Research Statement

XUANMING (HINS) LIU
M.Eng. student at Computer Science and Technology College, Zhejiang University
hinsliu@zju.edu.cn, https://hinsliu.com
December 1, 2024

I work in applied cryptography. My research interests primarily focus on the **efficiency and applications of cryptographic techniques**, especially *zero-knowledge proofs* [15]. The goal is to leverage these efficient cryptographic techniques to enhance fairness and privacy in real-world applications, such as blockchain [27].

Background and Motivation

My belief is that an increasing number of people today are paying attention to their privacy, security, and fairness. In modern real-world applications, developers are required to balance efficiency with ensuring user security and privacy. Advances in cryptographic techniques make this goal achievable. Since the concept of zero-knowledge proofs was introduced by Goldwasser, Micali, and Rackoff in [15], it has received extensive research attention in theoretical domains [14, 1, 26, 16]. A zero-knowledge proof allows an efficient prover to convince the correctness of a statement to a computation-limited verifier without revealing any private witness implied by the statement. Over the past decade, with the rise of blockchain technology, specialized proof systems known as SNARK have emerged as a prominent research topic [30, 17, 10] and have gradually matured into practical tools deployed in real-world blockchain systems [2, 42, 23] to enhance privacy and scalability.

However, a significant challenge lies in enhancing the efficiency and scalability of existing SNARK systems for practical applications. When applied to large-scale or complex scenarios, current SNARKs still encounter substantial bottlenecks in computation time and memory usage. Besides, another avenue is that, given the magic of proof systems, SNARKs are regarded as highly promising for addressing ”trusted party” dilemmas in real-world systems such as blockchain, thereby ensuring privacy and fairness. Consequently, enhancing the efficiency of SNARK-like applied cryptographic techniques and broadening their applicability to practical scenarios will form the central thrusts of my future research.

Boarder interests. I am also deeply passionate about exploring other applied cryptographic techniques beyond zero-knowledge proofs, such as *multiparty computation* [35] and *homomorphic encryption* [13, 3, 4]. My viewpoint is that these technologies should not exist in isolation but can be cross-applied to create more interesting application scenarios (See my research thrust I for an explanation), in areas such as privacy-preserving and fairness assurance. In my future research, I aim to combine these techniques together to provide more comprehensive solutions for real-world applications.

Thrust I: Scalable Proof Generation in Various Scenes

Research on SNARKs has achieved significant success. We now know how to construct SNARK systems with prover time scales linearly with the size of the application [41, 34, 7]. It is also possible to achieve constant-sized verifier time and proof by sacrificing some prover efficiency [17, 10, 8]. These constructions are typically obtained by combining a *Polynomial Interactive Oracle Protocol* (PIOP) with a *Polynomial Commitment Scheme* (PCS) [19, 29, 5, 45, 43] that possesses specific properties. Nevertheless, the concrete efficiency of these constructions is still not satisfactory when considering very large-scale statements and application scenarios. Therefore, a critical question is how to further enhance the efficiency of proof generation in SNARKs across different scenarios, enabling them to scale to larger and more complex applications?

Distributed Proof Generation. One possible direction is to explore distributed proof generation (DIZK), where the prover is divided into multiple machines, each responsible for a portion of the proof generation workload. This approach accelerates the proof generation process while reducing the memory usage on each machine. Research has already demonstrated the feasibility of this method, and the proof generation of many SNARKs [17, 45, 10, 20] have been successfully distributed [40, 42, 23, 33]. However, this approach, which directly distributes the prover's secret witness among multiple machines, may introduce risks of privacy leakage. Therefore, a valuable research question is, how to achieve distributed proof generation while preserving witness privacy?

Collaborative Proof Generation. To address this issue, another direction is to explore collaborative proof generation, where the prover's private witness is distributed to several parties using secret-sharing techniques. These parties then jointly run a collaborative zk-SNARK [28] to generate a proof for a given statement. This method actually leverages multiparty computation to compute the proof generation algorithm, thereby preserving the prover's witness privacy. However, the main challenge lies in designing an efficient collaborative proof generation protocol that can be applied to large-scale applications. For instance, [11] implemented collaborative proof generation for several SNARKs [17, 10], but their protocol requires a powerful leader to handle most of the computational workload, therefore cannot be considered scalable.

To overcome this limitation, our previous work [25] shows that scalable collaborative proof generation can actually be achieved. Two SNARKs, Libra [41] and HyperPlonk [7], are extended to support collaborative proofs. In our approach, each party, without access to the (full) private witness, only needs to perform an equal and minimal computational workload. Moreover, we show that if the circuit is data-parallel, the communication cost between parties can be sublinear with respect to the circuit size, which is a highly desirable property. This design enables proof generation to be conducted in a privacy-preserving and distributed manner while maintaining excellent scalability, capable of proving circuits with more than 2^{30} gates. Despite this success, there are still many open questions in this area. For example, our scheme currently only applies to SNARKs without FFT operation [41, 7]. Considering that FFT is hard to be distributed, how can our result be extended to SNARKs with FFT [17, 10] to directly address the bottlenecks identified in [11]? Furthermore, both our work and the results in [11] require $O(\mathcal{C})$ communication overhead for general circuits, where $\mathcal{C}$ is the circuit size. Given that communication overhead is often a bottleneck in multiparty computation, how can we reduce this communication cost effectively?

Fusion of techniques. One can view collaborative zk-SNARKs, such as our work [25], as a fusion of zero-knowledge proofs and multiparty computation, where privacy-preserving yet efficient protocols are designed leveraging the strengths of both techniques. I believe this recipe holds significant potential for future research and could achieve unexpected outcomes in certain scenarios. Furthermore, we think this fusion can be extended to other cryptographic techniques, such as combining homomorphic encryption with zero-knowledge proofs. Exploring this direction of integration will be one of the key focuses of my future research.

Future Direction: Oblivious Proof Generation. Inspired by the above considerations and previous research questions, a worthwhile question to ask is whether the private witness can be sent to a server via homomorphic encryption, allowing the server to generate a proof using the encrypted witness. The prover could then decrypt the (ciphertext-based) proof, fulfilling an oblivious proof generation. A similar theoretical concept was proposed in [12], but no practical implementation or concrete solution exists yet.

I believe this idea promising, and my key insight is that the proof generation in existing SNARKs can be modeled as a circuit with very shallow multiplicative depth. While nowadays fully homomorphic encryption is still impractical for use, some leveled homomorphic encryption schemes [4] have been demonstrated to achieve acceptable efficiency for circuits that require only a limited number of ciphertext multiplications. Therefore, a future research direction is to validate this concept for a specific SNARK. Furthermore, since this approach avoids the communication inherent in multiparty computation, there is reason to believe that it could be more efficient than collaborative proof generation in certain scenarios.

Potential Applications. The above schemes enable the possibility of *Proof Delegation*, where a resource-constrained client can delegate the computation of a proof (for a certain price) to parties with strong computational capabilities, all while ensuring that the client's sensitive information remains private. This could give rise to a *proof market*, which inspires further research possibilities, such as how to ensure the fairness of auctions [21].

Thrust II: Secure & Fair Systems against Trusted-party

In many real-world systems, users are required to trust other (central) parties to ensure fairness and privacy. However, this trust is often misplaced, leading to various privacy breaches and unfair practices. I believe in designing specific protocols, leveraging various cryptographic techniques to get trustless systems.

Fair Data Exchange. One example scenario is fair data exchange, where a seller wants to sell data—often required to satisfy certain properties—to a buyer for a specified price. The key is to ensure the atomicity and fairness of the protocol, i.e., "payment upon delivery". Commonly, solutions rely on blockchain to facilitate the transaction. For example, a classic protocol called ZKCP [6] uses zero-knowledge proofs to ensure that the seller provides valid data and uses hash time-lock contracts (HTLC) to guarantee that the buyer pays the agreed-upon fee. In our work [24], we identified three issues with this protocol, referred to in our paper as the Eavesdropper Attack, DoS Attack, and Reputation Attack, which could severely compromise the fairness of the seller in the protocol. Finally, we proposed a new protocol, SmartZKCP (Fig. 1), which addresses these issues.

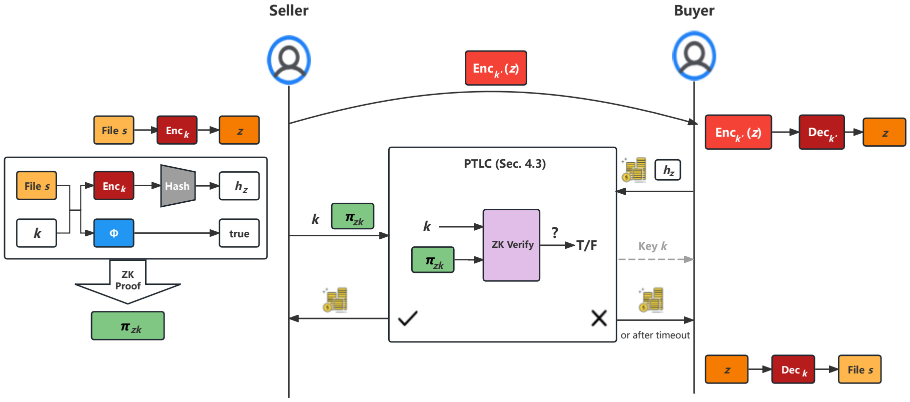

Figure 1: The SmartZKCP protocol.

In fact, there are still many open problems worth exploring in this field. For example, in the above protocol, we use smart contracts on the blockchain as an impartial third party. However, if a party colludes with miners to cheat, the protocol may fail. Therefore, an interesting question is to design a protocol that ensures fairness even in the presence of collusion between a party and miners. For instance, [37, 38] proposed some game-theory based approaches to address this issue.

Future Directions. In my future studies, I hope to explore other scenarios where fairness and privacy need to be guaranteed, believing that cryptographic techniques and other methods (e.g., game theory) can help address these challenges. For instance, some directions I am interested in exploring include private statistics [9, 32], private auctions [21], and fair transaction fee mechanisms [36].

Other Research Directions

Apart from the aforementioned areas, I have enthusiasm and knowledge on a number of other aspects of applied cryptography and their applications.

Efficiency of PCS. The efficiency of Polynomial Commitment Schemes (PCS) is crucial for SNARKs. Together with Guo et al., we construct a Reed-Solomon code-based multilinear PCS called DeepFold [18], which is highly efficient in prover time, verifier time, and proof size. This scheme is an improvement over a PCS named BaseFold [43] and is well-suited for integration with other cryptographic components, enabling applications in various scenarios.

Verifiable Machine Learning (zkML). Using zero-knowledge proofs to ensure the correctness of machine learning inference, known as zkML, is an emerging research area [44, 22]. Together with Qu et al., our work [31] achieved the first verifiable inference for the GPT-2 model, enabling succinct proof generation within 30 seconds.

Verifiable Private Information Retrieval (vPIR). Verifiable Private Information Retrieval is another area of interest. Together with Wang et al., we proposed a verifiable PIR scheme [39] named Crust, designed for practical scenarios in a two-server model, where one server is assumed to be semi-honest while the other can be arbitrarily malicious. Through some simple yet effective techniques, we achieved a highly efficient vPIR scheme that enables users to verify the server's responses without revealing the content of their queries.

References

[1] L. Babai, L. Fortnow, and C. Lund. Non-deterministic exponential time has two-prover interactive protocols. In *31st FOCS*, pages 16–25. IEEE Computer Society Press, Oct. 1990. doi: 10.1109/FSCS.1990.89520.

[2] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014. doi: 10.1109/SP.2014.36.

[3] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In R. Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, Oct. 2011. doi: 10.1109/FOCS.2011.12.

[4] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In S. Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, Jan. 2012. doi: 10.1145/2090236.2090262.

[5] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018. doi: 10.1109/SP.2018.00020.

[6] M. Campanelli, R. Gennaro, S. Goldfeder, and L. Nizzardo. Zero-knowledge contingent payments revisited: Attacks and payments for services. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM CCS 2017*, pages 229–243. ACM Press, Oct. / Nov. 2017. doi: 10.1145/3133956.3134060.

[7] B. Chen, B. Bünz, D. Boneh, and Z. Zhang. HyperPlonk: Plonk with linear-time prover and high-degree custom gates. In C. Hazay and M. Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 499–530. Springer, Cham, Apr. 2023. doi: 10.1007/978-3-031-30617-4_17.

[8] A. Chiesa, Y. Hu, M. Maller, P. Mishra, P. Vesely, and N. P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Cham, May 2020. doi: 10.1007/978-3-030-45721-1_26.

[9] H. Corrigan-Gibbs and D. Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In *14th USENIX symposium on networked systems design and implementation (NSDI 17)*, pages 259–282, 2017.

[10] A. Gabizon, Z. J. Williamson, and O. Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. URL https://eprint.iacr.org/2019/953.

[11] S. Garg, A. Goel, A. Jain, G.-V. Policharla, and S. Sekar. zkSaaS: Zero-knowledge SNARKs as a service. In J. A. Calandrino and C. Troncoso, editors, *USENIX Security 2023*, pages 4427–4444. USENIX Association, Aug. 2023.

[12] S. Garg, A. Goel, and M. Wang. How to prove statements obliviously? In L. Reyzin and D. Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 449–487. Springer, Cham, Aug. 2024. doi: 10.1007/978-3-031-68403-6_14.

[13] C. Gentry. Fully homomorphic encryption using ideal lattices. In M. Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. doi: 10.1145/1536414.1536440.

[14] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, Oct. 1986. doi: 10.1109/SFCS.1986.47.

[15] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985. doi: 10.1145/22145.22178.

[16] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum. Delegating computation: interactive proofs for muggles. In R. E. Ladner and C. Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008. doi: 10.1145/1374376.1374396.

[17] J. Groth. On the size of pairing-based non-interactive arguments. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Berlin, Heidelberg, May 2016. doi: 10.1007/978-3-662-49896-5_11.

[18] Y. Guo, X. Liu, K. Huang, W. Qu, T. Tao, and J. Zhang. DeepFold: Efficient multilinear polynomial commitment from reed-solomon code and its application to zero-knowledge proofs. Cryptology ePrint Archive, Paper 2024/1595, 2024. URL https://eprint.iacr.org/2024/1595.

[19] A. Kate, G. M. Zaverucha, and I. Goldberg. Constant-size commitments to polynomials and their applications. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Berlin, Heidelberg, Dec. 2010. doi: 10.1007/978-3-642-17373-8_11.

[20] A. E. Kosba, D. Papadopoulos, C. Papamanthou, and D. Song. MIRAGE: Succinct arguments for randomized algorithms with applications to universal zk-SNARKs. In

S. Capkun and F. Roesner, editors, *USENIX Security 2020*, pages 2129–2146. USENIX Association, Aug. 2020.

[21] A. Lazzaretti, C. Papamanthou, and I. Hishon-Rezaizadeh. Robust double auctions for resource allocation. Cryptology ePrint Archive, Paper 2024/1750, 2024. URL https://eprint.iacr.org/2024/1750.

[22] T. Liu, X. Xie, and Y. Zhang. zkCNN: Zero knowledge proofs for convolutional neural network predictions and accuracy. In G. Vigna and E. Shi, editors, *ACM CCS 2021*, pages 2968–2985. ACM Press, Nov. 2021. doi: 10.1145/3460120.3485379.

[23] T. Liu, T. Xie, J. Zhang, D. Song, and Y. Zhang. Pianist: Scalable zkRollups via fully distributed zero-knowledge proofs. In *2024 IEEE Symposium on Security and Privacy*, pages 1777–1793. IEEE Computer Society Press, May 2024. doi: 10.1109/SP54263.2024.00035.

[24] X. Liu, J. Zhang, Y. Wang, X. Yang, and X. Yang. SmartZKCP: Towards practical data exchange marketplace against active attacks. Cryptology ePrint Archive, Report 2024/941, 2024. URL https://eprint.iacr.org/2024/941.

[25] X. Liu, Z. Zhou, Y. Wang, J. He, B. Zhang, X. Yang, and J. Zhang. Scalable collaborative zk-SNARK and its application to efficient proof outsourcing. Cryptology ePrint Archive, Report 2024/940, 2024. URL https://eprint.iacr.org/2024/940.

[26] C. Lund, L. Fortnow, H. J. Karloff, and N. Nisan. Algebraic methods for interactive proof systems. In *31st FOCS*, pages 2–10. IEEE Computer Society Press, Oct. 1990. doi: 10.1109/FSCS.1990.89518.

[27] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Satoshi Nakamoto*, 2008.

[28] A. Ozdemir and D. Boneh. Experimenting with collaborative zk-SNARKs: Zero-knowledge proofs for distributed secrets. In K. R. B. Butler and K. Thomas, editors, *USENIX Security 2022*, pages 4291–4308. USENIX Association, Aug. 2022.

[29] C. Papamanthou, E. Shi, and R. Tamassia. Signatures of correct computation. In A. Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 222–242. Springer, Berlin, Heidelberg, Mar. 2013. doi: 10.1007/978-3-642-36594-2_13.

[30] B. Parno, J. Howell, C. Gentry, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013. doi: 10.1109/SP.2013.47.

[31] W. Qu, Y. Sun, X. Liu, T. Lu, Y. Guo, K. Chen, and J. Zhang. Zk-gpt: An efficient non-interactive zero-knowledge proof framework for llm inference. To be released.

[32] M. Rathee, Y. Zhang, H. Corrigan-Gibbs, and R. A. Popa. Private analytics via streaming, sketching, and silently verifiable proofs. In *2024 IEEE Symposium on Security and Privacy*, pages 3072–3090. IEEE Computer Society Press, May 2024. doi: 10.1109/SP54263.2024.00245.

[33] M. Rosenberg, T. Mopuri, H. Hafezi, I. Miers, and P. Mishra. Hekaton: Horizontally-scalable zkSNARKs via proof aggregation. Cryptology ePrint Archive, Report 2024/1208, 2024. URL https://eprint.iacr.org/2024/1208.

[34] S. Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In D. Micciancio and T. Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 704–737. Springer, Cham, Aug. 2020. doi: 10.1007/978-3-030-56877-1_25.

[35] A. Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, Nov. 1979. doi: 10.1145/359168.359176.

[36] E. Shi, H. Chung, and K. Wu. What can cryptography do for decentralized mechanism design. *arXiv preprint arXiv:2209.14462*, 2022.

[37] I. Tsabary, M. Yechieli, A. Manuskin, and I. Eyal. MAD-HTLC: Because HTLC is crazy-cheap to attack. In *2021 IEEE Symposium on Security and Privacy*, pages 1230–1248. IEEE Computer Society Press, May 2021. doi: 10.1109/SP40001.2021.00080.

[38] S. Wadhwa, J. Stoeter, F. Zhang, and K. Nayak. He-HTLC: Revisiting incentives in HTLC. In *NDSS 2023*. The Internet Society, Feb. 2023.

[39] Y. Wang, X. Liu, J. Zhang, J. Liu, and X. Yang. Crust: Verifiable and efficient private information retrieval with sublinear online time. Cryptology ePrint Archive, Report 2023/1607, 2023. URL https://eprint.iacr.org/2023/1607.

[40] H. Wu, W. Zheng, A. Chiesa, R. A. Popa, and I. Stoica. DIZK: A distributed zero knowledge proof system. In W. Enck and A. P. Felt, editors, *USENIX Security 2018*, pages 675–692. USENIX Association, Aug. 2018.

[41] T. Xie, J. Zhang, Y. Zhang, C. Papamanthou, and D. Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 733–764. Springer, Cham, Aug. 2019. doi: 10.1007/978-3-030-26954-8_24.

[42] T. Xie, J. Zhang, Z. Cheng, F. Zhang, Y. Zhang, Y. Jia, D. Boneh, and D. Song. zkBridge: Trustless cross-chain bridges made practical. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *ACM CCS 2022*, pages 3003–3017. ACM Press, Nov. 2022. doi: 10.1145/3548606.3560652.

[43] H. Zeilberger, B. Chen, and B. Fisch. BaseFold: Efficient field-agnostic polynomial commitment schemes from foldable codes. In L. Reyzin and D. Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 138–169. Springer, Cham, Aug. 2024. doi: 10.1007/978-3-031-68403-6_5.

[44] J. Zhang, Z. Fang, Y. Zhang, and D. Song. Zero knowledge proofs for decision tree predictions and accuracy. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *ACM CCS 2020*, pages 2039–2053. ACM Press, Nov. 2020. doi: 10.1145/3372297.3417278.

[45] J. Zhang, T. Xie, Y. Zhang, and D. Song. Transparent polynomial delegation and its applications to zero knowledge proof. In *2020 IEEE Symposium on Security and Privacy*, pages 859–876. IEEE Computer Society Press, May 2020. doi: 10.1109/SP40000.2020.00052.