Research Statement

F. Betül Durak

Microsoft Research, Redmond, USA
betul.durak@microsoft.com

My name is Betül Durak, and I am an applied cryptographer with a broad interest in the intricate balance between security and privacy. My journey began as a cryptanalyst, where I honed my skills in identifying weaknesses, exposing vulnerabilities, and employing sophisticated techniques to fortify cryptographic designs. Over time, I have evolved, and now most of the time, I leverage my cryptanalyst skills in crafting secure systems.

About me. After completing my Ph.D., I found myself at a crossroads, unsure whether to pursue a career in academia or in an industrial lab. Throughout my Ph.D., I conducted impactful research across various domains and interned in both industry and academic labs, but I needed time to discover what I truly enjoyed. I joined EPFL as a post-doc, eager to explore multiple domains in a large lab like LASEC. In my second year, I had the chance to interact with industry professionals, including executives, architects, and developers. This experience convinced me to gain industry experience, adopt a different perspective on research, find impactful opportunities, and collaborate with a large team of engineers. I believe I have successfully gained this unique industry perspective, contributed to research in industry, and standardized solutions while continuing to publish in top-ranked academic conferences. In this short statement, I will share my research journey to provide insight into my thinking and approach to the research process, concluding with my current stance on research.

The Security of Property-Preserving Encryption. I hold my Ph.D from Rutgers University. During my PhD, I delved into Property-Preserving Encryption schemes, which play a crucial role in securing databases. These cryptographic primitives are designed to maintain the database's format post-encryption (Format-Preserving Encryption) and enable feasible search capabilities on encrypted data (Searchable Encryption, Order-Revealing Encryption) [11, 15, 17, 13, 23]. My initial research was primarily centered on the cryptanalysis of these primitives in real-world scenarios.

In my work with Order-Revealing Encryption (ORE), I uncovered fundamental issues with the primitive itself: the leakage from encrypted data, even in the most secure designs, is inevitably too significant to make the primitive viable for deployment. This work was published in Computer and Communications Security (CCS) 2016 [11]. This research, conducted during my early PhD years, marked my first real-world impact. It sharpened my critical perspective on the promises of structure-preserving encryption mechanisms and their actual efficacy in deployed systems, such as legacy database systems. Our work demonstrated that the fundamental security vulnerabilities of ORE outweigh the functionality benefits it offers. As a result, to the best of my knowledge, ORE has not been deployed in any real-world systems, thanks to the explicit issues our research highlighted.

My subsequent research had an even greater impact. I identified and repaired vulnerabilities in a Format-Preserving Encryption (FPE) design called FF3, which is standardized by the National Institute of Science and Technology (NIST) [15, 16]. FPE is crucial for securing payment systems end-to-end. Attackers target payment terminals to steal credit card information, leaving hundreds of employees' credit cards in the hands of attackers. FF3, designed by Ingenico, the major payment terminal leader, was compromised. My work garnered attention from the Center of Discrete Mathematics and Theoretical Computer Science (DIMACS), which published an article based on my findings [10]. I was invited to NIST's office in Washington, DC, to discuss my research internally and to present at the flagship conference Real World Cryptography (RWC) to a broader audience [16, 18]. As a result, FF3 was dropped from the standards when the NIST team realized the vulnerability we demonstrated undermined their confidence in the algorithm.

Inspired by my talk at RWC, I was approached by Comforte, very successful company specializing in FPE with a large customer base, including major financial institutions such as VISA, Mastercard, and many of the world's leading banks [2]. I collaborated with them to cryptanalize their design and recommend changes to enhance its security. Together, we published a new design that the company now implements for their customers [13].

Post-Quantum Cryptography. After completing my Ph.D., I embarked on a post-doctoral journey at Ecole Polytechnique Fédérale de Lausanne (EPFL). The Laboratory for Cryptologic Algorithms (LASEC) at EPFL provided an exceptional environment for addressing various student-driven problems and marked my initial foray into real-world challenges through collaboration with a start-up. During my time at EPFL, I adopted a breadth-first approach to explore various domains and understand the impact and possibilities of the technologies studied. Specifically, I delved into Post-Quantum Cryptography (PQC), secure end-to-end encryption (E2EE), and biometric access control.

My interest in PQC was sparked during student presentations at LASEC, where students expressed their curiosity about the future of internet security in the advent of quantum computers. I began reading academic papers considered under NIST's standardization process and noticed a recurring statement about the security guarantees of the proposed algorithms: the key exchange (KE) and public key cryptosystems (PKC) were only secure for ephemeral key settings. This implied that key generation was required for every single operation, which looked impractical given the cost of KE algorithms. It became clear to me that the "rule" of not reusing secret key material would not be followed by practitioners, who often prioritize performance over security. I proposed that our lab investigates the real harm caused by ephemeral key reuse. The outcome of our work, which involved various intern projects and student seminars, was a publication in Eurocrypt 2019 [3]. This work became my second most cited research, highlighting that academics studying secure designs for PQC are paying close attention to this result.

End-to-End Encryption. During Crypto 2016 conference, I observed that academics were keen to study the problem of E2EE across various platforms, from direct messaging with text to encrypted calls. The applied cryptography community was excited about the formal security guarantees of communication platforms. However, the initial formal guarantees were overly simplified, assuming unidirectional communication. I wondered how far we could extend these formal guarantees, which the industry seemed to be seeking. At LASEC, we explored the security guarantees for bidirectional E2EE and how to design such a protocol. Our first paper was published a month or two too late to compete in the major venues like Eurocrypt or Crypto, but it was published in a lower-ranked venue [19] and has since been cited by dozens of follow-up works. I saw this setback as an opportunity to do more: all current E2EE design proposals were rigid, leaving no room for flexibility based on the platform's needs. This meant that any deviation for efficiency and usability reasons would cause the formal guarantees to fail. I worked on this problem with a student, Andrea Caforio, and we published our findings in Public Key Cryptography (PKC) [6]. After this project, Andrea decided to pursue further research and later joined EPFL as a Ph.D. student.

Biometry. Last but certainly not least, EPFL provided me with a life-changing opportunity to engage with an Innovation Park startup to research biometric access control [4, 14]. We tackled numerous real-world problems faced by businesses and technologies, from prototyping our solutions for Jura Hospital to pitching our proposals in customer IT meetings and presenting at the Scientastic science fair at EPFL in 2018. Our ideas were patented and are now used by GlobalID.

Privacy of AI. Although it was bittersweet to leave the collaborative environment in academia, I was eager to transition into the real world. My first experience as a full-time researcher in a corporate technology and development lab was at the Robert Bosch Technology and Research Center in Pittsburgh, USA. During my first year, I immersed myself in Artificial Intelligence (AI) models, particularly focusing on Convolutional Neural Networks (CNNs). CNNs are powerful tools commonly used in image processing and recognition, consisting of two phases: training, which inputs data and outputs a model, and inference, which inputs data and outputs the processed "decision."

Given Bosch's investment in autonomous vehicles, it was crucial for the company to process images responsibly, ensuring privacy and security. As an expert in privacy and applied cryptography, I contributed to a project where CNNs were implemented in a Multi-Party Computation (MPC) model. MPC provides a way to process data in a privacy-preserving manner, based on distributed trust with the assumption that no servers collude to uncover the secret-shared data. Since processing is done with multiple servers, it inherently incurs performance challenges, making it critical to solve these issues at both the hardware and software levels for real-world deployment.

I authored my first solo patent on computation optimization techniques, which I tested in a large-scale code base supported by the MP-SPDZ framework, an academic initiative. Subsequently, I explored integrating well-known block ciphers, such as AES, into MPC to enhance performance for specific applications. During this work, I started looking at a recent block cipher called LowMC, designed specifically for computing block ciphers with MPC operations with better throughput. Later on, LowMC designers collaborated with the Microsoft Research team to integrate LowMC into a post-quantum digital signature scheme called Picnic. Given that LowMC was a very new design, I was cautious about proposing anything based on it without first conducting a thorough cryptanalysis, after all I am a cryptanalyst at heart. The Picnic team received numerous questions about the cipher's strength, and I shared the skepticism about the LowMC design. An opportunity arose when the core designers of LowMC announced a cryptanalyst challenge, partially sponsored by Microsoft. I decided to participate in the LowMC challenge, aiming to break the block cipher with a single plaintext-ciphertext pair, a task deemed very difficult. We successfully won the challenge and received the best paper award at the Fast Software Encryption (FSE) conference [1]. Our techniques inspired well-known researchers to invent new and generic cryptanalysis tools for block ciphers [9]. My research concluded that block ciphers like LowMC is much more efficient than AES for MPC operations, but they require further maturation for security [12, 5].

During my investigation of MPC and CNNs, the leadership made the decision to reallocate funding away from that research. After working on privacy-preserving technologies such as analyzing (breaking) the technology developed by a company called AirClock, I began exploring my next opportunities as a researcher. In 2021, I was hired by Microsoft Research. As a Senior Research Scientist at Microsoft Research, I have the privilege and liberty of selecting the most critical problems to tackle. Working in a large organization like Microsoft Research, I have access to renowned researchers and engineers from all imaginable fields. My personal approach to research is to work with interdisciplinary projects with various experts and to align with the company's goals so that my efforts can help shape future technologies through Microsoft.

Replacing Third Party Cookies: Digital Identities and Privacy. I will start with the "toolkit", as I like to call it, which I collected in my first two years at Microsoft Research. In my first year, I wanted to continue exploring MPC and its real-world applications within Microsoft. I was introduced to the Microsoft Edge team and their partners at Microsoft Bing to investigate new privacy-preserving technologies for online advertisement. It was a time when Google and many players in the domain were amending to remove third-party cookies. It was a compelling idea to be part of developing a new technology to replace third-party cookies. The new technology would enable the previously used functionalities with the priority of privacy. We designed a new algorithm that extended the functionalities and efficiency of Prio [8]. I found myself presenting our work [20, 22] in various meetings, from IETF standardization meetings to internal product team meetings. Our ideas were so simple and adaptable that Microsoft Edge sought my expertise to help standardize the protocols for their use. While Google faced challenges in making significant progress on privacy initiatives, such as the removal of cookies and the adoption of MPC technologies, the collaboration with Microsoft Edge was invaluable. Although the project faced hurdles due to Microsoft Edge's reliance on Chromium, it provided me with a wealth of knowledge about the web ecosystem and advertisement technologies. This experience highlighted the real-world deployment, legal, and economic challenges of MPC technologies, but it also underscored the importance of continued innovation and collaboration in this field.

My collaboration with Microsoft Edge continued with privacy concerns on the web, where I studied anonymous tokens and their applications on the web with motivated applications [7, 21, 24]. Anonymous tokens are a lightweight version of anonymous credentials, enabling individuals to prove attributes such as age, gender, residency, and employment without revealing more information than necessary. These tools utilize zero-knowledge proofs, and making them efficient while finding the right architecture for their operation is a significant challenge: a central Certificate Authority (CA) should issue the credential after verifying the user's statements. Academic works often assumed that the CA would be the Division of Motor Vehicle (DMV) or government. However, creating an infrastructure to take people's offline identities, verify them, convert them into online identities, and allow anyone to verify them in a privacy-preserving way is not easy.

Nevertheless, we academics are smart and creative. Envisioning online identities with anonymous credentials led us to start with a smaller scope: what if a server becomes both the issuer and verifier? The first deployment of anonymous tokens where the issuer and verifier are the same entity was Privacy

Pass, as we know it today. Invented and deployed by Cloudflare, Privacy Pass allows VPN users to access the internet privately while preventing DDoS attacks. This approach brilliantly protects privacy while maintaining security.

However, this approach was not enough. The fraud team at Microsoft realized that when attackers noticed they were not issued an anonymous access token, they changed their behavior, which disrupted Microsoft's detection methods. To address this, we ensured that attackers were not immediately aware that their request for an access token was denied until at a later time. I published my work in Crypto 2023 and CCS 2024 [7, 21].

Where do I stand today? I am genuinely concerned about the technologies and their intended uses versus the side effects on societies and vulnerable populations. It would be easy to isolate myself and my research from emerging technologies such as AI and LLMs, as I see various researchers doing. However, as a security, privacy, and cryptography researcher, I want to be part of the sociotechnical change where I study the AI systems with security, reliability, and safety for the end users.

References

[1] LowMC Challenge. `https://lowmcchallenge.github.io/`.

[2] Comforte AG. `https://www.comforte.com/company`, 2024.

[3] Ciprian Băetu, F. Betül Durak, Loïs Huguenin-Dumittan, Abdullah Talayhan, and Serge Vaudenay. Misuse Attacks on Post-quantum Cryptosystems. In *Advances in Cryptology – EUROCRYPT 2019*, pages 747–776. Springer International Publishing, 2019.

[4] Fatih Balli, F. Betül Durak, and Serge Vaudenay. BioID: A Privacy-Friendly Identity Document. In *Security and Trust Management*, pages 53–70. Springer International Publishing, 2019.

[5] Subhadeep Banik, Khashayar Barooti, F. Betül Durak, and Serge Vaudenay. Cryptanalysis of LowMC instances using single plaintext/ciphertext pair. *IACR Transactions on Symmetric Cryptology*, 2020, Issue 4:130–146, 2020.

[6] Andrea Caforio, F. Betül Durak, and Serge Vaudenay. Beyond Security and Efficiency: On-Demand Ratcheting with Security Awareness. In *Public-Key Cryptography – PKC 2021*, pages 649–677, 2021.

[7] Melissa Chase, F. Betül Durak, and Serge Vaudenay. Anonymous Tokens with Hidden Metadata Bit From Algebraic MACs. In *Advances in Cryptology – CRYPTO 2023*, volume 14082, pages 418–449. Springer, 2023.

[8] Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, Robust, and Scalable Computation of Aggregate Statistics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation*, page 259–282. USENIX Association, 2017.

[9] Itai Dinur. Cryptanalytic Applications of the Polynomial Method for Solving Multivariate Equation Systems over GF(2). In *Advances in Cryptology – EUROCRYPT 2021*, pages 374–403. Springer International Publishing, 2021.

[10] F. Betül Durak. Betül Durak and Her Research. `http://dimacs.rutgers.edu/news_archive/durak-research`, 2017.

[11] F. Betül Durak, Thomas M. DuBuisson, and David Cash. What Else is Revealed by Order-Revealing Encryption? In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS 2016, page 1155–1166. Association for Computing Machinery, 2016.

[12] F. Betül Durak and Jorge Guajardo. Improving the Efficiency of AES Protocols in Multi-Party Computation. In *Financial Cryptography and Data Security*, pages 229–248, 2021.

[13] F. Betül Durak, Henning Horst, Michael Horst, and Serge Vaudenay. Fast: Secure and High Performance Format-Preserving Encryption and Tokenization. In *Advances in Cryptology – ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part III*, page 465–489, 2021.

[14] F. Betül Durak, Loïs Huguenin-Dumittan, and Serge Vaudenay. BioLocker: A Practical Biometric Authentication Mechanism Based on 3D Fingervein. In *Applied Cryptography and Network Security: 18th International Conference, ACNS 2020, Rome, Italy, October 19–22, 2020, Proceedings, Part II*, page 62–80, 2020.

[15] F. Betül Durak and Serge Vaudenay. Breaking the FF3 Format-Preserving Encryption Standard over Small Domains. In *Advances in Cryptology – CRYPTO 2017*, pages 679–707. Springer International Publishing, 2017.

[16] F. Betül Durak and Serge Vaudenay. NIST Announcement on Breaking the FF3 Format-Preserving Encryption Standard over Small Domains. `https://csrc.nist.gov/news/2017/recent-cryptanalysis-of-ff3`, 2017.

[17] F. Betül Durak and Serge Vaudenay. Generic Round-Function-Recovery Attacks for Feistel Networks over Small Domains. In *Applied Cryptography and Network Security*, pages 440–458. Springer International Publishing, 2018.

[18] F. Betül Durak and Serge Vaudenay. Real World Cryptography (RWC) talk: Breaking the FF3 Format-Preserving Encryption Standard over Small Domains. `https://www.youtube.com/watch?v=_UIcJ1Bnxa0`, 2018.

[19] F. Betül Durak and Serge Vaudenay. Bidirectional Asynchronous Ratcheted Key Agreement with Linear Complexity. In *Advances in Information and Computer Security: 14th International Workshop on Security, IWSEC 2019, Tokyo, Japan, August 28–30, 2019, Proceedings*, page 343–362, 2019.

[20] F. Betül Durak, Chenkai Weng, Erik Anderson, Kim Laine, and Melissa Chase. Precio: Private Aggregate Measurement via Oblivious Shuffling. Cryptology ePrint Archive, Paper 2021/1490, 2021.

[21] F. Betül Durak, Laurane Marco, Abdullah Talayhan, and Serge Vaudenay. Non-Transferable Anonymous Tokens by Secret Binding. ACM CCS 2024, to appear.

[22] F. Betül Durak, Chenkai Weng, Erik Anderson, Kim Laine, and Melissa Chase. Precio: Private Aggregate Measurement via Oblivious Shuffling. ACM CCS 2024, to appear.

[23] Thang Hoang, Attila A. Yavuz, F. Betül Durak, and Jorge Guajardo. Oblivious Dynamic Searchable Encryption on Distributed Cloud Systems. In *Data and Applications Security and Privacy XXXII*, pages 113–130. Springer International Publishing, 2018.

[24] Rachel McAmis, Betül Durak, Melissa Chase, Kim Laine, Franziska Roesner, and Tadayoshi Kohno. Handling Identity and Fraud in the Metaverse . *IEEE Security & Privacy*, pages 2–12.