Privacy has been a critical concern since the beginning of the information age. Data-driven services incorporate every search, view, and click of users into analytics. Different entities exchange their data to improve products further. However, this also exposes sensitive information to potential risks and threats. Recent years have witnessed the enhancement of privacy laws worldwide and the recognition of privacy as a fundamental human right. Unrestricted sharing of data is no longer presumed to be a default option. This poses challenges for service providers and users: (1) The quality of personalized services highly relies on the demographic profile of users, but privacy concerns obstruct the acquisition of such information. (2) It is difficult to verify the correctness of remote execution when the data is stored and processed remotely by other parties. (3) Existing information security principles focus on end-to-end encryption and ad-hoc anonymization but do not guarantee complete protection in data sharing.

I develop privacy-preserving techniques with a solid cryptographic foundation, which allows Internet and data services to thrive without exposing sensitive information. My research focuses on designing and improving advanced secure computing algorithms and protocols with the goal of applying them to tackle real-world problems. I build systems based on customized, secure, and robust cryptographic primitives that cater to the demands of information assurance. My research portfolio is enriched by working closely with researchers from databases, networking, formal methods, machine learning, finance, and healthcare industries. Systems I build are already in the pipeline for real-world deployment by tech companies, including Meta, Google, Chainlink, and JPMorgan Chase.

1 Computing and Proving Secrets: From Theory to Practice

My work centers on secure multi-party computation (MPC) and zero-knowledge proof (ZKP). MPC enables multiple parties to jointly compute a function while each party's input is not shared with other parties. ZKP allows a prover to demonstrate the properties of its secret without revealing it. Despite the rapid evolution of these techniques, many essential aspects are usually overlooked. Examples include (1) Protocol design and performance measurement concerning the cumulative effect of software stack, hardware stack, and network communication. (2) Rigorous security analysis. (3) Understanding and utilization of cryptographic assumptions. To address these challenges, I design concretely efficient protocols with strong security guarantees and build prototype implementations to turn them into almost production-ready privacy-enhancement tools.

Computation and communication-efficient MPC protocols. Vector oblivious linear evaluation (VOLE) is a common building block in many MPC and ZK protocols. Pioneering work on pseudorandom correlation generators (PCG) promises a significant reduction in communication but still incurs substantial computational overhead. As a result, it only outperforms its previous work in low-bandwidth settings. My work Ferret achieves up to 12× improvement over the previous PCG and up to 238× improvement over non-PCG protocols in terms of the end-to-end performance [YWL+20, WYKW21]. Its efficiency is maintained in all network conditions. Previous PCG concluded that protocols based on the primal-LPN assumption are deemed inefficient; therefore, it uses its variant dual-LPN. My work incorporated the concept of VOLE extension in the setup phase of PCG, which avoids the excessive communication cost of primal-LPN-based constructions while benefiting from its lightweight computation. *Ferret is currently integrated into the Meta Private Computation Framework, and used for ads measurement in its Private Lift Measurement project*[1].

An open problem in cryptography is how to design an MPC protocol that scales to a large number of parties while still being secure against a majority of corrupted parties. Previous work shows the feasibility of reducing the total traffic to a constant per program unit, but there have yet to be any practical protocols introduced. In collaboration with JPMorgan researchers, our work SuperPack [EGP+23] significantly improves the responsiveness of the protocol to near practical. The benchmark shows up to 6.4× improvement over the previous best result for 80-party computation. The advantage is more prominent if the proportion

[1] Private Computation Framework 2.0, Meta, https://research.facebook.com/publications/private-computation-framework-2-0/; BBC Becomes first UK organization to use Private Lift from Meta to measure marketing success, Meta, https://www.facebook.com/business/measurement/case-studies/BBC.

of honest parties increases. One of the core techniques is a novel sharing transformation that converts two types of secret sharing so that the protocol benefits from both: one is easier to achieve security against active adversaries, and another introduces less communication overhead. This is achieved by exploring the linearity and polynomial structures of these sharing schemes.

Proving trillion-size computation in commodity hardware. Most existing ZKPs are not scalable regarding the prover efficiency and memory consumption. Their current deployments involve only small-scale statements ($2^{10} \sim 2^{25}$ gates), and the provers usually need computing-optimized servers with hundreds of gigabytes of memory. To empower large-scale ZKP applications, I initiated the study of a designated-verifier proof system named VOLE-ZK [YSWW21, WYKW21]. The main innovation is efficiently committing to secrets by leveraging VOLE, followed by a constant-size proof of consistency. After introducing a round of interaction, it supports dynamic deallocation of committed values, which is challenging in existing ZKPs. Due to high scalability and efficiency, VOLE-ZK easily handles massive circuits of $2^{30} \sim 2^{40}$ gates. Its lightweight feature also allows the provers to be deployed in commodity hardware even with only 1 CPU and 1 gigabytes RAM. *Quicksilver [YSWW21] is among the runner-ups for CCS 2021 Best Paper Awards. Chainlink oracles use it to provide proof of provenance to blockchains.*[2] *A recent submission for the NIST post-quantum digital signature standardization*[3] *is also built upon the proof technique from Quicksilver.* As follow-up works, I developed VOLE-ZK that supports mixed-mode statements [WYX+21], sublinear-communication [WYY+22], and RAM programs [FKL+21].

Concrete security of cryptographic protocols. The security of most cryptographic protocols is proven in asymptotic notations for ease of composition. However, real-world applications require concrete security during deployment. My work [GKW+20] was among the first to study the concrete security of a popular MPC protocol, which has been used by several companies to build commercial applications. I showed an attack to prove that the security of a major MPC protocol is much lower than people's impression. Some popular implementations could be completely broken by a cluster of machines in a month. In previous designs, the hardness assumptions either do not have efficient instantiation (e.g., random oracle model) or can not achieve desired security (e.g., random permutation model). I proposed a design based on the ideal cipher model and applied instruction-level optimization, which provides optimal security with almost no reduction in efficiency. *This new design is used in the IETF standardization of verifiable distributed aggregation functions proposed by Google, Cloudflare, and Cisco*[4].

2 Cryptography Enabled Privacy-Preserving Systems

I design secure, efficient, and scalable cryptographic primitives and apply them for the privacy-enhancement of digital systems. I am interested in patching vulnerable systems and exploring novel applications emphasizing information security. Such interdisciplinary research requires domain-specific knowledge from multiple areas to capture privacy risks, model adversarial behaviors, and optimize defense mechanisms.

Zero-knowledge proof of correct (ML and DB) executions. The integrity of machine learning is a critical security concern of ML applications. It is essential to the users that they are appropriately served by an ML model. However, the models usually cannot be revealed for inspection as digital assets. By extending my Boolean-friendly VOLE-ZK for mixed computation, the Mystique framework [WYX+21] generates the proof of correct ML inference while hiding model parameters. Existing proof techniques cannot efficiently prove mixed statements containing linear components (e.g., convolutions) and non-linear components (e.g., ReLU, Sigmoid, and Batch Normalization). I proposed the zk-edaBits protocol that relates the cryptographically committed Boolean values and large field elements, enabling the ZKP to switch between the proof of convolution layers and other non-linear layers. Furthermore, the Boolean-friendly feature of my protocol enables the support of IEEE standard floating-point arithmetics. The accuracy loss is only 0.02% compared

[2] VOLE-Based ZK, Chainlink Labs Research, https://blog.chain.link/vole-based-zk/.

[3] FAEST Signature Algorithm, Baum et al., https://faest.info/authors.html.

[4] VDAF, Barnes et al., https://datatracker.ietf.org/doc/draft-irtf-cfrg-vdaf/.

to the plaintext evaluation, which is hard to achieve by other ZKPs.

ZKP can also be applied to the cloud database so that data users can verify the correctness and soundness of SQL results concerning the integrity of database contents. Meanwhile, the zero-knowledge property prevents the leaking of irrelevant information. However, the high complexity of query processing algorithms poses a challenge to the design of ZKPs. Instead of directly applying general-purpose frameworks, my work ZKSQL is specially tailored to leverage the fact that some results are easier to "verify" than "compute" [LWX+23]. It first reduces the database operators to lower-level operations, such as set-based and circuit-based operations. By combining VOLE-ZK and mathematical proofs, I designed customized protocols of set relations, reducing the verification complexity from superlinear to linear. The benchmark on the TPC-H workload shows $> 10^2$ improvement over pure circuit-based implementations in terms of performance and financial cost.

Privacy-enhancement in computer networks. Regular expression (regex) pattern matching is a highly accurate and expressive tool used in network packet filtering. However, inspecting the packet contents usually requires firewalls to break the enforcement of end-to-end encryption. My work on private regex matching [LWS+23] provides privacy enhancements for such applications. It includes 1) a ZKP that allows a packet sender to convince an inspector whether a publicly-known regex matches its private strings. 2) MPC protocols that take the private string and regex from different parties and only output the result of matching. Both are based on the design of an efficient simulation algorithm for Thompson nondeterministic finite automata. It achieves linear overhead by exploring the graph sparsity. Also, its memory access pattern obviates the need to hide the input string by oblivious data structures, which usually dominate the cost.

Internet users access the domain name system through distributed recursive resolvers (ReR), which collect a considerable amount of individual access patterns and profits from them. Such abuse cannot be avoided by end-to-end encryption or proxies. My work, PDNS, takes one more step toward a fully private DNS by utilizing a powerful cryptographic primitive called private information retrieval (PIR) [XWY+23]. It forces ReR to process DNS queries and deliver responses in encrypted mode without learning their contents. However, simply plugging in PIR is insufficient. ReR relies on caching DNS records to reduce the response time. Blind operations prevent caches from being updated. PDNS leverages the trust assumption in authoritative DNS servers to populate caches without ReR correlating the queries with users. Instead of entirely subverting the modern DNS architecture, my work is compatible with its hierarchical construction and allows gradual deployment. *Based on the preliminary results from this project, I co-lead the development of an NSF grant: Privacy-Preserving and Censorship-Resistant Domain Name System (Award #2310927).*

Privacy-preserving data sharing and aggregation. I worked with researchers from Microsoft, Alibaba, and medicine schools in Chicago to develop privacy-preserving data analysis systems.

- The Microsoft Precio system [ACD+21] is a private online conversion measurement protocol used to collect users' reactions to ad campaigns, analyze their effectiveness, and publish the results to advertisers. *I designed its cryptographic protocol and implemented the prototype, which is further developed and officially open-sourced by Microsoft.*[5] In Precio, multiple servers maintained by different entities collect the secretly shared user data and jointly compute a histogram to reflect the performance of ad campaigns. By utilizing oblivious shuffling and differential privacy, servers and advertisers cannot infer the information of individual users from noisy aggregates.
- Collaborative filtering is a popular tool used in recommendation systems. An entity with a large amount of high-dimension user data is often paid for helping small entities (with fewer users) fulfill their missing knowledge, e.g., the preferences of their users. In cases when none of them is willing to expose their data, *my work with Alibaba* provides a privacy-preserving solution for an unbalanced recommendation system based on collaborative filtering [HHW+23]. It is a high-performance parallel homomorphic encryption (HE) system built from the MPI distributed computing framework. In contrast to common attitudes re-

[5] Precio, Microsoft, https://github.com/Microsoft/Precio.

garding the bad performance of HE, it only takes 3 minutes for a benchmark with 16 million user records.

- *In collaboration with medicine schools around Chicago*, I developed a high-performance MPC protocol for the patient risk stratification [DRW+21]. It is a process to predict the amount of resources that should be allocated in each medical center, and it usually can only be done by cross-institute data analytics. With my protocol, the hospitals jointly analyze data without exposing confidential patient information. By applying Cuckoo hash, I reduce the complexity of patient matching from quadratic to linear.

3 Future Research

Current cryptography-based privacy-enhancing mechanisms face the following struggle. In data-sharing scenarios, they provide straightforward and reliable solutions to mitigate privacy risks and protect sensitive information. Major tech companies and government organizations already integrated them into products and public services. However, these advanced cryptographic protocols consume tremendous computing and network resources and delay the delivery of services. Organizations and individuals are not ready to trade off their profits and user experience for privacy. With that in mind, my research agenda focuses on improving the security and usability of cryptographic protocols. By understanding the privacy risk imposed in digital services and applications, I build customized privacy-preserving systems that effectively alleviate the conflicts between privacy and functionality.

Approachable implementation for secure computation. One critical issue that prevents the deployment of secure computation is the lack of apprehensive interfaces. The current design and implementation of privacy-preserving applications rely on crypto theory and engineering expertise. Meanwhile, existing MPC libraries either target too specific use cases or expose all low-level cryptographic operations to users who cannot correctly configure and implement. Moving forward, my research will focus on building both high-quality implementations of secure computation primitives as well as designing domain-specific languages for general privacy-preserving tasks. A particular challenge is the development of a compiler that precisely translates the security requirements, maps high-level computing objectives to low-level cryptographic operations, and maintains the performance by protocol-specific optimizations. I will address these issues by combining optimization strategies of cryptography, programming language, and formal verification.

Zero-knowledge with applications to the legal system. Digital objects in the legal system are easily fabricated or modified, so they need thorough examinations. Meanwhile, the data owner may be concerned whether its information is inappropriately accessed since it often contains sensitive information. ZKP is a technical solution that enables inspectors to verify the authenticity of the data source and provides reliable proof that the investigative software does not attain access to irrelevant information. However, current studies on ZKP do not provide immediate support for software analysis and data auditing. Moreover, such infrastructure should be affordable and approachable to the general public in order not to put an extra burden on digital forensics. My previous research on efficient interactive proof and private pattern matching is a starting point for this challenge. In the future, I seek to design customized proof systems and authenticated data structures that will promote the reliability and privacy of evidence inspection in the legal process.

Privacy protection in financial services and data science. Financial institutions gather and store extensive amounts of sensitive and valuable data, including the identifiable and behavioral information of their clients. Banks can acquire sensitive strategic information from other entities when operating business-to-business platforms. Unfortunately, their privacy is extensively compromised in existing financial systems. MPC and ZKP are two techniques that have proven to help detect money laundering across banks, optimizing bank transaction settlements and distributing stock bids without compromising valuable data from individuals. *During my long-term collaboration with the cryptography researchers in JPMorgan Chase*, we came across countless opportunities both internally and externally aiming for enhancing the security, privacy, and robustness of current financial services. Shortly, I will focus on projects related to private transactions in over-the-counter markets and joint portfolio optimization between banks and hedge funds.

References

[ACD+21] Erik Anderson, Melissa Chase, F. Betul Durak, Esha Ghosh, Kim Laine, and Chenkai Weng. Aggregate measurement via oblivious shuffling. Cryptology ePrint Archive, Report 2021/1490, 2021. https://eprint.iacr.org/2021/1490.

[DRW+21] Xiao Dong, David Randolph, Chenkai Weng, Abel Kho, Jennie Rogers, and Xiao Wang. Developing high performance secure multi-party computation protocols in healthcare: A case study of patient risk stratification. In *AMIA 2021 Informatics Summit*, 2021.

[EGP+23] Daniel Escudero, Vipul Goyal, Antigoni Polychroniadou, Yifan Song, and Chenkai Weng. SuperPack: Dishonest majority MPC with constant online communication. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 220–250. Springer, Heidelberg, April 2023.

[FKL+21] Nicholas Franzese, Jonathan Katz, Steve Lu, Rafail Ostrovsky, Xiao Wang, and Chenkai Weng. Constant-overhead zero-knowledge for RAM programs. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 178–191. ACM Press, November 2021.

[GKW+20] Chun Guo, Jonathan Katz, Xiao Wang, Chenkai Weng, and Yu Yu. Better concrete security for half-gates garbling (in the multi-instance setting). In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 793–822. Springer, Heidelberg, August 2020.

[HHW+23] Z. Huang, C. Hong, C. Weng, W. Lu, and H. Qu. More efficient secure matrix multiplication for unbalanced recommender systems. *IEEE Transactions on Dependable and Secure Computing*, 20(01):551–562, jan 2023.

[LWS+23] Ning Luo, Chenkai Weng, Jaspal Singh, Gefei Tan, Ruzica Piskac, and Mariana Raykova. Privacy-preserving regular expression matching using nondeterministic finite automata. Cryptology ePrint Archive, Paper 2023/643, 2023. https://eprint.iacr.org/2023/643.

[LWX+23] Xiling Li, Chenkai Weng, Yongxin Xu, Xiao Wang, and Jennie Rogers. Zksql: Verifiable and efficient query evaluation with zero-knowledge proofs. *Proceedings of the VLDB Endowment*, 16(8):1804–1816, 2023.

[WYKW21] Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *2021 IEEE Symposium on Security and Privacy*, pages 1074–1091. IEEE Computer Society Press, May 2021.

[WYX+21] Chenkai Weng, Kang Yang, Xiang Xie, Jonathan Katz, and Xiao Wang. Mystique: Efficient conversions for zero-knowledge proofs with applications to machine learning. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 501–518. USENIX Association, August 2021.

[WYY+22] Chenkai Weng, Kang Yang, Zhaomin Yang, Xiang Xie, and Xiao Wang. AntMan: Interactive zero-knowledge proofs with sublinear communication. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2901–2914. ACM Press, November 2022.

[XWY+23] Yunming Xiao, Chenkai Weng, Ruijie Yu, Peizhi Liu, Matteo Varvello, and Aleksandar Kuzmanovic. Demo: Pdns: A fully privacy-preserving dns. In *Proceedings of the ACM SIGCOMM 2023 Conference*, ACM SIGCOMM '23, page 1182–1184, New York, NY, USA, 2023. Association for Computing Machinery.

[YSWW21] Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2986–3001. ACM Press, November 2021.

[YWL+20] Kang Yang, Chenkai Weng, Xiao Lan, Jiang Zhang, and Xiao Wang. Ferret: Fast extension for correlated OT with small communication. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1607–1626. ACM Press, November 2020.