CAREER: Pushing the Practicality of Secure Multiparty Computation

Overview

Recent advances in networking, hardware, data science, and artificial intelligence have improved our life by leveraging personal information. However, it has raised a critical need in data security and privacy with the emergence of regulations like the General Data Protection Regulation. Secure multi-party computation (MPC), which allows a set of mutually untrusted parties to compute any function jointly and reveals nothing except for the output, is crucial in enabling computation on datasets from multiple sources without sharing them. Research in the past decade has gradually pushed MPC towards practice, but vital challenges still exist in various aspects of MPC research that prevent MPC from being widely deployed. As a result, the need for practical MPC systems is still much higher than the current deployment progress.

This CAREER proposal aims to push the practicality of MPC by advancing the scientific foundation of efficient MPC protocols and software infrastructures for applying MPC in the real world. To this end, the proposed research will focus on three important aspects. First, we focus on developing maliciously secure two-party computation protocols as efficient and scalable as their semi-honest counterparts. Second, we will improve MPC efficiency when used for big data by proposing a new deployment model with huge practical potential and designing improved RAM-based MPC protocols for real-world workloads. The last research aims at developing new classes of protocols that can scale MPC to hundreds of parties and compute huge statements even when the adversary can corrupt the vast majority of the participants. This CAREER proposal fits perfectly with the PI's long-term career goal of designing practical cryptographic protocols to enhance the privacy and security of real systems and making MPC technologies accessible to educators and industry practitioners. Building on top of the PI's long history of open-source efforts, the research outcomes from this proposal will be applied and evaluated via academic collaboration and industry deployment. The project also contains educational and outreach activities to enhance MPC education for students at different stages, improve social welfare, and promote diversity of participation.

Intellectual Merit

The proposed project will significantly advance the frontier of MPC research. For malicious two-party computation, the first thrust will study two distinct paradigms, the authenticated garbling paradigm and the GMW-compiler with zero-knowledge (ZK) proofs tailored for proving non-black-box garbling, which can reduce the overhead due to malicious security to the minimal level. For MPC on big data, the second thrust will study a new deployment model that can utilize existing content distribution networks and thus provides enormous opportunities for improved performance based on variations of pseudorandom correlation generators. It will also advance the state-of-the-art RAM-based MPC protocols. Finally, for the multi-parting setting, the third thrust will study generic compilers to obtain malicious security efficiently with system-aware protocol designs; it will also lead to further research to design efficient protocols (e.g., multi-verifier ZK and multi-party distributed point function) under an all-but-a-few corruption.

Broader Impacts

The proposed project will lead to a significant impact on broader members of society. 1) The project will advance the practicality of MPC via academic research and industry collaboration to accelerate the deployment of privacy-preserving technologies. 2) The project will produce a set of high-quality educational materials that will help the researchers and practitioners at different stages to learn, deploy, and advocate MPC. 3) The project will support the PI in hosting cryptographers' mentoring workshops to enhance individual advising with an emphasis on promoting the participation of females and minorities; it will particularly focus on one-on-one interaction, something significantly lacking due to the pandemic. 4) The project will produce a solid submission to the NIST threshold cryptography standardization to accelerate the broad deployment of MPC technologies. 5) In collaboration with Northwestern Medicine, the project will design and deploy advanced MPC systems to privately aggregate health record data in the Chicago area. 6) In collaboration with Northwestern Law, the project will design, implement, and seek deployment of an MPC system that can identify repeated sex offenders in confidential settlements.

CAREER: Pushing the Practicality of Secure Multiparty Computation

1 Introduction

Recent advances in networking, hardware, data science, and artificial intelligence have improved our life by leveraging personal information. Companies harvest data from their users but are often prohibited from sharing the data with others, although data sharing could lead to even more convenience in our life. It becomes particularly challenging with recent regulations like the General Data Protection Regulation and the California Consumer Privacy Act. Can we resolve the irreconcilables between privacy and utility needs? Secure multi-party computation (MPC), which allows a set of mutually untrusted parties to compute any function jointly without revealing anything except for the output [179, 74, 25], is crucial in enabling computation on datasets from multiple sources without sharing them. It allows parties to obtain useful information without revealing anything unnecessary, thus a perfect tool to solve this dilemma.

Although MPC was initially proposed merely as a theoretical solution, recent advances have gradually pushed MPC towards practice. Prominent real-world deployment examples include Boston University and the Boston Women's Workforce Council using MPC to study the gender wage gap [1], Northwestern University and the Chicago Area Patient-Centered Outcomes Research Network using MPC to aggregate hypertension patient information [2], and numerous startups using MPC for threshold signature wallets (see [8] for a list of companies). On the other hand, the need for practical MPC systems is even more significant than the current deployment progress. Just in the current (117th) US Congress, MPC deployment has been proposed in numerous bills for secure research data networks [156], private student-loan aggregation [157], and promoting digital privacy technologies [146, 154, 155]. However, despite the recent advances in MPC and the surging needs, several vital challenges still exist that prevent MPC from being widely deployed.

1. **Challenges in efficiency.** MPC protocols provide high security but come with a high cost. Although the practicality of MPC has significantly improved in the past decade, it is still insufficient for many applications, especially under demanding settings. In particular, security against malicious adversaries, scalability to computing on big data, and support of massive participation all lead to orders of magnitude overhead in protocol efficiency.
2. **Disconnection from systems.** Much MPC research is still disconnected from systems. For example, most protocol design does not consider memory usage, cache-friendliness, or ease of parallelization. As a result, we miss tremendous opportunities to bring together system-level and protocol-level optimizations.
3. **Lack of real-world deployments.** Deployments of MPC in the real world are still rare; even among those deployed examples, they usually assume a weak security model and only scale to simple tasks and small data. Many unexpected hurdles could hamper real-world adoption. For example, building public trust, regulations restrictions, and lack of cryptography background are among the greatest challenges.

1.1 Executive Summary

With a highly integrated plan in research, education, and outreach, this proposal aims to solve the most pressing challenges in pushing MPC to practice. Research outcomes from this proposal will significantly reduce the cost of MPC for malicious security, big-data computation, and massive participants. Through industry (Google, Meta, and [redacted]) and academic (Northwestern Medicine, Northwestern Law, and Vector Institute for Artificial Intelligence) collaboration, the PI will evaluate the proposed research at the frontiers of MPC applications. In addition, the proposed education and outreach activities will provide comprehensive educational materials for students/researchers at different stages to understand and deploy MPC. Various proposed efforts will also enhance the participation of minority groups and social welfare. See Figure 1 for our main research tasks and their relation to system building, educational activities, and other broader impact activities.

Research Thrust 1: Next-generation malicious 2PC. Two-party computation (2PC) secure against malicious adversaries is one of the most fundamental problems with a long line of prior research. Unfortunately, even for the latest authenticated garbling framework initiated by the PI, the performance overhead compared to a semi-honest 2PC is still high. The PI will explore two approaches to minimize this gap: 1) designing customized silent preprocessing to reduce the triple generation overhead for authenticated garbling. 2)

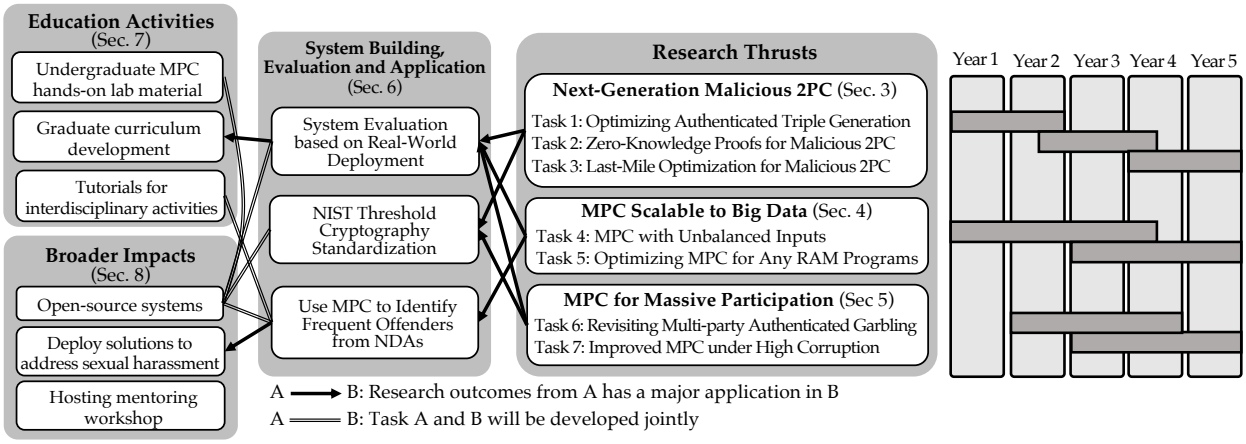

Figure 1: **Overview of this proposal.** This figure shows the main tasks in research and other activities. Their relationships are shown in arrows and double lines.

designing sublinear designated-verifier zero-knowledge (ZK) proofs to show the correctness of any garbling procedure. Both approaches bring minimum communication overhead (a few bits per gate), high computational efficiency, and circuit scalability. Finally, the PI will squeeze the last bit of the performance by designing protocols that can lead to application-specific optimizations. This is achieved by considering mix-visibility computation and malicious gadget computation for common building blocks.

Research Thrust 2: MPC scalable to big data. Secure computation on big data is a challenging task. When the data is huge, just processing the input already incurs a high cost. What's worse, many such tasks require private RAM accesses in MPC, of which performance is still not sufficient for large-scale computation even after massive progress in recent years. The PI will explore new models of computation and novel protocols to solve these issues. First, the PI will focus on a setting where a server uses the same big input for secure computation with multiple clients and will design protocols that can amortize the cost of big-data processing across multiple clients. Then, in collaboration with Google, the PI will extend the protocol for unbalanced private set intersection that can be deployed by Google Password Checkup in this setting. To optimize RAM-based MPC protocols, the PI will propose improved building blocks (e.g., better distributed point functions and read-only oblivious RAM) and distributed oblivious RAMs friendly to parallelization and batching.

Research Thrust 3: MPC for massive participation. MPC tolerating very high corruption can effectively defend against Sybil attacks even when the number of participants is huge; however, the cost increases significantly as the number of parties increases. To scale MPC with a massive number of parties tolerating all-but-one corruption, the PI will study a framework generalized from the authenticated garbling, which can compile many protocols to malicious security with high efficiency and reduced computational resource consumption. The PI will also explore security-efficiency trade-offs in a setting with all-but-a-few corruption and will design more efficient multi-party pseudorandom correlation generators and multi-verifier zero-knowledge proofs in this setting.

System Building, Evaluation, and Application. System building plays a central role in this proposal since the goal of this proposal is to push the practicality of MPC. The PI will build systems based on our research outcomes and evaluate them in real use cases (see letters of collaboration).

- The PI will work with collaborators from academia (Jennie Rogers from Northwestern, Abel Kho from Northwestern Medicine, and Nicolas Papernot from University of Toronto) and industry ([redacted] from Google, [redacted] from Meta, and [redacted] from [redacted]) to integrate and evaluate our protocols in real systems and products.
- The PI will prepare a competitive submission to NIST future standardization effort on threshold cryptography based on the research outcome in this proposal.
- The PI will explore the intersection between law and MPC with Alyson Carrel (Northwestern Law) and the PI's Ph.D. student (Peter Chan, PhD+JD joint program). We will design, develop, and seek deployment of a system that uses MPC to discover from confidential settlements important information like

frequent sex offenders who are otherwise undetectable.

The PI will use both generic metrics (like execution time) and application-dependent metrics (like usability) to evaluate our research and systems.

Education Activities and Broader Impacts. The main theme of the proposed educational efforts is to produce high-quality materials for MPC education at different levels and types of audiences. The PI will create hands-on MPC lab project materials like the SEED labs (targeting undergraduate students), a new curriculum focusing on the interplay between cryptography and systems (targeting research-focused students), and finally, video tutorials for non-cryptography researchers and industry deployment.

In addition, the PI will continue making progress in open souring MPC protocols based on research outcomes from this proposal. The PI will improve EMP-toolkit with better programmability between different security settings, richer data structures, and tools to facilitate MPC prototype over massive participation. Based on the PI's prior experience in hosting a mentoring workshop at Crypto 2021, the PI will host another mentoring workshop focusing on individual connection and mentorship. Finally, the PI, with collaborators and the PI's student in the joint JD degree, will reach out to the local lawyers and courts to explore the possibility and legal hurdles of deploying our systems for confidential settlement.

1.2 PI's Qualifications

The PI has a unique strength in bringing together theory and practice in cryptography; his prior works have significantly advanced the state of the art of many cryptographic protocols. For example, the PI has extensive experience in designing MPC with malicious security [166, 167, 165, 108, 77, 93, 176], MPC in the RAM model [171, 169, 135, 122, 163, 183, 168, 76], ZK proofs [106, 172, 175, 64, 123], and other important cryptographic building blocks [113, 40, 82, 81, 177, 41]. Furthermore, the PI has designed and implemented privacy-preserving systems for machine learning [135, 48, 173] and medical and health applications [170, 161, 62]. The PI's work on malicious MPC received a best-paper award in ACM CCS (2017); his work on designated-verifier ZK received a best-paper award runner-up in ACM CCS (2021).

The PI is also one of the leading forces in pushing novel cryptographic protocols to practice through open-source systems, most prominently, EMP-toolkit [164]. EMP-toolkit was initiated by the PI in 2016 and has been used widely in academic and industry research (examples since 2020 include [45, 39, 90, 185, 153, 91, 178, 5, 101, 84, 97]), as well as real-world systems readily deployable or being deployed. The PI's work on zero-knowledge proofs [106] is a key ingredient in the post-quantum signature scheme Picnic [184], a round-3 candidate in the NIST Post-Quantum Cryptography Standardization. Recently, a work by the PI [177] was deployed by Meta for private advertisement measurement [158].

2 Related Works

Due to the large volume of MPC research, technical works related to each research thrust will be discussed in each thrust separately.

Relationship to other ongoing projects. Among recent NSF-supported projects related to the applied aspects of MPC, [180] focuses specifically on approximation algorithms for MPC and differential privacy, while this proposal focuses on a broader range of protocols and scenarios. [129] focuses on combining TEE and MPC in a provable way, while this proposal is purely improving the practicality of cryptography-based approaches. [78] focuses on large-scale MPC protocols but mostly for the honest majority, while the related research thrust in this proposal focuses on the dishonest majority.

Relation to PI's other supports. The PI is currently being supported by the DARPA SIEVE (Securing Information for Encrypted Verification and Evaluation) project to study ZK proofs of program vulnerabilities, which leads to many of the PI's preliminary works [175, 174] for this proposal. This proposal is mainly about pushing the practicality of MPC, and ZK is used as a tool to obtain malicious security efficiently.

3 Thrust 1: Next-Generation Malicious 2PC

As one of the most canonical protocols, maliciously secure two-party computation (2PC) has attracted a lot of research effort to study its concrete efficiency since 2007 [118]. Different methods have been proposed that have reduced the protocol overhead significantly, including cut-and-choose [118, 159, 160, 115, 119, 117,

36, 11, 165] bucketing [137, 66, 96, 120, 121, 65, 138, 139, 112, 189, 188, 187], dual execution [131, 94, 95, 148], authenticated garbling [166, 87, 108, 176, 57], etc.

3.1 Task 1: Optimizing Authenticated Triple Generation for Malicious 2PC

The state-of-the-art malicious 2PC protocol is the authenticated garbling paradigm [166] proposed by the PI. In this paradigm, an authenticated garbled gate is computed with a cost similar to the semi-honest half-gate garbling plus one authenticated Beaver triple; thus, the overhead due to malicious security comes from generating authenticated Beaver triples [108]. The most concretely efficient triple generation protocol [176] needs to communicate roughly $(\frac{4\rho}{\log C}+1)\kappa$ bits per triple for a circuit with C gates, ρ bits of statistical security parameter, and κ bits of computational security parameter. This can be further reduced to around $\frac{\rho}{\log C}\kappa$ bits per triple by incorporating optimizations from the silent oblivious transfer (OT) [31, 30, 177, 49]. General pseudorandom correlation generators (PCG) [31] can generate authenticated beaver triples, but the concrete cost is very high in this context. The latest advancement [57] observed that authenticated garbling does not need perfect Beaver triples: even triples that are vulnerable to a selective failure attack from the garbler can potentially be used. Based on this observation, they designed a triple-generation protocol that only communicates $5\rho + O(1)$ bits per gate. Although an asymptotic improvement, the concrete communication cost is similar to the above under realistic parameters (e.g., $C = 2^{20}, \rho = 40, \kappa = 128$). To minimize and eventually eliminate the overhead of 2PC due to malicious security, it is crucial to further reduce this cost.

$\Longrightarrow$ **Research Challenge:** *How to design a concretely efficient protocol that generates authenticated triples with communication of $O(1)$ bits per triple?*

Recall that an authenticated bit known to P_1 works by having P_1 hold $(x, M) \in \mathbb{F}_2 \times \mathbb{F}_{2^\kappa}$ and P_2 hold $(K, \Delta_2) \in \mathbb{F}_{2^\kappa} \times \mathbb{F}_{2^\kappa}$ such that $M = K \oplus x\Delta_2$. An authenticated share consists of two authenticated bits, one known to P_1 and one known to P_2; an authenticated triple consists of three authenticated shares such that the underlying shared values have an AND relationship.

To generate triples with just a few bits of communication, we focus on authenticated triples with selective failure attacks and improve the work by Dittmer et al. [57]. In this case, the triple shares from the evaluator can essentially be compressed to a short vector of size m that is sublinear to the circuit size n. We observe that, by incorporating observations from recent optimizations for garbling schemes [91], it is possible to compute an authenticated bit under multiple global keys with a cost almost the same as authenticating to one global key. In detail, this works by exploiting the recent silent OT paradigm [31, 30, 177, 49]. These protocols usually contain multiple calls to a subprotocol called single-point correlated OT that outputs $(\vec{m}, \vec{x}) \in (\mathbb{F}_{2^\kappa}^n \times \mathbb{F}_2^n)$ to one party and $(\vec{k}, \Delta) \in (\mathbb{F}_{2^\kappa}^n, \mathbb{F}_{2^\kappa})$ to the other party, where $\vec{x}$ is a zero vector except for one location and $\vec{m} = \vec{k} \oplus \vec{x} \cdot \Delta$. To authenticate the same bit vector $\vec{x}$ over multiple global keys (say $\Delta_1, \ldots, \Delta_m$), it is wasteful to repeat the above protocol for m times naively. Instead, we can run it for once, use symmetric-key-based local expansion and then send correction words of size only $O(m\kappa)$ bits. When used for computing authenticated triples, the total communication is $O((\log n + m)\kappa)$ bits, which is sublinear in n. By extending this idea to work on authenticated values rather than the keys, we can compute the authenticated cross terms with sublinear communication.

We will design a full triple generation based on the high-level blueprint of TinyOT [136]. The above observation would enable us to compute the authenticated cross terms, while QuickSilver [175] is sufficient to prove local terms in a triple. However, it is still challenging to stitch them together without introducing selective failure attacks from the evaluator (attacks from the garbler do not hurt, as shown by Dittmer et al. [57]). We will explore the interplay between different share formats (e.g., BDOZ [26] and SPDZ [51]) and possibly design new authentication formats for efficient connections.

3.2 Task 2: Zero-Knowledge Proofs for Malicious 2PC

The above effort would bring authenticated garbling to an almost optimal cost compared to its semi-honest counterpart. However, it is unclear how to generalize this approach for the recent three-halves garbling [151] or future protocols. To this end, we will study a completely different way for malicious 2PC, specifically, the GMW compiler [74]. The GMW compiler, in summary, can compile any semi-honest protocol to malicious security by proving the correctness of all messages in zero-knowledge. Boyle et al. [33] applied it to secret-

sharing based protocols. When adapted to semi-honest 2PC based on garbled circuits, the major challenge is proving the correctness of garbling. Abascal et al. [9] explored this paradigm by combining Ligero [14] and half-gate garbling. Still, the overall running time is 1 to 2 orders of magnitude slower than the latest authenticated garbling when full AES is used for garbling.

$\Longrightarrow$ **Research Challenge:** *How to customize garbling and zero-knowledge proofs to make GMW-based malicious 2PC practical?*

Most garbled circuits work in the gate-by-gate paradigm, meaning that the computation of each gate is the same. As a result, the computation of garbling can be viewed as a single-instruction multiple-data (SIMD) circuit. In the PI's preliminary work, we designed a designated-verifier ZK for SIMD circuit [174], which has an asymptotic communication complexity of $O(n+m)$ field elements to prove a statement containing a batch of n subcircuits each of size m gates. Briefly speaking, this protocol is based on a new way to authenticate multiple values with small communication. For P_1 to authenticate a vector $\vec{x}$ to P_2, P_1 gets $M \in \mathbb{F}$ while P_2 gets $K \in \mathbb{F}$ and two global keys $(\Delta, \Omega) \in \mathbb{F}^2$, such that $M = K + f_{\vec{x}}(\Omega) \cdot \Delta$, where $f_{\vec{x}}(\cdot)$ is a polynomial interpolated over points $(i, \vec{x}[i])$. This compact authentication can be efficiently computed by P_2 sending additive-homomorphic encryption of powers of Ω. A SIMD ZK can be built on top of such authentication with high efficiency. Our preliminary work (Figure 2) shows that SIMD ZK can be highly efficient in computation and communication. This protocol can be used for the GMW compiler to achieve optimal asymptotic complexity, but without optimization, the actual cost could be high. Nevertheless, this approach is widely applicable to many garbling algorithms. The PI will study how to reduce the cost of this approach in the following directions.

Batch size	Running time (ns)	Communication (field elements)
512	56	0.0256
1024	56	0.0127
2048	56	0.0064
4096	56	0.0032

Figure 2: **The performance per gate of our preliminary SIMD-ZK over a 59-bit prime field.** Based on 16 threads and 1 Gbps bandwidth proving different batches of circuits with size 2^{20}.

1. The only non-linear operation in AES is SBox, which operates over $\mathbb{F}_{2^8}$. However, there is no efficient, if any, additive homomorphic encryption over this field. One could use a protocol over a prime field to simulate operations on $\mathbb{F}_{2^8}$, but the extra overhead would be intolerable. We will explore the noisy encoding assumption as an alternative way to compute polynomial-based compact authentication. The noisy encoding assumption has been used to construct oblivious linear evaluation (OLE) [133, 100], a closely related concept to our polynomial-based compact authentication. We will: 1) perform cryptanalysis on this assumption and possible variations (e.g., with regular noise) for better security and efficiency, 2) design a new compact authentication protocol for arbitrary fields from the noisy encoding assumption, and 3) enable homomorphic rotations to further reduce the concrete cost.
2. We will also investigate ways to reduce the circuit size for garbling a gate. The AES SBox is based on multiplicative inverse over $\mathbb{F}_{2^8}$ except for 0, which is mapped to itself. The inverse can be proven efficiently by the prover providing the inverse and proving in ZK that the multiplication with the input is 1, but it does not work for 0. A prior work [53] in the context of digital signature resolves this issue by letting the prover pick the input to AES to ensure no SBox will take zero as inputs. Although the garbler can control the inputs for some garbling schemes (e.g., the classical Yao's garbled circuit), it is not compatible with the latest schemes (e.g., half-gates [182] and three-halves [151]). In these schemes, the output wire label of a gate is a deterministic function of the input wire labels, leaving no degree of freedom to the garbler. In the proposed work, we will explore ways to express SBox verification as a low-degree polynomial without any modification. We will also explore SIMD-ZK protocols tailored for subcircuits that can be represented as low-degree polynomials.
3. The SIMD ZK protocol [174] is also proof of knowledge (PoK), which leads to some communication linear to the input size. When applied to the GMW compiler, this ZK comes with an extractor that can extract all garbled labels of all gates. The PoK property seems to be overkill because, for malicious 2PC, the overall simulator only needs to extract the inputs to 2PC. We will formulate this idea as a ZK protocol that only allows extraction of partial witness (e.g., labels of the input gates) in trade for higher efficiency.

3.3 Task 3: Last-Mile Optimization for Malicious 2PC

The first two tasks focus on improving the efficiency of malicious 2PC when the computation is a monolithic circuit, but realistic applications often have extra features that can lead to various optimizations.

$\Longrightarrow$ **Research Challenge:** *How to customize malicious 2PC for various applications with different features?*

To easily incorporate various optimizations, we will develop a unified framework that can capture various customization of the computation in 2PC. In particular, we will explore the following aspects.

1. **Visibility mix-mode circuits.** In general, 2PC ensures all internal wire values are private to both parties. Prior works [107, 19, 12] proposed protocols allowing improved efficiency if some internal wires can be made visible to one or both parties. However, these protocols are not directly compatible with the latest innovations of malicious 2PC, or proposed protocols in the first two tasks. We will explore protocols that can efficiently switch between our newly proposed malicious 2PC and various ZK protocols including VOLE-base ZK [172] and MPC-in-the-head-based ZK [99].
2. **Special gadgets in malicious 2PC.** Many customized gadgets have been proposed in the semi-honest setting that do not have efficient counterparts with malicious security, including one-hot garbling [91] (for computing outer products), arithmetic garbled circuits [18, 17] (for more efficient arithmetic computation), stacked garbling [89] (for better conditional branches), etc. We will explore ways to make them maliciously secure and more importantly, protocols to compose them with our generic malicious 2PC in this research thrust. While using ZK could provide a solution with reasonable performance, we will focus on using more efficient techniques. For example, one-hot garbling could potentially be made malicious by the consistency check in Ferret OT [177], arithmetic garbling could potentially be authenticated using information-theoretic MACs in arithmetic fields.

3.4 Evaluation Plan

We will implement all protocols mentioned in this section and make them openly available following the data management plan. We will benchmark computation, communication, and peak memory consumption efficiency. Furthermore, we will also evaluate the simplicity of the protocol, something highly relevant in practical deployment. We will compare the benchmark result with prior works as well as the state-of-the-art semi-honest 2PC to demonstrate how close we can bring down the cost of malicious 2PC.

Section 3.1 has the potential to bring down the cost difference between semi-honest 2PC and malicious 2PC to less than 10 bits. Section 3.2 can be very cheap in computation and only needs one scan of the circuit, which means the memory consumption is optimal. With further improvements in Section 3.3, outcomes from this research thrust will make malicious 2PC almost as fast as semi-honest 2PC for a wide range of hardware-network configurations.

4 Thrust 2: MPC scalable to Big Data

Computation on big data has become indispensable for enterprise-level computation and is becoming more and more accessible for individual users. In this thrust, we design new paradigms and protocols that can scale MPC to big data efficiently.

4.1 Task 4: MPC with Unbalanced Inputs

Big-data scalability is difficult and even more when it is done in MPC. Therefore, we first explore a more restricted setting, which we believe is realistic and has features to allow significant improvement. In particular, we consider a setting including a server with a huge input x and many clients, each with a small input y_i such that $|x| \gg |y_i|$. As in many X-as-a-service models, the server computes with each client so that it learns $f(x, y_i)$, where $f(\cdot, \cdot)$ depends on the provided service. This model is ubiquitous in cloud services and the workflow can be summarized in two phases (Figure 3).

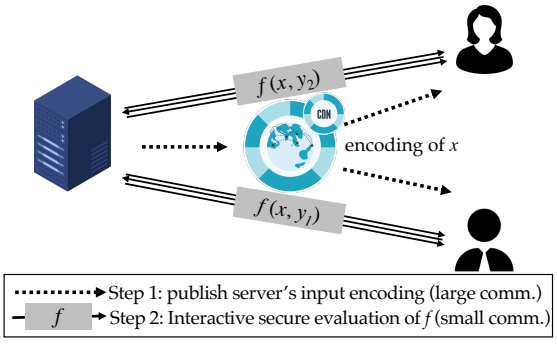

Figure 3: **Our framework for MPC with unbalanced inputs.** The server publishes an encoding of its input x; each client (with input y_i), with the encoding, uses a lightweight protocol with the server to get $f(x, y_i)$.

1. The server publishes an encoding of its large input x, which can be pre-distributed, e.g., via content distribution network (CDN) cheaply. CDN is viable because many clients download the same encoding, in which situation the network cache can accelerate the transmission.
2. For a client to use this service, it first downloads the encoding. Then a protocol between a client and the server is executed to let the client learn $f(x, y_i)$. The communication complexity after getting the encoding should be sublinear to $|x|$.

Some non-interactive secure computation protocols [98, 11, 44, 132] could theoretically work, but almost all of them focus only on feasibility and may require large communication in the second step.

$\Longrightarrow$ **Research Challenge:** *What kind of tasks have efficient protocols in the new unbalanced-input setting? How to design concretely efficient protocols in this setting?*

We first explore a use case where $f(\cdot, \cdot)$ represents the set intersection. This preliminary investigation is in collaboration with [redacted], a research scientist at Google (see letter of collaboration). State-of-the-art private set intersection (PSI) protocols [149, 67, 147, 38] all require communication linear to both sets. Although there are protocols with communication quasilinear to the small set using fully homomorphic encryption [46] or techniques from laconic OT [13], they are at least two orders of magnitude slower in computation. Bui and Couteau [38] proposed a concept called non-interactive PSI, but they require a one-time setup between two parties. As a result, the messages cannot be reused with multiple clients.

Our starting point is a recent PSI protocol based on vector OLE (VOLE) [149]. In this protocol, VOLE is used as a building block for an oblivious PRF from which there are well-known methods to build a PSI [111, 143]. In more detail, two parties need to run a VOLE, such that the receiver uses an encoding of its set as the input. After this chosen-input VOLE, the rest part of the protocol only needs to communicate linear to the other set size. This is a perfect candidate to start with since now the task boils down to constructing a VOLE protocol where the first message encodes the receiver's input.

The most efficient VOLE protocol with malicious security is by the PI [172] and follows the pseudorandom correlation generator paradigm [29, 31, 30, 32], which is based on the learning parity with noise (LPN) assumption. Under a dual-LPN assumption, we essentially need a VOLE where the simulator can extract the LPN noise from the first message. Because the amount of noise is sublinear, it can be very efficient when the input is large. In existing VOLE protocols, the value of the noise is generated from another base-VOLE protocol, which suffers from a selective failure attack but can be eliminated by random linear combination due to the leftover hash lemma [172], similar to MASCOT [109]. However, this process is interactive and thus the first message can not encode the noise value which also depends on the randomness in the linear combination. A selective failure attack would normally be fine under a stronger LPN with leakage assumption, but in our setting, the attacker could control a set of clients and thus gradually learn all the noise values via a Sybil attack. One potential idea is to regularly refresh the encoding once a sufficient number of selective failure attacks are observed but it is still vulnerable to denial of service attacks. Built on top of this preliminary idea, the PI will explore the following aspects.

1. To prevent the selective failure attack in base-VOLE, the PI will seek alternative approaches for computing chosen-input VOLE. Inspired by recent VOLE-based ZK, we propose to run the random-input VOLE interactively (so that selective failure attack can be eliminated without leaking information from the encoding) and then let the server fix the random input to the encoded input. To prevent the server from using inconsistent values, we let the server include a commitment of the all noise values and prove in zero-knowledge that they are consistent with the VOLE values after the fix. We believe that an additive-homomorphic commitment should be helpful, but most of them are not extractable, let alone UC-secure. Lysyanskaya [125] proposed a sigma protocol for Pedersen commitment but it is not clear how to make it UC-secure efficiently. The PI will explore how to perform such a consistency check; we will also design a UC-secure NIZK to prove the knowledge of the committed value for commonly used additive-homomorphic commitment schemes. Due to the flexibility of the MPC-in-the-head approach [99, 106], we believe it has the potential to work for both lattice-based commitment and Pedersen commitment.
2. Collaborating with Google, we will evaluate the practical application requirements of Google Password Checkup (GPC) [144], which allows any user to check if their passwords have been compromised without revealing the whole database. Our preliminary discussion with relevant teams from Google is promising:

the overall workflows works for GPC but tweaks are needed to 1) only allow the clients to learn the results and 2) explore leakage-efficiency trade-off to reduce the amount of encoding that the clients needs to download. We will explore these adaptation to our protocol to match the workflow and leakage model currently used by GPC.

3. The PI will extend the protocol beyond set intersection. In more detail, we will explore 1) how to allow the server updating the set without recomputing everything, 2) how to support more complicated set operations; and 3) we will explore how to extend the idea beyond set operations. For the first two directions, we will explore how to put dynamic set structures in the public encodings as well as how to make connections, e.g., via ZK, between encodings before and after the update. For the third direction, we will particularly focus on linear-algebra operations because there are also VOLE-based protocols. We envision that this could lead to a protocol where the server can publish encodings of an ML model and perform private inference with small communication.

4.2 Task 5: Optimizing MPC for Any RAM Programs

MPC in the RAM model [141, 75] allows secure array accesses with cost sublinear to the array size. While this line of work initially focused on optimizing oblivious RAM (ORAM) [71, 163, 169, 110, 63] and oblivious data structures [181, 171, 27, 130, 59, 135, 140] for MPC, recent works [183, 60, 92, 142] started designing customized protocols to aim for even higher efficiency. To scale MPC to large inputs for any program, it is crucial to improve the efficiency of MPC protocols in the RAM model.

$\Longrightarrow$ **Research Challenge:** *How to improve the efficiency of state-of-the-art RAM-model MPC protocols to be capable of handling real-world workload?*

For an array of size 4096, the current state-of-the-art RAM-based protocol [60] requires about 100KB communication and three rounds of interaction for each access. On the other hand, Garbled RAM enables constant-round RAM-based protocol, and the most efficient protocols [92, 142] still require about 4MB of communication per access. For this proposal, we will focus on non-constant-round RAM-based protocols that provide more design flexibility and are more efficient under low-latency networks.

Improving two-party distributed point function. The two-party distributed point function (DPF) takes secret shares of an index $i \in \{0,1\}^{\log n}$ as input and outputs two Boolean vectors $x_1, x_2 \in \mathbb{F}_2^n$, one to each party, such that $x_1 \oplus x_2$ is a zero vector except for the i-th location. This is an important building block for recent RAM-based protocols [60, 76], and its efficiency is crucial since it is needed for every RAM access. At the heart of the DPF protocol is the GGM tree [73] used as puncture PRF [28], which is constructed using a length doubling PRG. The PI's prior work [82] shows a length-doubling PRG in the random permutation model, which can efficient instantiated via fixed-key AES [23]. In our preliminary exploration, we observe that it is possible to further improve the efficiency by instantiating the PRG as $G(x) = H \| H \oplus x$, where $H = \pi(\sigma(x)) \oplus \sigma(x)$ and $\sigma(x)$ is a linear orthomorphism. Although this expansion function is no longer a PRG, we believe the whole protocol is still secure because of the exact use case of DPF, where only the index i needs to be protected. This construction requires only half of the computation and can still enjoy provably secure instruction-level optimization. We will formally prove its security and design a compatible distributed DPF-key generation. We will also explore its applications to accelerate other primitives as well.

Improved read-only oblivious RAM. An important intermediate milestone for RAM-based MPC is to build an oblivious-read, public-write protocol. As an overview, two parties hold secret sharing of a secret key, and both have encryptions of an array of blocks. Two parties, holding secret shares of the index i, use the above DPF-based protocol to privately select and obtain secret sharing of the encryption of the i-th block. To obtain the secret sharing of the content, previous work [60] performs private decryption inside of 2PC (and thus non-black-box use of the encryption scheme), which leads to high overhead. We will study the use of lattice-based encryption schemes to improve efficiency. A useful feature of many lattice-based encryptions [35, 34, 72] is that the decryption algorithm is linear to the secret key and thus can be potentially used for cheap homomorphic decryption [47]. However, homomorphic decryption requires that the ciphertext is known to both parties but the revealing ciphertext will leak the index. Therefore, we will design a secure ciphertext refresh protocol so that the updated ciphertext can be revealed without revealing the index. This would avoid heavy decryption in 2PC and thus significantly improve the efficiency.

Reusable ORAM Initialization. A highly desirable feature would be allowing $f(\cdot,\cdot)$ in Section 4.1 to be a RAM program. This would require a reusable ORAM initialization. Existing ORAMs do not support this because randomness used during the initialization will be partially revealed during the access (e.g., in tree-based ORAMs, the path of the block being accessed will be revealed), and thus the protocol is not secure if the same ORAM initialization is used twice. The only exception is PI's prior work [76], where the randomness in the initialization is never revealed. We will combine this idea with the research from Section 4.1 for an efficient instantiation. We will also explore applications, e.g., private database queries on committed databases, secure decision-tree learning on committed ML, where the initialization of databases/ML models only needs to be done only once.

Customizing ORAM algorithm for MPC. A major drawback of the current fastest RAM protocol [60] is its poor asymptotic complexity (square-root rather than polylogarithmic). It particular impacts the performance when the array contains millions or even billions of elements. To improve the efficiency especially for very large arrays, we will focus on improving a prior work by the PI [76], that also uses the building blocks we propose to improve. [76] uses a tree-based ORAM but avoids the recursive structure, a $O(\log n)$ factor improvements for n-sized arrays. Due to the latest improvements in OT [177, 49], secret-share based 2PC uses $50\times$ less communication than garbled circuits but requires more round trips. We will customize the ORAM protocol to perform more parallelizable batch evictions with the goal to reduce the frequency of eviction and the depth of the eviction circuit. We will particularly explore to seek a combination of Onion ORAM [56] and [76].

4.3 Evaluation Plan

We will implement all protocols proposed in this section and benchmark their performance in realistic configurations. Since scalability to big data is the main focus, the evaluation will focus on how the protocol performs on large datasets. For the protocols in Section 4.1, we will evaluate the encoding time as well as the subsequent interaction. Since the interaction is between the server with multiple clients, we will also evaluate the server throughput. We anticipate at least two orders of magnitude improvement compared to prior works when the big set is much larger than the small set. For protocols in Section 4.2, we will pay extra attention to the cost of each building block since they are of independent interest. We will evaluate the performance using common RAM programs, including binary search, stable marriage, shortest path, etc. Here, we anticipate one to two orders of magnitude improvement, especially when the array size is large and taking reusable initialization into consideration.

5 Thrust 3: MPC for Massive Participation

While the previous two thrusts primarily focus on a two-party setting, the focus here is the multi-party setting. In particular, we will explore protocols that can scale to hundreds to thousands of participants efficiently. At this scale, honest-majority protocols have been improved with high efficiency in recent years [79, 80]. Still, a stronger dishonest majority is often desired to prevent the adversary from performing Sybil attacks and pretending to be multiple participants. This thrust focuses on MPC protocol secure in a highly adversarial corruption model: given n parties, we assume that up to $t = n - c$ parties could be corrupted, where c is a small constant. When $c = 1$, it corresponds to the all-but-one corruption model that has been long studied (starting from [136, 51, 26, 74, 22]).

5.1 Task 6: Revisiting Multi-Party Authenticated Garbling

For MPC with all-but-one corruption, authenticated garbling is one of the most scalable approaches for general circuits, but even for this protocol, the computation cannot scale to large circuits when the number of parties is huge. When computing between 128 parties, the largest circuit only has less than 100K gates. Prior works [116, 186] attempted to address it from a system perspective, but to bring scalability to a new level, we also need novel designs of protocols.

$\Longrightarrow$ **Research Challenge:** *How to scale MPC for a massive number of participants with all-but-one corruption at the same time?*

Authenticated garbling can be viewed as authenticating a particular semi-honest protocol efficiently. We

first generalize it to a broader range of protocols. Suppose we are given an MPC protocol (namely Π_1) and an authenticated-triple-generation protocol (namely Π_2), both of which are maliciously secure up to additive attacks [69]. To securely evaluate function f, we generalize authenticated garbling as below.

1. Use Π_1 on f to obtain a secret share of each wire value in the circuit and then authenticate them.
2. Use Π_2 to generate ℓ number of authenticated triples that are secure up to additive errors.
3. Use the authenticated triples to check against the authenticated wire sharings via a probabilistic triple-checking process. Reveal the output if the checking passes.

Protocols secure up to additive attacks were first studied for feasibility results [69, 68, 70]. More recently, Boyle et al. [33] proposed an improved compiler, but it assumes a trusted dealer which, when instantiated using a distributed protocol, requires overhead linear to the circuit size. On the other hand, our compiler uses efficiently instantiable protocols and is geared specifically towards concrete efficiency. In the proposed project, the PI will focus on instantiating this compiler for scalable MPC.

1. **Distributed garbling of gadgets.** we can use distributed garbling as Π_1 but even the latest works [24, 167] require a four-row garbling with free-XOR [114] and point-and-permute [22]. In general, garbled row-reduction (GRR) [134] appears to be incompatible to the distributed setting because the output labels need to be picked by the parties. The PI will seek to bypass this difficulty by designing distributed garbling of gadgets as a whole. When we go beyond the gate-by-gate paradigm, we could potentially apply GRR to intermediate wires in the gadget and thus reduce the garbling size in the distributed setting. We will incorporate prior works on garbling formulas and gadgets [18, 17] in the central setting.
2. **Protocol switch under additive attacks.** Our compiler does not limit Π_1 to any specific type of protocol and thus can potentially be applied to mixed protocols, e.g., ones combining distributed garbling and GMW-type protocols based on secret shares. Although both types of protocol are secure up to additive attacks, it is unclear if the conversion is. The PI will explore efficient conversions of different types of protocols with security up to additive attacks.
3. **Reducing memory overhead.** Memory overhead in authenticated garbling is caused by the BDOZ sharings, which occupy $O(n\kappa)$ bits to authenticate one bit. Our compiler is not tied with any particular format of authentication and thus could potentially benefit from improved representation. Immediately, by using SPDZ shares, one could reduce the memory overhead of $O(\kappa)$. The PI will explore more authentication methods that can further reduce the memory overhead. For example, we will further explore authentication based on MiniMac [50, 52] and reverse multiplication-friendly embeddings [42, 43] (which has been shown to be practical in MPC-in-the-head signatures [105]) and how to generalize them to the multi-party setting.
4. **Improving cache/memory access locality.** The probabilistic triple-checking procedure [136], even after optimization [15], requires randomly shuffling all triples. This procedure is inherently cache-inefficient because accesses are random over a large range, which becomes a big overhead. These random-permutation-based approaches use significantly more entropy than the provided statistical security, which means that a simpler probabilistic checking could exist. Indeed, we believe an approach based on cyclic shifting could provide ρ bits of statistical security while using ρ bits of randomness. The PI will formally prove its soundness and also the level of cache friendliness. We will study the normal setting where the cache size is known to the protocol as a parameter and also a more general case where the protocol does not need to know the cache size (a.k.a. cache-obliviousness).

5.2 Task 7: Improved MPC under High Corruption

When the number of participants is massive, it could be reasonable to assume a small constant number of honest parties as a trade-off between efficiency and security. Protocol in this security model was studied by Hazay et al. [86, 85], who observed that when there is more than one honest party, the protocol can be designed so that one only needs to ensure that the amount of entropy contributed by all honest parties is high enough. This often leads to short keys, but the efficiency is no longer competitive considering the latest improvements in OT [31, 30, 177, 49].

$\Longrightarrow$ **Research Challenge:** *How to utilize a constant number of honest parties to improve the efficiency of MPC with high corruption?*

Pseudorandom correlation generator for massive participants. In the preprocessing model, MPC is fairly efficient even under a massive amount of participants [51], but the instantiation of the preprocessing could be inefficient. To improve the overall efficiency, the PI will study efficient multi-party PCG under high corruption. PCG-style protocols [29, 31, 32] have led to great efficiency improvements in the two-party settings. In the multi-party setting, the only concretely efficient work [10] only generates a special type of correlation for threshold signature and supports up to three parties, far from supporting massive participants. The most important building block in PCG is distributed point function (DPF), which only has efficient construction between two and three parties (with two corruptions). The PI will explore DPF protocols in the multi-party setting with more than one honest party. We will particularly explore DPF protocols that can benefit from observations from TinyKeys. The PI will also explore applying the TinyKey intuition to reduce the amount of LPN secrets in PCG protocols. We will design protocols to distribute the LPN secret among honest parties to improve the efficiency of popular building blocks from pseudorandom correlation generators. Recall that the state-of-the-art multi-party PCG [31] works by decomposing a constant-degree relationship (e.g., beaver triple is a bilinear relationship) into multiple pairwise two-party relationships; two-party relationships can then be computed via PCG protocols where parties can specify inputs. By ensuring that LPN secret from all honest parties has sufficient entropy, each party can use less noise, and the saving increases when the number of honest parties increases.

Multi-verifier ZK under high corruption. Generally speaking, a multi-verifier zero-knowledge (MVZK) protocol allows a prover to prove to a set of verifiers. By allowing it to be potentially interactive, the protocol could be more efficient. MVZK is very useful in malicious MPC: it can be used to validate parties' inputs without full-fledged MPC and could also be used to prove parties' local computation, e.g., for lifting the security from semi-honest to malicious security. In PI's preliminary work, an MVZK protocol was designed assuming an honest majority of verifiers. The protocol can be made information-theoretic or constant-round in the random oracle model. We will explore different methods to design MVZK in the dishonest majority setting focusing particularly on ones from the MPC-in-the-head approach [99]. MPC-in-the-head is a popular approach for constructing non-interactive zero-knowledge (NIZK), which is more flexible to use than MVZK. However, if only an MVZK is needed, then we could potentially design better protocols.

5.3 Evaluation Plan

The implementation and evaluation of research results from this thrust will focus on settings with very large number of parties. In the multi-party setting, we will focus on the total communication complexity as well as the maximum communication that a party needs to send. While the first metric is reflective of the monetary cost, the second metric represents the running time better. We will also focus on the scalability of the protocol, in particular, the amount of resources needed to double the number of participants and the size of the computation. In collaboration with Nicolas Papernot (University of Toronto, see letter of collaboration), we will also benchmark their performance in secure distributed and robust machine learning tasks. We expect that the improved malicious protocol from Section 5.1 could improve the running time and scalability by at least one order of magnitude and even more under limited hardware resources. Results from Section 5.2 would lead to significant improvement in asymptotic complexity and concrete efficiency improvement when the number of parties is very high.

6 System Building and Deployment

6.1 System Evaluation based on Real-World Deployment

We will implement protocols from our research thrusts and evaluate their practicality based on real-world deployments with our collaborators in academia and industry. In addition to the collaboration on Google Password Checkup described in Section 4.1, we describe three more deployment scenarios.

Federated relational databases. The PI has been collaborating with Jennie Rogers (associate professor at Northwestern University) and Abel Kho (professor at Northwestern Medicine) in deploying efficient federated relational databases that use MPC to execute SQL queries over multiple data owners securely. Our current implementation uses EMP-toolkit as the backend. In our preliminary effort [150], we have successfully deployed an end-to-end MPC system between Northwestern Medicine, Rush University Medical

Center, and Alliance Chicago, covering more than 500,000 patients' health record data. The PI will continue this collaborative effort to deploy research outcomes from this proposal for faster federated relational databases (See letters of collaboration). In particular, the research results from Section 4 can significantly improve the efficiency of important data structures when running SQL in MPC. We will implement, integrate, and evaluate our RT2 protocols in this context. Implementation and evaluation within the scope of this project do not involve any human subject research or personal information.

Very-large-scale MPC. The PI has been collaborating with [redacted], a research scientist at Meta, to deploy the PI's research in practical deployment. So far, they have built a semi-honest system on top of our recent work on efficient OT [177] as part of the Facebook Private Computation Framework (FBPCF) for privately computing the effectiveness of advertising [158]. In our future collaborations (letter of collaboration included), we will integrate improvements from Section 4.2 into FBPCF for improved performance. Furthermore, we also plan to deploy malicious protocols designed in Section 3 for enhanced security.

Efficient Decentralized Oracle Networks. Efficient malicious 2PC and designated-verifier ZK protocols also have applications to bring digital data to decentralized oracle networks. In particular, they enable decentralized oracle services for TLS sessions [185] and digital identity [126]. The PI has been collaborating with [redacted] and deployed a version of the DECO system [185] based on the PI's preliminary works [172, 175] using EMP-toolkit. The PI will work with [redacted] to deploy new protocols from this proposal, in particular, ones from Section 3.2, to the decentralized oracle networks hosted by [redacted]. See letter of collaboration from [redacted], [redacted] [redacted].

6.2 Submitting to NIST Standardization of Threshold Cryptography

Standardization is a crucial step in pushing MPC to practice as it can increase public confidence in MPC technologies. The PI participated in the NIST standardization of post-quantum cryptography as a member of the Picnic team. The submission advanced to the third round of the standardization process and has inspired a long line of research [53, 21, 104, 103, 102, 20, 83, 54, 58]. NIST is currently planning on standardizing threshold cryptography, where MPC plays a key role in executing public-key and symmetric-key operations over multiple parties [37]. Results from this proposal (Section 3 and Section 5) will enable us to make a strong submission with outstanding efficiency and security. Based on the most recent report on the roadmap of this standardization [37], we plan to make a submission to the multi-party track for computing symmetric-key primitives and explore potential extensions to support public-key primitives. To promote open research and deployment, the PI is committed to not filing patents on technologies related to this standardization.

The NIST standardization process involves high scrutiny in protocol design and system implementation. We will make extra effort to develop reference and optimized implementations to meet the requirements, notably making sure the code is const-time. We will make our implementation available across multiple architectures, including x86 and ARM architectures like Cortex-M.

6.3 Use MPC to Identify Frequent Offenders from Confidential Settlements

The PI is collaborating with Prof. Alyson Carrel from Northwestern Law School (See letter of collaboration) using MPC to solve dilemmas in confidential settlements like non-disclosure agreements (NDA). Professor Carrel has extensive experience in judicial negotiation and mediation.

As some background, mediation and settlement outside of the court date back to 1976 [152]. Since then, the debate has never stopped on its impact (See [127] and references therein). On the one hand, it provides justice to the victims: it allows individual autonomy over the outcome of the case and provides a voice for those who may otherwise not be able to use the court systems due to limited financial resources. On the other hand, it has also received criticism that it could harm society as criminals can get away undetected. The recent revelations of sexual harassment from celebrities have demonstrated this debate concretely. We have witnessed multiple occurrences where victims risk breaking NDAs to speak publicly about repeated sex offenders. These victims are normally restricted from speaking publicly due to the NDAs they signed. From this perspective, NDAs prevent sex offenders from being identified in the early stages. However, NDAs indeed provide the victims with some justice and interests, especially when there is a huge imbalance in social status between the offenders and the victims.

We plan to use MPC to solve such a dilemma. In particular, we will use an MPC protocol executed among a set of law firms, where they submit the NDAs that they executed. The MPC protocol can output, e.g., a list of offenders that appears in more than a certain number of cases. Although the use of MPC technology seems simplistic, a lot of practical factors must be considered, which in turn calls for new requirements for MPC. Callisto [145] and SATE [16] are also cryptography-enabled systems for detecting serial sexual predators. However, our emphasis here is not on individual reporting but on identification from signed NDAs. The goal of our system is complementary to them: NDAs might prevent the victim from using their systems. Our preliminary study discovered new problems that we will study and overcome in this proposal.

1. **Incentives.** Law firms representing the interest of the defendants (that is, those who are accused of sexual harassment) may not be willing to participate in an MPC execution. Based on our preliminary discussion with attorneys in the local area, we will add more features to motivate participation better. For example, our preliminary exploration shows that we could also identify frequent plaintiffs who appear in many NDAs, which could potentially be linked to fraudulent accusations. This is something useful to the defendants, particularly those with high social status. Other long-term solutions include making the validity of NDA dependent on the frequency of the committed sexual harassment by the defendants. We will explore their legal implications in this proposal.
2. **MPC with enhanced properties.** Due to the sensitivity of the execution, we may consider extra properties, including guaranteed output delivery, public auditability, computation on committed inputs, and so on. We will develop these add-ons when needed.

The PI is advising a graduate student, Peter Chan, who is currently in the PhD+JD program at Northwestern University (details see Section 7). With knowledge in both cryptography and law, Peter will be partially supported by this proposal to bridge the effort made by the two fields. Unlike other directions in this proposal, the MPC system here will be used by people in law firms without any background in computer science. Therefore, we will also develop a web-based solution, aiming to bring the whole code infrastructure to popular browsers. System design and building in this project do not involve any human subjects or sensitive information.

6.4 Evaluation Plan

The ultimate metric of success of activities in this section is the actual deployment of systems running on real data on a regular basis. The development of all protocols will be evaluated based on real-world scenarios. Evaluation metrics include the concrete security of the implemented protocol, end-to-end execution time, and simplicity. We will benchmark the systems in an isolated environment as well as in practical settings with potentially unreliable networks. For the effort toward NIST standardization, we will also evaluate the protocol based on other metrics, including the number of lines of code, memory consumption, and energy consumption. For applying MPC for private NDAs, we will additionally evaluate the usability of our system from the perspective of users without cryptography backgrounds.

7 Integrated Education Plan

The PI's education plan is highly integrated with the proposed research and outreach activities in Section 8. The main goal is to produce high-quality materials for every stage of MPC education and deployment.

Task 1: Undergraduate hands-on lab project on MPC. In the past years, the PI has developed two new courses at Northwestern: introduction to cryptography (undergraduate) and advanced cryptographic protocols (graduate). In the undergraduate curriculum, the PI allocated one lecture to introduce the concept of MPC with the help of the five-card trick for privately computing an AND gate [55]. Poker cards were distributed to the students so they could try the protocol in class.

To prepare for a future with blooming cryptographic applications, the PI will design a set of hands-on lab projects to expose students to cryptographic protocols and systems. The PI will use the SEED Labs [4] as a role model and make the materials widely usable by other institutions for free. The PI will start with the EMP-toolkit for MPC protocols and extend to other libraries (e.g., SCALE-MAMBA [3]) and protocols (e.g., zkSNARK, private information retrieval). The goal is twofold. 1) By letting our undergraduate students develop "toy" applications, we foster a better environment for students to think about start-up ideas to build

useful applications based on MPC and other cryptographic tools. 2) By allowing students only with basic backgrounds in cryptography to use our system, we will be able to evaluate the usability of our framework and understand how to improve our systems for a broader deployment.

Task 2: New curriculum development. The PI will develop a new graduate-level course focusing on the interplay between applied cryptography and system building. While the previous task focuses on the education of *basic concepts of MPC and how to use MPC as a black box*, this task is about teaching students how to *design and implement* new cryptographic systems. The design perspective will cover examples of the real-world security requirement in complex systems and principled trade-offs between security, simplicity, and efficiency. From the system perspective, it will cover principles to implement cryptographic protocols efficiently on various CPU/memory architectures and write to implement constant-time programs efficiently. The PI will invite industry collaborators and other researchers with MPC deployment experiences to give talks. All curriculum materials will be made publicly available.

Task 3: Videos and tutorials for interdisciplinary activities. In the PI's effort to apply MPC for confidential settlements and NDA, we will need to explain the concept and workflow of MPC to lawyers and judges. Therefore, the PI will make short videos and examples to facilitate the explanation to people without a background in computer science. We will also make a set of tutorials on using our open-source framework to enable rapid development of MPC-enabled prototypes by programmers without any cryptography background. The development of these tutorials will go hand-in-hand with the hands-on lab project mentioned above with different emphases.

Task 4: Student advising. The PI believes every student can excel with their unique strength and has been sticking to this philosophy when advising students. Chenkai Weng has an extensive background in system building and has worked on practical MPC/ZK protocols with fruitful results. Nicholas Franzese has done research in machine learning before and is leading PI's collaboration between cryptography and machine learning. Finally, Peter Chan is a PhD-JD student joint with the Northwestern Law School. Peter's expertise in judiciary and cryptography gives him a unique angle on pushing MPC to practical applications in the legal domain, especially the research collaboration described in Section 6.3.

The PI has been very active in promoting the participation of female and minority students at Northwestern. The PI serves on the graduate recruitment committee of the CS department and is actively promoting the hiring of female and minority students. Indeed, by the 2022 fall, three out of seven Ph.D. students in the PI's group are either female students or from gender and sexual minority communities. With their help, the PI will continue promoting academic and research participation from female and minority communities.

Evaluation Plan. Activities in the educational plan will be evaluated based on feedback from students, developers, and lawyers. The lab materials will also be evaluated by the quality of students' project submissions. The PI plan to iteratively improve the quality by incorporating their feedback. Concrete success metrics include the number of individual downloads of the MPC educational materials in the public domain and the number of other universities adopting our MPC lab projects.

8 Broader Impacts

In addition to the above educational impacts, the PI commits to the following activities to bring broader benefits to the research community and the general public.

Task 1: Open-source systems. The main objective of this proposal is to push MPC to practice; thus, trustful system building is indispensable. As the maintainer of EMP-toolkit [164], the PI has a strong record of system building and open sourcing (also see Section 1.2). The quality and usability of EMP are being carefully maintained. One industry collaborator has hired third-party security auditing on the backbone repositories of EMP for its deployment. The PI will continue the open-source effort to make it easier to deploy MPC in the real world. Our open-source plan will be coordinated with other parts of this proposal, including system building and benchmarking of main research thrusts, hands-on lab projects for MPC (Section 7), and a tool for our academic and industry collaboration (Section 6). The PI plan to accomplish the following open-source milestones, each corresponding to a major research thrust.

- **Formalizing EMP intermediate representation.** To enable fast deployment of malicious 2PC in real products, the PI will formalize an EMP intermediate representation compatible with both semi-honest and

malicious protocols. The goal is to allow the deployment of semi-honest protocols at first and a seamless transition to malicious security without rewriting the application.

- **Flexible RAM structures.** We will build a library of advanced secure data structures in the RAM model based on the research from Section 4.2. To make our library easy to use, we will follow the conventions from the C++ Standard Template Library (STL) so that programmers' prior experience can be brought to MPC programming seamlessly.
- **Deployment tools for large-scale MPC.** One pain point in experimenting with large-scale MPC is to test it on many machines. When it comes to hundreds of machines, manually setting up each cloud server is not feasible. The PI will develop tools and examples for deploying MPC with many parties over the cloud using popular toolchains link Terraform.

Task 2: Deploy solutions to address sexual harassment. The proposed research collaboration described in Section 6.3 will produce a system that uses MPC to enhance private settlement, which has a major application in protecting victims of sexual harassment. To make impacts on real matters, the PI plans to make real runs and potentially deploy this MPC-based solution. The main goal is to protect individual victims' interests and rights that come with NDA while at the same time protecting social justice by identifying frequent sex offenders who can do huge harm to our society. With the help of local courts, the PI and the collaborators will work with local law firms to demonstrate its potential impact and promote the system to law practitioners. We believe our tool has the potential to significantly change the status quo on how sexual harassment settlements are handled and thus make a long-term impact on our society.

Task 3: Host mentoring workshops. The PI co-hosted a mentoring workshop at Crypto 2021 with Peihan Miao (assistant professor at Brown) and Tal Rabin (professor at UPenn). Although there have been mentoring workshops in security [6] and program language [7], there was no mentoring workshop specifically in cryptography before our effort. The workshop consists of three panels for sharing experiences between junior faculties, new students, and female researchers. As part of this workshop, we also interviewed more than 40 cryptographers all over the world and made a list of 24 videos [128] that cover topics on research, academic career, work-life balance, and challenges to female researchers. At the time of submission of this proposal, these videos have attracted more than 9000 views and 800 hours of cumulative watch time.

The PI will host another mentoring workshop emphasizing one-on-one communication between junior students and senior researchers. This is particularly beneficial to junior students who have missed much in-person interaction with other researchers due to COVID. We plan to adopt the individual mentoring model by iMentor [6], where each senior researcher is paired with one or two junior students and is responsible for guiding them throughout the conference.

Evaluation Plan. Our activities will be evaluated based on the feedback from the open-source community, participants of our mentoring workshops, and the lawyers specializing in NDAs. The success of the plan will be measured by the amount and the depth of adoption of our open-sourced code, the number of people participating in our advising activities, and the progress in deploying our system against sexual harassment.

9 Results from Prior NSF Support

The PI is a co-PI of the NSF award #2016240, entitled *"Collaborative Research: SaTC: CORE: Medium: Quicksilver: A Write-oriented, Private, Outsourced Database Management System."* (10/01/2020-09/30/2024; $599,993). **Intellectual Merit**: This project is supporting the PI to explore secure outsourcing of database computation without trusting the outsourcing servers. Work supported by this award has resulted in conference publications in database [150], cryptography [175, 64, 124, 123], and security [162, 61]; among them, [175] was a CCS best-paper runner-up. **Broader Impacts**: This project has supported the PIs' effort to 1) deploy a real-life MPC clinical trial between three hospitals in the Chicago area [150], 2) host a tutorial session in SIGMOD 2021 on topics at the intersection of security and database [88], and 3) host a mentoring workshop at Crypto 2021 with an emphasis on helping junior and female researchers [128].

References

[1] Boston Women's Workforce Council. https://thebwwc.org/mpc, Accessed in June 2022.

[2] Chicago Area Patient-Centered Outcomes Research Network. https://www.capricorncdrn.org/projects/vault-db/, Accessed in June 2022.

[3] Scale and mamba. https://github.com/KULeuven-COSIC/SCALE-MAMBA, access in June 2022.

[4] Seed security labs. https://seedsecuritylabs.org accessed in June 2022.

[5] Cheetah: Lean and fast secure Two-Party deep neural network inference. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, August 2022. USENIX Association.

[6] iMentor workshop at CCS 2022, 2022. https://sites.google.com/vt.edu/imentor.

[7] Programming Language Mentoring Workshop at PLDI 2022, 2022. https://pldi22.sigplan.org/track/PLMW-PLDI-2022.

[8] Mpc alliance, Accessed in July, 2020. https://www.mpcalliance.org/.

[9] Jackson Abascal, Mohammad Hossein Faghihi Sereshgi, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Is the classical GMW paradigm practical? The case of non-interactive actively secure 2PC. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1591–1605, Virtual Event, USA, November 9–13, 2020. ACM Press.

[10] Damiano Abram, Ariel Nof, Claudio Orlandi, Peter Scholl, and Omer Shlomovits. Low-bandwidth threshold ECDSA via pseudorandom correlation generators. Cryptology ePrint Archive, Report 2021/1587, 2021. https://eprint.iacr.org/2021/1587.

[11] Arash Afshar, Payman Mohassel, Benny Pinkas, and Ben Riva. Non-interactive secure computation based on cut-and-choose. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 387–404, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.

[12] Arash Afshar, Payman Mohassel, and Mike Rosulek. Efficient maliciously secure two party computation for mixed programs. Cryptology ePrint Archive, Report 2017/062, 2017. https://eprint.iacr.org/2017/062.

[13] Navid Alamati, Pedro Branco, Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Sihang Pu. Laconic private set intersection and applications. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part III*, volume 13044 of *LNCS*, pages 94–125, Raleigh, NC, USA, November 8–11, 2021. Springer, Heidelberg, Germany.

[14] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkitasubramaniam. Ligero: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[15] Toshinori Araki, Assi Barak, Jun Furukawa, Tamar Lichter, Yehuda Lindell, Ariel Nof, Kazuma Ohara, Adi Watzman, and Or Weinstein. Optimized honest-majority MPC for malicious adversaries - breaking the 1 billion-gate per second barrier. In *2017 IEEE Symposium on Security and Privacy*, pages 843–862, San Jose, CA, USA, May 22–26, 2017. IEEE Computer Society Press.

[16] Venkat Arun, Aniket Kate, Deepak Garg, Peter Druschel, and Bobby Bhattacharjee. Finding safety in numbers with secure allegation escrows. In *NDSS 2020*, San Diego, CA, USA, February 23-26, 2020. The Internet Society.

[17] Marshall Ball, Brent Carmer, Tal Malkin, Mike Rosulek, and Nichole Schimanski. Garbled neural networks are practical. Cryptology ePrint Archive, Report 2019/338, 2019. https://eprint.iacr.org/2019/338.

[18] Marshall Ball, Tal Malkin, and Mike Rosulek. Garbling gadgets for Boolean and arithmetic circuits. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 565–577, Vienna, Austria, October 24–28, 2016. ACM Press.

[19] Carsten Baum. On garbling schemes with and without privacy. Cryptology ePrint Archive, Report 2016/150, 2016. https://eprint.iacr.org/2016/150.

[20] Carsten Baum, Cyprien de Saint Guilhem, Daniel Kales, Emmanuela Orsini, Peter Scholl, and Greg Zaverucha. Banquet: Short and fast signatures from AES. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 266–297, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany.

[21] Carsten Baum and Ariel Nof. Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 495–526, Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany.

[22] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *22nd ACM STOC*, pages 503–513, Baltimore, MD, USA, May 14–16, 1990. ACM Press.

[23] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *2013 IEEE Symposium on Security and Privacy*, pages 478–492, Berkeley, CA, USA, May 19–22, 2013. IEEE Computer Society Press.

[24] Aner Ben-Efraim, Yehuda Lindell, and Eran Omri. Optimizing semi-honest secure multiparty computation for the internet. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 578–590, Vienna, Austria, October 24–28, 2016. ACM Press.

[25] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10, Chicago, IL, USA, May 2–4, 1988. ACM Press.

[26] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 169–188, Tallinn, Estonia, May 15–19, 2011. Springer, Heidelberg, Germany.

[27] Marina Blanton, Aaron Steele, and Mehrdad Aliasgari. Data-oblivious graph algorithms for secure computation and outsourcing. In Kefei Chen, Qi Xie, Weidong Qiu, Ninghui Li, and Wen-Guey Tzeng, editors, *ASIACCS 13*, pages 207–218, Hangzhou, China, May 8–10, 2013. ACM Press.

[28] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300, Bengalore, India, December 1–5, 2013. Springer, Heidelberg, Germany.

[29] Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 896–912, Toronto, ON, Canada, October 15–19, 2018. ACM Press.

[30] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 291–308. ACM Press, November 11–15, 2019.

[31] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 489–518, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.

[32] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators from ring-LPN. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 387–416, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.

[33] Elette Boyle, Niv Gilboa, Yuval Ishai, and Ariel Nof. Sublinear GMW-style compiler for MPC with preprocessing. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 457–485, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[34] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

[35] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325, Cambridge, MA, USA, January 8–10, 2012. ACM.

[36] Luís T. A. N. Brandão. Secure two-party computation with reusable bit-commitments, via a cut-and-choose with forge-and-lose technique - (extended abstract). In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 441–463, Bengalore, India, December 1–5, 2013. Springer, Heidelberg, Germany.

[37] Luís T. A. N. Brandão, Michael Davidson, and Apostol Vassilev. Nist roadmap toward criteria for thresholdschemes for cryptographic primitives, June 2020. https://doi.org/10.6028/NIST.IR.8214A.

[38] Dung Bui and Geoffroy Couteau. Private set intersection from pseudorandom correlation generators. Cryptology ePrint Archive, Report 2022/334, 2022. https://eprint.iacr.org/2022/334.

[39] Niklas Büscher, Daniel Demmler, Nikolaos P. Karvelas, Stefan Katzenbeisser, Juliane Krämer, Deevashwer Rathee, Thomas Schneider, and Patrick Struck. Secure two-party computation in a quantum world. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20, Part I*, volume 12146 of *LNCS*, pages 461–480, Rome, Italy, October 19–22, 2020. Springer, Heidelberg, Germany.

[40] Ran Canetti, Pratik Sarkar, and Xiao Wang. Blazing fast OT for three-round UC OT extension. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 299–327, Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany.

[41] Ran Canetti, Pratik Sarkar, and Xiao Wang. Efficient and round-optimal oblivious transfer and commitment with adaptive security. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 277–308, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.

[42] Ignacio Cascudo, Ronald Cramer, Chaoping Xing, and Chen Yuan. Amortized complexity of information-theoretically secure MPC revisited. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 395–426, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.

[43] Ignacio Cascudo and Jaron Skovsted Gundersen. A secret-sharing based MPC protocol for boolean circuits with good amortized complexity. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 652–682, Durham, NC, USA, November 16–19, 2020. Springer, Heidelberg, Germany.

[44] Melissa Chase, Yevgeniy Dodis, Yuval Ishai, Daniel Kraschewski, Tianren Liu, Rafail Ostrovsky, and Vinod Vaikuntanathan. Reusable non-interactive secure computation. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 462–488, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.

[45] Melissa Chase, Esha Ghosh, and Oxana Poburinnaya. Secret-shared shuffle. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part III*, volume 12493 of *LNCS*, pages 342–372, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.

[46] Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1243–1255, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[47] Ilaria Chillotti, Emmanuela Orsini, Peter Scholl, Nigel Paul Smart, and Barry Van Leeuwen. Scooby: Improved multi-party homomorphic secret sharing based on fhe. Cryptology ePrint Archive, Paper 2022/862, 2022. https://eprint.iacr.org/2022/862.

[48] Christopher A. Choquette-Choo, Natalie Dullerud, Adam Dziedzic, Yunxiang Zhang, Somesh Jha, Nicolas Papernot, and Xiao Wang. Capc learning: Confidential and private collaborative learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.

[49] Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 502–534, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[50] Ivan Damgård, Rasmus Lauritsen, and Tomas Toft. An empirical study and some improvements of the MiniMac protocol for secure computation. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 398–415, Amalfi, Italy, September 3–5, 2014. Springer, Heidelberg, Germany.

[51] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

[52] Ivan Damgård and Rasmus Winther Zakarias. Fast oblivious AES A dedicated application of the MiniMac protocol. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 16*, volume 9646 of *LNCS*, pages 245–264, Fes, Morocco, April 13–15, 2016. Springer, Heidelberg, Germany.

[53] Cyprien de Saint Guilhem, Lauren De Meyer, Emmanuela Orsini, and Nigel P. Smart. BBQ: Using AES in picnic signatures. In Kenneth G. Paterson and Douglas Stebila, editors, *SAC 2019*, volume 11959 of *LNCS*, pages 669–692, Waterloo, ON, Canada, August 12–16, 2019. Springer, Heidelberg, Germany.

[54] Cyprien de Saint Guilhem, Emmanuela Orsini, and Titouan Tanguy. Limbo: Efficient zero-knowledge MPCitH-based arguments. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 3022–3036, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[55] Bert den Boer. More efficient match-making and satisfiability: The five card trick. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 208–217, Houthalen, Belgium, April 10–13, 1990. Springer, Heidelberg, Germany.

[56] Srinivas Devadas, Marten van Dijk, Christopher W. Fletcher, Ling Ren, Elaine Shi, and Daniel Wichs. Onion ORAM: A constant bandwidth blowup oblivious RAM. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 145–174, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.

[57] Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Authenticated garbling from simple correlations. In *CRYPTO 2022*, LNCS, Santa Barbara, CA, USA, August 19–23, 2022. Springer, Heidelberg, Germany.

[58] Christoph Dobraunig, Daniel Kales, Christian Rechberger, Markus Schofnegger, and Greg Zaverucha. Shorter signatures based on tailor-made minimalist symmetric-key crypto. Cryptology ePrint Archive, Report 2021/692, 2021. https://eprint.iacr.org/2021/692.

[59] Jack Doerner, David Evans, and abhi shelat. Secure stable matching at scale. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1602–1613, Vienna, Austria, October 24–28, 2016. ACM Press.

[60] Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 523–535, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[61] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10164–10173, June 2022.

[62] Xiao Dong, David A Randolph, Chenkai Weng, Abel N Kho, Jennie M Rogers, and Xiao Wang. Developing high performance secure multi-party computation protocols in healthcare: A case study of patient risk stratification. *AMIA Summits on Translational Science Proceedings*, page 200, 2021.

[63] Sky Faber, Stanislaw Jarecki, Sotirios Kentros, and Boyang Wei. Three-party ORAM for secure computation. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 360–385, Auckland, New Zealand, November 30 – December 3, 2015. Springer, Heidelberg, Germany.

[64] 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, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[65] Tore Kasper Frederiksen, Thomas P. Jakobsen, Jesper Buus Nielsen, and Roberto Trifiletti. TinyLEGO: An interactive garbling scheme for maliciously secure two-party computation. Cryptology ePrint Archive, Report 2015/309, 2015. https://eprint.iacr.org/2015/309.

[66] Tore Kasper Frederiksen, Thomas Pelle Jakobsen, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. MiniLEGO: Efficient secure two-party computation from general assumptions. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 537–556, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.

[67] Gayathri Garimella, Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. Oblivious key-value stores and amplification for private set intersection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 395–425, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[68] Daniel Genkin, Yuval Ishai, and Antigoni Polychroniadou. Efficient multi-party computation: From passive to active security via secure SIMD circuits. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 721–741, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.

[69] Daniel Genkin, Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and Eran Tromer. Circuits resilient to additive attacks with applications to secure computation. In David B. Shmoys, editor, *46th ACM STOC*, pages 495–504, New York, NY, USA, May 31 – June 3, 2014. ACM Press.

[70] Daniel Genkin, Yuval Ishai, and Mor Weiss. Binary AMD circuits from secure multiparty computation. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 336–366, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.

[71] Craig Gentry, Kenny A. Goldman, Shai Halevi, Charanjit S. Jutla, Mariana Raykova, and Daniel Wichs. Optimizing ORAM and using it efficiently for secure computation. In Emiliano De Cristofaro and Matthew K. Wright, editors, *PETS 2013*, volume 7981 of *LNCS*, pages 1–18, Bloomington, IN, USA, July 10–12, 2013. Springer, Heidelberg, Germany.

[72] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

[73] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *25th FOCS*, pages 464–479, Singer Island, Florida, October 24–26, 1984. IEEE Computer Society Press.

[74] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229, New York City, NY, USA, May 25–27, 1987. ACM Press.

[75] S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 513–524, Raleigh, NC, USA, October 16–18, 2012. ACM Press.

[76] S. Dov Gordon, Jonathan Katz, and Xiao Wang. Simple and efficient two-server ORAM. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 141–157, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany.

[77] S. Dov Gordon, Samuel Ranellucci, and Xiao Wang. Secure computation with low communication from cross-checking. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 59–85, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany.

[78] Samuel Gordon and Arkady Yerukhimovich. Satc: Core: Medium: Collaborative: New approaches for large scale secure computation. https://www.nsf.gov/awardsearch/showAward?AWD_ID=1955264, access in June 2022.

[79] Vipul Goyal, Hanjun Li, Rafail Ostrovsky, Antigoni Polychroniadou, and Yifan Song. ATLAS: Efficient and scalable MPC in the honest majority setting. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 244–274, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[80] Vipul Goyal, Antigoni Polychroniadou, and Yifan Song. Unconditional communication-efficient MPC via hall's marriage theorem. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 275–304, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[81] 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, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.

[82] Chun Guo, Jonathan Katz, Xiao Wang, and Yu Yu. Efficient and secure multiparty computation from fixed-key block ciphers. In *2020 IEEE Symposium on Security and Privacy*, pages 825–841, San Francisco, CA, USA, May 18–21, 2020. IEEE Computer Society Press.

[83] Yaron Gvili, Julie Ha, Sarah Scheffler, Mayank Varia, Ziling Yang, and Xinyuan Zhang. TurboIKOS: Improved non-interactive zero knowledge and post-quantum signatures. In Kazue Sako and Nils Ole Tippenhauer, editors, *ACNS 21, Part II*, volume 12727 of *LNCS*, pages 365–395, Kamakura, Japan, June 21–24, 2021. Springer, Heidelberg, Germany.

[84] Kyoohyung Han, Dukjae Moon, and Yongha Son. Improved circuit-based PSI via equality preserving compression. Cryptology ePrint Archive, Report 2021/1440, 2021. https://eprint.iacr.org/2021/1440.

[85] Carmit Hazay, Emmanuela Orsini, Peter Scholl, and Eduardo Soria-Vazquez. Concretely efficient large-scale MPC with active security (or, TinyKeys for TinyOT). In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 86–117, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany.

[86] Carmit Hazay, Emmanuela Orsini, Peter Scholl, and Eduardo Soria-Vazquez. TinyKeys: A new approach to efficient multi-party computation. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 3–33, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.

[87] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. Low cost constant round MPC combining BMR and oblivious transfer. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 598–628, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.

[88] Xi He, Jennie Rogers, Johes Bater, Ashwin Machanavajjhala, Chenghong Wang, and Xiao Wang. Practical security and privacy for database systems. In *Proceedings of the 2021 International Conference on Management of Data*, SIGMOD '21, page 2839–2845, New York, NY, USA, 2021. Association for Computing Machinery.

[89] David Heath and Vladimir Kolesnikov. Stacked garbling - garbled circuit proportional to longest execution path. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 763–792, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.

[90] David Heath and Vladimir Kolesnikov. Stacked garbling for disjunctive zero-knowledge proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 569–598, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.

[91] David Heath and Vladimir Kolesnikov. One hot garbling. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 574–593, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[92] David Heath, Vladimir Kolesnikov, and Rafail Ostrovsky. EpiGRAM: Practical garbled RAM. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 3–33, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany.

[93] Cheng Hong, Jonathan Katz, Vladimir Kolesnikov, Wen-jie Lu, and Xiao Wang. Covert security with public verifiability: Faster, leaner, and simpler. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 97–121, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.

[94] Yan Huang, Jonathan Katz, and David Evans. Quid-Pro-Quo-tocols: Strengthening semi-honest protocols with dual execution. In *2012 IEEE Symposium on Security and Privacy*, pages 272–284, San Francisco, CA, USA, May 21–23, 2012. IEEE Computer Society Press.

[95] Yan Huang, Jonathan Katz, and David Evans. Efficient secure two-party computation using symmetric cut-and-choose. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 18–35, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

[96] Yan Huang, Jonathan Katz, Vladimir Kolesnikov, Ranjit Kumaresan, and Alex J. Malozemoff. Amortizing garbled circuits. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 458–475, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[97] Siam Umar Hussain, Mojan Javaheripi, Mohammad Samragh, and Farinaz Koushanfar. COINN: Crypto/ML codesign for oblivious inference via neural networks. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 3266–3281, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[98] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 406–425, Tallinn, Estonia, May 15–19, 2011. Springer, Heidelberg, Germany.

[99] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30, San Diego, CA, USA, June 11–13, 2007. ACM Press.

[100] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Secure arithmetic computation with no honest majority. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 294–314. Springer, Heidelberg, Germany, March 15–17, 2009.

[101] Keyu Ji, Bingsheng Zhang, Tianpei Lu, Lichun Li, and Kui Ren. UC secure private branching program and decision tree evaluation. Cryptology ePrint Archive, Report 2021/1390, 2021. https://eprint.iacr.org/2021/1390.

[102] Daniel Kales, Sebastian Ramacher, Christian Rechberger, Roman Walch, and Mario Werner. Efficient FPGA implementations of LowMC and Picnic. In Stanislaw Jarecki, editor, *CT-RSA 2020*, volume 12006 of *LNCS*, pages 417–441, San Francisco, CA, USA, February 24–28, 2020. Springer, Heidelberg, Germany.

[103] Daniel Kales and Greg Zaverucha. An attack on some signature schemes constructed from five-pass identification schemes. In Stephan Krenn, Haya Shulman, and Serge Vaudenay, editors, *CANS 20*, volume 12579 of *LNCS*, pages 3–22, Vienna, Austria, December 14–16, 2020. Springer, Heidelberg, Germany.

[104] Daniel Kales and Greg Zaverucha. Improving the performance of the Picnic signature scheme. *IACR TCHES*, 2020(4):154–188, 2020. https://tches.iacr.org/index.php/TCHES/article/view/8680.

[105] Daniel Kales and Greg Zaverucha. Efficient lifting for shorter zero-knowledge proofs and post-quantum signatures. Cryptology ePrint Archive, Paper 2022/588, 2022. https://eprint.iacr.org/2022/588.

[106] Jonathan Katz, Vladimir Kolesnikov, and Xiao Wang. Improved non-interactive zero knowledge with applications to post-quantum signatures. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 525–537, Toronto, ON, Canada, October 15–19, 2018. ACM Press.

[107] Jonathan Katz, Alex J. Malozemoff, and Xiao Wang. Efficiently enforcing input validity in secure two-party computation. Cryptology ePrint Archive, Report 2016/184, 2016. https://eprint.iacr.org/2016/184.

[108] Jonathan Katz, Samuel Ranellucci, Mike Rosulek, and Xiao Wang. Optimizing authenticated garbling for faster secure two-party computation. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 365–391, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.

[109] Marcel Keller, Emmanuela Orsini, and Peter Scholl. MASCOT: Faster malicious arithmetic secure computation with oblivious transfer. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 830–842, Vienna, Austria, October 24–28, 2016. ACM Press.

[110] Marcel Keller and Peter Scholl. Efficient, oblivious data structures for MPC. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 506–525, Kaoshiung, Taiwan, R.O.C., December 7–11, 2014. Springer, Heidelberg, Germany.

[111] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 818–829, Vienna, Austria, October 24–28, 2016. ACM Press.

[112] Vladimir Kolesnikov, Jesper Buus Nielsen, Mike Rosulek, Ni Trieu, and Roberto Trifiletti. DUPLO: Unifying cut-and-choose for garbled circuits. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 3–20, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[113] Vladimir Kolesnikov, Mike Rosulek, Ni Trieu, and Xiao Wang. Scalable private set union from symmetric-key techniques. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 636–666, Kobe, Japan, December 8–12, 2019. Springer, Heidelberg, Germany.

[114] Vladimir Kolesnikov and Thomas Schneider. Improved garbled circuit: Free XOR gates and applications. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 486–498, Reykjavik, Iceland, July 7–11, 2008. Springer, Heidelberg, Germany.

[115] Benjamin Kreuter, abhi shelat, and Chih-Hao Shen. Billion-gate secure computation with malicious adversaries. In Tadayoshi Kohno, editor, *USENIX Security 2012*, pages 285–300, Bellevue, WA, USA, August 8–10, 2012. USENIX Association.

[116] Sam Kumar, David E. Culler, and Raluca Ada Popa. MAGE: Nearly zero-cost virtual memory for secure computation. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*, pages 367–385. USENIX Association, July 2021.

[117] Yehuda Lindell. Fast cut-and-choose based protocols for malicious and covert adversaries. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 1–17, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

[118] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 52–78, Barcelona, Spain, May 20–24, 2007. Springer, Heidelberg, Germany.

[119] Yehuda Lindell and Benny Pinkas. Secure two-party computation via cut-and-choose oblivious transfer. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 329–346, Providence, RI, USA, March 28–30, 2011. Springer, Heidelberg, Germany.

[120] Yehuda Lindell and Ben Riva. Cut-and-choose Yao-based secure computation in the online/offline and batch settings. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 476–494, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.

[121] Yehuda Lindell and Ben Riva. Blazing fast 2PC in the offline/online setting with security for malicious adversaries. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 579–590, Denver, CO, USA, October 12–16, 2015. ACM Press.

[122] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. ObliVM: A programming framework for secure computation. In *2015 IEEE Symposium on Security and Privacy*, pages 359–376, San Jose, CA, USA, May 17–21, 2015. IEEE Computer Society Press.

[123] Ning Luo, Timos Antonopoulos, William Harris, Ruzica Piskac, Eran Tromer, and Xiao Wang. Proving UNSAT in zero knowledge. In *ACM CCS 2022*. ACM Press, November 15–19, 2022.

[124] Ning Luo, Samuel Judson, Timos Antonopoulos, Ruzica Piskac, and Xiao Wang. ppSAT: Towards two-party private SAT solving. In *USENIX Security 2022*. USENIX Association, August 11–13, 2022.

[125] Anna Lysyanskaya. *Signature schemes and applications to cryptographic protocol design*. PhD thesis, Massachusetts Institute of Technology, 2002.

[126] Deepak Maram, Harjasleen Malvai, Fan Zhang, Nerla Jean-Louis, Alexander Frolov, Tyler Kell, Tyrone Lobban, Christine Moy, Ari Juels, and Andrew Miller. CanDID: Can-do decentralized identity with legacy compatibility, sybil-resistance, and accountability. In *2021 IEEE Symposium on Security and Privacy*, pages 1348–1366, San Francisco, CA, USA, May 24–27, 2021. IEEE Computer Society Press.

[127] Carrie Menkel-Meadow. Whose dispute is it anyway: A philosophical and democratic defense of settlement (in some cases). *Geo LJ*, 83:2663, 1994.

[128] Peihan Miao, Xiao Wang, and Tal Rabin. Mentoring Workshop and Videos, Affiliated workshop at CRYPTO, 2021. https://mentor-crypto-2021.github.io.

[129] Andrew Miller and Ari Juels. Collaborative research: Satc: Core: Medium: Hybridizing trusted execution environments and secure multiparty computation. https://www.nsf.gov/awardsearch/showAward?AWD_ID=2112726, access in June 2022.

[130] John C Mitchell and Joe Zimmerman. Data-oblivious data structures. In *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.

[131] Payman Mohassel and Matthew Franklin. Efficiency tradeoffs for malicious two-party computation. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 458–473, New York, NY, USA, April 24–26, 2006. Springer, Heidelberg, Germany.

[132] Andrew Morgan, Rafael Pass, and Antigoni Polychroniadou. Succinct non-interactive secure computation. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 216–245, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.

[133] Moni Naor and Benny Pinkas. Oblivious transfer and polynomial evaluation. In *31st ACM STOC*, pages 245–254, Atlanta, GA, USA, May 1–4, 1999. ACM Press.

[134] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, page 129–139, New York, NY, USA, 1999. Association for Computing Machinery.

[135] Kartik Nayak, Xiao Shaun Wang, Stratis Ioannidis, Udi Weinsberg, Nina Taft, and Elaine Shi. GraphSC: Parallel secure computation made easy. In *2015 IEEE Symposium on Security and Privacy*, pages 377–394, San Jose, CA, USA, May 17–21, 2015. IEEE Computer Society Press.

[136] Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 681–700, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

[137] Jesper Buus Nielsen and Claudio Orlandi. LEGO for two-party secure computation. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 368–386. Springer, Heidelberg, Germany, March 15–17, 2009.

[138] Jesper Buus Nielsen and Claudio Orlandi. Cross and clean: Amortized garbled circuits with constant overhead. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 582–603, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.

[139] Jesper Buus Nielsen, Thomas Schneider, and Roberto Trifiletti. Constant round maliciously secure 2PC with function-independent preprocessing using LEGO. In *NDSS 2017*, San Diego, CA, USA, February 26 – March 1, 2017. The Internet Society.

[140] Valeria Nikolaenko, Stratis Ioannidis, Udi Weinsberg, Marc Joye, Nina Taft, and Dan Boneh. Privacy-preserving matrix factorization. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 801–812, Berlin, Germany, November 4–8, 2013. ACM Press.

[141] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *29th ACM STOC*, pages 294–303, El Paso, TX, USA, May 4–6, 1997. ACM Press.

[142] Andrew Park, Wei-Kai Lin, and Elaine Shi. NanoGRAM: Garbled RAM with $\widetilde{O}(\log N)$ overhead. Cryptology ePrint Archive, Report 2022/191, 2022. https://eprint.iacr.org/2022/191.

[143] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. SpOT-light: Lightweight private set intersection from sparse OT extension. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 401–431, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.

[144] Google Security& Privacy. To stay secure online, password checkup has your back. https://blog.google/technology/safety-security/password-checkup/, Accessed in June 2022.

[145] Anjana Rajan. Callisto: A cryptographic approach to #metoo. Burlingame, CA, January 2019. USENIX Association. https://mycallisto.org/assets/docs/callisto-cryptographic-approach.pdf.

[146] Rep. Stevens, Haley and Rep. Gonzalez, Anthony. H.R.847 - Promoting Digital Privacy Technologies Act. https://www.congress.gov/bill/117th-congress/house-bill/847, Accessed in June 2022.

[147] Peter Rindal and Srinivasan Raghuraman. Blazing fast PSI from improved OKVS and subfield VOLE. Cryptology ePrint Archive, Report 2022/320, 2022. https://eprint.iacr.org/2022/320.

[148] Peter Rindal and Mike Rosulek. Faster malicious 2-party secure computation with online/offline dual execution. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 297–314, Austin, TX, USA, August 10–12, 2016. USENIX Association.

[149] Peter Rindal and Phillipp Schoppmann. VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 901–930, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.

[150] Jennie Rogers, Elizabeth Adetoro, Johes Bater, Talia Canter, Dong Fu, Andrew Hamilton, Amro Hassan, Ashley Martinez, Erick Michalski, Vesna Mitrovic, Fred Rachman, Raj Shah, Matt Sterling, Kyra VanDoren, Theresa L. Walunas, Xiao Wang, and Abel Kho. Vaultdb: A real-world pilot of secure multi-party computation within a clinical research network, 2022.

[151] Mike Rosulek and Lawrence Roy. Three halves make a whole? Beating the half-gates lower bound for garbled circuits. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 94–124, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

[152] F.E.A. Sander. *Varieties of Dispute Processing*. West Publishing Company, 1976.

[153] Phillipp Schoppmann, Lennart Vogelsang, Adrià Gascón, and Borja Balle. Secure and scalable document similarity on distributed databases: Differential privacy to the rescue. *PoPETs*, 2020(2):209–229, April 2020.

[154] Sen. Cortez Masto and Sen. Fischer, Deb. S.224 - Promoting Digital Privacy Technologies Act. https://www.congress.gov/bill/117th-congress/senate-bill/224, Accessed in June 2022.

[155] Sen. McSally, Martha and Sen. Sinema, Kyrsten. H. Rept. 117-305 - PROMOTING DIGITAL PRIVACY TECHNOLOGIES ACT. https://www.congress.gov/congressional-report/117th-congress/house-report/305, Accessed in June 2022.

[156] Sen. Wyden, Ron and Sen. Portman, Rob and Sen. Boozman, John and Sen. Gillibrand, Kirsten. S.4200 - Secure Research Data Network Act. https://www.congress.gov/bill/117th-congress/senate-bill/4200, Accessed in June 2022.

[157] Sen. Wyden, Ron and Sen. Rubio, Marco and Sen. Warner, Mark R. S.3952 - Student Right to Know Before You Go Act of 2022. https://www.congress.gov/bill/117th-congress/senate-bill/3952, Accessed in June 2022.

[158] Aditya Shastri, David Wu, Lynn Chua, Ruiyu Zhu, Tal Davidi, Yu Sheng, and Josh Mintz. Private computation framework 2.0, Accessed in June, 2022. White paper at https://research.facebook.com/publications/private-computation-framework-2-0/.

[159] abhi shelat and Chih-Hao Shen. Two-output secure computation with malicious adversaries. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 386–405, Tallinn, Estonia, May 15–19, 2011. Springer, Heidelberg, Germany.

[160] abhi shelat and Chih-Hao Shen. Fast two-party secure computation with minimal assumptions. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 523–534, Berlin, Germany, November 4–8, 2013. ACM Press.

[161] Justin Wagner, Joseph N Paulson, Xiao Wang, Bobby Bhattacharjee, and Hector Corrada Bravo. Privacy-Preserving Microbiome Analysis Using Secure Computation. In *Bioinformatics*, 2016.

[162] Lixu Wang, Shichao Xu, Ruiqi Xu, Xiao Wang, and Qi Zhu. Non-transferable learning: A new approach for model ownership verification and applicability authorization. In *International Conference on Learning Representations*, 2022.

[163] Xiao Wang, T.-H. Hubert Chan, and Elaine Shi. Circuit ORAM: On tightness of the Goldreich-Ostrovsky lower bound. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 850–861, Denver, CO, USA, October 12–16, 2015. ACM Press.

[164] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient MultiParty computation toolkit. https://github.com/emp-toolkit, 2016.

[165] Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. Faster secure two-party computation in the single-execution setting. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 399–424, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.

[166] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 21–37, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[167] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 39–56, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.

[168] Xiao Shaun Wang, S. Dov Gordon, Allen McIntosh, and Jonathan Katz. Secure computation of MIPS machine code. In Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows, editors, *ESORICS 2016, Part II*, volume 9879 of *LNCS*, pages 99–117, Heraklion, Greece, September 26–30, 2016. Springer, Heidelberg, Germany.

[169] Xiao Shaun Wang, Yan Huang, T.-H. Hubert Chan, abhi shelat, and Elaine Shi. SCORAM: Oblivious RAM for secure computation. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014*, pages 191–202, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.

[170] Xiao Shaun Wang, Yan Huang, Yongan Zhao, Haixu Tang, XiaoFeng Wang, and Diyue Bu. Efficient genome-wide, privacy-preserving similar patient query based on private edit distance. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 492–503, Denver, CO, USA, October 12–16, 2015. ACM Press.

[171] Xiao Shaun Wang, Kartik Nayak, Chang Liu, T.-H. Hubert Chan, Elaine Shi, Emil Stefanov, and Yan Huang. Oblivious data structures. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014*, pages 215–226, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.

[172] 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, San Francisco, CA, USA, May 24–27, 2021. IEEE Computer Society Press.

[173] 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 11–13, 2021.

[174] Chenkai Weng, Kang Yang, Zhaomin Yang, Xiang Xie, and Xiao Wang. Antman: Interactive zero-knowledge proofs with sublinear communication. In *ACM CCS 2022*. ACM Press, November 15–19, 2022.

[175] 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, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.

[176] Kang Yang, Xiao Wang, and Jiang Zhang. More efficient MPC from improved triple generation and authenticated garbling. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1627–1646, Virtual Event, USA, November 9–13, 2020. ACM Press.

[177] 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, Virtual Event, USA, November 9–13, 2020. ACM Press.

[178] Yibin Yang, David Heath, Vladimir Kolesnikov, and David Devecsery. Ezee: Epoch parallel zero knowledge for ansi c. In *2022 IEEE 7th European Symposium on Security and Privacy* (*EuroS&P*), pages 109–123, 2022.

[179] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167, Toronto, Ontario, Canada, October 27–29, 1986. IEEE Computer Society Press.

[180] Arkady Yerukhimovich. Career: Sketching for secure computation on large inputs. https://www.nsf.gov/awardsearch/showAward?AWD_ID=2144798, access in June 2022.

[181] Samee Zahur and David Evans. Circuit structures for improving efficiency of security and privacy tools. In *2013 IEEE Symposium on Security and Privacy*, pages 493–507, Berkeley, CA, USA, May 19–22, 2013. IEEE Computer Society Press.

[182] Samee Zahur, Mike Rosulek, and David Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 220–250, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.

[183] Samee Zahur, Xiao Shaun Wang, Mariana Raykova, Adria Gascón, Jack Doerner, David Evans, and Jonathan Katz. Revisiting square-root ORAM: Efficient random access in multi-party computation. In *2016 IEEE Symposium on Security and Privacy*, pages 218–234, San Jose, CA, USA, May 22–26, 2016. IEEE Computer Society Press.

[184] Greg Zaverucha, Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, Jonathan Katz, Xiao Wang, Vladmir Kolesnikov, and Daniel Kales. Picnic. Technical report, National Institute of Standards and Technology, 2020. available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions.

[185] Fan Zhang, Deepak Maram, Harjasleen Malvai, Steven Goldfeder, and Ari Juels. DECO: Liberating web data using decentralized oracles for TLS. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1919–1938, Virtual Event, USA, November 9–13, 2020. ACM Press.

[186] Wenting Zheng, Ryan Deng, Weikeng Chen, Raluca Ada Popa, Aurojit Panda, and Ion Stoica. Cerebro: A platform for multi-party cryptographic collaborative learning. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 2723–2740. USENIX Association, August 11–13, 2021.

[187] Ruiyu Zhu, Darion Cassel, Amr Sabry, and Yan Huang. NANOPI: Extreme-scale actively-secure multi-party computation. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 862–879, Toronto, ON, Canada, October 15–19, 2018. ACM Press.

[188] Ruiyu Zhu and Yan Huang. JIMU: Faster LEGO-based secure computation using additive homomorphic hashes. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 529–572, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.

[189] Ruiyu Zhu, Yan Huang, and Darion Cassel. Pool: Scalable on-demand secure computation service against malicious adversaries. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 245–257, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.