

Multi Party Computation: From Theory to Practice

Nigel P. Smart

Department of Computer Science,
University Of Bristol,
Merchant Venturers Building,
Woodland Road,
Bristol, BS8 1UB.

January 8, 2013

What if?

Take two drug companies.

Each has a database of molecules and toxicology test results.

They want to combine their results

Without revealing what molecules are in the databases.

What if?

A government wants to search network traffic for a specific anomolous behaviour.

But the network operator does not want to give access to the network to the government.

And the government does not want to reveal exactly what behaviour it is searching for.

Computing on Encrypted Data

There are two main ways of performing such **Computations On Encrypted Data**:

Fully Homomorphic Encryption

- ▶ First scheme developed in 2009
- ▶ Party *A* sends encrypted data to party *B*.
- ▶ Party *B* does some computation and returns the encrypted result to party *A*
- ▶ Party *A* now decrypts to find out the answer.

Multi-Party Computation

- ▶ First schemes developed in mid 1980's.
- ▶ Parties jointly compute a function on their inputs using a protocol
- ▶ No information is revealed about the parties inputs.

Theory

In theory both such technologies can compute anything.

In FHE one has a huge computational cost, but zero communication.

In MPC one has virtually no computational cost, but huge communication.

In theory we can make either technology error tolerant

- ▶ Even against malicious players.

Practice

FHE is currently impractical for all but the simplest functions

- ▶ Although you can do some useful things with it.

MPC has been deployed for some operations

- ▶ Mainly against semi-honest adversaries.
- ▶ Tolerating only one baddie out of exactly three players.

We will show how to **combine** FHE and MPC to get something much better and practical.

Set up

Assume n parties of which $n - 1$ can be malicious.

Assume a global (secret) key $\alpha \in \mathbb{F}_p$ is determined

Each party i holds α_i with

$$\alpha = \alpha_1 + \dots + \alpha_n.$$

Secret Sharing

All data is represented by elements in \mathbb{F}_p .

A secret value $x \in \mathbb{F}_p$ is shared between the parties as follows

- ▶ Party i holds a data share x_i
- ▶ Party i holds a “MAC” share $\gamma_i(x)$

such that

$$x = x_1 + \dots + x_n \quad \text{and} \quad \alpha \cdot x = \gamma_1(x) + \dots + \gamma_n(x).$$

Note we can share a public constant v by

- ▶ Party 1 sets $x_1 = v$
- ▶ Party $i \neq 1$ sets $x_i = 0$
- ▶ Party i sets $\gamma_i(v) = \alpha_i \cdot v$.

Preprocessing Model

Such a sharing of x is denoted by $[x]$.

Our protocol works in the preprocessing model.

We (overnight say) generate a lot of data which is independent of the function to be computed, or its inputs.

In its basic form the data consists of triples of shared values

$$[a], [b], [c]$$

such that

$$c = a \cdot b.$$

We discuss how to produce these triples later.

The Computation

To perform the computation we utilize the following idea

Any computation can be represented by a series of additions and multiplications of elements in \mathbb{F}_p .

In other words $+$ and \times are a set of Universal Gates over \mathbb{F}_p .

We assume the players inputs are shared first using the above sharing

- ▶ Will not explain how to do this, but it is easy

So all we need do is working out how to **add** and **multiply** shared values.

Addition will be easy, multiplication will be hard.

Addition

Suppose we have two shared values $[x]$ and $[y]$.

To compute the result $[z]$ of an addition gate the parties individually execute

- ▶ $z_i = x_i + y_i$
- ▶ $\gamma_i(z) = \gamma_i(x) + \gamma_i(y)$

Note this is a local operation and that we end up with

$$\begin{aligned}z &= \sum z_i = \sum (x_i + y_i) = \left(\sum x_i\right) + \left(\sum y_i\right) \\ &= x + y, \\ \alpha \cdot z &= \sum \gamma_i(z) = \sum (\gamma_i(x) + \gamma_i(y)) = \alpha \cdot x + \alpha \cdot y \\ &= \alpha \cdot (x + y).\end{aligned}$$

Linear Secret Sharing

The addition trick works because we have a Linear Secret Sharing Scheme.

We can **locally** compute **any** linear function of shared values

i.e. given constants v_1 , v_2 and v_3 and shared values $[x]$ and $[y]$ we can compute

$$v_1 \cdot [x] + v_2 \cdot [y] + v_3 = [v_1 \cdot x + v_2 \cdot y + v_3].$$

We will now use this in our method to perform multiplication.

Note: In what follows “partially opening” a share $[x]$ means revealing x_i but not the MAC share.

Multiplication

To multiply $[x]$ and $[y]$ to obtain $[z]$ we work as follows:

- ▶ Take a new triple $([a], [b], [c])$ off the precomputed list.
- ▶ Partially open $[x] - [a]$ to obtain $\epsilon = x - a$.
- ▶ Partially open $[y] - [b]$ to obtain $\rho = y - b$.
- ▶ Locally compute the linear function

$$[z] = [c] + \epsilon \cdot [b] + \rho \cdot [a] + \epsilon \cdot \rho.$$

Note

- ▶ Each multiplication requires interaction
- ▶ If a (resp. b) is random then ϵ (resp. ρ) is a one-time pad encryption of x (resp. y).

We get the correct result because

$$\begin{aligned}c + \epsilon \cdot b + \rho \cdot a + \epsilon \cdot \rho \\&= a \cdot b + (x - a) \cdot b + (y - b) \cdot a + (x - a) \cdot (y - b) \\&= (a \cdot b) + (x \cdot b - a \cdot b) + (y \cdot a - a \cdot b) + (x \cdot y - x \cdot b - y \cdot a + a \cdot b) \\&= x \cdot y.\end{aligned}$$

Verifying Correctness

So given we can add and multiply we can compute anything

At the end of the computation we check **correctness** by interactively checking the MAC values are all correct.

Each player i has an agreed set of partially open values

$$a_j, \quad 1 \leq j \leq t$$

and each one has a sharing of the associated MAC value

$$\gamma(a_j)_i, \quad 1 \leq j \leq t.$$

Each player i also has a share of the MAC key

$$\alpha_j.$$

Verifying Correctness

We generate an agreed set of random values

$$r_j, \quad 1 \leq j \leq t$$

Each player i computes

$$a = \sum_{j=1}^t r_j \cdot a_j.$$

They also compute their share of the MAC on a

$$\gamma_i = \sum_{j=1}^t r_j \cdot \gamma(a_j)_t$$

and then

$$\sigma_i = \gamma_i - \alpha_i \cdot a.$$

Note, if all is correct then σ_i is a sharing of zero.

- ▶ So players broadcast σ_i
- ▶ Then all check whether $\sigma_1 + \dots + \sigma_n = 0$.

Preprocessing and FHE

We return to the preprocessing, which we do using FHE

- ▶ Following is a naive version, the real version has lots of bells and whistles.

We assume an FHE scheme with keys (pk, sk) whose plaintext is \mathbb{F}_p

- ▶ In practice for efficiency work on vectors of such elements in a SIMD fashion

Given $ct_1 = \text{Enc}_{pk}(m_1)$ and $ct_2 = \text{Enc}_{pk}(m_2)$ we have

$$\text{Dec}_{sk}(ct_1 + ct_2) = m_1 + m_2$$

and

$$\text{Dec}_{sk}(ct_1 \cdot ct_2) = m_1 \cdot m_2.$$

We only need to evaluate circuits of multiplicative depth one.

Preprocessing and FHE

We require a little more of our FHE scheme though

We assume a shared FHE public key pk for an FHE scheme.

- ▶ Party i holds a share sk_i
- ▶ Together they can decrypt a ciphertext ct via $\text{Dec}_{sk_1, \dots, sk_n}(ct)$.
- ▶ Each party computes $\text{Enc}_{pk}(\alpha_i)$ and broadcasts this.

Last step needed so that each party has $\text{Enc}_{pk}(\alpha)$.

Reshare

Given a ciphertext ct encrypting a value m we can make each party obtain

- ▶ An additive share m_i , s.t. $m = \sum m_i$
- ▶ And (if needed) a new fresh ciphertext ct' encrypting m .

Reshare(ct)

- ▶ Party i generates a random f_i and transmits $ct_{f_i} = \text{Enc}_{pk}(f_i)$.
- ▶ All compute $ct_{m+f} = ct + \sum ct_{f_i}$.
- ▶ Execute $\text{Dec}_{sk_1, \dots, sk_n}(ct_{m+f})$ to obtain $m + f$.
- ▶ Party 1 sets $m_1 = (m + f) - f_1$.
- ▶ Party $i \neq 1$ sets $m_i = -f_i$.
- ▶ Set $ct' = \text{Enc}_{pk}(m + f) - \sum ct_{f_i}$.

Use some “default” randomness for the last encryption.

Generating $[a]$ and $[b]$

We can generate our sharing $[a]$ as follows

- ▶ Party i generates a random a_i and transmits $ct_{a_i} = \text{Enc}_{pk}(a_i)$.
- ▶ All compute $ct_a = \sum ct_{a_i}$.
- ▶ All compute $ct_{\alpha \cdot a} = ct_\alpha \cdot ct_a$.
- ▶ Execute Reshare on $ct_{\alpha \cdot a}$ so party i obtains $\gamma_i(a)$.

Note this can also be executed to obtain $[b]$.

Generating $[c]$

This is also easy

- ▶ We have ct_a and ct_b .
- ▶ All compute $ct_c = ct_{a \cdot b}$ from $ct_a \cdot ct_b$.
- ▶ Get shares c_i via executing Reshare on ct_c ; also obtaining a fresh ciphertext ct'_c .
- ▶ All compute $ct_{\alpha \cdot c} = ct_\alpha \cdot ct'_c$.
- ▶ Execute Reshare on $ct_{\alpha \cdot c}$ so party i obtains $\gamma_i(c)$.

This is efficient despite using FHE technology because we only compute with depth one circuits.

Similar tricks with FHE allow us to perform other preprocessing making the computation phase even faster.

SPDZ and NNOS

The above is called the SPDZ protocol

Very efficient and practical for some applications.

Better security properties than other MPC implementations

More flexible in terms of parameters than other MPC implementations.

- ▶ Not quite suited to evaluating binary circuits, or circuits over small finite fields.
- ▶ A variant of the above protocol can do this.
- ▶ Or another protocol related to SPDZ due to Nielsen, Nordholt, Orlandi and Sheshank (NNOS).

Performance SPDZ

There has been a lot of work on protocols to perform various higher level operations via MPC

- ▶ Without evaluating “circuits”.
- ▶ Using ability to open data.

Many of these protocols can be improved by offloading function independent processing into the Offline phase.

- ▶ We expect more impressive results to come out in the next few months

On the next slide we present timings for our current Online phase for various functions

- ▶ Again we expect these to improve **dramatically** in the coming months

SPDZ Timing (Large Finite Field)

	Latency (sec)	Throughput (ops per sec)	Notes
32-bit Integer Mult	0.0001	800000	\approx 386 performance
32-bit Integer Comparison	0.001	2500	An Intel 4004 did 46000 per sec
Floating point Addition	0.02	50	ENIAC did 384 FLOPS
Floating point Mult	0.01	100	ENIAC did 384 FLOPS

All timings are for two players

Floating point operations are for single precision operations

- ▶ ENIAC timing from <http://knology.net/johnfjarvis/HistCompNotes.html>

Small Finite Field Example

As an example we present timings for evaluating the AES functionality for two players and active security.

	Latency (sec)	Throughput (blocks per sec)
SPDZ	0.236	4
NNOS	3.000	33

Note, we expect both of these latencies could be significantly improved by more efficient programming techniques.

AES Functionality

Why do we care about the AES functionality?

Many attacks against stored password systems in recent past by people breaking into individual servers

EMC/RSA have the following solution for **static** passwords:

- ▶ User splits password up $p = p_1^U \oplus p_2^U$
- ▶ Sends p_1 to server one, and p_2 to server two.
- ▶ Servers have another share of the password $p = p_1^S \oplus p_2^S$.
- ▶ Servers compute $t_i = p_i^U \oplus p_i^S$
- ▶ Servers sent t_i to each other.
- ▶ Accept password if

$$t_1 = p_1^U \oplus p_1^S = (p \oplus p_2^U) \oplus (p \oplus p_2^S) = p_2^U \oplus p_2^S = t_2.$$

Attacker needs to compromise both servers to get the password.

AES Functionality

The EMC/RSA solution does not work with **dynamic** passwords

- ▶ SecureID tokens.
- ▶ e-banking applications using CAP/EMV
- ▶ etc

Password is typically

$$p = \text{AES}_k(m)$$

where m is a counter, or a challenge from the server and k is the master “password”.

AES Functionality

In joint work with Bar Ilan University (Lindell) and Partisia (Damgård and Nielsen) we have a proposed system to verify such dynamic passwords using multiple servers

- ▶ Need to compromise all servers to break the system.

Basically just apply MPC to the dynamic password situation.

Now working on system to do this for real

- ▶ Extend to other dynamic password methods, e.g. DES, MD5, SHA-1 etc.
- ▶ Improve run-times

Any Questions ?