UNIVERSITY OF CALGARY

PRIVACY PRESERVING CRYPTOGRAPHY

Teaching Aide

Group Members:

Pradip Ghimire
Azeezat Lawal
Dami Ogunnupebi
Tarun Sidhu
Abdul Osuwa

What is Privacy Preserving Cryptography?

Privacy Preserving Cryptography, or PPC, refers to the cryptographic methods that allows data to be shared and processed without revealing sensitive information to other unauthorized parties. The general concept of PPC lies in being able to use data or information without full, unauthorized access to the original data.

In order to achieve this, PPC employs a number of techniques, that will be explored further, that share its goal, but reach this goal in different ways.

Image Source

PPC Techniques

- Zero-Knowledge Proofs (ZKP)
- Homomorphic Encryption
- Secure Multi-Party Computation (MPC)
- Differential Privacy
- Anonymous Credentials and Commitment Schemes

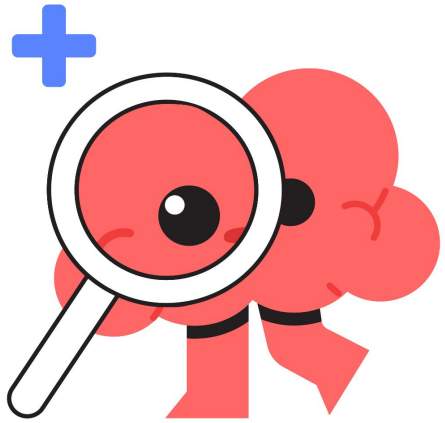

Image Source

Zero-Knowledge Proofs (ZKP)

UNIVERSITY OF CALGARY

- Cryptographic method used to prove knowledge about a piece of data, without revealing the data itself.
- First appeared in the 1985 paper "**The knowledge complexity of interactive proof systems [GMR85]**"
- Zero-knowledge proofs must satisfy three properties:
  - **Completeness:** if the statement is true, an honest verifier will be convinced by an honest prover.
  - **Soundness:** if the statement is false, no dishonest prover can convince the honest verifier. The proof systems are truthful and do not allow cheating.
  - **Zero-Knowledge:** if the statement is true, no verifier learns anything other than the fact that the statement is true

Zero-Knowledge Proofs (ZKP)

Example: The Magic Door Puzzle (Classic Example)

- Alice knows the secret word that opens a magic door inside a cave. Bob wants to be sure she knows it, but Alice doesn't want to tell him the word.
  - The cave has two tunnels, **A** and **B**, connected by the magic door.
  - Alice walks into the cave while Bob waits outside. Bob doesn't know which tunnel she took.
  - Bob then asks her to come out through **A** or **B**.
  - If Alice knows the word, she can open the door and come out wherever Bob asks.
  - If she doesn't, she can only guess—and will eventually get caught.

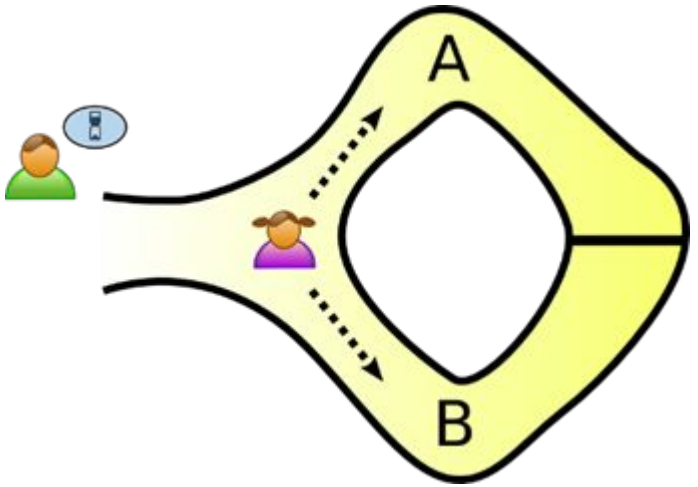

Figure: Magic Gate Example (Source)

XKCDon't

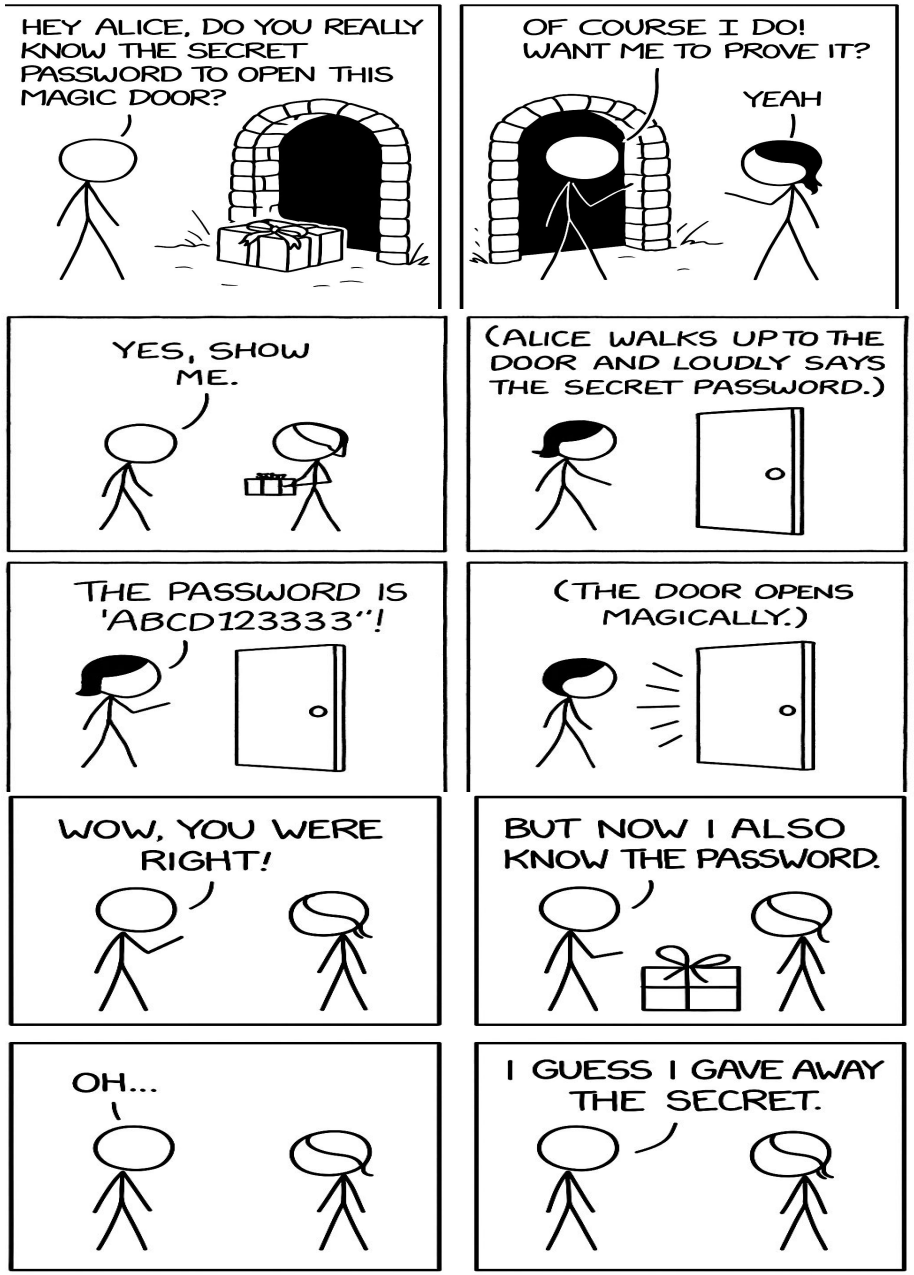

UNIVERSITY OF CALGARY

In this comic, Alice wants to prove she knows the secret password to open a magic door. To prove it, she says the password aloud and opens the door. While this shows she knows the secret, it also reveals it to Bob. Therefore, it's not a zero-knowledge proof because the secret itself is exposed during the demonstration.

All comic images generated by AI

XKCDo

This comic demonstrates zero-knowledge proofs with a cave. The prover claims to have a password that unlocks a magical door. The verifier randomly asks if they want to exit left or right. If the prover succeeds repeatedly, they are likely to know the password but not reveal it. It's proof through consistency, not disclosure, humorously compared to modern dating.

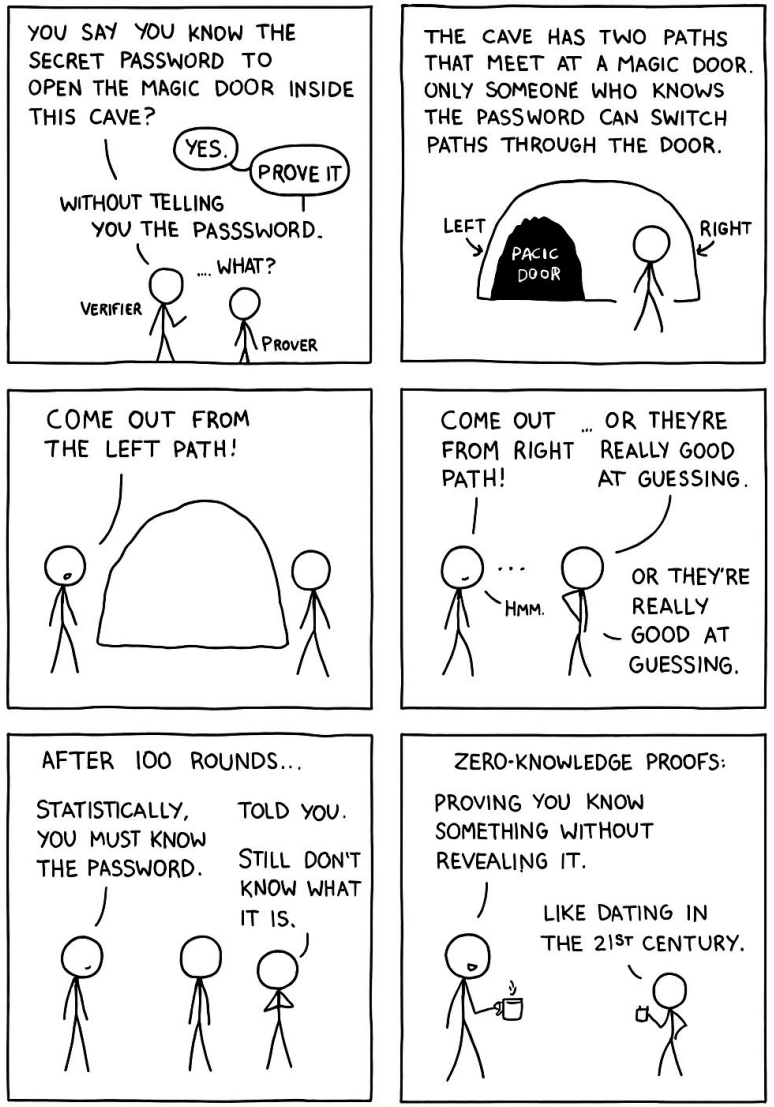

All comic images generated by AI

Homomorphic Encryption

Homomorphic encryption is a cryptographic method that allows computations, calculations and analytics to be performed on encrypted data without having to decrypt it first.

This means that data can then be shared with third parties in various industries who need to run the data through various algorithms, analyze or otherwise manipulate the data, without sharing the contents of the data.

For example, doing addition or multiplication on values inside a locked box and when the box is opened, the result is still correct.

Traditional Encryption vs. Homomorphic Encryption

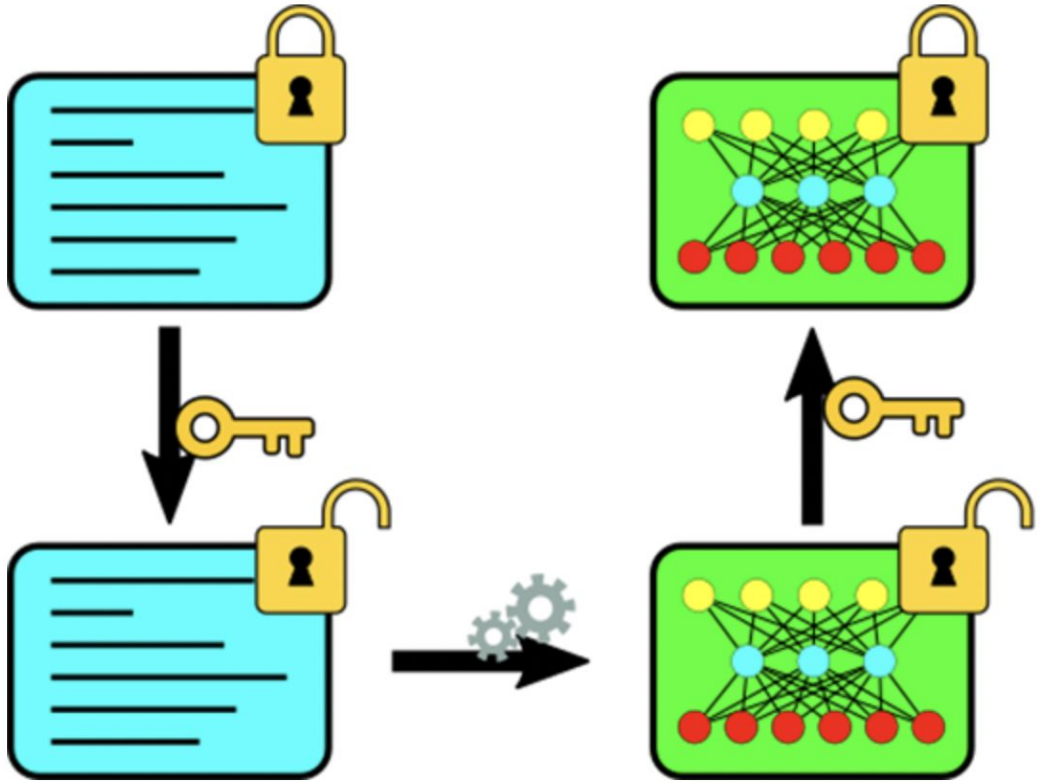

This is the traditional encryption method, where to apply analytics, the data must first be decrypted, then the analytics are applied, then re-encrypted.

This image shows homomorphic encryption in action - the data does not need to be decrypted before any calculations can be performed. The data remains protected.

Source

XKCDon't

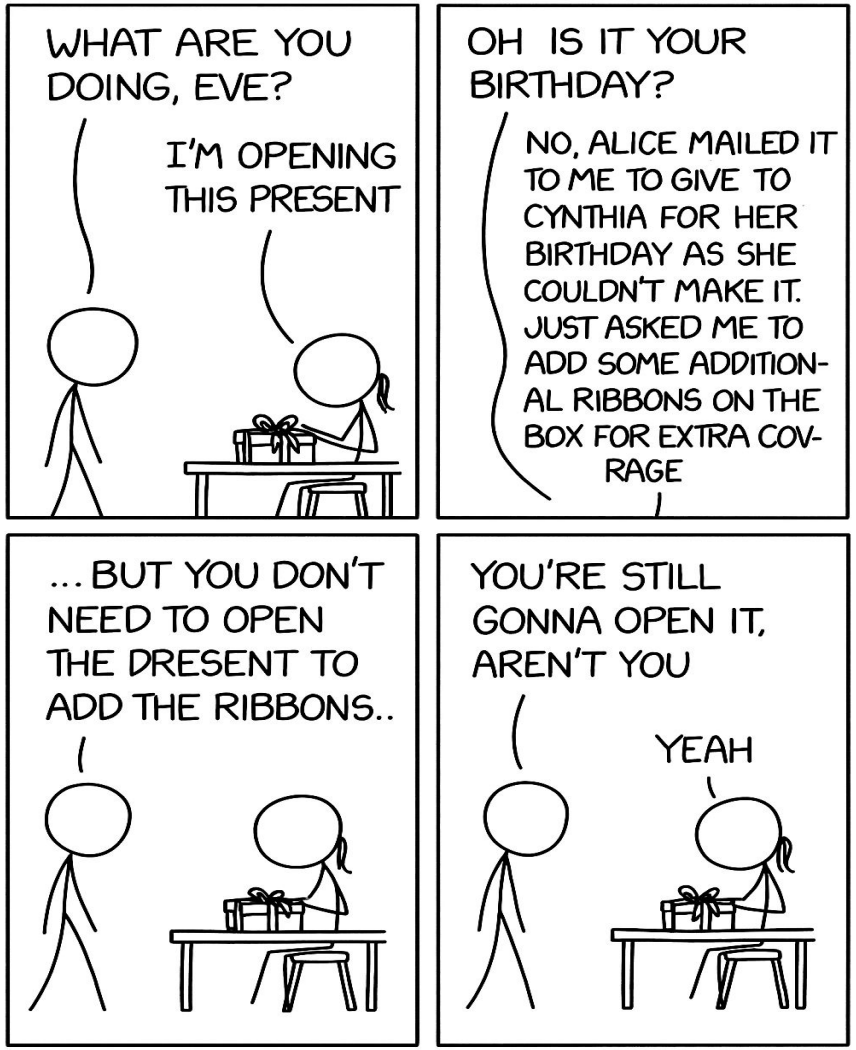

All comic images generated by AI

Here, we have Eve opening a present, even though she doesn't need to know what's in it to add the ribbons, the additional computations in this case.

Eve is very nosy.

XKCDo

But here, even though Eve is still nosy, she cannot open open the present. Homomorphic encryption has been applied, and the present is essentially encrypted and cannot be opened.

The lesson here is to always encrypt your gift before handing them to potentially nosy people like Eve.

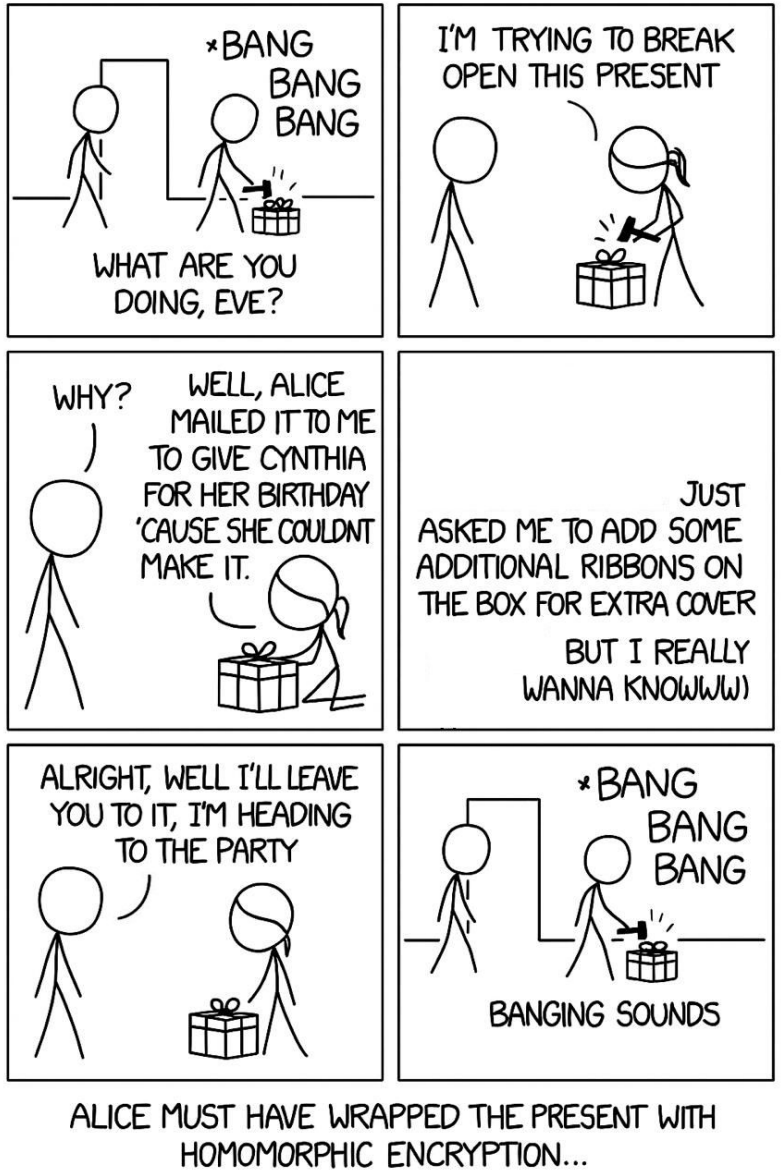

All comic images generated by AI

Secure Multi-Party Computation (MPC)

Sometimes called Secure Multi-Party Computation (SMPC), MPC is a cryptographic and mathematical concept that ensures that everyone in a function is able to keep their raw data unknown to other parties of the group.

The computation that makes up this technique mandates that parties first protect their raw data through a method like encryption before then submitting them as input for the MPC function to be combined with that of other parties. While the output is known and glaring for the group, that value cannot be used to trace back the original unencrypted input.

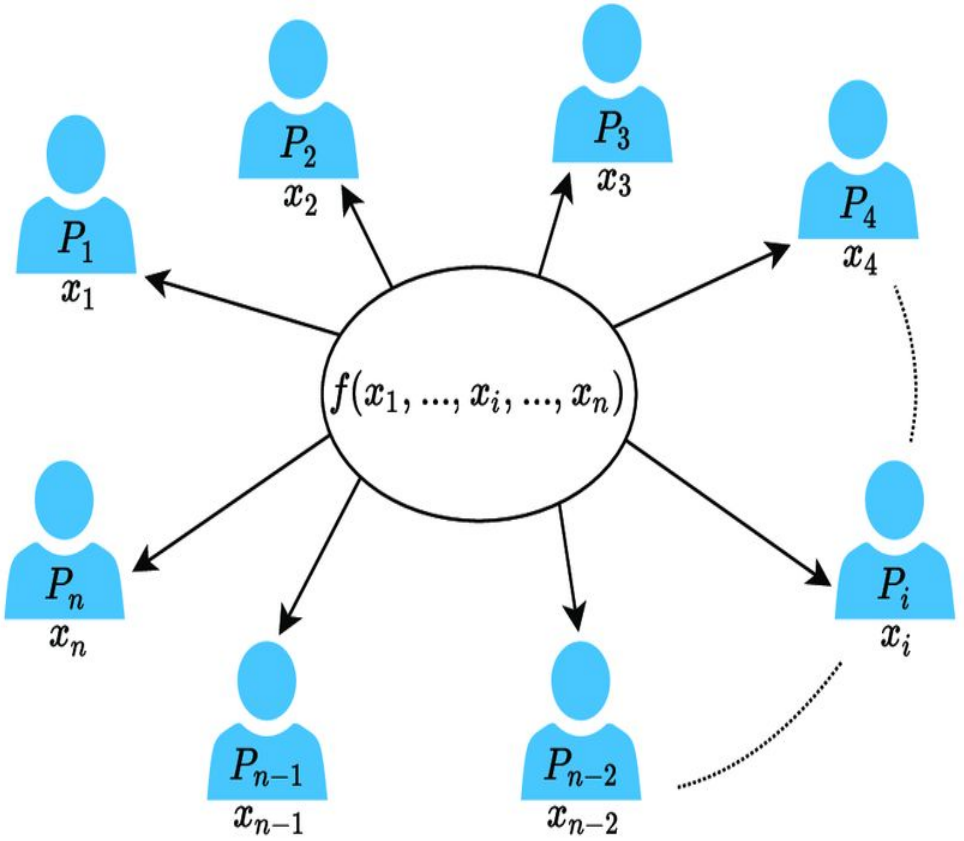

Image Source

Secure Multi-Party Computation (MPC)

Imagine that three coworkers Alice, Bob, and Cynthia, want to know their average hourly wage but don't want to share their own salary with each other.

First, they break their wage into four amounts that add up to their hourly earnings. Next, they keep one of those figures, and share one each with the other coworkers along with a trusted third party. Now, each party calculates the average of the numbers they received. Finally, these averages are then shared and summed to provide the average hourly wage. While they all know the average, they don't know each other's individual earnings.

While this example uses a relatively simple additive secret-sharing technique, you can explore a more **advanced mathematical example here**.

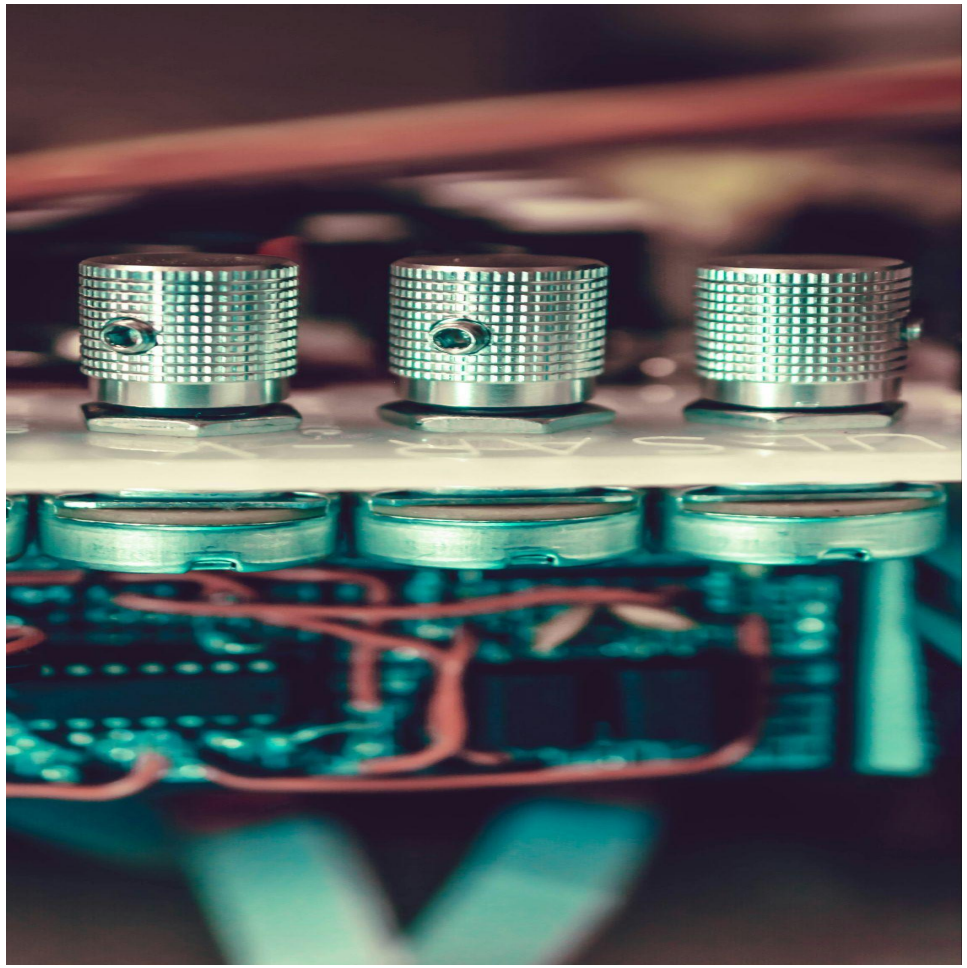

Image Source

XKCDon't

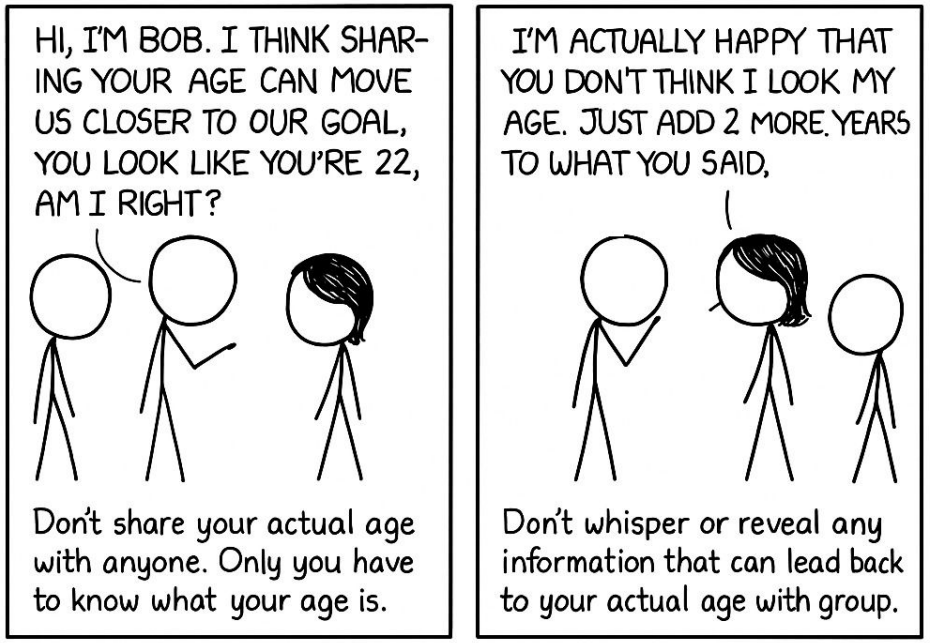

All comic images generated by AI

Here, Bob unknowingly uses a social engineering technique to try and guess Shannon's age, knowing fully well that neither he or the rest of the group can have any direct details about the group's ages.

Shannon should keep information about her age secure regardless of how cute Bob finds her.

XKCDo

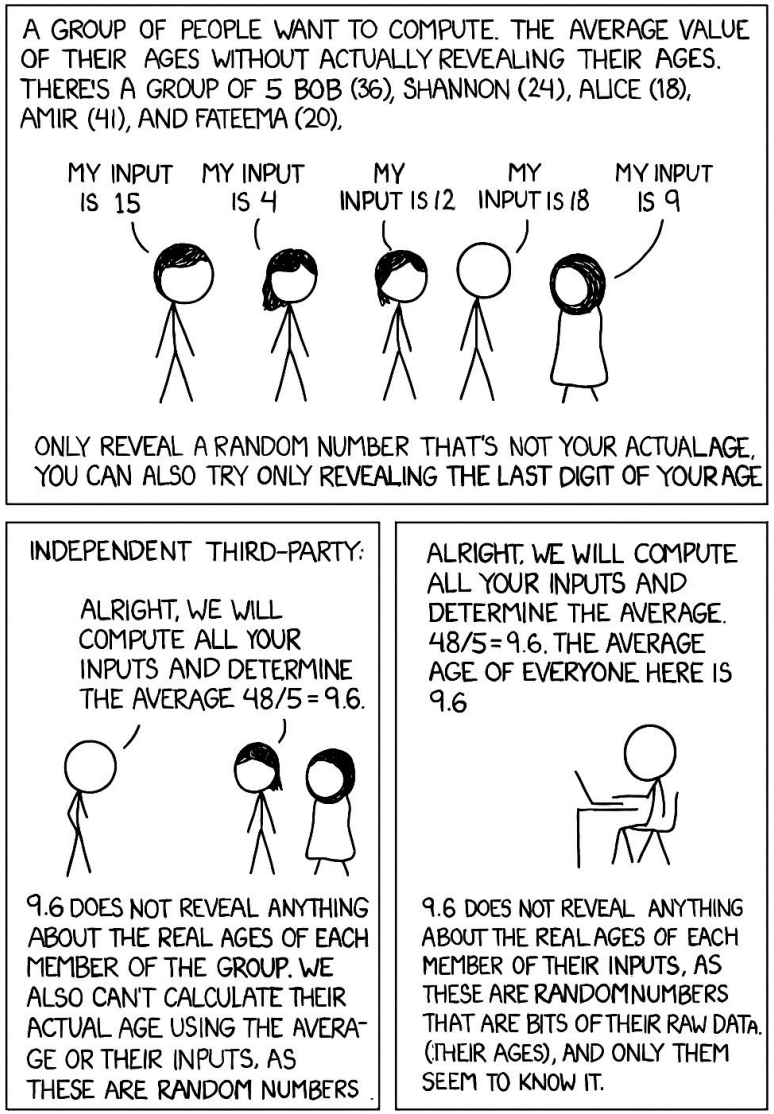

All comic images generated by AI

Here, we see that the evaluator was able to calculate the group's age average with only bits of the information provided by the group. We also see that the output (9.6) does not in any way reveal anyone's actual age. This is a simple illustration of what MPC seeks to achieve.

Place priority on the security of the data rather than the goal. Be like Apple. Protect until the end!

Differential Privacy

- Differential Privacy (DP) establishes a mathematical framework which enables protection of personal information when data undergoes analysis or sharing. It ensures that the output of a data analysis or algorithm does not reveal sensitive information about any single individual, even if an attacker has additional background knowledge.
- Provides measurable privacy guarantees using noise addition.
- Recognized by NIST (2025) as the most effective privacy-preserving framework.

How Differential Privacy Works

A database containing individual medical information exists as an example: Differential Privacy ensures that the results of any analysis will stay identical regardless of whether your data is analyzed or not.

The data remains anonymous so no one can determine if your information exists within the dataset.

The implementation of differential privacy requires researchers to add particular random noise measurements to the data:

1. The data itself (local differential privacy), or
2. The outputs of queries (global differential privacy).

Example of DP

- The system introduces a small random “fuzziness” to the exact number when an analyst requests the average age of people in the dataset.
- The noise level becomes so high that individual data points disappear from view yet the system can still extract meaningful information from the complete dataset.

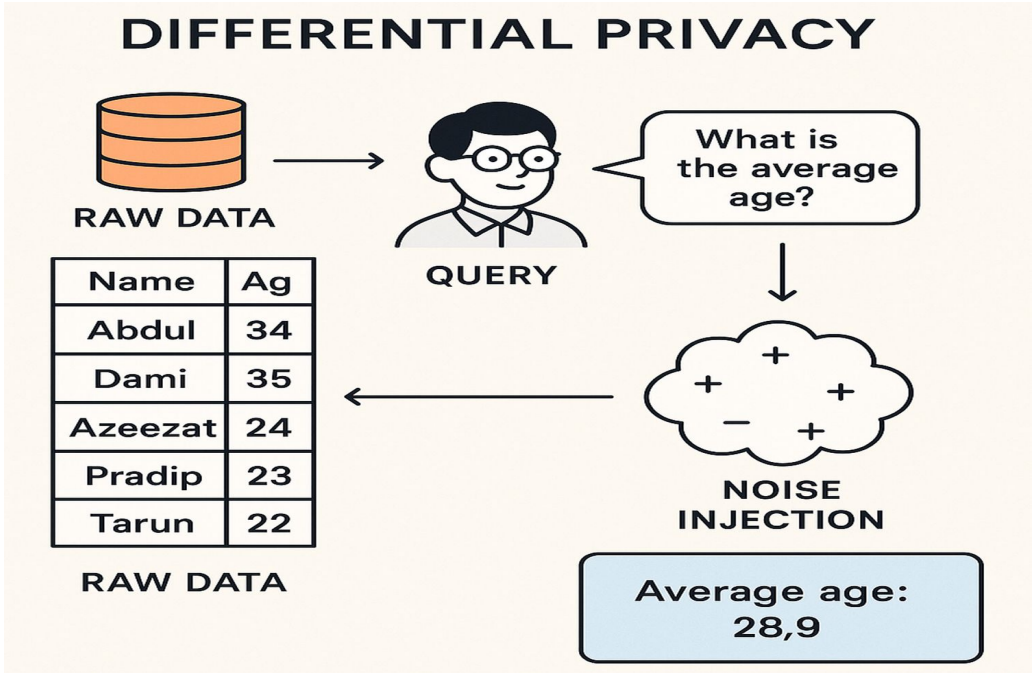

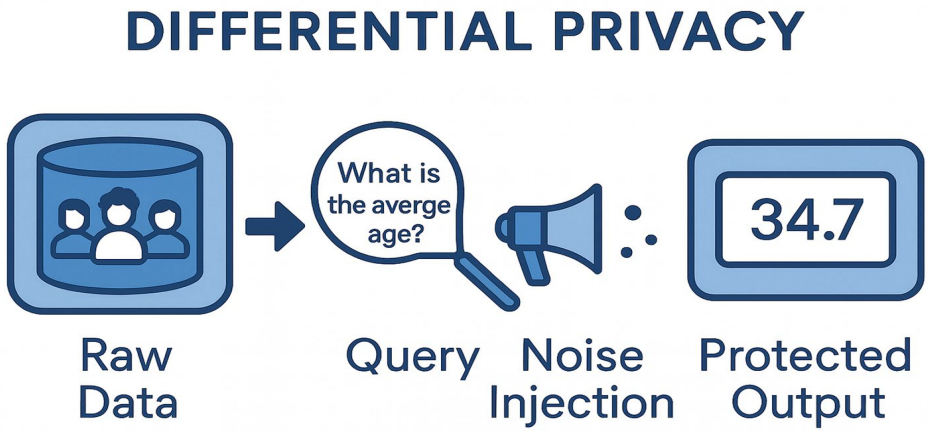

Images generated by AI

Applications of Differential Privacy

UNIVERSITY OF CALGARY

- **Census & Government Stats** : U.S. Census 2020, economic & population data.
- **Tech Companies :** Apple (iOS/macOS analytics), Google (Chrome, RAPPOR).
- **Healthcare:** Sharing medical data for research safely.
- **Location Services:** Traffic, ride-sharing, city planning with GPS data.
- **AI & Machine Learning :** Training models (DP-SGD) on sensitive data.
- **Finance :** Analyzing transactions without exposing individuals.
- **Social Media:** Studying user behavior while protecting identities.

XKCDon't

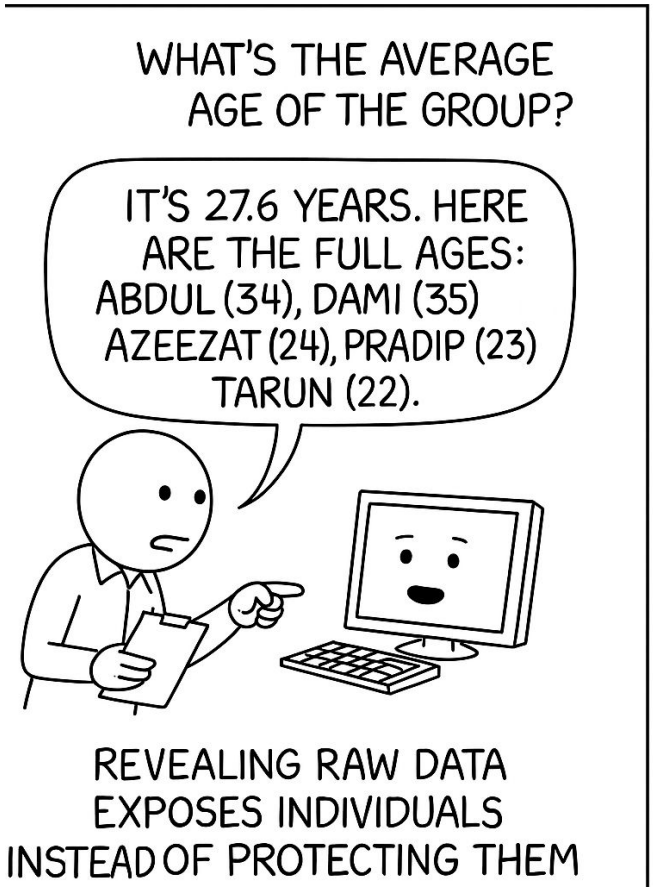

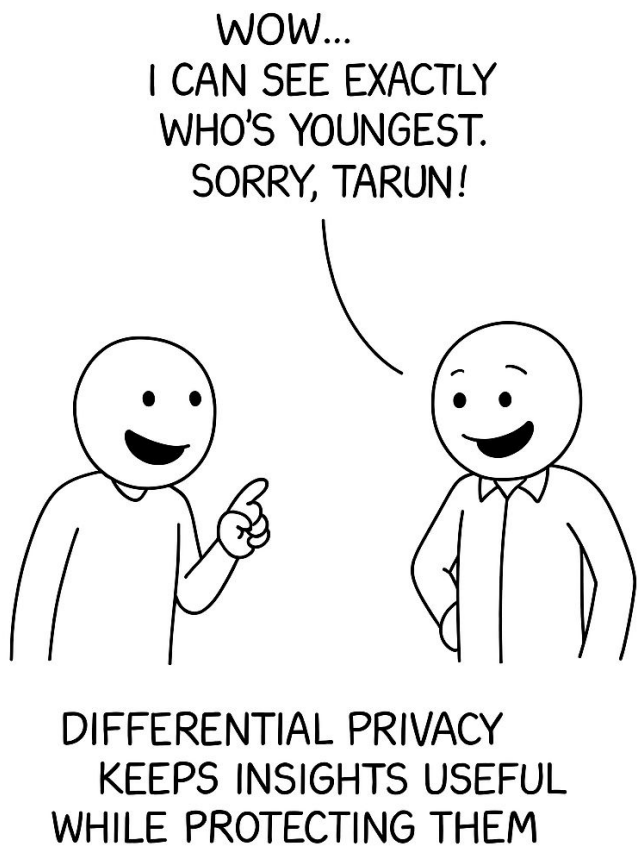

All comic images generated by AI

In this version, Analyst Ahmed asks for the group's average age. The system provides both the total average age and shows the precise age of each individual. The system enables simple identification of people through the example of Tarun who is the youngest member. The comic demonstrates that revealing unprocessed data results in privacy issues because it reveals sensitive information which needs to remain protected.

XKCDo

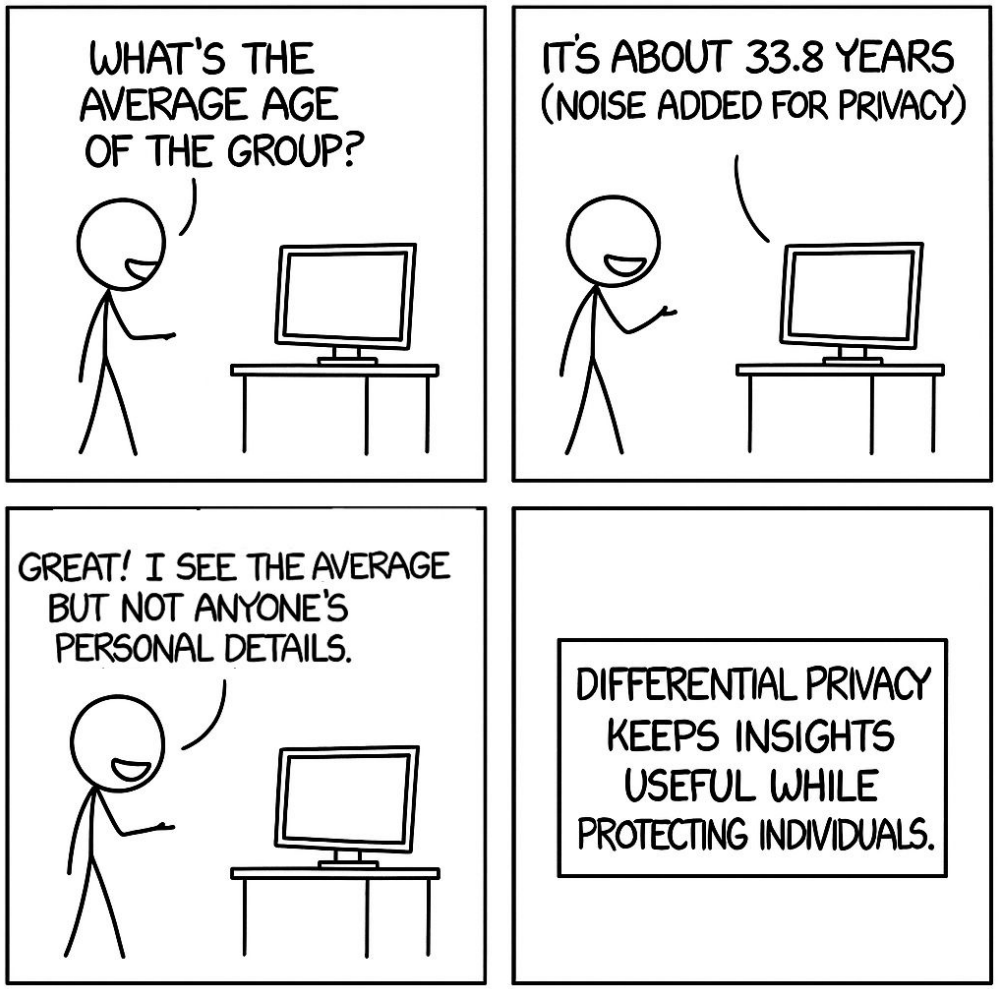

All comic images generated by AI

The system responds to Ahmed's inquiry through differential privacy. Instead of showing individual ages, it only provides an approximate average (with noise added). Ahmed still learns the useful trend but without anyone's personal information being revealed. The method shows how differential privacy achieves both data protection and information access.

Anonymous Credentials and Commitment Schemes

UNIVERSITY OF CALGARY

Anonymous credentials & commitment schemes are powerful cryptographic tools that allows users to prove attributes about themselves without giving away any extra or unnecessary information. Users can verify certain rights while keeping their identity a secret and hence remain anonymous. A major building block to this is commitment schemes, they work like a digital promise. The user locks in a value such that it cannot be changed, but it remains a secret until revealed. The user basically commits to a value. Together, these tools have helped us design systems that preserve our privacy. They are the backbone of today's modern applications where trust and privacy coexist.

Image Source

Anonymous Credentials and Commitment Schemes

The second attribute to focus on is how they enable privacy and security to exist together in a way that makes sense and is pragmatic. It also enforces accountability and correctness, so that while the users remain anonymous, the system can still guarantee that their actions are accurate. Selective disclosure plays a major role here, only information that is required is shared, while all other information remains a secret.

This balance grants users control over their data, making it possible for them to engage with digital systems securely without constantly oversharing. Since concerns regarding privacy are increasing, anonymous credentials are becoming more and more important to create systems that users can trust for privacy.

Image Source

XKCDon't

In this comic, Alice tries to join a gaming tournament. The verifier asks for proof of age that she's over 18. Instead of showing only that, Alice dumps way too much information, such as her full ID, social security number, even a list of her favorite snacks. The verifier is caught off guard, since none of the extra details were actually needed. It might be funny but focuses on a very important point - we often overshare, way more than necessary. The message is simple, only share what's absolutely necessary, because giving away more can be awkward at best and unsafe at worst.

All comic images generated by AI

XKCDo

In this comic, it's the same scenario as the last one, Alice once again wants to enter the gaming tournament, but this time she handles it differently. Instead of dumping all her personal details, she just proves that she's over 18 while keeping her identity a secret. The verifier is verifies her as it was exactly what was needed. The takeaway is simple: You can prove what's needed without giving away your whole identity.

All comic images generated by AI

Why is Privacy Preserving Cryptography Important?

UNIVERSITY OF CALGARY

Secures Personal Data: Safeguards private, economic or health details.

Ensure Trust: Encourages faith in the digital platform by guaranteeing privacy rights of the products.

Prevents Misuse: Protects against the threats of profiling, watching and data selling.

Enables Compliance: Complies with current policy industry regulations (GDPR, HIPAA, etc.).

Future-Proof Security: Implies that modern systems can carry out new functions and provide answers to challenging problems while retaining anonymity and confidentiality.

Challenges of Privacy Preserving Cryptography

Sometimes, even the best mathematical computations can be found to have some flaws, either in their design or application. Below are a few challenges facing the adoption of PPC as a whole:

Performance and Scalability: computational overhead makes it quite challenging to scale effectively especially for situations where the use of large datasets are needed.

Implementation Issues: though algorithms thrive better with constant updates and fixes, a bug-infested implementation does expose the software to vulnerabilities.

Key Management Risk: In most cases, keys may just be as sensitive as the information they're used to secure and protect, so poor key management practices make the entire system vulnerable to security breaches.

Challenges of Privacy Preserving Cryptography (Cont'd)

UNIVERSITY OF CALGARY

Multi-Company Data Exchange: The complexity required to process varying data between companies can complicate the design and deployment of privacy-preserving solutions.

Decryption in Servers: In many applications, data must be decrypted for server-side processing, creating a window of vulnerability where the data is exposed to threats on the third-party server, according to ScienceDirect.

Quantum Computing: Though a quality of PPC is that it's quantum safe, the advancement of powerful quantum computers could still threaten to break even the most powerful cryptography algorithms.

System Integration Complexity: Assimilating advanced privacy-preserving cryptographic tools into large-scale, real-world operations poses a complex technical problem.

Conclusion

Privacy transcends companies, it's woven into data and information, and from storage, to access, to usage, and to the general application of data/information, people are at its heart. Privacy continues to be a tricky subject to truly manage, and Privacy Preserving Cryptography is only one of the steps to having privacy truly **be** private. With its many faces (techniques), as well as its privacy preserving makeup, PPC is another proof that cryptographic applications and its evolutions are great choices for cybersecurity and privacy as a whole. There's a lot of hope yet for individuals and companies moving closer to secure systems and practices that foster efficient privacy and security.

Image Source

Bibliography

Sheybani, Nojan, et al. “Zero-knowledge Proof Framework: A Systematic Survey”, 27 Apr. 2015. arxiv.org/pdf/2502.07063

Jayodya Methmal, “Zero Knowledge Proofs: A Comprehensive Review of Applications, Protocols, and Future Directions in Cybersecurity”, August 2023. DOI:10.13140/RG.2.2.11606.22080

A. Pathak, T. Patil, et al. "Secure Authentication using Zero Knowledge Proof," 2021 Asian Conference on Innovation in Technology (ASIANCON), PUNE, India, 2021, pp. 1-8. DOI: 10.1109/ASIANCON51346.2021.9544807

Lu Zhou, Abebe Diro, et al. “Leveraging zero knowledge proofs for blockchain-based identity sharing: A survey of advancements, challenges and opportunities”, Journal of Information Security and Applications, Volume 80, 2024, 103678, ISSN 2214-2126. DOI: 10.1016/j.jisa.2023.103678

Zanussi, Zachary. “Privacy Preserving Technologies Part Two: Introduction to Homomorphic Encryption”. Statistics Canada, 01 March 2022. www.statcan.gc.ca/en/data-science/network/homomorphic-encryption

“What is homomorphic encryption?” IBM. www.ibm.com/think/topics/homomorphic-encryption. Accessed 16 September 2025.

Bibliography

“Homomorphic Encryption”. Chainlink, 6 June 2025. www.chain.link/education-hub/homomorphic-encryption

Paine, Kirsty. “Homomorphic Encryption: How It Works”. Splunk, 5 February 2024. https://www.splunk.com/en_us/blog/learn/homomorphic-encryption.html

Bryanton, Betty Ann. “Introduction to Privacy Enhancing Cryptographic Techniques: Secure Multiparty Computation”. Statistics Canada, 15 March 2024. www.statcan.gc.ca/en/data-science/network/multiparty-computation

“Secure Multi-Party Computation”. Chainlink, 14 August 2024.

www.chain.link/education-hub/secure-multiparty-computation-mcp

Lindel, Yehudal. “Secure Multiparty Computation (MPC)”. Unbound Tech and Bar-Ilan University. eprint.iacr.org/2020/300.pdf. Accessed 16 September 2025.

Evans, David, et al. "A Pragmatic Introduction to Secure Multi-Party Computation". now Publishers Inc, 19 December 2018. www.nowpublishers.com/article/Details/SEC-019

Maurer, Ueli. “Secure Multi-Party Computation Made Simple”. ScienceDirect, 3 October 2005. www.sciencedirect.com/science/article/pii/S0166218X05002428

Bibliography

Volgushev, Nikolaj, et al. “Conclave: Secure Multi-Party Computation on Big Data”. ACM Digital Library, 25 March 2019. www.dl.acm.org/doi/abs/10.1145/3302424.3303982

Alborch Escobar, Ferran, et al. “Computational Differential Privacy for Encrypted Databases Supporting Linear Queries.” Proceedings on Privacy Enhancing Technologies, 2024. petsymposium.org/popets/2024/popets-2024-0131.pdf

Dwork, Cynthia. “Differential Privacy: A Survey of Results.” 2008. www.cs.ucdavis.edu/~franklin/ecs289/2010/dwork_2008.pdf

Mironov, Ilya. “On Significance of the Least Significant Bits for Differential Privacy.” University of Waterloo, crysp.uwaterloo.ca/courses/pet/F18/cache/Mironov.pdf

Chaum, David. “Security without Identification: Transaction Systems to Make Big Brother Obsolete“, October 1985. https://www.cs.ru.nl/~jhh/pub/secsem/chaum1985bigbrother.pdf

Damgård, Ivan & Nielsen, Jesper. “Commitment Schemes and Zero-Knowledge Protocols“, 2011. https://cs.au.dk/%7Eivan/ComZK06.pdf

Lysyanskaya, Anna & Rivest, Ronald & Sahai, Amit & Wolf, Stefan. “Pseudonym Systems“, 1999. https://crypto.ethz.ch/publications/files/LRSW99.pdf

Bibliography

Camenisch, Jan & Lysyanskaya, Anna. “An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation“, 2001. https://cs.brown.edu/people/alysyans/papers/cl01a.pdf

Camenisch, Jan & Herreweghen, Els. “Design and Implementation of the idemix Anonymous Credential System“. IBM Research. https://www.freehaven.net/anonbib/cache/idemix.pdf

Paquin, Christian. “U-Prove Technology Overview V1.1“, Microsoft, April 2013 https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Technology20Overview20V1.120Revision202.pdf?msockid=19ca7afbd4c1646b36bd6c9cd54b6500

Camenisch, Jan & Lysyanskaya, Anna. “Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials“, February 2002. https://cs.brown.edu/people/alysyans/papers/camlys02.pdf

Alborch Escobar, F., Canard, S., Laguillaumie, F., & Phan, D. H. (2024). Computational Differential Privacy for Encrypted Databases Supporting Linear Queries. Proceedings on Privacy Enhancing Technologies, 2024(4), 583–604.

Belorgey, M. G., & Carpov, S. (2024). Combining Cryptography and Other Techniques for Various Privacy-Preserving Applications. NIST Crypto Reading Club, May 15, 2024.

Dwork, C. (2008). Differential Privacy: A Survey of Results. Lecture Notes in Computer Science, 4978, 1–19.

Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.

Bibliography

Movsowitz Davidow, D., Manevich, Y., & Toch, E. (2023). Privacy-Preserving Transactions with Verifiable Local Differential Privacy. Tel Aviv University & IBM Research.

Near, J., Darais, D., & Boeckl, K. (2020, July 27). Differential privacy for privacy-preserving data analysis: An introduction to our blog series. Cybersecurity Insights (NIST blog). https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

NIST. (2025). Guidelines for Evaluating Differential Privacy Guarantees. NIST Special Publication 800-226. U.S. Department of Commerce.

Morris, Dana. “How Encryption Works to Preserve Data Privacy”. Dataversity, January 3, 2023. https://www.dataversity.net/articles/how-encryption-works-to-preserve-data-privacy/#:~:text=to%20have%20both.-,Privacy%2DPreserving%20Cryptography,-Privacy%2Dpreserving%20cryptography

Craddock, Mark et. al. “UN Handbook on Privacy-Preserving Computation Techniques”. Big Data UN Global Working Group, Accessed September 30, 2025. https://unstats.un.org/bigdata/task-teams/privacy/UN%20Handbook%20for%20Privacy-Preserving%20Techniques.pdf

“Privacy-Preserving Collaboration Using Cryptography”. Digital.gov, Accessed September 30, 2025. https://digital.gov/resources/privacy-preserving-collaboration-using-cryptography#:~:text=Secure%20multi%2Dparty%20computation%20(MPC)%20is%20a%20type%20of,optimize%20MPC%20for%20complex%20functions.