UNBOUND

A BASIC INTRODUCTION TO SECURE MULTIPARTY COMPUTATION (MPC)

Prof. Yehuda Lindell
CEO & Co-Founder | Unbound Tech

The Data Economy, the Privacy Challenge and MPC

Data has been called the modern oil, fueling our economic growth. Whether or not we agree with this statement, there is no doubt that the "monetization of data" is the primary business model of some of the largest global companies and is a high priority for many others.

In order to achieve this goal, it is necessary to gather large amounts of data. Furthermore, the utility of data is significantly enhanced by correlating multiple data sources. This results in severe privacy concerns since this places a great deal of sensitive information about individuals in the hands of corporations. Imagine if it were possible to utilize distinct data sources without ever bringing them together, with the guarantee that no raw data is revealed by the computation. This is the promise of secure multiparty computation, or MPC.

In this brief introduction we will describe what MPC is and what it can be used for. Among the many use cases of MPC, we will also show how it can be specifically used to solve the difficult problem of key management and protection in a virtual (software only) world.

What Is Multiparty Computation

Consider two people who wish to compare their DNA in order to determine if they are related. There are powerful algorithms for carrying out this computation today. However, a standard computation would involve the parties sharing their DNA with each other or with a third party; needless to say, revealing one's DNA can be very dangerous (e.g., it could contain genetic dispositions that insurance companies will then use to deny the person health insurance).

With the aid of MPC, these two people could compare their DNA without revealing it, and even if one of them behaved maliciously in an attempt to learn the other person's DNA.

Generally speaking, in the setting of secure multiparty computation (MPC), a set of distinct parties or devices wish to collaborate to carry out some computation on their private data. The challenge in this setting is that the parties do not trust each other – either because they fear that the other collaborating parties themselves are corrupted, or they fear that some of the other parties' networks and machines may have been breached.

As such, no party is willing to simply send its data to another party, because if that other party has been compromised then their private data will be lost. (It is one thing to trust that another party is not themselves corrupt, but a completely different thing to trust that no attacker can breach their network. For example, consider a situation of two parties: a major cloud service provider and its customer. The customer is unlikely to be concerned that the cloud provider itself will behave maliciously to steal information about the customer from their cloud. However, attackers breaching the cloud provider's services will have access to a wealth of information, and the silent subpoena issue is one that cannot be ignored.)

Loosely speaking, an MPC protocol should guarantee the following security properties, even if some of the participating parties are corrupted and may attempt to cheat:

- **Input privacy:** No corrupt party or subset of corrupted parties should be able to derive any information about the private data belonging to the other parties (except for what is revealed by the result itself). In the example of DNA comparison, privacy would mean that the only information revealed is whether or not the two people are related.
- **Correctness:** No corrupt party or subset of corrupted parties should be able to cause any honest party to output an incorrect result. In the example of DNA comparison, correctness would mean that a party behaving maliciously cannot make the result be that the input DNA sequences are of related people, if they are not.

In order to further illustrate this, we consider another example of an auction being run in MPC. Then, input privacy would guarantee that only the winning bid is revealed, and correctness would guarantee that the highest bid is indeed the one that is chosen as the result.

MPC has a long history of research in the cryptographic academic community, starting in the mid to late 1980s. There have been literally thousands of papers published on the topic, from its theoretical foundations and all the way to how to construct practical MPC protocols that can be used in practice. Thus, MPC is a new technology in the sense that it has only recently begun to be used in real-world settings, but MPC is actually based on deep science developed over decades of intensive research.

Key Protection Using MPC

One important application of secure multiparty computation is that it enables the use of cryptographic keys without ever having them in a single place, thereby eliminating the secret key as a single point of failure. In order to achieve this, it is possible to split the secret key into two or more pieces so that all pieces are needed in order to get any information about the key. Then, the different pieces are placed on different servers and devices, so that an attacker would have to breach them all in order to steal the key. By ensuring strong separation between these devices (e.g., different administrator credentials, environments, and so on), it is possible to achieve a very high level of key protection.

The above strategy of key splitting alone makes sense for preventing unauthorized access to key material at rest. However, if it is necessary to reconstruct the key on a single machine in order to use it, then an attacker could just wait for that to happen and then steal the key from the memory of that machine.

In order to prevent the aforementioned attack, MPC can be used to compute cryptographic operations with the keys, such as decryption or digital signing, while the keys are distributed among multiple devices. In particular, the privacy requirement of MPC guarantees that the pieces of the key are not revealed, even during the cryptographic computation that uses that key. Thus, even if an adversary has breached a subset of the machines, it cannot steal the key – even while it is in use.

This is a powerful paradigm, enabling organizations to transition away from legacy key protection methods based on secure hardware to solutions that better suit today's virtualized software-only environments. Specifically, since MPC is just a protocol run between any two or more machines in order to obtain the result, it is possible to run it on standard hardware and in any environment, while maintaining strong and clear security guarantees.

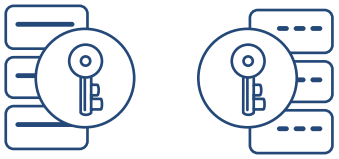

Each private key exists as two separate random shares stored on separate locations & refreshed constantly

Key shares are never combined at any point in time - not even when used or when created

Key material never exists in the clear at any point of its lifecycle

This method can be used to efficiently protect all standard cryptographic keys and algorithms today, including RSA, Elliptic curve cryptography, and symmetric cryptography, for encryption/decryption, authentication and signing.

Encrypted Data in Use

MPC can also be used to obtain a new paradigm of security: encryption of data while in use. We are all familiar with the two basic paradigms of encryption: data-at-rest and data-in-motion. However, encryption of data-in-use seems like an oxymoron: if data is encrypted, how is it possible to use it?

With secure multiparty computation this can be achieved. Consider the case that a user encrypts her data with a key that she holds, and then uploads the encrypted data to one cloud provider and the private decryption key to another cloud provider. Neither cloud provider can learn anything about the data, solving one of the big challenges of uploading highly confidential data to the cloud. Since the data is encrypted, it cannot be processed as is. However, using MPC, it is possible for the two cloud providers to run a joint protocol to process the data (e.g., search for keywords) without ever decrypting it.

This is a very powerful paradigm, and one that reduces the risk of exporting data to the cloud. We stress, however, that this type of MPC computation can be computationally intensive, and so the specific application needs to be examined.

Summary

Secure multiparty computation enables parties to compute over data without ever revealing it. Though it can sound like magic, MPC is based on deep science and is the result of decades of research in academia.

Today, MPC is a mature technology that can be used to carry out private computations in practice, and is being used to solve important security problems like key protection. Importantly, all MPC protocols are mathematically proven secure and therefore provide very clear guarantees that data that should remain private (if there is no such proof, then the protocol should not be considered valid).

For the specific application of key protection, such a proof guarantees that an attacker must breach all devices holding key shares in order to learn anything about the key material. The use of MPC-backed technology to protect information provides a high level of security in an architecture aligned with today's modern digital computing environments.