Efficient Implementation of the Orlandi Protocol

Thomas P. Jakobsen[1], Marc X. Makkes[2], and Janus Dam Nielsen[1]

[1] The Alexandra Institute
Aabogade 34, 8200 Aarhus N
Denmark
{thomas.jakobsen,janus.nielsen}@alexandra.dk
[2] Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven
Netherlands
m.x.makkes@student.tue.nl

Abstract. We present an efficient implementation of the Orlandi protocol which is the first implementation of a protocol for multiparty computation on arithmetic circuits, which is secure against up to $n-1$ static, active adversaries. An efficient implementation of an actively secure self-trust protocol enables a number of multiparty computation where one or more of the parties only trust himself. Examples includes auctions, negotiations, and online gaming. The efficiency of the implementation is largely obtained through an efficient implementation of the Paillier cryptosystem, also described in this paper.

Keywords: Secure multiparty computation, MPC, homomorphic encryption, protocols.

1 Introduction

Secure multiparty computation is a cryptographic technique allowing n parties to jointly compute the result of a function $f(x_1, x_2, ..., x_n)$ while ensuring that the input x_i of each party P_i is kept private, even with a number t of the parties acting maliciously. The only information that is allowed to be revealed is the result of the function.

In the 80s it was proved that secure multiparty computation could in fact be applied to any computable function, making it an extremely general and useful technique, at least in theory. This was first done by Yao [34] in the restricted case of two parties, but soon followed similar results for the general case of n parties [4,13]. These results were, however, mostly of theoretical interest due to the complexity of the protocols.

Since then a large number of results have been obtained using different security and adversary models, underlying network assumptions, and improvements of previously known results.

In recent years, the theory has advanced enough to allow practical implementations of secure multiparty computation. Examples of practical systems which support evaluation of general multiparty computation are the FairPlay [23],

J. Zhou and M. Yung (Eds.): ACNS 2010, LNCS 6123, pp. 255–272, 2010.
© Springer-Verlag Berlin Heidelberg 2010

VIFF [32], ShareMind [7], and SIMAP [9] systems. However, many applications are still infeasible in practice, especially those that rely on quick response times like online auctions. Also, in order to be practical, the aforementioned systems tend to either be restrited to a limited number of parties or to loosen up the security model. Some examples of the latter could be assuming that the corrupted parties do not deviate from the protocol (the passive security model) or that at most a certain threshold t of parties gets corrupted (threshold security model). Especially the active security model have until recently been regarded as too complex for practical implementations. However, recently Lindell, Pinkas, and Smart showed that active security in the two-party case is indeed practical [22], and Damgård, Geisler, Krøigaard, and Nielsen showed that active security can be practical if less than $n/3$ parties out of n are corrupted [14].

In this paper we go one step further and document an implementation of the Orlandi protocol [28] for secure multiparty computation which is both actively secure and tolerates up to $n-1$ corruptions. We further describe our benchmarks of this implementation, compare it to benchmarks of related protocols and argue that the protocol is indeed practical.

The rest of the paper is organized as follows. Section 2 gives an introduction to the Orlandi protocol. Section 3 describes our implementation, and we introduce our benchmarks, their setup, and discuss our results in Section 4. In Section 5 we describe and discuss how we removed the main performance bottlenecks using a high-performance implementation of the Paillier cryptosystem. A description of related work is given in Section 6, and we conclude in Section 7 along with discussing future work.

2 The Orlandi Protocol

The Orlandi protocol [28] is a protocol for secure multiparty computation on arithmetic circuits, which is secure against up to $n-1$ static, active adversaries. We will introduce the protocol in this section by giving a high-level description of the protocol and a detailed account of the parts of the protocol which have been the target of our optimizations. The Orlandi protocol is based on a Verifiable Secret Sharing (VSS) scheme, secure against a dishonest majority, augmented with a protocol for generating random shared multiplicative triples, based on a homomorphic cryptosystem. The Orlandi protocol needs a group $\mathbb{G}$ of some prime order p which is specified by the generator $g \in \mathbb{G}$. The order p and the generator g are part of the public parameters. A secret $x \in \mathbb{Z}_p$ is shared in the Orlandi protocol using additive secret sharing. Every party of the computation holds a share x_i of the secret, two uniformly randomly chosen additively secret shared elements $\rho_{i,1}$ and $\rho_{i,2}$ in $\mathbb{Z}_p$ and a public commitment C. The two random elements $\rho_{i,1}$ and $\rho_{i,2}$ are needed in order to compute the commitment to the secret, and the commitment is used when reconstructing the secret to check that no party contributed a wrong share. The commitment is computed using a double trapdoor Pedersen commitment scheme [31] based on the hardness of the discrete logarithm in the group $\mathbb{G}$. The commitment C is computed as

$C = \mathrm{Com}(x, \rho_{i,1}, \rho_{i,2}) = g^x h_1^{\rho_1} h_2^{\rho_2}$ where $h_i = g^{t_i}$ for $i \in \{1,2\}$ and t_i is a trapdoor. We denote h_1, h_2 as the public key of the commitment scheme. A share in the Orlandi protocol is a four-tuple $(\mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathcal{C})$, consisting of the share of the secret, $x_i \in \mathbb{Z}_p$, two uniformly randomly chosen numbers $\rho_1, \rho_2 \in \mathbb{Z}_p$, and a commitment $C \in \mathcal{C}$ to the secret. We write a share of the secret x as $[x]$. The protocol is secure in the Common Reference String (CRS) Model [12], and a proof of the security is sketched in Orlandi's PhD progress report [28] under the assumption of the hardness of the discrete logarithm problem in $\mathbb{G}$, the availability of a secure broadcast protocol, and the semantic security of the homomorphic cryptosystem. The security of the protocol holds, up to the security level 2^{-s} if λ and d are chosen such that:

$$s < d\ \log_2(M) + (d+1)\log_2(\ln(1+\lambda)) + 2 \tag{1}$$

where M is the number of multiplicative triples needed for a given computation. We refer the reader to Orlandi's PhD progress report [28] for the intuition behind the above expression. The parameters λ and d are used in the definition of the commands below.

The protocol can be divided into two parts: a preprocessing part where multiplicative triples are generated and an online part where arithmetic expressions are evaluated. The online part provides the commands one would usually expect from a VSS scheme such as commands for sharing a given value $\texttt{Input}(x)$, reconstructing a secret $\texttt{Open}([x])$, creating a random secret $\texttt{Rand}()$, addition, subtraction, and multiplication ($\texttt{Mul}([x],[y],[a],[b],[c])$) of shared numbers. We will not explain these commands further, except for the multiplication command. We instead refer the reader to Orlandi's PhD progress report [28]. The preprocessing part is divided into a number of building blocks (*leak-tolerant multiplication*, *triple generation*, and *triple test*), which are composed into the final triple generating (*random triple generation*) functionality which produces a list of triples. We will describe online multiplication, triple generation, and random triple generation below.

Basic Multiplication. We define the multiplication of the shares $[x]$ and $[y]$ as $[z] = \texttt{Mul}([x],[y],[a],[b],[c])$ where we assume that the parties are given a random triple $([a],[b],[c])$ s.t. $c = a \cdot b$ from a honest dealer. The multiplication is realized as follows:
1. $d = \texttt{Open}([x] - [a])$ and $e = \texttt{Open}([y] - [b])$
2. $[z] = e[x] + d[y] - de + [c]$

The basic multiplication is used both as a building block in the preprocessing phase and also for performing online multiplications. This is the main reason why multiplicative triples are generated in the preprocessing, so that they can be used in online multiplications. It also indicates that one multiplication requires one triple.

The leak-tolerant multiplication of shares $[x]$ and $[y]$ is defined as $[z] = \texttt{LTMul}([x],[y],\mathcal{M})$ where $\mathcal{M} = \{([a_i],[b_i],[c_i])\}_{i \in \{1,\ldots,2d+1\}}$ is a set of multiplicative triples. Leak-tolerant multiplication is an extension of the basic multiplication with the property that if $d+1$ triples (a_i, b_i, c_i) are uniformly random in

view of the adversary then the protocol leaks no information about x, y, and $x \cdot y$.

`TripleGen`() generates a triple by having each party first choose random shares $[a]$ and $[b]$ including the needed randomness and the commitments. Second, each party encrypts ($\mathrm{Enc}_{\mathrm{ek}_i}(a_i)$) his share a_i using his public key ek_i and a homomorphic cryptosystem. Then he broadcasts the encrypted share, the corresponding commitment, and the commitment for b_j. The share of the product $[c] = [a] \cdot [b]$ is computed by using the homomorphic property of the received encrypted values to multiply the shares $[a_i]$ and $[b_j]$. The product is then masked with some randomness $d_{i,j}$ and sent. The share c_i is then computed by decrypting $\mathrm{Dec}_{\mathrm{sk}_i}(\gamma_{i,j})$ the product shares, adding them up and subtracting the randomness. The private key of party i is sk_i.

Triple Generation. The triple generation command `TripleGen`() creates a multiplicative triple which is shared among the parties. The triple generation is realized as follows:

1. Every party P_i chooses $a_i, r_{i,1}, r_{i,2} \in_R \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p$, computes $\alpha_i = \mathrm{Enc}_{\mathrm{ek}_i}(a_i)$, $A_i = \mathrm{Com}(a_i, r_{i,1}, r_{i,2})$, and broadcasts them
2. Every party P_j does:
   (a) choose $b_j, s_{j,1}, s_{j,2} \in_R \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p$, compute $B_j = \mathrm{Com}(b_j, s_{j,1}, s_{j,2})$ and broadcast B_j
   (b) Party P_j does, for every other party P_i: choose $d_{i,j} \in_R \mathbb{Z}_{p^3}$, compute and send $\gamma_{i,j} = \alpha_i^{b_j} \mathrm{Enc}_{\mathrm{ek}_i}(1;1)^{d_{i,j}}$ to P_i
3. Every party P_i does:
   (a) compute $c_i = \sum_j \mathrm{Dec}_{\mathrm{sk}_i}(\gamma_{i,j}) - \sum_j d_{i,j} \bmod p$
   (b) pick $t_{i,1}, t_{i,2} \in_R \mathbb{Z}_p \times \mathbb{Z}_p$, compute and broadcast $C_i = \mathrm{Com}(c_i, t_{i,1}, t_{i,2})$
4. Everyone computes $(A, B, C) = (\prod_i A_i, \prod_i B_i, \prod_i C_i)$
5. Every party P_i outputs:
   $([a]_i, [b]_i, [c]_i) = ((a_i, r_{i,1}, r_{i,2}, A_i), (b_i, s_{i,1}, s_{i,2}, B_i), (c_i, t_{i,1}, t_{i,2}, C_i))$

The computation inside encrypted values gives rise to the requirement that the modulus of the cryptosystem N must be much larger than the modulus of the shares and the commitment scheme p. This is not an issue in practice because the key size of a factorization based cryptosystem is usually much bigger than the order of the group of points on an elliptic curve, if the same level of security is to be obtained.

The triple test command `TripleTest`() creates one multiplicative triple from two. The first triple is used to check the correctness of the second triple. This removes the risk of overflow in the encrypted computation in `TripleGen`() with overwhelming probability. The overflow may occur due to the differrence in the modulus of the cryptosystem and the shares and the commitment scheme.

Random triple generation `RandomTriple`() creates a set of multiplicative triples $\mathcal{M}$ of size M, which we call the *result set*. The result set is created by first generating a larger *distillation set* $\mathcal{D}$ of triples using `TripleTest`(). The size of the distillation set depends on the security parameter and M. The result set is distilled from the distillation set by first choosing a uniformly random subset called the *test set*

$\mathcal{T} \subset \mathcal{D}$ of size $\lambda(2d+1)M$. The triples in the test set are checked for correctness, and if any inconsistency is detected the protocol is aborted. Second the remaining triples $\mathcal{D} \backslash \mathcal{T}$ are partitioned into M random sets of size $(2d+1)$. The result set is generated using these sets and the $\mathcal{F}_{\text{PP}}(\textbf{rand}, \ldots)$ functionality.

Random Triple Generation. The implementation of the random triple generation command `RandomTriple`() creates a set $\mathcal{M}$ of multiplicative triples of size M which is shared among the parties. The random triple generation is realized as follows:

1. $\mathcal{D} = \emptyset$. For $i = 1, \ldots, (1+\lambda)(2d+1)M$ do: $\mathcal{D} = \mathcal{D} \cup$ `TripleTest`()
2. Coin-flip a subset $\mathcal{T} \subset \mathcal{D}$ of size $\lambda(2d+1)M$
3. For all $i \in \mathcal{T}$ the parties reveal the randomness used for `TripleTest`()
4. Check that the randomness is consistent with the view. Check that $a_i, b_i < p$ and $d_{i,j} < p^3$. Abort otherwise.
5. Partition $\mathcal{D} \backslash \mathcal{T}$ in M random subsets $\mathcal{D}_i$ of size (2d + 1)
6. For $i = 1, \ldots, M$ do:
   (a) $[a] = \mathcal{F}_{\text{PP}}(\textbf{rand}, \ldots)$, $[b] = \mathcal{F}_{\text{PP}}(\textbf{rand}, \ldots)$, $[r] = \mathcal{F}_{\text{PP}}(\textbf{rand}, \ldots)$
   (b) $[c] =$ `LTMul`$([a], [b], \mathcal{D}_i)$ and `Open`$([c] + [r])$
   (c) Add $([a], [b], [c])$ to $\mathcal{M}$

The $\mathcal{F}_{\text{PP}}(\textbf{rand}, \ldots)$ functionality used in the random triple generation creates a random share using the $\prod_{\text{comm}}$ protocol described in Chapter 4 of Orlandi's PhD progress report [28]. The difference between using the `Rand`() function and the $\mathcal{F}_{\text{PP}}$ functionality is twofold. First, they differ in the security model of the overall protocol. If one uses the `Rand`() function then the protocol provides stand-alone security [11] whereas if one uses the $\mathcal{F}_{\text{PP}}$ functionality then it is secure in the CRS model. Second, they differ in speed. The `Rand`() function is faster than the $\mathcal{F}_{\text{PP}}$ functionality because the latter generates random shares using Universal Composable commitments whereas the first does not. In the implementation we use the `Rand`() function and thus achieve stand-alone security.

In the original protocol it is assumed that the public key for the commitment scheme is provided to the parties by a trusted third party (TTP), so that the key is randomly chosen. However in a real world setting, where the parties don't trust each other, it might not be the case that there is a single TTP that all parties trust. Other ways of generating the public key might include: measuring some physical random quantity, running a coin-flip protocol, or modeling a hash function with a random oracle (e.g. the first party can choose a random string r and publish $(r, H(r))$ and everyone parses $H(r)$ as the public key). The security of the whole protocol will reflect the security of the method used to generate the public key.

3 Implementation of the Orlandi Protocol

In this section we describe how we implemented the Orlandi protocol using VIFF, the Virtual Ideal Functionality Framework [14,32]. VIFF is an open source framework implemented in Python for executing general multiparty computations. It

is possible to extend VIFF with new protocols for evaluation of arithmetic circuits. Such protocols are called *runtimes* in VIFF lingo and are materialized by the `Runtime` class, which new runtimes must subclass. A share in VIFF is represented by instances of the `Share` class. A `Share` instance represents a value to be computed in the future, and one can attach callbacks which will be executed once the share gets a concrete value. A share in the Orlandi protocol is represented using the `OrlandiShare` class which extends `Share`. The concrete value held by an `OrlandiShare` forms a tuple as described in Section 2.

The Orlandi protocol is implemented as the `OrlandiRuntime`, a subclass of `Runtime` and as such overloading the usual addition, subtraction and multiplication operators. It also provides some further methods largely corresponding to the commands described in Section 2. The implementation of the various commands follows the protocol closely, except that we combine steps and/or schedule them in parallel whenever possible. An example where we combine steps is step 1, 2.a, and 2.b of `TripleGen`() where we save one broadcast operation. An example of scheduling operations in parallel is the `TripleTest`() command where two `TripleGen`() commands are scheduled in parallel with one `Open`() and a `Rand`() command.

To speed-up the computation it can be observed that in step 2.c of the `TripleGen`() function that $\mathrm{Enc}_{\mathrm{ek}i}(1;1)$ will result in g^{N+1} when encrypting with the Paillier system. Hence, γ_{ij} can be computed by using a simultaneous multi-exponentiation method as described in Section 5, i.e. $\gamma_{ij} = \alpha_i^{b_j}(g^{N+1})^{d_{ij}}$. In addition, when using homomorphic properties of the Paillier cryptosystem, step 3.a can be rewritten to $c_{ij} = \mathrm{Dec}_{\mathrm{sk}}(\prod_j \gamma_{ij} \mod N^2) - \sum_j d_{ij}$, which results in just doing one exponentiation in total instead of one per party.

The security of the Orlandi protocol is based on the assumption of the hardness of the discrete logarithm of the group used and the presence of a broadcast channel. We satisfy the hardness assumption of the discrete logarithm by computing the commitments in a group defined by an elliptic curve over the field $\mathbb{F}_p$ with prime p of 192-bits with the generator g and the public key h_1, h_2 which consists of points on the curve. We have implemented the commitment scheme as a Python C extension using the industry strength PrimeInk ECC library v. 6.4.0 [1]. The main obstacle was the conversion from integers in base 2^{15}, which is used as the internal representation of arbitrary precision integers in Python, to base 2^{32} which is the representation used by PrimeInk ECC. The broadcast channel assumed by the Orlandi protocol is implemented using an instance of the *weak-crusader* broadcast. The weak-crusader broadcast is a variant of the *crusader* broadcast [16] where we allow a malicious adversary to make some honest parties output a message while others abort. The crusader broadcast is not needed in the Orlandi case since the protocol is already vulnerable to denial of service attacks, e.g. an adversary can just refrain from sending messages at all. By relaxing the requirements on the broadcast protocol we also get a more efficient implementation since we do not need a signature scheme. The protocol consists of two rounds. In the first round the *senders* send a value to each of the receivers, who then computes a collision resistant hash of the received value, and

sends it to the other receivers in the second round, who check the correctness. We generally use the broadcast protocol in the implementation for broadcasting from a set of parties to all parties, except for the share reconstruction command in the case where only some subset of the parties should learn the output.

Broadcast. $ls = \text{broadcast}(value, senders, receivers)$, where the result ls is a list of received values.

1. Each party $P_j \in senders$ sends *value* to every party $P_i \in receivers$
2. Every party P_i in *receivers* computes a collision resistant hash on the received value and sends the hash to every other party in *receivers*
3. Each party in *receivers* checks that the received hash is equal to the hash computed by the party in the previous step, and returns *value* if true, or aborts if not

4 Benchmarks

In this section we describe how we have benchmarked our implementation with various levels of optimization, and discuss the results. We have chosen to benchmark the three commands `Mult`, `TripleGen`, and `RandomTriple`, because the other commands are not much different than the commands in a standard additive secret sharing scheme. `Mult` and `TripleGen` are straightforward to benchmark since they do not depend on the security parameter. The execution of `RandomTriple` on the other hand depends on the security parameter and the needed number of triples.

The `RandomTriple` command generates a set of triples which is distilled into a smaller set that is the result of the command. The total number of triples generated is $(1+\lambda)(2d+1)M$ where M is the size of the result set, and λ and d have to satisfy Equation 1. The overhead of distilling M triples is $(1+\lambda)(2d+1)-1$, and it is clear from Equation 1 that the overhead increases as the security parameter goes up, but also that it decreases as the number of needed triples M increases. This gives two interesting dimensions along which to investigate the execution time.

It is infeasible to benchmark every possible combination of security parameter and number of triples so we chose the security parameter values 1 (covert security [2]), 16, and 21, and 5, 10, and 30 triples, because they are representative and feasible for the interval of interesting security parameters $[1, 32]$. They are feasible in the sense that they can be computed in a reasonable amount of time. We have chosen λ and d such that the overhead is minimal.

The benchmarks are created using the VIFFBench Framework [33], which automates benchmarking of VIFF protocols. The benchmarks are defined as a small program parametrized with the number of parties and the VIFF repository revision. The results are automatically stored in a relational database. We have chosen to hardwire two numbers t_1, t_2 into the implementation, in order to avoid unnecessary complexity. The numbers are used to compute the public keys (as g^{t_1}, g^{t_2}) for the commitment scheme. This breaks the security of the implementation if one is to use the implementation for practical applications.

It does, however, not influence the efficiency of the commands, because the key can be computed in a setup phase, before the preprocessing phase. We have performed three benchmarks which were executed for each of the VIFF revisions containing significant improvements to the commands. Except for random triple generation which is only benchmarked for revision 1435. The online multiplication benchmark consist of 100 multiplications run in parallel. If we only executed one multiplication we would get too close to the resolution of the system clock that it would affect the precision of our measurements. The triple generation and the random triple generation benchmarks, on the other hand, only execute one invocation of the corresponding commands, because the execution time is much longer. For each revision we have repeated each benchmark 50, 50, and 1 times for online multiplication, triple generation, and random triple generation, respectively, in order to eliminate random noise. Executing the random triple generation 50 times for each revision would be prohibitively time consuming. All the benchmarks are performed using 1024-bits key size for the Paillier cryptosystem. We do not investigate how the implementation behaves as the latency on the network changes. The benchmarks were performed by using up to 10 identical computers equipped with 1 GHz dual-core AMD Opteron 2216 processors with 2x1 Mb level 2 cache and 2 Gb RAM each. The hosts are running Red Hat Enterprise Linux 5.2 on a 64-bit x86 architecture and were connected using gigabit Ethernet with a round-trip latency of 0.104 ms. One of the machines was chosen as the coordinator, whose responsibility it was to distribute and execute the benchmarks on the needed subset of the nine other machines. VIFFBench chooses the subset randomly.

4.1 Benchmark Results

The results of the basic multiplication benchmarks are shown in Table 1 where the average execution time for one multiplication is presented for two to nine parties along with the standard deviation. We clearly see that the implementation is efficient and achieves 15.9 ms per multiplication for three parties. The figure also shows that the average execution time increases linearly as the number of parties increases which is as expected due to the broadcast. A multiplication basically consists of two `Open`() operations with execution time linear in the number of parties. The execution time for two parties is not as expected. Based on the protocol we would expect it to be faster than for three parties, but the measurements shows that it is slower than for three, four, five, six, and even seven parties. We contemplate that the cause of this anomaly is that for two parties the implementation is CPU bound and not network bound as we have observed for three parties. The standard deviations are large compared to the measurements, and indicates some variation in our measurements. However the timings are meaningful and the basic multiplication is useful in practice even if we take the standard deviation into account.

Table 2 shows the average execution time of triple generation for two, three, and nine parties and the data is visualized as a graph in Figure 1. We only show a subset of our measurements, please see the extended version of this

Table 1. The average *execution time* in ms. of selected Basic Multiplication benchmarks as function of the number of *parties*

parties	2	3	4	5	6	7	8	9
time (ms)	27.4	15.9	19.7	22.8	25.6	26.7	28.2	35.9
stdvar (ms)	0.1	3.5	4.7	6.7	7.4	6.8	8.1	8.3

paper [18] for the full set of measurements. We have benchmarked different revisions of our implementation corresponding to the various optimizations we have performed. Revision 1231 is the initial unoptimized implementation which uses the implementation of Paillier in VIFF, revision 1355 in-lined step 1, 2.a, and 2.b of `TripleGen`(), 1370 uses our efficient implementation of the Paillier cryptosystem, 1393 moves step 2.c into C, 1399 moves step 3.a into C, and 1400 is a minor technical optimization.

The performance of the final revision is below 200 ms for all but two and four parties. This is encouraging for practical uses of the protocol. Based on the definition of `TripleGen`() we would expect to see the execution time increase linearly ($\mathcal{O}(n)$) in the number of players n. This is also the case until revision 1393. One explanation is that random noise is more dominant when the measured time is small. We again see that two parties are slower than even nine parties, but it seems like the anomaly is introduced in revision 1370, where we use a more efficient implementation of Paillier. This is consistent with our earlier observation that the two party case is CPU bound and the other are network bound. It is clear from Figure 1 that the use of an efficient implementation of Paillier gives a substantial improvement of the execution time and is the main contributor to the efficiency of the Orlandi implementation. The Figure shows that rewriting step 2.c in C gives a larger performance increase than rewriting step 3.a does, which is as we would expect. Step 2.c is more computational intensive. The improvements we have done in the Python code in revision 1355 and 1400 are dwarfed by the other improvements.

Table 3 shows the average execution time of random triple generation defined in revision 1435 for two, three, and nine parties. The full set of measurements can be found in the extended version of this paper [18]. One would expect the benchmarks to show that the execution time per triple increases as the security

Table 2. The average *execution time* in ms. of triple generation as a function of *number of parties*

parties		1231	1355	1370	1393	1399	1400
2	time	3519.6	3519.6	894.6	243.8	226.5	224.2
2	stdvar	1.0	0.8	3.2	0.9	0.7	0.7
3	time	3972.7	4012.1	376.3	155.0	168.3	170.9
3	stdvar	94.8	157.4	72.1	59.2	35.9	38.2
9	time	8937.4	8849.7	846.9	237.0	188.9	188.4
9	stdvar	460.2	281.2	27.0	36.5	20.7	29.0

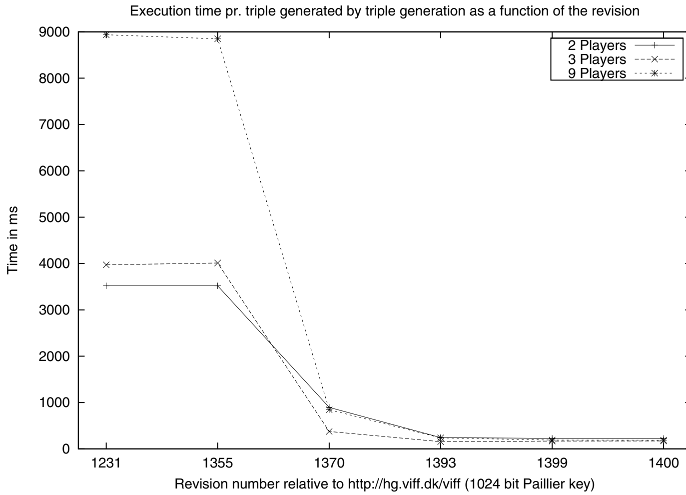

Fig. 1. The average *execution time* in ms. of selected Triple Generation benchmarks as a function of the changes to the implementation

parameter increases ($\mathcal{O}(d \log \ln(\lambda))$) and decreases as the number of triples increases. The measurements shows that the execution time increases as the security parameter increases. And, in most cases the execution time also decreases as the number of triples increases. Random noise may explain the cases where we do not see the a decrease in execution time. We have only run the benchmarks once for each combination of security parameter and number of triples due to time considerations.

4.2 Performance Comparison

We are aware of two other implementations of secure multiparty computation protocols with active security: A protocol by Damgård, Geisler, Krøigaard, and Nielsen (DGKN) which has been implemented in VIFF [14], and a protocol by Lindell, Pinkas, and Smart (LPS) [22]. The performance of these implementations cannot be directly compared to the performance of the Orlandi protocol since they rely on different security models and the benchmarks have been executed

Table 3. The average execution time in seconds of random triple generation as a function of *parties* (2, 5, and 9), *security parameter*, and *number of triples*

s	1	1	1	16	16	16	21	21	21
t	5	10	30	5	10	30	5	10	30
2	1.872	1.511	1.370	15.879	15.157	11.641	20.959	16.560	16.453
3	1.598	0.952	1.059	11.796	11.883	10.944	16.931	15.981	15.269
9	2.238	1.799	1.794	25.931	24.444	25.638	31.901	32.572	37.545

on different hardware. However, in this section we will try to elaborate on the difference between the performance of the systems.

The DGKN protocol provides evaluation of arithmetic circuits and is secure against an adaptive active adversary up to a threshold of $n/3$ corrupted parties. An adversary may halt the computation up to a synchronization point, not after - in which case termination is guaranteed. The LPS protocol is a 2-party protocol for evaluation of boolean circuits. The Orlandi protocol is a full-threshold multiparty protocol. Both the Orlandi and the LPS protocol are secure against a static active adversary, the security is based on cryptographic assumptions, and they are "unfair" in the sense that a corrupt party can prevent honest parties from getting any result while the corrupt party get results himself.

It is difficult to make a direct comparison between our results and those reported for the LPS protocol, since they do not benchmark multiplications, but rather comparisons of 16-bit integers.

Results have been reported for the DGKN protocol for 4, 7, 10, 13, 16, 19, 22, and 25 parties. If we compare the numbers for 4 and 7 parties, which is the setups we have numbers for, then DGKN takes 4 and 6 ms which is only a factor 5 (roughly) better than our results for the online case. Whereas in the preprocessing case the DGKN uses 5 ms and 22 ms in the best case, which is a factor 80-240 better then our implementation even with the smallest security parameter. Based on these numbers our implementation may seem inferior, but remember that the Orlandi protocol provides full threshold. And the benchmarks show that it is possible to use the Orlandi protocol in practice.

5 High-Performance Paillier

Various non-deterministic cryptosystems have been proposed based on randomized encryption schemes which encrypt a message m by raising a base g to the power m and suitably randomizing this result [5,15,17,26,27,29]. The security of these systems is based on the intractability of various "residuosity" problems. As an important consequence of this encryption technique, those schemes have homomorphic properties. These homomorphic properties enable computation on ciphertexts without knowing the context. This allows for a wide spread of applications such as secure multiparty computations, like Orlandi's protocol. In this section we discuss the Paillier cryptographic system and its implementation and optimization.

5.1 Description of the Paillier Schemes

Paillier has presented multiple closely related cryptosystems [29,30]. We will focus on the main- and subgroup variants of these cryptosystems shown in Figure 2 and 3, respectively. The subgroup variant is slightly different as it computes residues in a subgroup of order $\lambda(N)$.

Key Generation Let N be a RSA modulus $N = pq$, where p and q are large prime integers. Let $g \in \mathbb{Z}^*_{N^2}$ be chosen such that its order is a multiple of N. Let $\lambda(N) = \text{lcm}(p-1, q-1)$. The public key is (g, N), and the private key is $\lambda(N)$

Encryption To encrypt a message $m \in \mathbb{Z}_N$, randomly chose $r \in \mathbb{Z}^*_N$ and compute the ciphertext $c = g^m r^N \mod N^2$.

Decryption The decryption of c is defined by $\frac{L(c^\lambda \mod n^2)}{L(g^\lambda \mod N^2)} \mod N$. Where the $L(\mu)$ function is defined as $\frac{\mu-1}{N}$ and takes inputs of $S_N = \{u < N^2 \mid u = 1 \mod N\}$.

Fig. 2. The main variant

5.2 Paillier Performance Evaluation

A common task in implementations of many public-key cryptosystems is multi-exponentiation in commutative groups. This is also the case for the Paillier cryptosystem, namely computing $g^m r^N \mod N^2$ for the main variant and $g^m (g^N)^r \mod N^2$ for the subgroup. Many algorithms have been proposed to speed-up the computation a single exponentiation [10,24,25,35,21]. These algorithms can be modified to compute a product over of multiple exponentiations in such a way that it is faster than a product of single exponentiations. In the following subsections we show how to reduce this overhead with different simultaneous multi-exponentiation algorithms.

The simultaneous 2^k-ary method was first introduced by Brauer [10], the idea behind the method is slicing the binary representation of an exponent into pieces using a windows of length k and processing the exponent in a larger basis. For each evaluation of the exponent the intermediate results get raised by power of 2^k and multiplied by its base raised to the power of the evaluated bits in the exponent. The powers $\{0, 1, 2, \ldots, 2^k - 1\}$ of base g are precomputed in an auxiliary table.

In order to make the 2^k-ary method evaluate two powers at the same time (i.e. $g_1^{e_1} g_2^{e_2} \mod n$). Two separate auxiliary tables with there powers of g_1 and g_2 are required. Each time both exponents get evaluated at the same time. First the intermediate result is raised to 2^k and is multiplied by each separate base raised to the power of the evaluated bits from the corresponding exponent. This saves a squaring for every bit that is evaluated.

Key Generation Let N be a RSA modulus $N = pq$, where p and q are large prime integers. Let $\lambda(N) = \text{lcm}(p-1, q-1)$ and choose α such that it divides $\lambda(N)$. Let $h \in \mathbb{Z}^*_{N^2}$ such that is has maximal order of $n\lambda(N)$, and $g = h^{\lambda/\alpha} \mod N^2$. The public key is (g, N), and the private key is α

Encryption To encrypt a message $m \in \mathbb{Z}_N$, randomly chose $r \in \mathbb{Z}^*_N$ and compute the ciphertext $c = g^{m+r \cdot N} \mod N^2$.

Decryption The decryption of c is defined by $m = \frac{L(c^\alpha \mod N^2)}{L(g^\alpha \mod N^2)} \mod N$. Where the $L(\mu)$ function is defined as $\frac{\mu-1}{N}$ and takes inputs of $S_N = \{u < N^2 \mid u = 1 \mod N\}$.

Fig. 3. The subgroup variant

Algorithm 1. 2^k-ary Method

Require: $aux_a, aux_b, b = 2^k - 1$, e_1, e_2
Ensure: $g_1^{e_1} \cdot g_2^{e_2}$

```
A ← 1
for j = ⌊(b − 1)/w⌋w down to 0 do
A ← A^(2^k)
  if (e1[j + w − 1 . . . j])
    A ← A · aux_g1[e_{1,j+1}, e_{1,j+2}, . . . e_{1,j+w−1}]
  if (e2[j + w − 1 . . . j])
    A ← A · aux_g2[e_{2,j+1}, e_{2,j+2}, . . . e_{2,j+w−1}]
```

The simultaneous 2^k-ary matrix method is a slight modification of the simultaneous 2^k-ary method. The main difference is the computation of the auxiliary table which consists of $k \times k$ table entries which holds for $0 < i, j < 2^k - 1$ the product of $g_1^i g_2^j$ in entry $aux[i][j]$. Building such a table requires more pre-computation, but gives one less multiplication per evaluated window of length k.

The simultaneous sliding window exponentiation method of Yen, Laih, and Lenstra [35] is an improvement of the 2^k-ary method. Just like the 2^k-ary method the sliding window method consists of slicing the binary representation of e_i into pieces using a window of length ω and processing the part one by one. The addition of letting the window slide allows us to skip consecutive zeros in e_i while squaring the intermediate result. As a result, evaluation of two even exponents are avoided, and computation of the entries of these entries in the auxiliary table can be avoided. This results in a generally faster algorithm for evaluating exponents.

To evaluate two exponents simultaneously we apply the same trick as for simultaneous 2^k-ary method. But with the change that we check if both evaluated bits are zero. Additional bookkeeping is needed to keep track of the bits.

When comparing both decryption version of the Paillier scheme, the basic computation consists of one fixed exponentiation and a multiplication in $\mathbb{Z}_{N^2}^*$ and a multiplication in $\mathbb{Z}_N^*$. The subgroup variant requires the same operation, except the size of the exponent α is smaller, which makes the subgroup variant faster. Paillier has suggested an alternative decryption method by means of the Chinese Remainder Theorem (CRT). By defining $L_p = \frac{\mu-1}{p}$ and $L_q = \frac{\mu-1}{q}$ we can decrypt by separately computing the message modulo p and q and combining m_p and m_q with CRT. First compute $h_p = L_p(g^{p-1} \mod p^2)$ and $h_q = L_q(g^{q-1} \mod q^2)$ then $m_p = L_p(c^{p-1} \mod p^2)h_p \mod p$ and $m_q = L_q(c^{q-1} \mod q^2)h_q \mod q$ and finally recombine using CRT. Additional speed-up can be found by computing $L(\mu)$ with $\mu \cdot n^{-1} \mod 2^{|N|}$. This is just a multiplication and a logical `AND`. Another way to make computations more efficient is a careful choice of parameters. For instance if one chooses $g = 1 + n$ then the exponentiation g^m can be executed using only one multiplication, namely $g^m = (1+n)^m \equiv (1+mn) \mod n^2$. This only works for the encryption in the main variant of the Paillier scheme. Such optimizations can provide substantial speed-ups as we show in the next subsection.

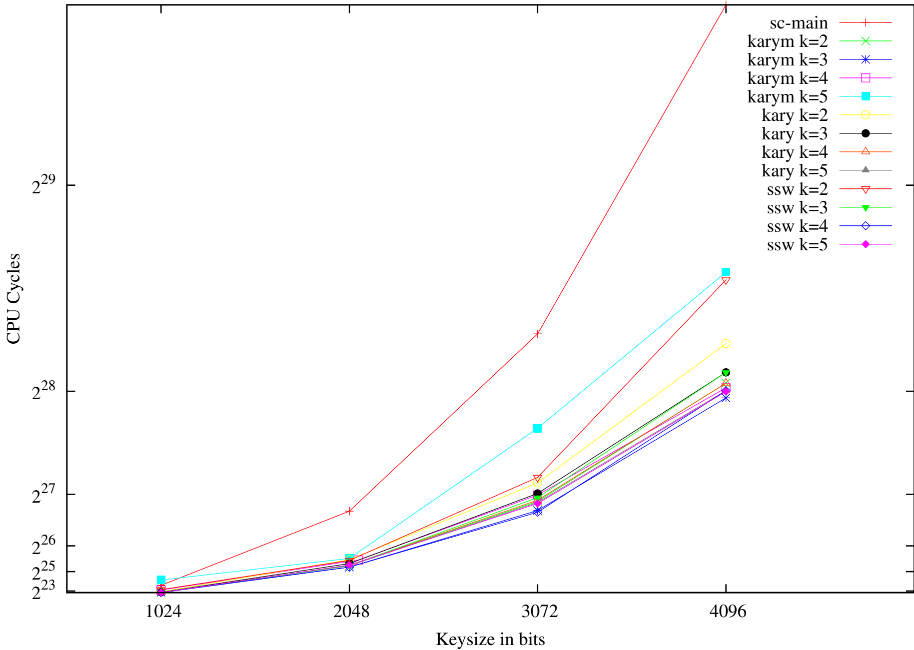

Fig. 4. The *execution time* in CPU cycles with different keys sizes with *simultaneous sliding window method* (ssw), *simultaneous* 2^k*-ary method* (kary) and *simultaneous* 2^k*-ary method* (karym) Parameter k is the size of the window in bits and the *main variant* (sc-main) with $g = n + 1$

5.3 Results

In this section we describe how we benchmarked our implementation with various optimizations and discuss the results. We have chosen to benchmark the three above simultaneous multi-exponentiation algorithms and their parameters for key sizes N ranging from 1024-bit to 4096-bit with increments of 1024-bits.

To benchmark the speed of encryption of the Paillier cryptosystem we ran the main variant with $g = N + 1$ and the subgroup variant with the 3 different simultaneous multi-exponentiation algorithms, with window size k as parameter. For windows size k we select $1 < k \leq 5$, as choosing k higher than 5 will result in longer pre-computation for these bit sizes.

The benchmarks can be found in Figure 4. The speed of the algorithms becomes clear as the key-sizes increase. As it is infeasible to test all combinations of random element r and message m, we have chosen m and r to have the same size as the sub-group length in bits this is approximately $1/4$ of N. This is due to the requirement of the Paillier cryptosystem. The benchmark includes the generation of auxiliary tables.

It is clear that the main variant of the Paillier with optimized parameters is slowest, this is mostly due to the computation of r^N with N having $\{1024, \ldots 4096\}$-bits. The simultaneous 2^k-ary normal and matrix variant as well as simultaneous

sliding window method perform better with a greater k if the key length gets bigger. Also, we can see that 2^k-ary matrix performs better than 2^k-ary method as the key size grows. The 2^k-ary matrix has larger pre-computation but has one multiplication less for evaluating one bit of the exponent. The simultaneous sliding window method is in most cases the best performing algorithm, this is due to the fact that it skips consequent zeros.

The benchmarks were performed by using a 2 GHz Intel Pentium E2180 dual core with 1024KB cache per core and 2 GB Ram. The system is running Fedora release 8 with kernel 2.6.26.8-57.fc8 in 64-bit mode. For benchmarking we used *cpucycles* with is part of eBACS [6] to measure the amount of CPU cycles used by the execution.

6 Related Work

Several practical systems for general multiparty computation have been implemented during the recent years. FairPlay [23] is the earliest implementation that the authors are aware of. In the system one can specify computations in a high-level, procedural programming language. Using the FairPlay compiler, the high-level programs are then compiled to low-level representations of one-pass boolean circuits. These circuits are then used for secure computation as described by Yao [34]. The timings reported on FairPlay show that FairPlay is efficient, but it should be noted that it only supports two-party computation in the passive security model. FairPlay has later been supplemented by FairPlayMP [3] which is capable of handling the case with more than two parties in the passive security model assuming less than $n/2$ corrupted parties. A two-party protocol for secure computation which is secure against a static active adversary has recently been implemented [22]. Like the protocol used in FairPlay, this protocol is also based on boolean circuits.

Another practical system for general multiparty computation was created by Bogetoft et al. [9] in the SIMAP project. The system was used for the first known large-scale commercial application of secure multiparty computation [8]. It supports general multiparty computation in a passive threshold security model assuming less than $n/2$ corrupted parties. Like FairPlay it lets users express programs in a high level language, but contrary to FairPlay, it evaluates the programs as arithmetic rather than boolean circuits. The downside of this strategy is that comparison of integers becomes more complex and time consuming. The protocol used in the SIMAP system has also been implemented in the VIFF framework [32]. In addition, VIFF contains a passively secure two-party protocol based on the Paillier cryptosystem as well as an implementation of a multiparty protocol described in Section 4. The ShareMind system [7] represents yet another efficient approach to practical multiparty computation based on arithmetic circuits and additive sharing. It only supports three parties in the passive model and assumes that at most one party gets corrupted. None of the above implementations, though, support the combination of active security and self-trust, that is available with the implementation of the Orlandi protocol described in this paper.

7 Conclusion and Future Work

In this paper we presented an implementation of the Orlandi protocol, which is the first implementation of a MPC protocol based on arithmetic circuits, which is secure against up to $n-1$ static, active adversaries. We showed that the protocol can be implemented efficiently in the presence of an efficient implementation of a double trapdoor Petersen commitment scheme and a homomorphic cryptosystem. We also described an efficient implementation of the Paillier cryptosystem.

Practical uses is an interesting direction of future work e.g. auctions, benchmarks, and online games. Also it would be interesting to implement and benchmark the $\mathcal{F}_{\mathrm{PP}}(\mathbf{rand}, \ldots)$ functionality and a suitable setup phase.

A further direction of future work would be to implement and benchmark the Lim/Lee [21] and the fractional window exponentiation [24] algorithms, which we expect would provide further speed up.

Acknowledgements

The authors would like to sincerely thank Ivan Damgård, Claudio Orlandi, and Jesper Buus Nielsen for answering our questions about the Orlandi protocol. In addition many thanks to Tanja Lange, Daniel J. Bernstein, and Peter Schwabe for comments and suggestions on the Paillier cryptosystem. We thank the VIFF and VIFFBench Development Teams for creating VIFF and VIFFBench, respectively. Thanks to our partners in the CACE project for making this collaboration possible. Also thanks to Cryptomatic A/S for letting us use the PrimeInk ECC library, and to the anonymous reviewers for suggestions on improving the paper.

References

1. Cryptomatic A/S. PrimeInk ECC library v. 6.4.0, `http://www.cryptomatic.com`
2. Aumann, Y., Lindell, Y.: Security against covert adversaries: Efficient protocols for realistic adversaries. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 137–156. Springer, Heidelberg (2007)
3. Ben-David, A., Nisan, N., Pinkas, B.: FairplayMP: a system for secure multi-party computation. In: Ning, P., Syverson, P.F., Jha, S. (eds.) ACM Conference on Computer and Communications Security, pp. 257–266. ACM, New York (2008)
4. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: Simon, J. (ed.) [19], pp. 1–10
5. Benaloh, J.D.C.: Verifiable Secret-Ballot Elections. PhD thesis, Yale University (1978)
6. Bernstein, D.J., Lange, T.: eBACS: ECRYPT benchmarking of cryptographic systems, `http://bench.cr.yp.to`
7. Bogdanov, D., Laur, S., Willemson, J.: Sharemind: A framework for fast privacy-preserving computations. In: Jajodia, S., Lopez, J. (eds.) ESORICS 2008. LNCS, vol. 5283, pp. 192–206. Springer, Heidelberg (2008)

8. Bogetoft, P., Christensen, D.L., Damgård, I., Geisler, M., Jakobsen, T., Krøigaard, M., Nielsen, J.D., Nielsen, J.B., Nielsen, K., Pagter, J., Schwartzbach, M.I., Toft, T.: Secure multiparty computation goes live. In: Dingledine, R., Golle, P. (eds.) FC 2009. LNCS, vol. 5628, pp. 325–343. Springer, Heidelberg (2009)
9. Bogetoft, P., Damgård, I., Jakobsen, T.P., Nielsen, K., Pagter, J., Toft, T.: A practical implementation of secure auctions based on multiparty integer computation. In: Di Crescenzo, G., Rubin, A. (eds.) FC 2006. LNCS, vol. 4107, pp. 142–147. Springer, Heidelberg (2006)
10. Brauer, A.: On addition chains. Bulletin of the American Mathematical Society 45(10), 736–739 (1939)
11. Canetti, R.: Security and composition of multiparty cryptographic protocols. Journal of Cryptology 13(1), 143–202 (2000)
12. Canetti, R., Fischlin, M.: Universally composable commitments. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 19–40. Springer, Heidelberg (2001)
13. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (extended abstract). In: Simon, J. (ed.) [19], pp. 11–19
14. Damgård, I., Geisler, M., Krøigaard, M., Nielsen, J.B.: Asynchronous multiparty computation: Theory and implementation. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009. LNCS, vol. 5443, pp. 160–179. Springer, Heidelberg (2009)
15. Damgård, I., Geisler, M., Krøigard, M.: Homomorphic encryption and secure comparison. International Journal of Applied Cryptography 1(1), 22–31 (2008)
16. Dolev, D.: The byzantine generals strike again. Technical report, Stanford University, Stanford, CA, USA (1981)
17. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences 28(2), 270–299 (1984)
18. Jakobsen, T.P., Makkes, M.X., Nielsen, J.D.: Efficient Implementation of the Orlandi Protocol Extended Version. Cryptology ePrint Archive, Report 2010/224 (2010), `http://eprint.iacr.org/`
19. In: STOC '88: Proceedings of the twentieth annual ACM symposium on Theory of computing, May 1988. ACM, New York (1988)
20. Lee, P.J., Lim, C.H. (eds.): ICISC 2002. LNCS, vol. 2587. Springer, Heidelberg (2003)
21. Lim, C.H., Lee, P.J.: More flexible exponentiation with precomputation. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 95–107. Springer, Heidelberg (1994)
22. Lindell, Y., Pinkas, B., Smart, N.P.: Implementing two-party computation efficiently with security against malicious adversaries. In: Ostrovsky, R., De Prisco, R., Visconti, I. (eds.) SCN 2008. LNCS, vol. 5229, pp. 2–20. Springer, Heidelberg (2008)
23. Malkhi, D., Nisan, N., Pinkas, B., Sella, Y.: Fairplay - Secure Two-Party Computation System. In: USENIX Security Symposium, pp. 287–302. USENIX (2004)
24. Möller, B.: Improved techniques for fast exponentiation. In: Lee, Lim (eds.) [20], pp. 298–312
25. Montgomery, P.L.: Modular multiplication without trial division. Mathematics of computation 44(170), 519–521 (1985)
26. Naccache, D., Stern, J.: A new public-key cryptosystem. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 27–36. Springer, Heidelberg (1997)
27. Okamoto, T., Uchiyama, S.: A new public-key cryptosystem as secure as factoring. In: Nyberg, K. (ed.) EUROCRYPT 1998. LNCS, vol. 1403, pp. 308–318. Springer, Heidelberg (1998)

28. Orlandi, C.: LEGO and Other Cryptographic Constructions - PhD Progress Report (March 2009), http://www.cs.au.dk/~orlandi/
29. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999)
30. Paillier, P., Pointcheval, D.: Efficient public-key cryptosystems provably secure against active adversaries. In: Lam, K.-Y., Okamoto, E., Xing, C. (eds.) ASIACRYPT 1999. LNCS, vol. 1716, pp. 165–179. Springer, Heidelberg (1999)
31. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992)
32. VIFF - The Virtual Ideal Functionality Framework, http://viff.dk
33. VIFFBench Framework, http://bitbucket.org/tpj/viffbench
34. Yao, A.C.-C.: How to generate and exchange secrets (extended abstract). In: Foundations of Computer Science, pp. 162–167. IEEE, Los Alamitos (1986)
35. Yen, S.M., Laih, C.S., Lenstra, A.K.: Multi-exponentiation. Computers and Digital Techniques 141(6), 325–326 (1994)