are big data veracity challenges (a), current approaches (b), and longer term vision (c).

This work is sponsored by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, recommendations and conclusions are those of the authors and are not necessarily endorsed by the United States Government.

There are many security challenges in such an IoT data processing system (see Figure 1a): external denial of service, credential stealing, cross virtual machine (VM) side channels, VM hypervisor privilege escalation, remote code injection, data integrity attacks, data loss/exfiltration, insider threats, internal network resource attacks, and supply chain attacks [Evans 2013].

These attacks threaten to degrade the availability of the IoT processing system, compromise the confidentiality of the data and analytics used, and violate the integrity of both the original data and the analytic results. There are several approaches to mitigating these security challenges. Data centric protections are particularly useful for preserving the confidentiality of the data. Typical defenses of this type include (see Figure 1b): encrypting the links between users and the IoT processing system, encrypting the links between the data sources and the IoT processing system, encrypting the data in the file system, and encrypting data in the database when it is at rest. These approaches are all significant steps forward in improving the security of an IoT processing system. However, all of these approaches require that the data be decrypted for it to be used inside the IoT processing system, which requires that the keys to the data be available to the IoT processing system thus exposing the IoT processing to any attacker able to breach the boundaries of the system.

One vision (Figure 1c) for an IoT processing system is to have data sources encrypt data prior to transmitting it to the system, have the IoT processing system operate on the data in encrypted form, and only allow authorized users the keys to decrypt the answer for their specific result. Such a system makes the underlying IoT processing technologies oblivious to the details of the data and would go a long way towards mitigating the IoT processing security challenges described above. As a result, the data and processing can be outsourced to an untrusted cloud while preserving the confidentiality of the data and results.

There are several cryptographic tools that one could use to build a system like the one shown in Figure 1c. Fully homomorphic encryption allows for arbitrary analytic computations to be performed on encrypted data without decrypting it while preserving its semantic security (i.e., no information about the data is leaked other than its length). Fully homomorphic encryption has been an active topic of research since its discovery [Gentry 2009]. Nevertheless, the best currently available schemes [Perl 2011, Halevi 2014] have an overhead of 10^5 or more, making them too slow for use in many systems [Varia et al 2015a].

If one is willing to allow a small amount of information about the encrypted data to be revealed, a higher speed approach is to design protocols that leverage more traditional cryptographic techniques to carry out queries on encrypted data. One example of such a protocol is CryptDB [Popa 2011], which constructs a practical database system capable of handling most types of SQL queries on encrypted data. It uses deterministic encryption, which always encrypts the same data to the same ciphertext, to enable equality queries; order-preserving encryption, which encrypts data in a way that preserves the original order of the data, to enable range queries; and additively homomorphic encryption, which enables summing values directly on encrypted data, to perform basic

analytics. Several other protocols achieving alternative trade-offs between leakage and efficiency have been proposed by [Raykova 2012, Pal 2012, Cash 2013, Varia et al 2015b]. Additional solutions are also possible using techniques for secure multi-party computation [Yao 1982, Ben-Or 1988] but these require further improvement to achieve the required speed.

Incorporation of the aforementioned cryptographic techniques into secure algorithms and analytics requires a co-design with encrypted computation and encrypted query. Currently these areas have few common design processes or tools and employ different design criteria, mathematics, and programming.

A typical sensor algorithm involves filtering the raw data, correlating with known signatures, detecting significant correlations, and summarizing the precise location of the detections in the data. Throughout the sensor algorithm development process the designer seeks to balance a variety of system characteristics such as signal-to-noise ratio, probability of detection, and probability of false alarm. The primary mathematics of sensor algorithms is linear algebra, and the most common algorithm development environment is MATLAB.

A typical database analytic might first select records from a database, perform a database join on the records, construct a graph out of the records, and cluster the graph to find records of most interest. During the design of a database analytic, the designer will simultaneously balance database system characteristics such as ingest rate and query latency. The primary mathematics used is set theory, and SQL is the default language of database environments.

Encrypted computation utilizes a variety of techniques such as homomorphic computation, fully homomorphic computation, multi-party computation, and computing on masked data. The design of encrypted computation focused on developing adversary models and understanding the information leakage of various encryption techniques with respect to different adversary models. The mathematics of encryption relies on number theory and in particular Galois Fields. There are a number of environments for developing encrypted computation, and Python is widely used.

Encrypted query relies on various approaches for obscuring the data while still allowing some search to occur: for example, deterministic and order preserving encryption. These are often layered in an “onion” so that the strongest encryption is on the outside of the data and can be “peeled” to the weaker levels as needed. Furthermore, the actual patterns of query can also be designed in a privacy-preserving manner so that the query itself reveals as little information as is possible. Encrypted query is similar to encrypted computation and focuses on minimizing information leakage with respect to specific adversary models. Likewise similar Galois Field mathematics are also used. The dominant programming environment is the same SQL used for standard database analytics.

An important step towards enabling the co-design of secure algorithms and analytics is to provide common environments where co-design can occur. This requires both mathematical and technological unification of the environments. Associative array provides a common mathematics for sensor algorithms

and database analytics [Kepner 2012, Chaidez 2014, Gadepally 2015a, Kepner 2015, Kepner & Jansen 2016]. Computing on Masked Data has demonstrated encrypted computation and query on associative arrays using external encryption libraries [Kepner 2014, Gadepally 2015b]. The next step is to integrate encryption mathematics (Galois Fields) directly into these associative array environments. The advent of high performance Galois Field mathematics in high level programming environments such as MATLAB and Magma [Agullo 2009] now make it feasible to incorporate encryption directly into sensor processing and database analytic design environments. This paper demonstrates this capability with the Advanced Encryption Standard (AES) algorithm and shows that fast encryption can be performed in a co-design environment suitable for sensor algorithm development and database analytic design.

We caution here that MATLAB is primarily an algorithm prototyping tool. As such, our implementation is not intended directly for use in practice. Any production implementation of AES should take into account side channel attacks that can exploit variations in the software's running time [Kocher 1996] or power consumed [Kocher 1999, Chari 2002, Peeters 2007] in order to learn the secret key. Such attacks are very effective in practice, often requiring just tens or hundreds of measurements to learn the key completely [DPA contest, Osvik 2006]. Production implementations of AES require careful handling to provide proper defenses against side channel attacks [Messerges 2001].

The outline of the rest of this paper is as follows. Section II gives a brief introduction to the mathematics of AES. Section III describes the implementation of AES using the Galois Field primitives found in the MATLAB Communications Toolbox. Section IV presents the comparative results of this implementation versus other implementations. Section V presents our conclusions and plans for future work.

II. ADVANCED ENCRYPTION STANDARD

AES is a specification established by U.S. National Institute of Standards and Technology (NIST) in 2001. AES is based on the Rijndael algorithm developed by Joan Daemen and Vincent Rijmen [Daemen & Rijmen 2001]. Rijndael is a family of ciphers with different key sizes, block sizes, and modes. The standard blocksize for AES is 128 bits (8 bytes), which means that unencrypted data (plaintext) is transformed into encrypted data (ciphertext) into blocks 8 bytes at a time. The encryption can be performed with keys of varying lengths: 128 bit, 192 bit, or 256 bit. For this work, the focus will be on 128 bit encryption. Finally, the encryption scheme can be utilized in a variety of modes. Electronic Codebook (ECB) mode encrypts each 8 byte block of plaintext independently of all of the others. Cipher Block Chaining (CBC) mode combines the plaintext with an Initialization Vector (IV) and then encrypts the data. The ciphertext of the 8 byte block is then used as the IV for the next 8 byte block. For this work, CBC mode is implemented.

The core mathematics of AES encryption is a Galois Field arithmetic [Artin 2011], where the bits in each byte represent the coefficients of a polynomial. The decimal, binary, and Galois Field representations of various numbers are as follows:

$$\begin{aligned} 1 &= 00000001 = 1 \\ 3 &= 00000011 = x + 1 \\ 7 &= 00000111 = x^2 + x + 1 \\ 15 &= 00001111 = x^3 + x^2 + x + 1 \\ 31 &= 00011111 = x^4 + x^3 + x^2 + x + 1 \\ 63 &= 00111111 = x^5 + x^4 + x^3 + x^2 + x + 1 \\ 127 &= 01111111 = x^6 + x^5 + x^4 + x^3 + x^2 + x + 1 \\ 255 &= 11111111 = x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1 \end{aligned}$$

Addition and multiplication of numbers in Galois Field are performed using the normal rules of polynomial addition and multiplication of binary coefficients

$$\begin{aligned} 1 + 3 &= 00000001 + 00000011 \\ &= 1 + (x + 1) \\ &= x + 2 \\ &= x + 0 \\ &= 00000010 \\ &= 2 \end{aligned}$$

In some cases this would result in a higher order polynomial that cannot be represented by 8 bits

$$(x^7)(x^5) = x^{12}$$

To solve this problem, the results of all calculations are computed modulo the AES polynomial

$$100011011 = x^8 + x^4 + x^3 + x + 1$$

Hence

$$(x^7)(x^5) \bmod (x^8 + x^4 + x^3 + x + 1) = x^7 + x^5 + x^3 + x + 1$$

Letting addition, multiplication, and matrix multiplication be defined by the aforementioned Galois Field arithmetic, the AES encryption algorithm can be written succinctly as follows

$$\begin{aligned} \mathbf{P}, \mathbf{C} &: \mathbb{GF}_{8b}^{N \times M} && \text{Matrices of input/output in 8bit Galois Field} \\ \mathbf{V} &: \mathbb{GF}_{8b}^{N \times 16} && \text{Matrix of IVs in 8bit Galois Field} \\ \mathbf{K} &: \mathbb{GF}_{8b}^{R \times N \times 16} && \text{Tensor of key schedules in 8bit Galois Field} \\ \mathbf{s} &: \mathbb{GF}_{8b}^{256} && \text{s-box in 8bit Galois Field} \\ \mathbf{M} &: \mathbb{GF}_{8b}^{4 \times 4} && \text{Mix matrix in 8bit Galois Field} \\ \mathbf{i} &= [1:16] && \text{Set of 16 rows} \end{aligned}$$

for each block of \mathbf{i} rows

$$\mathbf{C}(:, \mathbf{i}) = \mathbf{P}(:, \mathbf{i}) + \mathbf{V} + \mathbf{K}(1, :, \mathbf{i}) \quad \text{Add IV and round key}$$

for each round r

$$\begin{aligned} \mathbf{C}(:, \mathbf{i}) &= \mathbf{s}(\mathbf{C}(:, \mathbf{i}) + 1) && \text{Apply s-box} \\ \mathbf{C}(:, \mathbf{i}_1) &= \mathbf{C}(:, \mathbf{i}_2) && \text{Shift rows} \\ \mathbf{C}(:, 1:4) &= (\mathbf{M} \mathbf{C}(:, 1:4))^T && \text{Multiply by M} \\ \mathbf{C}(:, 5:8) &= (\mathbf{M} \mathbf{C}(:, 5:8))^T && \text{Multiply by M} \\ \mathbf{C}(:, 9:12) &= (\mathbf{M} \mathbf{C}(:, 9:12))^T && \text{Multiply by M} \\ \mathbf{C}(:, 13:16) &= (\mathbf{M} \mathbf{C}(:, 13:16))^T && \text{Multiply by M} \\ \mathbf{C}(:, \mathbf{i}) &= \mathbf{C}(:, \mathbf{i}) + \mathbf{K}(r, :, \mathbf{i}) && \text{Add round key} \end{aligned}$$

$$\mathbf{V} = \mathbf{C}(:, \mathbf{i}) \quad \text{Copy back to IV}$$

$$\mathbf{i} = \mathbf{i} + 16 \quad \text{Increment rows}$$

where T denotes the transpose of the matrix, and for 128 bit AES encryption there are 10 rounds. The above algorithm will encrypt N plaintexts of length M simultaneously with the same key. The most computationally expensive part of the algorithm is the four mix matrix multiplications. Encrypting many

plaintexts at the same time allows these operations to be performed on all the messages using efficient matrix multiplication operations. This technique of combining many items together so they can be performed simultaneously is often referred to as *vectorization* [Birkbeck 2007].

AES decryption can be described in much the same way by reversing the steps of AES encryption. Likewise, the IVs, key schedules, and s-box can also be calculated using simple algorithms.

III. MATLAB IMPLEMENTATION

The MATLAB implementation of AES has three main goals. The first goal is to make the cryptographic technique easy to understand. The second goal is to provide relatively fast encryption/decryption of many short messages of the type that would be used in a sensor algorithm or a database analytic. The third goal is to provide reasonable speed so that a sensor algorithm or database analytic can be quickly tested on sufficiently large data to determine its effectiveness.

Easy-to-understand implementations allow sensor algorithms, database analytics, and cryptographic techniques to be combined in a way that they can be adapted to a variety of applications. In a MATLAB context, easy-to-understand means that there is a strong correspondence between the MATLAB code and the corresponding mathematics. To meet this goal, the high level Galois Field mathematics found in the MATLAB Communication Toolbox is used that allows the mathematics of AES to be implemented in manner that is very similar to the written mathematics.

Sensor algorithms and database analytics perform their operations on many relatively small data elements such as a sensor measurements or a short strings of characters. Thus cryptographic techniques applied to this data should be optimized for many short messages. Many standard cryptographic implementations are implemented to run well on long messages. Vectorizing cryptographic algorithms so that they perform on many messages at once allows the overall speed of the algorithm on short messages to be increased.

The design process of sensor algorithms and database analytics is an iterative, trial-and-error process whereby the designer implements mathematical techniques to achieve their goals. Throughout this process the implementations are tested on ever increasing datasets to validate the mathematics. Extremely slow implementations can be a significant impediment to the design process and requires the implementation must have good overall speed. A simple way to deliver good overall speed is to employ parallel computing. Fortunately, many vectorized implementations are relatively easy to run in parallel. In this case, since the encryption and decryption of each message is independent, this provides a natural approach to running the algorithm in parallel. Likewise, parallel MATLAB [Kepner 2009] provides the necessary software infrastructure for easily running a MATLAB program on a parallel computer.

Employing the aforementioned techniques, a MATLAB implementation of the AES encryption algorithm can be written succinctly as shown in Appendix A.

IV. COMPARATIVE RESULTS

The MATLAB implementation listed in Appendix A is compared to three other AES implementations. The two mostly widely recognized AES MATLAB implementations ([Bucholz 2001] and [Matejka 2011]) are very similar to each other. They both focus on providing an AES implementation in MATLAB primarily for illustrative purposes. Neither of these implementations uses the Galois Field primitives found in the Communications Toolbox. The third implementation is a MATLAB binding to the Open Secure Socket Layer [OpenSSL] implementation of AES written in C [Gadepally 2015]. OpenSSL is among the most widely used encryption libraries.

The three goals of the MATLAB Communication Toolbox implementation were for it to be easy to understand, have good speed on many short messages, and have good overall speed. Easy to understand in this context refers to the correspondence of the implementation to the mathematics. Visual inspection of mathematics and the code shows a strong correspondence. A proxy for the ease of understanding is the relative volume of the code measured in software lines of code (SLOC) using a code counting tool [SCLC]. Figure 2 shows the code volume for all four implementation of AES in MATLAB. The Communications Toolbox implementation is 3x smaller than the other MATLAB implementations, 10x smaller than the MATLAB binding to OpenSSL, and 100x smaller if the underlying OpenSSL C code is included.

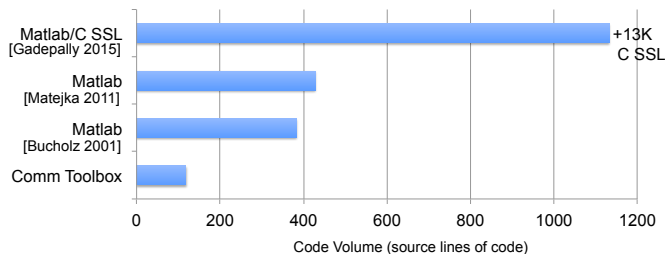


Figure 2. Code volume of different implementations of AES in MATLAB. The Communications Toolbox implementation is 3x smaller than the other MATLAB implementations, 10x smaller than MATLAB binding to OpenSSL, and 100x smaller if the underlying OpenSSL C code is included.

The speed of the Communications Toolbox implementation for messages of different size and different number of messages are shown in Figure 3. The top curves show the rate for 16 byte messages with varying numbers of messages. The bottom curves show the rate for 128 messages with varying message lengths. As expected, the vectorized Communications Toolbox implementation performs well for many short messages.

Figure 4 is the same as Figure 3 but also includes the rate of the other three implementations on many 16 byte messages. As expected, the OpenSSL implementation is 10x faster. The other MATLAB implementations are 100x slower. Thus, the Communications Toolbox implementation achieves nearly 10% of the rate of a highly optimized C library with only 1% of the code.

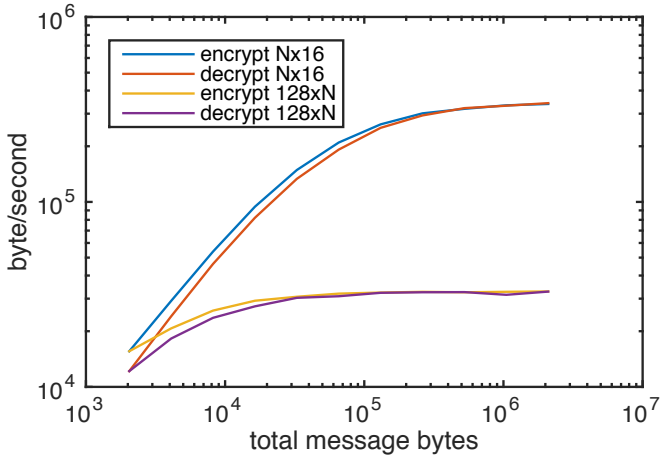


Figure 3. Encryption and decryption rates of the MATLAB Communications Toolbox implementation as a function of total message size. The top curves show rate for 16 byte messages with varying numbers of messages. The bottom curves show the rate for 128 messages with varying message lengths. As expected, the vectorized Communications Toolbox implementation performs well for many short messages.

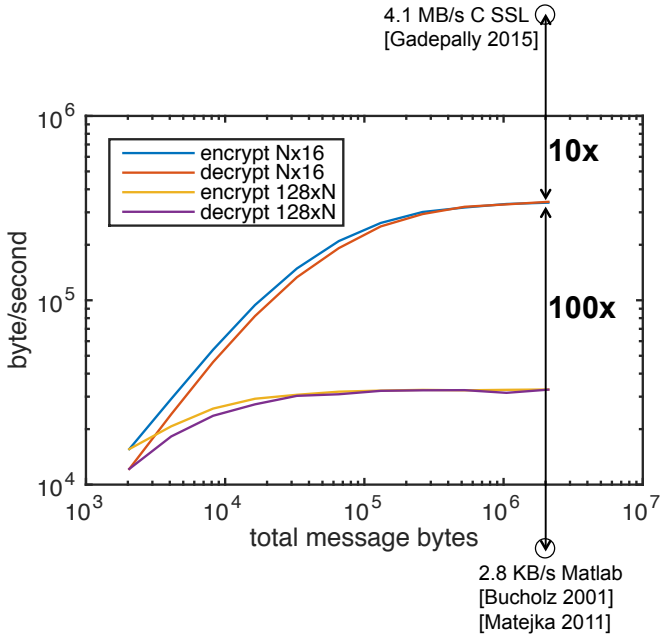


Figure 4. Comparison of encryption and decryption rates of the MATLAB Communications Toolbox implementation with other implementations. OpenSSL implementation is 10x faster. The other MATLAB implementations are 100x slower.

The overall speed of the Communications Toolbox implementation can be improved by running it in parallel. Because this implementation is designed to run well on many small messages, this provides an easy mechanism for running the code in parallel. Different groups of messages can be processed independently on different processors. Figure 5 shows that a speedup of 20x was achieved using 24 processor cores on a single node of a compute cluster.

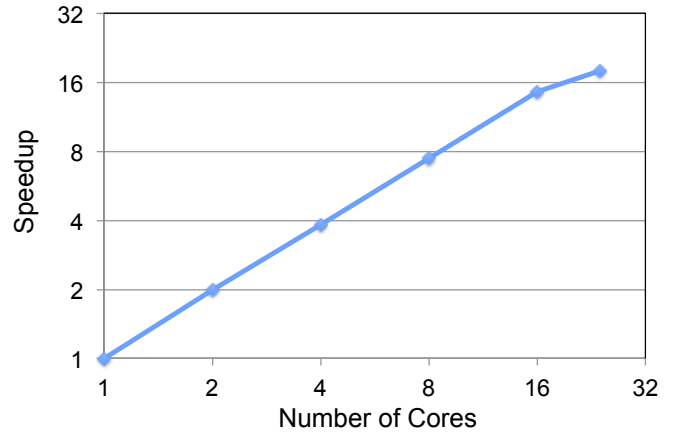


Figure 5. Speedup versus number of cores used for MATLAB Communication Toolbox implementation. A speedup of 20x was achieved using 24 processor cores on a single node of a compute cluster.

V. CONCLUSIONS & FUTURE WORK

The increasing use of networked sensor systems and networked databases has led to an increased interest in incorporating encryption directly into sensor processing and database algorithms. MATLAB is the dominant tool for rapid prototyping of sensor processing algorithms and has extensive database analytics capabilities. The advent of high level and high performance Galois Field mathematical environments allows encryption algorithms to be expressed succinctly and efficiently. This work leverages the Galois Field primitives found in the MATLAB Communication Toolbox to implement a mode of the Advanced Encrypted Standard (AES) based on first principals mathematics. The resulting implementation requires 100x less code than standard AES implementations and delivers speed that is effective for most algorithm development purposes. The parallel version achieves speed comparable to OpenSSL on a single node and is sufficient for real-time prototyping of many sensor processing algorithms and database analytics.

Although the speed of this implementation is good, it is limited by the Galois Field matrix multiplication in the MATLAB Communications Toolbox. Other computational algebra systems have demonstrated extremely impressive speed on equivalent computations. For example, the Magma computational algebra system can perform the same Galois Field matrix multiplication 50x faster than MATLAB [Sutherland 2015]. Thus, additional optimization of the MATLAB Communications Toolbox could allow this implementation to equal or even surpass the speed of the OpenSSL version.

Other limitations in the MATLAB Communications Toolbox limit the ability to easily implement more complex cryptographic operations. For example, the widely used Galois Counter Mode (GCM) is difficult to implement without using 128 bit Galois Fields. Currently the MATLAB Communications Toolbox only supports 16 bit Galois Fields. As the MATLAB Communications Toolbox expands its support for Galois Fields, these more complex cryptographic modes should be easy to implement.

REFERENCES

- [Agullo 2009] E. Agullo, J. Demmel, J. Dongarra, B. Hadri, J. Kurzak, J. Langou, H. Ltaief, P. Luszczek, S. Tomov, "Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects," *Journal of Physics: Conference Series*, Vol. 180, 2009.
- [Artin 2011] M. Artin, *Algebra*, 2nd edition, Prentice Hall, 2011.
- [Atzori 2014] L. Atzori, A. Iera & G. Morabito, "The internet of things: A survey," *Computer networks* 54.15 (2010): 2787-2805.
- [Ben-Or 1988] M. Ben-Or, S. Goldwasser, & A. Wigderson, "Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation (Extended Abstract)," *STOC* 1988.
- [Birkbeck 2007] N. Birkbeck, J. Levesque & J. Amaral, "A dimension abstraction approach to vectorization in MATLAB," *IEEE International Symposium on Code Generation and Optimization*, 2007.
- [Cash 2013] D. Cash, S. Jarecki, C.S. Jutla, H. Krawczyk, M.-C. Rosu, & M. Steiner, "Highly-Scalable Searchable Symmetric Encryption with Support for Boolean Queries," *CRYPTO* 2013.
- [Chaidez 2014] J. Chaidez & J. Kepner, "The Abstract Algebra of Big Data and Associative Arrays," *SIAM Meeting on Discrete Math*, Jun 2014.
- [Chari 2002] S. Chari, J.R. Rao, P. Rohatgim, "Template attacks," *CHES*, LNCS vol. 2523, pp. 13–28, Springer, 2002.
- [DPA contest] http://www.dpacontest.org/v4/rsm_hall_of_fame.php.
- [Daemen & Rijmen 2001] J. Daemen & V. Rijmen, *The design of Rijndael: AES-the advanced encryption standard*, Springer Science & Business Media, 2002.
- [Evans 2013] E. Evans & R. Grossman, "Cyber Security and Reliability in a Digital Cloud," *U.S. Department of Defense Science Board Study*, Jan 2013.
- [Gadepally 2014] V. Gadepally, A. Krishnamurthy, & U. Ozguner, "A framework for estimating driver decisions near intersections," *IEEE Transactions on Intelligent Transportation Systems*, 15.2 (2014): 637-646.
- [Gadepally 2015a] V. Gadepally, S. Wu, J. Kepner & S. Madden, "Sifter: A Generalized, Efficient, and Scalable Big Data Corpus Generator," *New England Database Day*, Jan 2015.
- [Gadepally 2015b] V. Gadepally, B. Hancock, B. Kaiser, J. Kepner, P. Michaleas, M. Varia, A. Yerukhimovich, "Improving the Veracity of Homeland Security Big Data Through Computing on Masked Data," *IEEE HST*, Apr 2015.
- [Gentry 2009] C. Gentry, "Fully Homomorphic Encryption Using Ideal Lattices," *41st ACM Symposium on Theory of Computing (STOC)*, 2009.
- [Halevi 2014] S. Halevi & V. Shoup, "Algorithms in HElib." *Cryptology ePrint Archive*, Report 2014/106. <http://eprint.iacr.org/2014/106>.
- [Hubbell 2012] M. Hubbell & J. Kepner, "Large Scale Network Situational Awareness Via 3D Gaming Technology," *IEEE HPEC*, 2012.
- [Kepner 2009] J. Kepner, *Parallel MATLAB for Multicore and Multinode Computers*, SIAM Press, 2009.
- [Kepner 2012] J. Kepner, W. Arcand, W. Bergeron, N. Bliss, R. Bond, C. Byun, G. Condon, K. Gregson, M. Hubbell, J. Kurz, A. McCabe, P. Michaleas, A. Prout, A. Reuther, A. Rosa & C. Yee, "Dynamic Distributed Dimensional Data Model (D4M) Database and Computation System," *ICASSP*, Mar 25-30, 2012.
- [Kepner 2014] J. Kepner, V. Gadepally, P. Michaleas, N. Shear, M. Varia, A. Yerukhimovich, R. Cunningham, "Computing on Masked Data: a High Performance Method for Improving Big Data Veracity," *IEEE HPEC*, Sep 2014.
- [Kepner 2015] J. Kepner, "Associative Arrays: Unified Mathematics for Spreadsheets, Databases, Matrices, and Graphs," *New England Database Day*, Jan 2015.
- [Kepner & Jansen 2016] J. Kepner & H. Jansen, *Mathematics of Big Data: Spreadsheets, Databases, Matrices, and Graphs*, SIAM Press & MIT Press, in preparation.
- [Kocher 1996] P.C. Kocher, "Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems," *CRYPTO*, Springer Berlin Heidelberg, 1996.
- [Kocher 1999] P.C. Kocher, J. Jaffe, B. Jun. "Differential power analysis," *CRYPTO*, pp. 388–397, Springer-Verlag, London, 1999.
- [Messerges 2001] T.S. Messerges, "Securing the AES finalists against power analysis attacks," *Fast Software Encryption*, pp. 150-164, Springer, 2001.
- [Osvik 2006] D.A. Osvik, A. Shamir, E. Tromer, "Cache attacks and countermeasures: the case of AES," *CT-RSA*, pp. 1-20, Springer Berlin Heidelberg, 2006.
- [Pal 2012] P. Pal, G. Lauer, J. Khoury, N. Hoff, J. Loyall: "P3S: A privacy preserving publish-subscribe middleware." *In: Middleware 2012*. Narasimhan, P., Triantafillou, P. (eds.) LNCS, vol. 7662, pp. 476–495. Springer, Heidelberg, 2012.
- [Peeters 2007] E. Peeters, F.X. Standaert, J.J. Quisquater. "Power and electromagnetic analysis: Improved model, consequences and comparisons." *Integration* 40 (1), pp. 52–60. 2007.
- [Perl 2011] H. Perl, M. Brenner, & M. Smith, "Poster: an implementation of the fully homomorphic Smart-Vercauteren crypto-system." *CCS* 2011.
- [Popa 2011] R.A. Popa, C.M.S. Redfield, N. Zeldovich, & H. Balakrishnan, "CryptDB: Protecting Confidentiality with Encrypted Query Processing," *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, Oct 2011.
- [Raykova 2012] M. Raykova, A. Cui, B. Vo, B. Liu, T. Malkin, S.M. Bellovin, & S.J. Stolfo, "Usable, Secure, Private Search," *IEEE Security & Privacy* 10(5), 2012.
- [SCLC] <http://sclc.googlecode.com/svn/trunk/perldoc/sclc.html>.
- [Starner 2014] T. Starner, "How Wearables Worked their Way into the Mainstream," *IEEE Pervasive Computing* 4 (2014): 10-15.
- [Sutherland 2015] A. Sutherland, personal communication
- [Varia et al 2015a] M. Varia, S. Yakoubov & Y. Yang, "HEtest - A Homomorphic Encryption Testing Framework," *3rd Workshop on Encrypted Computing and Applied Homomorphic Cryptography*, Jan 2015.
- [Varia et al 2015b] M. Varia, B. Price, N. Hwang, A. Hamlin, J. Herzog, J. Poland, M. Reschly, S. Yakoubov & R.K. Cunningham, "Automated Assessment of Secure Search Systems," *ACM SIGOPS Operating Systems Review - Special Issue on Repeatability and Sharing of Experimental Artifacts archive*, Volume 49 Issue 1, January 2015.
- [Yao 1982] A.C.-C. Yao, "Protocols for Secure Computations (Extended Abstract)," *FOCS* 1982.

APPENDIX A1: MATLAB AES 128 BIT CBC ENCRYPTION IMPLEMENTATION

```

function ct = AESCBCencrypt(key,IV,pt)
    Nblock = ceil(size(pt,2)/16);           % Number of 16 byte blocks for plaintext.
    pt(:,size(pt,2)+1:Nblock*16) = 0;     % Pad plaintext to end of the last block.
    ct = uint8(pt); ct(:) = 0;            % Allocate ciphertext.
    Nround = 10;                          % Standard for 128 bit.
    ob_gf8 = repmat(AESfield(uint8(IV)), size(pt,1), 1); % Set up IV.
    sbox = genSbox();                      % Compute sbox from first principles.
    row = [2 3 1 1];                      % Define mix matrix.
    mixmat_gf8 = gallery('circul', row);
    key_schedule = genKeys(key, Nround, sbox); % Generate the key schedule.
    idx = 1:16;
    for iblock=1:Nblock
        pt_gf8 = AESfield(pt(:,idx)) + ob_gf8; % XOR plaintext with the IV.
        ctrnd00_gf8 = pt_gf8 + repmat(key_schedule(1,:), size(pt_gf8,1),1); % Add plaintext and key.
        for iround=1:Nround
            ctrnd01_gf8 = AESfield(sbox(uint32(ctrnd00_gf8)+1)); % SubBytes
            ctrnd01_gf8(:,[2 3 4 6 7 8 10 11 12 14 15 16]) = ctrnd01_gf8(:,[6 11 16 10 15 4 14 3 8 2 7 12]); % ShiftRows
            if (iround < Nround) % MixColumns
                ctrnd01_gf8(:, 1:4) = (mixmat_gf8 * ctrnd01_gf8(:, 1:4)).';
                ctrnd01_gf8(:, 5:8) = (mixmat_gf8 * ctrnd01_gf8(:, 5:8)).';
                ctrnd01_gf8(:, 9:12) = (mixmat_gf8 * ctrnd01_gf8(:, 9:12)).';
                ctrnd01_gf8(:,13:16) = (mixmat_gf8 * ctrnd01_gf8(:,13:16)).';
            end
            ctrnd01_gf8 = ctrnd01_gf8 + repmat(key_schedule(iround+1,:),size(pt_gf8,1),1); % AddRoundKey
            ctrnd00_gf8 = ctrnd01_gf8; % Copy 01 variable to 00 variable for CBC mode.
        end
        ob_gf8 = ctrnd01_gf8; % Last ciphertext becomes the IV for the next block.
        ct(:,idx) = reshape(uint8(uint32(ctrnd01_gf8)), size(pt,1), []); % Save current block of ciphertext.
        idx = idx + 16; % Increment to the next block.
    end

function x = uint32(g);
    x = g.x;
end

function f = AESfield(inp)
    f = gf(inp,8,hex2dec('11B')); % Standard primitive polynomial for AES
end

```

APPENDIX A2: MATLAB AES 128 BIT CBC GENERATE S-BOX IMPLEMENTATION

```

function sbox = genSbox()
    pp = hex2dec('11B'); % primitive polynomial for AES .
    row = [1 1 1 1 1 0 0 0]; % set up constants A and b.
    A = gf(gallery('circul', row), 2);
    b = gf([0; 1; 1; 0; 0; 0; 1; 1], 2);
    xin = gf(0:255, 8, pp); % Compute inverses.
    xin(2:256) = 1./xin(2:256);
    y = gf(transpose(dec2bin(uint32(xinv),8) == '1'),2);
    z = A * y + repmat(b,1,256); % Affine transformation.
    sbox = transpose(uint8(bin2dec(num2str(transpose(uint32(z)))))); % Reformat to make sbox.
end

```

APPENDIX A3: MATLAB AES 128 BIT CBC GENERATE KEY SCHEDULE IMPLEMENTATION

```

function key_schedule = genKeys(key, Nround, sbox)
    if sbox == 0 % Encrypt can pass in the sbox, but decrypt uses the inverse sbox.
        sbox = genSbox();
    end
    key_schedule = AESfield(zeros(Nround+1,16)); % Allocate space for Nround+1 keys of size 4x4.
    key_schedule(1,:) = AESfield(uint8(key)); % Embed key into Galois Field.
    x = AESfield(2); % Create round constant for any round.
    rcon = AESfield(zeros(1,4));

```

```

for iround=1:Nround
    rcon(1) = x^(iround - 1); % Compute round constant.
    key_schedule(iround+1,1:4) = circshift(key_schedule(iround,13:16),-1,2); % Shift column 4 and place into column 1.
    key_schedule(iround+1,:) = sbox(uint8(uint32(key_schedule(iround+1,:))+1)); % Transform values with S-Box.
    key_schedule(iround+1,1:4) = key_schedule(iround+1,1:4) + rcon; % Add round constant.
    % Add and shift forward to remaining columns.
    key_schedule(iround+1, 1:4 ) = key_schedule(iround, 1:4 ) + key_schedule(iround+1,1:4 );
    key_schedule(iround+1, 5:8 ) = key_schedule(iround, 5:8 ) + key_schedule(iround+1,1:4 );
    key_schedule(iround+1, 9:12) = key_schedule(iround, 9:12) + key_schedule(iround+1,5:8 );
    key_schedule(iround+1,13:16) = key_schedule(iround,13:16) + key_schedule(iround+1,9:12);
end
end

```

APPENDIX A4: MATLAB AES 128 BIT CBC DECRYPTION IMPLEMENTATION

```

function pt = AESCBCdecrypt(key,IV,ct)
    Nblock = size(ct,2)/16; % Number of 16 byte blocks in ciphertext.
    pt = char(zeros(size(ct,1), 16*Nblock)); % Preallocate for the plaintext.
    Nround = 10; % Standard for 128 bit key.
    ob_gf8 = repmat(AESfield(uint8(IV)), size(ct,1), 1); % Set up IV.
    sbox = genSboxInv(); % Compute inverse sbox from first principles.
    row = [14 11 13 9]; % Define mix matrix.
    mixmat_gf8 = gallery('circul', row);
    key_schedule = genKeys(key, Nround, 0); % Generate the key schedule.
    idx = 1:16;
    for iblock=1:Nblock
        state00 = uint8(ct(:,idx));
        for iround=Nround:-1:1 % Intermediate rounds.
            state01 = state00 + repmat(key_schedule(iround+1,:), size(ct,1),1); % AddRoundKey
            if (iround < Nround) % MixColumns
                state01(:, 1:4 ) = (mixmat_gf8 * state01(:, 1:4 ).)';
                state01(:, 5:8 ) = (mixmat_gf8 * state01(:, 5:8 ).)';
                state01(:, 9:12) = (mixmat_gf8 * state01(:, 9:12).)';
                state01(:,13:16) = (mixmat_gf8 * state01(:,13:16).)';
            end
            state01(:, [2 3 4 6 7 8 10 11 12 14 15 16]) = state01(:, [14 11 8 2 15 12 6 3 16 10 7 4]); % ShiftRows
            state01 = AESfield(sbox(uint32(state01)+1)); % SubBytes
            state00 = state01; % Copy 01 variable to 00 variable for CBC mode.
        end
        state01 = state01 + repmat(key_schedule(1,:),size(ct,1),1); % Reverse round 0 (initial round). Add plaintext and key.
        state01 = state01 + ob_gf8; % XOR state with the IV.
        ob_gf8 = uint8(ct(:,idx)); % This ciphertext becomes the IV for the next block.
        pt(:,idx) = uint32(state01); % Save current block of plaintext.
        idx = idx + 16; % Increment to the next block.
    end
end

```

APPENDIX A5: MATLAB AES 128 BIT CBC GENERATE INVERSE S-BOX IMPLEMENTATION

```

function sinv = genSboxInv()
    pp = hex2dec('11B'); % primitive polynomial for AES.
    row = [0 1 0 1 0 0 1 0]; % set up constants A and b.
    A = gf(gallery('circul', row), 2);
    b = gf([0; 0; 0; 0; 0; 1; 0; 1], 2);
    y = gf(transpose(dec2bin(0:255) == '1'), 2); % Affine transformation.
    z = A * y + repmat(b,1,256);
    z = gf(bin2dec(num2str(transpose(uint32(z))))), 8, pp); % Reformat.
    zinv = gf(0:255, 8, pp); % Invert results.
    zinv(1:99) = 1./z(1:99);
    zinv(100) = 0;
    zinv(101:256) = 1./z(101:256);
    sinv = uint8(uint32(zinv)); % Reformat to make inverse sbox.
end

```