Homomorphic Encryption:
Enabling Secure Computations on Encrypted Data

Abstract

For storing data and outsourcing its computations, end users (enterprises and individuals) are increasingly getting cloud service providers (CSPs) on board, instead of handling the data themselves. This is because CSPs provide services at an affordable cost and with much ease of use. One major challenge for users of cloud services, however, is preservation of the privacy and confidentiality of the data stored on the cloud, which is why the data is encrypted before upload.

However, the CSP has to decrypt the data to be able to perform computations on it. This reveals the data to the CSP, causing significant privacy concerns for end users. This is especially relevant given the current background of privacy concerns about third-party data and regulations such as GDPR that are meant to ensure data privacy.

While cryptographic computation mechanisms such as encrypted databases, order-revealing encryptions and order-preserving encryptions are helpful in preserving privacy, these are implemented at the price of some data leakage. To avoid such leakages, one can preserve the privacy of data stored on the cloud with the help of homomorphic encryption (HE). Fully homomorphic encryption (FHE), is a type of HE that enables arbitrary computations on encrypted data without decrypting it. This paper talks about the basic notion of HE, its applications and limitations.

The need for HE

Recent data trends suggest that there is an exponential increase in the growth rate of data creation.[1] Often, this data is shared with multiple parties, such as a CSP or a third-party organization, for the purpose of storing and processing. Notably, CSPs are the de facto destination for most enterprises and individual users. Cisco's report on the global cloud index (GCI) suggests that 94% of workloads and compute instances are processed by cloud data centers, and of that, 70% are managed by public cloud.[2]

Alarmingly, users do not have control over their data and are naturally concerned about data privacy. Furthermore, data is often exposed to breaches, where sensitive customer information is accessed in an unauthorized manner.[3] Customers often risk privacy of their data in exchange for services from CSPs. And while users can encrypt data and store it on the cloud for the purpose of confidentiality, this limits any kind of data processing. Therefore, the usual encryption is limited to data storage alone and does not allow for any meaningful computation.

To enable computations while guaranteeing data privacy, researchers are focusing on privacy-enabled computations. One of the promising approaches in this direction is the use of a cryptographic paradigm called homomorphic encryption (HE).

What does HE allow?

BHE allows computations on encrypted data without decryption. It was first proposed as a multiplicative homomorphic encryption scheme based on RSA[4] by Rivest, Adleman and Dertouzos in 1978. A HE scheme can be of two types:

1. Partial homomorphic encryption (supports either multiplication/addition; not both)
2. Fully homomorphic encryption (supports both – multiplication and addition)

Partial homomorphic schemes such as RSA and Paillier crypto systems that support multiplicative and additive homomorphism, respectively, have been around for long. However, it was only in 2009 that Craig Gentry proposed an

FHE scheme based on lattices for the first time.[5] An FHE scheme usually supports addition and multiplication of ciphertexts as follows:

E(a)+E(b)=E(a+b) and E(a)*E(b)=E(a*b)

That is, addition/multiplication of two ciphertexts gives encryption of the sum/product of the underlying plaintexts. These primitive operations are used to realize any arbitrary computation, which provides an opportunity to realize privacy-enabled computations in any application.

Possible Applications

HE is often considered the holy grail of cryptography, as it enables privacy-preserving computations in almost all possible scenarios (see Figure 1).

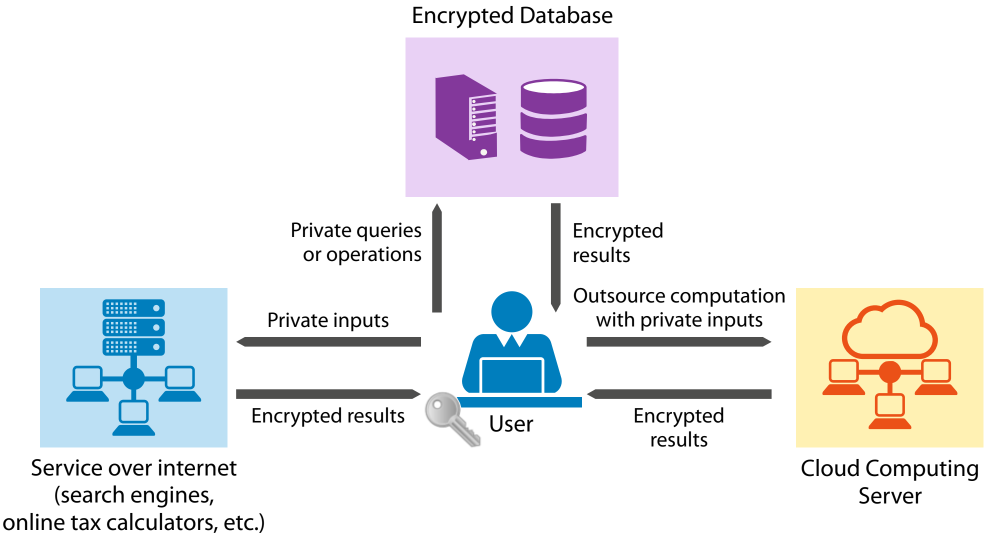
Encrypted Database
Private queries or operations
Encrypted results
Outsource computation with private inputs
Private inputs
Encrypted results
User
Encrypted results
Service over internet (search engines, online tax calculators, etc.)
Cloud Computing Server

Figure 1: Outsourcing Computation on Encrypted Data

- **Private Search:** The internet offers users an effective way to crawl through its universe – via search engines. These search engines index all available data in the digital universe and index them appropriately to serve customers with relevant responses to their search queries. However, the risk of users' data exposure to search giants has been a major concern for many years. Service providers rely on search queries to track and target users with advertisements. This is a significant privacy risk that can't be solved effectively with conventional mechanisms.

With HE to the rescue, it is possible to prevent service providers from learning queries while still allowing them to serve users with relevant results. The processed queries return encrypted results to the users, without revealing any information to the search engines.

- **Encrypted Databases:** Almost all information systems rely on databases for storing information. These databases form an important part of the IT infrastructure and solicit sound security guarantees in order to ensure security of data. It is often the case that in case of any security breach, attackers target the databases of applications. Database security is a widely acknowledged problem and hence several security mechanisms are employed. One such mechanism is the encrypted database, where the data is encrypted with the secret key of the client to prevent unauthorized disclosure.

  However, this solution makes the database unreadable and does not support search/range queries on encrypted records. To support such queries on encrypted databases, solutions such as deterministic encryption[6], order-preserving encryption[7] and order-revealing encryption[8] have been proposed. But these solutions are implemented at the cost of a certain degree of leakage[9], such as memory access patterns or search patterns.

  With HE, it is possible to encrypt data on the database for confidentiality, while still allowing processing on the data with minimal leakage. Only authorized users with the relevant secret key can access the data from database. Although HE is highly inefficient today, when made practical, it will solve the problem of data leakage in encrypted databases.

- **Outsourcing Computation on the Cloud:** The popularity of cloud computing is based on the fact that individual enterprises and users do not have to maintain the infrastructure for running their services: customers can lease cloud infrastructure on a need basis to run their applications. Also, the ability of CSPs to scale their services based on application traffic has made cloud computing a go-to option. However, since the whole cloud infrastructure is managed by the service provider, the current trust model relies on the CSP to be honest.

Enabling computation on encrypted data eliminates such trust requirements in this setting. The cloud can host applications that deal with encrypted data without the need for access to the plaintexts. This guarantees privacy to users as well as enterprises.

- **Other Applications:** In the recent past, new paradigms such as machine learning (ML) and artificial intelligence (AI) have been of interest to several enterprises looking to improve their services. These paradigms depend on building statistical models based on available data, and the models are then queried for intelligence on new requests or data. Enterprises use data collected from users to train these AI/ML models, which is a privacy risk for the users.

  To enable privacy-preserving computation, HE can be used to train the models with encrypted data. Over the past few years, there has been a growing interest in the research community to use FHE for private AI or machine learning. Furthermore, healthcare services in recent times are providing predictions based on human genomic data. HE can help preserve the privacy of such sensitive genomic data.

Drawbacks and Limitations

In general, HE computations are very slow and only a finite number of operations can be performed on encrypted data. Our research has shown that FHE-based computation is around 106 times slower than normal computation on plain data. Efforts are in place by the research community to speed up FHE-enabled computation.

Currently, it is not practically feasible to use FHE for any generic computation on encrypted data. However, it is manageable for a class of computations with fewer operations. For example, computing statistical functions such as mean, variance and standard deviations on encrypted data are feasible. Also, HE in its current stage can be applied to domain-specific limited private computations. Domain-specific knowledge provides scope for improving performance of applications that deal with encrypted data.

Open-source Implementations

Some of the prominent open-source implementations of HE schemes include HEAAN, Palisade, Homomorphic Encryption Library and Simple Encrypted Arithmetic Library by Microsoft.

What's Next

In its current state, FHE is computationally expensive and practically inefficient. But there has been consistent progress in designing efficient HE algorithms and researchers are working towards making them practical. Considering the importance of privacy in today's world and the role of HE in enabling privacy-preserving computations, this is set to be a space to watch for CSPs as well as end users.

References

[1] https://blog.westerndigital.com/2018-data-trends-today-tomorrow-change/

[2] https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html

[3] https://www.identityforce.com/blog/recent-data-breaches-december-2018

[4] https://pdfs.semanticscholar.org/3c87/22737ef9f37b7a1da6ab81b54224a3c64f72.pdf

[5] https://crypto.stanford.edu/craig/easy-fhe.pdf

[6] https://eprint.iacr.org/2006/108.pdf

[7] https://eprint.iacr.org/2012/624.pdf

[8] https://eprint.iacr.org/2016/612.pdf

[9] https://blog.cryptographyengineering.com/2019/02/11/attack-of-the-week-searchable-encryption-and-the-ever-expanding-leakage-function/

About The Authors

Nitesh Emmadi

Nitesh Emmadi is a researcher in the cybersecurity and privacy research area at TCS Research and Innovation (R&I). His areas of research include computations on encrypted data with a broader interest in application security, applied cryptography and blockchains. He looks closely into the practical side of novel systems and provides consulting services for evaluating and building products. Nitesh received his master's degree in information technology, with a specialization in information security, from IIIT, Hyderabad, India.

Harika Narumanchi

Harika Narumanchi is a researcher in the cybersecurity and privacy research area at TCS R&I. Her research broadly focuses on applying cryptography and blockchain solutions to business-oriented scenarios. She graduated from the Jawaharlal Nehru Technological University, Hyderabad, India, with a master's degree in information technology, with a specialization in information security.

Rajan MA

Rajan MA is a scientist at TCS R&I in the cybersecurity and privacy research area. He holds a PhD, MTech and BE in computer science, as well as a PhD, MPhil and MSc in mathematics. His research interests include applied cryptography, blockchain and graph theory. Prior to joining TCS, he worked at the ISRO Satellite Centre, Bangalore, India, as a senior scientist. He is also a senior member of IEEE.

Experience certainty.

IT Services
Business Solutions
Consulting

Contact

Visit the Research and Innovation page on www.tcs.com

Email: innovation.info@tcs.com

Blog: #Research and Innovation

Subscribe to TCS White Papers

TCS.com RSS: http://www.tcs.com/rss_feeds/Pages/feed.aspx?f=w
Feedburner: http://feeds2.feedburner.com/tcswhitepapers

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services is an IT services, consulting and business solutions organization that delivers real results to global business, ensuring a level of certainty no other firm can match. TCS offers a consulting-led, integrated portfolio of IT and IT-enabled, infrastructure, engineering and assurance services. This is delivered through its unique Global Network Delivery Model^{TM}, recognized as the benchmark of excellence in software development. A part of the Tata Group, India's largest industrial conglomerate, TCS has a global footprint and is listed on the National Stock Exchange and Bombay Stock Exchange in India.

For more information, visit us at www.tcs.com

TCS Design Services I M I 08 I 19