

# Secure Multiparty AES (short paper)

Ivan Damgård and Marcel Keller

Dept. of Computer Science, Aarhus University, Denmark  
{ivan,mkeller}@cs.au.dk

**Abstract** We propose several variants of a secure multiparty computation protocol for AES encryption. The best variant requires  $2200 + \frac{400}{255}$  expected elementary operations in expected  $70 + \frac{20}{255}$  rounds to encrypt one 128-bit block with a 128-bit key. We implemented the variants using VIFF, a software framework for implementing secure multiparty computation (MPC). Tests with three players (passive security against at most one corrupted player) in a local network showed that one block can be encrypted in 2 seconds. We also argue that this result could be improved by an optimized implementation.

## 1 Introduction and Motivation

In secure multiparty computation (MPC), a number of players each supply a private input and then compute an agreed function on these inputs *securely*, i.e., even if an adversary corrupts some of the players, honest players obtain correct results, and the intended outputs is the only new information released about the inputs. Several general feasibility results for MPC are known, for instance, given secure point to point channels, any function can be computed securely against an honest but curious adversary corrupting any minority of the players, and securely against a malicious adversary corrupting strictly less than one third of the players [1, 4].

Although MPC has been a topic in cryptographic research for many years, and despite the obvious potential for applications, implementations have evolved only recently [2, 6, 9]. Some of them have even been used to solve real-world tasks, such as privacy-preserving auctions [3].

In this paper, we present several variants of an MPC protocol for computing AES encryption [11]. We assume that key and plaintext are byte-wise secret shared among the players; the same holds for the outputted ciphertext.

Apart from the general motivation of investigating how far we can take MPC in practice, there is also a more direct motivation for looking at such a “threshold-approach” to symmetric encryption. An example: suppose a set of players hold some secret shared data and wish to communicate this data to an external party. A trivial solution is for each player to send his shares securely to the receiver, who can then reconstruct the data. But this will mean that the receiver must be aware of the fact that the data is secret shared and must apply a non-standard algorithm to get the data. In addition, his work is linear in the number of players. From this point of view it would be a more attractive solution if the players could

cooperate to generate a ciphertext for the receiver in standard form, which would typically be an encryption of an AES key  $K$  under the receiver's public key, followed by the data encrypted under  $K$ . Note that a similar solution used in the opposite direction could be used for a party to supply encrypted inputs to a multiparty computation, even if that party is not aware of the number of players, or the concrete MPC protocol they execute. He only needs to know a public key for the system, where the players share the private key. This could be useful in any application of MPC, e.g. for secure auctions, procurement or benchmarking. In practice, this would mean that parties submitting data to the system can use completely standard client software for sending data securely protected under a public key. Moreover, the back-end of the system can be updated with new MPC protocols or migrate to a new set of players with no change on the client side, as long as the public key remains the same.

Another application could be the following: Analogously to encrypted hard disks, one could imagine to store data encrypted in a place with weak security compliance (e.g., a cloud), whereas the key is secret shared between different secured machines. Those machines then can run multiparty AES to read and write data together with further MPC to process it. The secret sharing of the key reduces the risk of leakage, as well as the risk of losing the key. A more naive solution, where one reconstructs the key and encrypts/decrypts data in the normal way, would create a single point of attack from where the entire data-set can be stolen even if one only meant to read a small part.

Whereas choosing a random key  $K$  and encrypting it under a public key is easy using known techniques, there is virtually no previous work considering specifically MPC for symmetric encryption (except for an existing 2-player solution, see next section).

Our work on AES exploits the fact that AES is based on arithmetic in  $\text{GF}(2^8)$ . Therefore, our protocol can be based on any general MPC protocol that is based on Shamir secret sharing [12] and implements secure multiplication and addition in  $\text{GF}(2^8)$ . With respect to security threshold and type of adversary (passive/active), our protocol will be as secure as the underlying MPC protocol we use. We can, for instance, use the classic passively secure protocol from [1] tolerating a dishonest majority, or the actively secure protocol from [6] tolerating less than one third corrupted players.

The non-trivial problem we need solve is to implement the AES S-box efficiently, since this is the only non-linear part of the algorithm, and essentially requires us to securely compute a multiplicative inverse of an element in  $\text{GF}(2^8)$  where 0 should be mapped to 0. The naive solution to this is to raise to the power of 254. We propose several alternative solutions that improve on this by reducing the number of elementary operations, or the number of rounds, or both.

We have implemented our protocol in VIFF, a software framework for implementing secure multiparty computation [6]. Tests for three players running a passively secure protocol on a local network show that an AES block can be encrypted in 2 seconds, and tests also confirm that our methods for reducing the number of rounds lead to better performance when network delays are large

enough to influence speed. Since our implementation uses a general framework based on the high-level interpreted language Python, much better performance can certainly be obtained using a dedicated C implementation. We therefore believe our results demonstrate that MPC for symmetric encryption is definitely a possibility in practice.

## 2 Related Work

MPC protocols can be divided into two categories. The first consist of protocols computing an arithmetic circuit over a suitable field. These are usually related to a secret-sharing scheme [12]. Other protocols can be used to compute any binary circuit. These are mostly based on Yao’s garbled circuits [13]. An optimized implementation by Pinkas et al. was recently used for secure two-party AES [10]. Their protocol differs from ours, which is of the first category. Their protocol requires one party to know the key, the other to know the cleartext, and outputs the ciphertext to the latter one. Our protocol works for the multiparty case, it takes a secret shared key and cleartext as input and outputs a secret shared ciphertext. The communication complexity of our protocol is smaller, due to the utilization of the arithmetic properties of AES. Our implementation is also faster than that of [10], as detailed later. However, Yao’s garbled circuits lead to constant-round protocols, contrary to ours, in the sense that if one increases the number of AES rounds, our number of rounds increase as well.

Since the original proposal of MPC there have been several improvements to make it more efficient. One of those is pseudorandom secret sharing [5], which allows to generate a secret shared random number without any communication at all. Another improvement is an MPC protocol providing active security which allows preprocessing, i.e., performing some computations without knowing the input to reduce the online time [6]. We will use both techniques in the following.

## 3 Preliminaries

*The Finite Field  $\text{GF}(2^8)$*  AES treats bytes mostly as elements in  $\text{GF}(2^8)$  because there exists a bijective mapping from the set of bytes to the field:  $\{0, 1\}^8 \rightarrow \text{GF}(2^8) \cong \text{GF}(2)[x]/(p)$ ,  $a = a_7 \dots a_0 \mapsto \sum_{i=0}^7 a_i \cdot x^i$ , where  $\text{GF}(2)[x]/(p)$  is the field of polynomials over  $\text{GF}(2)$  modulo an irreducible polynomial  $p$ . Note that  $\text{GF}(2^8)$  has characteristic 2, i.e. subtraction is the same as addition.

*Secure Multiparty Computation* based on Shamir secret sharing over a field provides the following operations: Addition and multiplication can be done locally, multiplication in general and opening of shared values requires communication. We will refer to the latter two as elementary operations. Throughout the paper, we will use square brackets to denote secret shared values:  $[x]$ .

*Pseudorandom Secret Sharing* allows the distributed generation of random values without communication. For fields with characteristic 2, generation of random bits is also possible. We refer to Section 4.2 of [6] for details.

*Bit Decomposition* of an element in  $\text{GF}(2^8)$  is required for the S-box of AES. Since  $\text{GF}(2^8)$  has characteristic 2, this can be done by masking with a bit-wise random secret shared value. If the random bits are generated using PRSS, the communication cost is one opening. We refer to the full version for details [7].

## 4 The AES Protocol

AES encryption and decryption are round-based, with each round consisting of some operations on the internal state. This is a matrix of  $4 \times 4$  bytes, corresponding to a block size of 128 bits. Initial state is the input, final state the output. A typical encryption round looks as follows: SubBytes, ShiftRows, MixColumns, AddRoundKey. The only exceptions are an additional AddRoundKey at the beginning of encryption, and that MixColumns is skipped in the last round.

We now describe how to compute all operations using MPC. Both cleartext and key are assumed to be byte-wise secret shared over  $\text{GF}(2^8)$ . The internal state and so the output will be as well.

### 4.1 SubBytes

In SubBytes, an S-box is applied to every byte of the input. Because the S-box is defined arithmetically, we can compute it relatively efficiently with multiparty computation. This is the only part of the protocol requiring communication, everything else can be done locally. The S-box consists of two steps: an inversion on  $\text{GF}(2^8)$  and an affine linear transformation on  $\text{GF}(2)^8$ .

**Inversion** The field element represented by the byte is inverted in  $\text{GF}(2^8)$ , except 0, which is mapped to 0. There are several possibilities of doing this with multiparty computation.

*Square-and-multiply* We raise the field element to the power of  $\text{ord}(\text{GF}(2^8)^*) - 1 = 254$  using some square-and-multiply variant. This costs 11 multiplications in 9 rounds per byte, using the addition chain (1, 2, 4, 8, 9, 18, 19, 36, 55, 72, 127, 254). This is optimal regarding the number of multiplications. Since the number of rounds is the lowest possible for the number of multiplications, we will refer to this variant as *square-and-multiply with shortest addition chain and least number of rounds*. Another multiplication chain, (1, 2, 3, 4, 7, 8, 15, 16, 31, 32, 63, 64, 127, 254) requires 13 multiplications in 8 rounds, which is optimal regarding the number of rounds. We will refer to this as *square-and-multiply with least rounds*. Standard square-and-multiply costs 13 multiplications in 13 rounds.

*Masked Exponentiation* This method uses the fact that  $(x+y)^2 = x^2 + 2xy + y^2 = x^2 + y^2$  for fields with characteristic 2. By a simple induction, it follows that  $(x+y)^{2^i} = ((x+y)^{2^{i-1}})^2 = (x^{2^{i-1}} + y^{2^{i-1}})^2 = x^{2^i} + y^{2^i}$  for  $i \geq 2$ . We exploit this property to split up the computation in a preprocessing and an online phase,

which saves some rounds because the preprocessing operations for all S-boxes of the protocol can be executed in parallel.

In the preprocessing phase, we generate a random shared  $[r] \in \text{GF}(2^8)$  and square  $[r]$  7 times. This costs 7 multiplications in 7 rounds if pseudorandom secret sharing is used:  $[r^2] = [r] \cdot [r]$ ,  $[r^4] = [r^2] \cdot [r^2]$ ,  $\dots$ ,  $[r^{128}] = [r^{64}] \cdot [r^{64}]$ . To invert  $[x]$ , we mask  $[x]$  by  $[r]$ , open it, and exponentiate locally:

$$\text{open}([x] + [r]) = (x + r) \quad (x + r)^2, (x + r)^4, (x + r)^8, \dots, (x + r)^{128}.$$

Finally, we unmask the powers of  $[x]$  and multiply them to get  $[x^{254}]$ :

$$(x + r)^{2^i} + [r^{2^i}] = [x^{2^i}] \quad \forall i = 1, \dots, 7, \quad \prod_{i=1}^7 [x^{2^i}] = [x^{\sum_{i=1}^7 2^i}] = [x^{254}].$$

The online operations cost 7 elementary operations (1 opening and 6 multiplications) in 4 rounds.

*Masking* Here we exploit that the inversion is a homomorphism with respect to multiplication. The field element  $[a] \in \text{GF}(2^8)$  can be masked by multiplying with some shared random number  $[r] \in_R \text{GF}(2^8)$ . We open the masked value, invert it and unmask the result to get a sharing of the inverted element:

$$\text{open}([a] \cdot [r]) = ar, \quad (ar)^{-1} \cdot [r] = [a^{-1}r^{-1}r] = [a^{-1}].$$

This method would leak whether  $a$  is 0 because  $ar = 0$  for all  $r \in \text{GF}(2^8)$  if  $a = 0$ . Therefore, if  $a = 0$ , we add 1 before doing the inversion, and subtract 1 after it. This guarantees that 0 is mapped to 0 without leaking it. Let  $b$  be 0 if  $a \neq 0$ , and 1 if  $a = 0$ . It can be computed by decomposing  $a$  into bits  $a_i$ ,  $a = a_7 \dots a_0$ , and then letting  $[b] := 1 - \prod_{i=0}^7 (1 - [a_i])$ . The inversion is computed as follows:

$$\begin{aligned} \text{open}([(a] + [b]) \cdot [r]) &= (a + b) \cdot r, & ((a + b) \cdot r)^{-1} \cdot [r] &= [(a + b)^{-1}] \\ [(a + b)^{-1}] - [b] &= \begin{cases} [(a + 0)^{-1} - 0] = [a^{-1}], & a \neq 0, b = 0 \\ [(0 + 1)^{-1} - 1] = [0], & a = 0, b = 1. \end{cases} \end{aligned}$$

If  $(a + b)r$  is now 0, we know that  $r = 0$ . In that case, we just choose another random  $r$  and repeat the masking.

The computation of  $[b]$  costs 1 opening operation for bit decomposition, and 7 multiplications in 3 rounds afterwards. The computation of  $(a + b)r$  costs 1 multiplication and 1 opening, again assuming that the random shared number  $[r]$  can be generated locally. Since  $r$  might be zero, we require expected  $\frac{2}{1-1/256} = 2 + \frac{2}{255}$  elementary operations until  $(a+b)r$  is non-zero. The rest can be computed locally, so we get expected  $10 + \frac{2}{255}$  elementary operations in  $6 + \frac{2}{255}$  expected rounds overall.

**Affine Linear Transformation** Here, the byte  $a = \sum_{i=0}^7 a_i \cdot x^i$  is considered as a bit vector  $(a_0, \dots, a_7) \in \text{GF}(2)^8$ , which is multiplied with a fixed invertible matrix and then added to a constant vector. To do so, we decompose the input into bits (cost: 1 opening if PRSS is used), and then we compute the rest locally because the matrix and the vector are fixed.

**Communication Cost** The computation of one S-box costs at least 12 elementary operations in 10 rounds or 14 elementary operations in 9 rounds using square-and-multiply and expected  $11 + \frac{2}{255}$  elementary operations in  $7 + \frac{2}{255}$  rounds using masking. For masked exponentiation, preprocessing requires 7 elementary operations in 7 rounds, and online computation requires 8 elementary operations in 5 rounds, both per S-box.

## 4.2 Other operations

ShiftRows, MixColumns, and AddRoundKey consist only of byte permutations and linear operations on  $\text{GF}(2^8)$ , which can be executed locally. Key expansion uses the same S-box as SubBytes and local operations.

## 5 Security

The security of our protocol relies mainly on the security of the MPC scheme used. The only information that is revealed additionally to the leakage of the MPC scheme are openings of masked values, i.e. either of  $x + r$  or of  $y \cdot r$  for a random  $r$  and  $y \neq 0$ . It is easy to see that both openings do not reveal information about  $x$  and  $y$ , respectively.

It follows that a simulator, e.g., in the UC framework, can generate those values with the same distribution as in the real execution if there exists a simulator for the MPC scheme.

## 6 Analysis

Since the S-box is the only part which requires communication, it suffices to count the number of S-boxes computed. 16 S-boxes are computed in parallel in every SubBytes operation and thus in every AES round. The key expansion can be computed in parallel with the AES rounds. Putting all together, the total number of elementary operations is the number of S-boxes times the number of elementary operations per S-box, and the the total number of rounds is the number of AES rounds times the number of rounds per S-box.

Table 1 describes the different possibilities for inversion. From the AES specification, one can deduce that the encryption of one block in AES-128 requires 200 S-boxes (including key expansion). So, one can calculate that using inversion by masking, it takes  $2200 + \frac{400}{255}$  expected elementary operations in  $70 + \frac{20}{255}$  expected rounds. See the full version for the analysis of the other AES flavors [7].

Masked exponentiation only gives an advantage over the other methods if all the preprocessing is done before the encryption of a block. In this way, one can calculate with only  $7 / (\text{number of AES rounds})$  rounds per S-box, i.e., 0.7 rounds in the case of AES-128.

	El. operations	Rounds
Masking	$11 + \frac{2}{255}$	$7 + \frac{2}{255}$
Standard square-and-multiply	14	14
S-a-m with shortest add. chain and least rounds	12	10
S-a-m with least rounds	14	9
Masked exponentiation	15	$5 + \frac{7}{\#AES\ rounds}$

**Table 1.** One S-box in different inversion protocols.

## 7 Implementation

Our implementation is based on the Virtual Ideal Functionality Framework (VIFF), a Python-based framework for secure multiparty computation [8]. VIFF was developed to implement efficient MPC for asynchronous networks, i.e., every local computation is executed as soon as the needed input values are present. It provides protocols with passive security as well as protocols secure against active adversaries. Shamir secret sharing is used for protocols with at least three parties, and two-party MPC can be done based on the Paillier cryptosystem.

### 7.1 Benchmarks

The implementation of the encryption was tested on a local gigabit-network (ping 0.1 ms) with modern hardware: Dual-Core AMD Opteron Processor with 2.4 GHz per core, 2 GB RAM, Red Hat 5.2, Linux Kernel 2.6.18, Python 2.6.1. Using three machines and passive security against one opponent, the encryption of one block with AES-128 took about 2 seconds on average including key expansion when encrypting 10 blocks in parallel. This was achieved using inversion by exponentiation, which turned out to be faster than inversion by masking in the given setting. This is contradictory to our analysis. The reason is that masking needs more local computation (more pseudorandom secret sharing used for bit decomposition), which has a higher impact if the network latency is low and the bandwidth is high. However, the benchmarks behave as expected with a network delay of 40 ms or the bandwidth limited to 800 kbit/s, see the full version [7].

Our method is faster than two-party AES by Pinkas et al. [10] which takes 7 seconds with passive security. Note that their implementation is optimized, whereas ours uses Python, a high-level interpreted language. We observed that local computation is the main bottleneck in our implementation, so this gives the possibility for better results using an implementation in a low-level language with less overhead, such as C.

Moreover, in a setting with four players and active security against one malicious adversary our protocol takes about 7 seconds. This is considerably less than the solution by Pinkas et al. which requires 1148 seconds to encrypt one block with security against a malicious adversary. We used the PRSS-based variant of an actively secure MPC scheme by Damgård et al. [6]. The scheme allows

to generate so-called multiplication triples in a preprocessing phase, i.e., before knowing any input. By using that, the online time can be reduced to less than 4 seconds per block for masked exponentiation, which uses preprocessing also for AES inversion, as described in Section 4.1.

## 8 Conclusion

We have presented a secure multiparty computation protocol for AES together with benchmarking results of an implementation: roughly 2 seconds per block. Our results can not be applied directly to other algorithms (including ciphers) because we made use of the arithmetic properties of AES, namely of the fact that the S-box is not just a “random” substitution.

## References

1. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *STOC*, pages 1–10. ACM, 1988.
2. D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A framework for fast privacy-preserving computations. In S. Jajodia and J. López, editors, *ESORICS*, volume 5283 of *Lecture Notes in Computer Science*, pages 192–206. Springer, 2008.
3. P. Bogetoft, D. L. Christensen, I. Damgård, M. Geisler, T. Jakobsen, M. Krøigaard, J. D. Nielsen, J. B. Nielsen, K. Nielsen, J. Pagter, M. I. Schwartzbach, and T. Toft. Secure multiparty computation goes live. In R. Dingledine and P. Golle, editors, *Financial Cryptography*, volume 5628 of *Lecture Notes in Computer Science*, pages 325–343. Springer, 2009.
4. D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols (extended abstract). In *STOC*, pages 11–19. ACM, 1988.
5. R. Cramer, I. Damgård, and Y. Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In J. Kilian, editor, *TCC*, volume 3378 of *Lecture Notes in Computer Science*, pages 342–362. Springer, 2005.
6. I. Damgård, M. Geisler, M. Krøigaard, and J. B. Nielsen. Asynchronous multiparty computation: Theory and implementation. In S. Jarecki and G. Tsudik, editors, *Public Key Cryptography*, volume 5443 of *Lecture Notes in Computer Science*, pages 160–179. Springer, 2009.
7. I. Damgård and M. Keller. Secure multiparty AES (full paper). Cryptology ePrint Archive, Report 2009/614, 2009. <http://eprint.iacr.org/>.
8. M. Geisler. VIFF: Virtual ideal functionality framework. Homepage: <http://viff.dk/>, 2007.
9. D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella. Fairplay - secure two-party computation system. In *USENIX Security Symposium*, pages 287–302. USENIX, 2004.
10. B. Pinkas, T. Schneider, N. Smart, and S. Williams. Secure two-party computation is practical. Cryptology ePrint Archive, Report 2009/314, 2009. <http://eprint.iacr.org/>.
11. F. I. P. S. Publications. Advanced Encryption Standard. Technical Report FIPS PUB 197, National Institute of Standards and Technology, November 2001.
12. A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
13. A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167. IEEE, 1986.