

SoK: Privacy-Enhancing Technologies in Finance

Carsten Baum ✉

Technical University of Denmark, Lyngby, Denmark

James Hsin-yu Chiang ✉

Aarhus University, Denmark

Bernardo David ✉

IT University of Copenhagen, Denmark

Tore Kasper Frederiksen ✉

Zama, Paris, France

Abstract

Recent years have seen the emergence of practical advanced cryptographic tools that not only protect data privacy and authenticity, but also allow for jointly processing data from different institutions without sacrificing privacy. The ability to do so has enabled implementations of a number of traditional and decentralized financial applications that would have required sacrificing privacy or trusting a third party. The main catalyst of this revolution was the advent of decentralized cryptocurrencies that use public ledgers to register financial transactions, which must be verifiable by any third party, while keeping sensitive data private. Zero Knowledge (ZK) proofs rose to prominence as a solution to this challenge, allowing for the owner of sensitive data (*e.g.* the identities of users involved in an operation) to convince a third party verifier that a certain operation has been correctly executed without revealing said data. It quickly became clear that performing arbitrary computation on private data from multiple sources by means of secure Multiparty Computation (MPC) and related techniques allows for more powerful financial applications, also in traditional finance.

In this SoK, we categorize the main traditional and decentralized financial applications that can benefit from state-of-the-art Privacy-Enhancing Technologies (PETs) and identify design patterns commonly used when applying PETs in the context of these applications. In particular, we consider the following classes of applications: 1. Identity Management, KYC & AML; 2. Markets & Settlement; 3. Legal; and 4. Digital Asset Custody. We examine how ZK proofs, MPC and related PETs have been used to tackle the main security challenges in each of these applications. Moreover, we provide an assessment of the technological readiness of each PET in the context of different financial applications according to the availability of: theoretical feasibility results, preliminary benchmarks (in scientific papers) or benchmarks achieving real-world performance (in commercially deployed solutions). Finally, we propose future applications of PETs as Fintech solutions to currently unsolved issues. While we systematize financial applications of PETs at large, we focus mainly on those applications that require privacy preserving computation on data from multiple parties.

2012 ACM Subject Classification Security and privacy → Cryptography

Keywords and phrases DeFi, Anti-money laundering, MPC, FHE, identity management, PETs

Digital Object Identifier 10.4230/LIPIcs.AFT.2023.12

Related Version *Full Version:* <https://eprint.iacr.org/2023/122>

Funding *Carsten Baum:* Part of the work was carried out while the author was visiting Copenhagen University and supported by Partisia. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of Partisia.

James Hsin-yu Chiang: This work was supported by a DTU Compute scholarship.

Bernardo David: The project was supported by the Independent Research Fund Denmark (IRFD) grants number 9040-00399B (TrA²C), 9131-00075B (PUMA) and 0165-00079B, and by DIREC.

Tore Kasper Frederiksen: This work was carried out while working at Protocol Labs and the Alexandra Institute (supported by Copenhagen Fintech as part of as part of the “National Position of Strength programme for Finans & Fintech” funded by the Danish Ministry of Higher Education and Science).



© Carsten Baum, James Hsin-yu Chiang, Bernardo David, and Tore Kasper Frederiksen; licensed under Creative Commons License CC-BY 4.0

5th Conference on Advances in Financial Technologies (AFT 2023).

Editors: Joseph Bonneau and S. Matthew Weinberg; Article No. 12; pp. 12:1–12:30

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Modern Cryptography and Traditional Finance. Due to their sensitive nature, financial applications require strong security guarantees. Clearly, it is necessary to ensure authenticity and integrity of any financial operation, *i.e.* guaranteeing that the operation has been ordered by an entity authorized to do so and that this order has not been tampered with. Moreover, it is also necessary to achieve *privacy*, *i.e.* preventing attackers from obtaining sensitive information related to financial operations (*e.g.* *the identities of entities involved in a transaction and/or the value of that transaction*). In the digital realm, authenticity and privacy guarantees can be achieved against powerful adversaries who control communication networks (*e.g.* the Internet) by means of digital signatures and encryption, respectively.

A Decentralized Conundrum. The meteoric rise of decentralized financial applications based on cryptocurrencies and smart contracts hosted on blockchain platforms brought to light a whole new set of challenges. While traditional financial applications are hosted and executed by financial institutions in a centralized manner, the decentralized nature of blockchain-based applications requires all operations to be verifiable by third parties by means of publicly available records. If only simple cryptographic primitives are employed, this means that sensitive data that was once internally handled by financial institutions must now be exposed on the blockchain in order to perform a financial application. For example, Bitcoin requires revealing the sender and the receiver of a financial token that is transferred, so that a transfer transaction is considered valid if and only if the rightful owner of the token signs it.

Privacy-Enhancing Technologies (PETs) to the Rescue. While sacrificing privacy to achieve decentralization may be acceptable in some situations, most financial operations involving companies and private citizens cannot be conducted in this manner due to a number of reasons (*e.g.* protecting business interests and complying to regulations). In order to solve this issue, the cryptocurrency community turned to Privacy-Enhancing Technologies (PETs) that allow for achieving the same authenticity and privacy guarantees as in traditional centralized financial applications while providing the public verifiability guarantees needed in decentralized blockchain platforms. In particular, many of the first proposals towards this goal involved using a technology called Zero Knowledge (ZK) proof systems [67] : a method that allows the owner of sensitive data to prove a statement about this data without having to reveal it. For example, in the token transfer transaction example, a ZK proof allows a user to prove that an “encrypted” transfer transaction has been signed by the rightful owner of the token, without revealing neither the owner’s nor the receiver’s identity, *e.g.* as in [15].

From Decentralized to Traditional Finance. The vast usefulness of advanced PETs in blockchain applications also sparked an interest in deploying similar solutions for traditional financial applications. The aforementioned ZK proof technology has also been used in innovative solutions to the Know Your Client (KYC) and Anti Money Laundering (AML) problems commonly encountered in the banking industry, *e.g.*, as in [79, 118, 115, 110]. As in the case of privacy preserving cryptocurrency transactions, in many scenarios an entity wants to prove that they comply with KYC/AML regulations without revealing their identity nor their sensitive data. For example, a client can prove to a third party service provider that their identity has been verified by their bank and that they are authorized to use a certain service and perform operations up to a certain financial volume, while keeping their identity, and other attributes (*e.g.* the list of operations they are allowed to perform) private.

PETs for the Masses - or - From ZK to MPC. The ZK proof technology lends itself extremely well to applications that require a single entity to publicly prove a statement about its private data, *e.g.* the KYC/AML or cryptocurrency examples above. However, it is limited by the fact that the entity who generates a ZK proof must necessarily know all the private information about which the statement is proven. This is a serious limitation in two cases: 1. applications that must process sensitive data provided by multiple entities; 2. applications where certain data (*e.g.* cryptographic secret keys) are far too valuable to be stored on a single device, which leaks the data if its security is compromised. Fortunately, these limitations can be addressed by means of secure Multiparty Computation (MPC) [36, 66], which allows a set of entities to jointly execute an arbitrary program that computes on an “encrypted” version of their private data and only reveals the output of this computation.

For example, specific-purpose MPC protocols have long been used for sealed-bid auctions [70, 25] among entities who do not trust each other, nor a third party auctioneer. In this case, the parties provide as input “encrypted” versions of their bids and jointly compute a program that determines the winner of the auction, without revealing the value of the bids or any other information. In the context of blockchain-based cryptocurrencies, MPC protocols [77] have also been successfully deployed [42] for protecting secret signature keys used to authorize/authenticate transactions. In this case, many entities locally store a “share” of the signing key that does not reveal any information about the key itself unless all shares are united. When a transaction must be signed, all these entities use MPC to jointly execute a program that takes as input all the signing key shares, reconstructs the true key and computes the signature on the given transaction, while only revealing the resulting signature and nothing else. Since knowledge of the key is split among many entities, an attacker now has to compromise many, potentially all, entities instead of a single server.

Systematizing Privacy-enhancing Technologies in Finance

The goal of this SoK is to systematize financial applications that can benefit from PETs, as well as systematizing relevant PETs according to their respective applications and technological readiness. As summarized in Table 1, we consider 5 main classes of financial applications, which are each addressed in the specific section indicated next to the application class name. At a high level, these applications can be potentially facilitated by the PETs indicated in the **PET** column of Table 1, which are introduced in detail in Section 2. In particular, we focus on financial applications that require handling private data from multiple entities, as ZK proof technology for financial applications has been extensively addressed in previous works (*e.g.* [24, 3, 4, 91]). While we aim at providing a general overview of PETs for financial applications covering broad ranges of both PETs and applications, we do not intend to provide an exhaustive review of the PET literature. For each class of applications, we strive to survey the works that introduced the most relevant insights and groundbreaking results, since it would be infeasible to cover every single optimization of each PET that would be relevant for each application.

The financial applications we cover and the respective relevant PETs are summarized as follows:

Identity, KYC & AML (Section 3): Identity management is a classical problem that has the added challenges of Know Your Client (KYC) and Anti Money Laundering (AML) regulations in the financial sector. PETs can be used in these applications to provide robust identity management with privacy preserving methods for enforcing KYC & AML regulations in both decentralized and traditional scenarios.

■ **Table 1** PET stack for financial applications. TSS = Threshold Secret Sharing, DP = Differential Privacy, (F)HE=(Fully) Homomorphic Encryption, PSI = Private Set Intersection, MPC = Multiparty Computation, ZK = Zero Knowledge proofs. See Section 2 for a discussion of each concept.

	PET (§2)
(§3) Identity, KYC & AML	MPC, ZK
(§A, [10]) Legal	MPC, ZK
(§B, [10]) Digital Asset Custody	TSS, MPC
(§4) Markets & Settlement	(F)HE, MPC, ZK
(§5) Future applications	PSI, DP, MPC

Legal Procedures (Section A in [10]): Many legal procedures require evidence to be presented in court. However, in many cases the evidence or even the relevant law/regulation must be kept private. PETs allow for such legal procedures to be conducted without sacrificing neither privacy nor auditability (*i.e.* the ability of any entity to verify that a legal procedure has been properly executed). Furthermore, due to the novelty of PETs it is not always clear how they fit within existing legal frameworks and how they might help and provide utility while fulfilling privacy regulations such as GDPR.

Digital Asset Custody (Section B in [10]): Digital assets such as cryptocurrencies are usually transferred by means of a digital signature, which can only be generated given a secret key. Since storing this key in a single device poses a risk of key leakage, PETs can be employed to distribute the signing power (and thus the power to move the asset) among many entities, in such a way that the system is only compromised if all entities are compromised.

Markets & Settlements (Section 4): Both the traditional and decentralized financial markets use complex trading instruments that may be abused by entities who retain privileged information about trades. PETs provide a robust solution to this issue via distributed “Dark Pools” or privacy preserving DeFi mechanisms (*e.g.* Automated Market Makers) that process trades without revealing any sensitive information to the entities evolved.

Future Applications (Section 5): Besides financial applications that have already been addressed in previous work, we propose that several PETs can be potentially used to address other interesting challenges in finance. In particular, recent advances in PETs enable the execution of advanced machine learning (ML) algorithms on private data, allowing for detecting patterns (*e.g.* for fraud) without revealing neither the ML models nor the data.

2 Available Privacy-enhancing Technologies

Before we describe applications of *Privacy-Enhancing Technologies* (PETs) to finance, we will give a short overview over existing PETs and how mature they are.

Zero-Knowledge proofs. Zero-Knowledge proofs [67] are cryptographic algorithms which allows a prover to convince a distrusting verifier that a certain statement is true. While the statement (usually specified in the form of a program) is known to the verifier, the proof (*e.g.* a certain input that makes the program output 0) is never leaked to the verifier. There exists

a large variety of different ZK proof algorithms, and choosing the optimal proof depends largely on the application. Recently, efforts have been underway to standardize ZK proofs¹ to make them more accessible to practitioners.

Private Set Intersection. Private Set Intersection (PSI) [58] allows two (or more) distrusting parties with respective input sets S_1 and S_2 to securely learn their intersection, i.e. $S_1 \cap S_2$, without revealing the non-intersecting elements to the other party. For example, if party 1 has $S_1 = \{a, b, d\}$ and party 2 has $S_2 = \{b, c, e\}$ then both parties will learn that they have b in common in their sets. At the same time, party 2 will not learn that party 1 also had a, d in its input set and vice-versa for c, e . Highly efficient PSI protocols have been developed recently and some, such as the one developed by Chen *et al.* [39] found applications in industry.

Threshold Secret Sharing. Threshold Secret Sharing (TSS) allows a dealer to distribute [100] a secret x among n different parties, who each receive a *share* of the secret. Given a threshold $t < n$, TSS guarantees that if t or less parties pool their shares together, then they cannot reconstruct any information about x . If instead more than t parties cooperate (*i.e.* pool their shares), then x can be reconstructed. Multiple versions of secret sharing exist, for example with security against share-holders who don't act honestly during the reconstruction of the secret [41, 93]. Moreover, secret-sharing can be generalized so that not a threshold decides about the possibility of reconstruction, but instead any pattern can be used by the sender of the shares.

Multiparty Computation. Cryptographic protocols for Multiparty Computation (MPC) [14, 37, 65] allow 2 or more mutually distrusting parties who each have an input x_i to evaluate an arbitrary function $y = f(x_1, \dots, x_n)$ on their inputs. MPC guarantees that only the function output y and no other information about the inputs is revealed. One can see PSI as a special case of MPC where the computed function is the intersection of input sets. MPC can be made robust against parties who maliciously deviate from the protocol description, and security usually holds if less than a threshold t of the participants in the computation collaborate to undermine the security. Therefore, MPC can be seen as constructing a *distributed trusted entity*. It is worth highlighting that MPC can be carried out in many different security and communication models. In particular the communication model can have a big impact on the possible use-cases *and* security of the protocol. For example in a *synchronous* model any message sent can be assumed to arrive within a fixed amount of time, whereas in the *asynchronous* model messages can be arbitrarily, even infinitely, delayed. Hence the synchronous model is more suited for protocols being run on a stable network, like a LAN. Recent progress in MPC research has made practical use of MPC possible².

Fully-Homomorphic Encryption. Fully-Homomorphic Encryption (FHE) is a special type of encryption scheme first proposed in [95] and later realized in [63]. In FHE, everyone with a so-called public key can encrypt information, while only the holder of the private key can decrypt it later. In addition, given encrypted data as well as the public key, anyone can *perform computations* on the encrypted data and evaluate algorithms on secret inputs. For example. Given encryption $[x], [y], [z]$ of the values x, y, z , FHE allows to compute an

¹ See <https://zkproof.org/>.

² <https://www.mpcalliance.org/>

encryption $[x \cdot y + z]$ of $x \cdot y + z$ or *any other efficiently computable algorithm* on these inputs. The clue is that the decryptor who obtains $[x \cdot y + z]$ will only learn $x \cdot y + z$ but not the *inputs to the computation*. Although concrete FHE schemes are relatively new, the technology is already somewhat mature³ and powerful testing implementations⁴ are available.

Differential Privacy. Differential Privacy [55] (DP) is a technique to compute *add noise to outcomes of algorithms* such that leakage about the inputs of the computation is minimized. The level of the noise is calibrated such that mathematical guarantees about the privacy of the inputs can be given.

Readiness of the technologies. We next try to make a rough evaluation of the market readiness and usability of the different technologies. Our main metric for this is the *Technology Readiness Level* (TRL). Recall that the TRL scale is an estimation of technology maturity, originally developed by NASA and standardized by ISO [71]. The scale goes from 1 to 9, where level 1 indicates the initial studies of moving general research into applied research and level 9, means the system has been proven successful in the real-world. The systems we are considering in this paper are far up the scale so we outline what exactly we mean by the different TRL levels achieved for these systems:

7. Demonstration of the system in an operational environment. E.g. beginning of alpha-testing.
8. The system is fully developed and functioning under expected real-world conditions.
9. The system is actually deployed and used by end-users.

We summarize the main privacy enhancing technologies we consider in this SoK and our judgement of their Technology Readiness Level (TRL) in Table 2.

While the TRL level provides one metric in judging the market application of a technology; overhead (in particular in computation and communication) is another important factor, in particular when it comes to scalability and wideness of deployment.

Hence we have include overhead indicators in Table 2 in order to illustrate the *typical* overhead of the different technologies relate. An interesting observation on overhead is that DP can be applied with barely any computational overhead and TSS (*without* key generation) can in some situations have an overhead only linear in the amount of servers involved. Some ZK on the other hand (in particular SNARKs [17]) has the powerful feature that it can *reduce* the communication and computation overhead over native computation (for *some* of the involved parties only!). This feature has been used widely in practice, in particular in the blockchain setting for example to aggregate validation of multiple transactions through a *zk-rollup* [43].

A note on Trusted Execution Environments. Trusted Execution Environments (TEE) such as Intel’s SGX are special modes of modern processors. A processor in its trusted execution setting guarantees that programs and their data are shielded from every other program running on the computer - even the operating system or any user having full access. A secure TEE allows to build many of the aforementioned PETs such as ZK proofs, PSI, MPC etc. “cheaply” and without additional cryptographic tools. In practice, SGX and similar technologies from other vendors⁵ are regularly broken and do not offer the protection that they claim. We therefore do not consider it as a PET in this document.

³ <https://fhe.org/>

⁴ <https://www.openfhe.org/>

⁵ See e.g. the exhaustive list on <https://sgx.fail/>.

■ **Table 2** Technology Readiness Level (TRL) between 1-9 of different PETs along with how much overhead they *typically* add when used in the *most* suitable contexts. An empty circle means close to no overhead, and a full circle means multiple orders of magnitude. Finally a dot means *less* than native communication/computation. See Section 2 for a discussion of each concept.

PET (§2)	TRL	Comp. Overhead	Comm. Overhead	<i>Demonstrated by</i>
ZK	9	.	.	Ethereum [43], Filecoin [76]
TSS	9	○	○	Zengo [75]
DP	9	○	○	Apple [53]
PSI	8	○	○	Apple [16]
MPC	7	●	◐	JP Morgen [46], Meta [35]
FHE	7	●	○	Zama [96]

3 Identity, KYC, AML

A general issue facing the financial world is the validation of customer identities and attributes. Laws and regulations require financial institutions, both classical and decentralized, to employ Know Your Customer (KYC) rules, for example to prevent money laundering and to be able to aid in criminal cases - or even to lock accounts in case of sanctions. Being able to correctly determine a legal owner of an account can in itself help in preventing money laundering by precluding the use of fake accounts which could otherwise aid in *smurfing*, see Sec. 3. In the EU this is for example in place through the Anti-Money Laundering Directives and in the US through the Money Laundering Control Act of 1986. In the classical banking setting such validations are carried out through customers going to their physical bank and bringing required documents to prove their identity, residence or perhaps even criminal history, of which the bank would keep a copy. However, with the advent of online-only banks such as Lunar, Revolut and N26, along with crypto-currency exchanges like Binance and Coinbase, such validations become tricky, as there are no physical locations to validate identities.

Today KYC is instead carried out online, and in many cases through machine learning algorithms, where customers upload copies of their data which gets validated. When it comes to security this unfortunately has several disadvantages: i) it is easy to create a picture of a document or manufacture it, or even modify some data of a real document [105]; and ii) the leakage of legitimate documents online allows an adversary to steal identities. Simply considering how often a copy of ones passport is needed (e.g. basically any hotel or accommodation in any country), it is not hard to see that copies of legitimate documents will be easy to find on the dark market. Even though requirements can be made to include selfies or short videos to validate authenticity, this has turned into a race against Photoshop and deep fakes, which have shown tremendous advancement in the recent years [106]. While such attacks are also possible in physical space (i.e. creating fake documents and having them validated by a human), they are significantly more cumbersome due to human involvement and more expensive to mount, and therefore do not *scale* like digital-only attacks. Thus it is clear that the digital attack vector on KYC is the weakest link in account validation.

Identity management

One possible way of combating attacks when validating digital copies of physical identity documents is simply to move the documents into the digital realm. Digital signatures and revocation systems can ensure that digital documents are legitimate. This is done

by combining them with an identification scheme where a user needs to prove they know a password/key used in the construction of their digital identity. This can prevent theft by simply copying the digital document. Often a simple password or key is not deemed secure enough in financial applications by law [90], and a second factor is required. Thus the use of authentication apps is common in electronic ID (eID) solutions, like the Danish MitID. Such eID solutions validate user-identities by a trusted issuer during setup, allowing other applications to piggy-back on existing validation. This naturally comes with a risk of compromise or identity sharing through the eID provider, although it may arguably be harder than with their physical counterparts.

Single Sign-On. eIDs are typically validated by a centralized and trusted server that is able to perform relevant logging, and hence poses a risk to user privacy. Furthermore, such identity management is not exclusive to official or government identities, but can involve any kind of self-reported identity, which is the case for example for a Facebook or Google account. These platforms act as *federated identity management* services, allowing the sharing of the user's identity, along appropriate attributes of the user, to third-party websites. Thus facilitating a *single sign-on* (SSO) system. In this setting, the server validating the user's identity is known as the *identity provider* (IdP), which would be Facebook or Google in the above example. The third-party website is known as the *service provider*. This could for example be Netflix or Spotify. The idea of an SSO service goes beyond simply having an IdP facilitating a user authenticating towards a service provider. In fact an IdP may gather certified attributes about a user from multiple trusted issuers, and sign off on the user indeed being validated to have such attributes. This for example happens when Facebook validates that a given user has access to a specific email account, or phone number. While using an SSO makes things much simpler for a user, it also is a big privacy issue as an IdP now holds a large amount of the user's personal information, along with knowledge of whenever the user uses this information and towards which service provider. Furthermore, it also means that large amount of trust has to be put in the IdP as they would be able to impersonate any of their users towards any service provider. While such a thing is also possible for any attribute issuer (to a lesser extent) it becomes more of a problem for an IdP as they must be user-friendly enough that they can be used several times a day and since their only job is authentication. Hence becoming more exposed. Beyond this, simply using an SSO service can also lead to *traceability* and *linkability* of the user across the web. Traceability means that a user can be identified from the data resulting from using their eID. Whereas linkability means that it is possible for different service providers to find out if they have the same users. This can be an issue even if the user is authenticated using a pseudonym, since all it takes is one sharing of personal data, such as credit card information at a service provider, to de-anonymize the user. However, works like PASTA [2] and PESTO [12] use threshold cryptography to enhance the security of IdPs and limit traceability and linkability without reducing the usability. I.e. password based authentication can still be used and they remain compliant with solutions like OAuth and OpenID Connect. While they only focus on password-based authentication, they can be generalized to support multi-factor authentication [57] and thus be used when multi-factor authentication is required for financial compliance, as e.g. in Europe according to PSD2 [90].

Decentralized Identifiers. With the advent of blockchain technologies a lot of work has sprung up, trying to remove centralization from the management of eIDs. This is generally known as a Decentralized Identifier (DID) [102], but often also called Self-Sovereign Identity

(SSI), specifically when the user is in full control over the usage of their decentrally stored eIDs. The overall idea is that any kind of attribute provider issues a pseudonym to a user's blockchain account, reflecting a specific attribute. The user can then later use the pseudonym to prove certain attributes, or to simply get a reusable link to their pseudonymous identity at the same identity provider. However, it is clear to see that this basic construction is unfortunately not enough to ensure privacy, as again it is possible to link the user across the internet (or blockchain) through their pseudonym. For this reason DID systems are starting to incorporate more advanced cryptographic constructions allowing users to anonymously prove that they hold a certain pseudonym towards a service provider (in order to facilitate authentication). Such a construction is known as a cryptographic *credential*.

Camenisch and Lysyanskaya [29] were the first to show a fully self-managed solution allowing users to prove their identity has been certified by a trusted provider, in an anonymous manner. Their credential construction affords validation of issuance from a trusted authority, while allowing the user to anonymously use it and preventing anyone who does not know the user's key⁶ to impersonate it. However their construction did not allow the validation of arbitrary predicates on the attributes. Something which is needed in many financial situations. Consider for example the case for loan or insurance issuance, where the customer's financial situation or health status has to be validated. Classically these must be provided as signed physical documents from the customer's attribute provider (such as credit bureaus), but the line of work on credentials, known as *attribute based credentials* shows how to achieve this in the digital sphere [30, 98] with cryptographic security and privacy guarantees. It was later shown how to compute arbitrary predicates on the certificated attributes [27]. Further development of such schemes into commercial products have been done by both IBM with their Idemix framework [31] and by Microsoft through U-Prove [88]. The underlying primitives have even been taken up by standardization frameworks such as W3C [101]. Still, despite such commercial traction, widespread adoption is still lacking. Moreover, in the context of DeFi systems, decentralized versions of anonymous credentials [62, 26, 5, 48] have been proposed.

One could imagine that the requirement for self-managed private keys could be the reason that such approaches lack adoption since the regular news bulletins of people having lost their cryptocurrency keys, show that self-administered key management is not for the general public. However, multiple solutions based on threshold cryptography can be used to securely store keys under a client's password [28, 72]. A more likely explanation might be the need of existing attribute providers to completely change their work-flow and systems, without any direct financial, legal or customer requirements.

Deploying Privacy Preserving Identity Management. Fortunately, recent research have shown how PETs can be used to get certified attributes from issuers without modifying existing infrastructure, when such attributes can be retrieved from the provider through TLS-secured connections. The Town Crier system [117] shows how to construct certified attributes using secure hardware (like Intel SGX and using a TLS connection with a trusted provider). Concretely, they discussed how such certified attributes could be relayed to smart-contracts to allow more advanced decentralized user-attribute validation. Later, DECO [118] then showed how to remove the need for secure hardware and replace it with MPC while achieving the same goal. However, they extended their construction to also integrate with

⁶ Allowing the user to fully control the use of their credential through a single key can be conceptually advantageous.

12:10 SoK: Privacy-Enhancing Technologies in Finance

zero-knowledge proofs, allowing clients to construct certified proofs of arbitrary *predicates* on attributes from any provider, trusted through a TLS certificate which provides online access to the user's attributes. This could for example include a bank providing online banking access, where a user would then be able to construct a proof that they hold a bank account with e.g. at least \$20,000. If the user's government provides an online residency portal, then it could also be used for the users to prove that they legally reside in a given city in a given country without leaking their exact address.

The CanDID system [79] fully realizes a DID system with legacy support through either Town Crier or DECO. This is achieved through the usage of an MPC committee that validates legacy identity data and constructs a zero-knowledge friendly credential. Based on this credential, a user can prove arbitrary predicates on their attributes towards any provider.

Using attribute based credential allows the construction of fully private identity and attribute-based systems. However, in some situations full privacy might be undesirable, we would rather want to privately validate whether transactions are permissible based on attributes or identity, for example by ensuring that the identity of the credential holder is not on a deny-list. Kohlweiss *et al.* [74] showed that such a system can efficiently be constructed on top of credentials. The construction allows an auditor to specify any predicate on the attributes in a credential, where the identity of the credential holder gets leaked if the predicate is fulfilled. Such conditional privacy leakage could prove tremendously helpful in fighting money laundering as we discuss next.

Anti Money Laundering. Money laundering is the process of concealing the origins of money, such as financial gains from drug trafficking or other serious crimes, by changing its origin to a benevolent source. This is because criminals must acquire many services and goods in the regular economy: put simply, most luxury car dealers don't accept briefcases full of bills. Money laundering is a huge problem in the financial sector: the estimated amount of laundered money is at the level of 2-3% of the national GDP in the US alone, excluding tax evasion [94, Chap. 2].

Getting large amount of illegitimate cash into the financial system requires multiple steps and multiple accounts to avoid raising suspicion. Simply getting dirty money into the system is known as *placement*. A concrete and common approach for this is known as *smurfing*, where multiple legal people deposit small amounts of money for a criminal, with the promise of earning a small amount as a kick-back. After a period of time the smurfs move the money out of their accounts (minus their fee), by transferring to other accounts controlled by the criminal. If this process is done with small amounts, and the receiving party's account is not flagged, then smurfing is hard to identify⁷.

Once the money is in the legal financial system, it needs to be mixed with legitimate transfers, to counter suspicions caused by the initial transfers. This involves creating reasonable and justifiable transfers among multiple accounts of multiple entities in a process known as *layering*. By setting up a layering scheme through multiple banks, in different legal jurisdictions, using different legal entities, it becomes almost impossible to trace the flow of money. This is because the involved banks are (reasonably!) not allowed to communicate private account and customer information about the sender and recipient of a money transfer. After the layering, the money is finally moved out of financial institutions and into legitimate investments such as real estate or legitimate businesses. This last step is known as *integration*.

⁷ This step is sometimes also realized through other means, such as deposits from cash-driven businesses such as laundromats or food trucks.

What banks do to counter money laundering. As banks cannot share account and customer information with each other it is extremely hard for them to trace dirty money during layering. To address this, banks use multiple approaches usually subsumed as Anti-Money Laundering (AML) techniques. For example, banks internally use a *suspiciousness* score for customers. It is based on a *base* score, which is derived from the meta information about the account and its owner. The score may be derived from e.g. age of the account/holder, amount of money in the account, expected income and nationality of the owner. Through transfers, the base score is then updated, e.g. based on the score of the account a transfer goes to or comes from if both sender and receiver account are held by the same bank. If they instead are held by different banks, then metadata such as the amount of money going in/out and the frequency of the transfers is used in updates.

Finally, banks do have one common tool in measuring the suspiciousness of transfers, and that is a common, yet secret, grey list. This grey list contains accounts that have been deemed significantly suspicious, but for whom no provable money laundering has been identified yet. A transfer to or from a grey-listed account significantly increases the suspiciousness score. At certain time intervals, the suspiciousness score of an account is checked against a certain threshold and if the score is too high, then it gets flagged for manual⁸ inspection.

What can banks do? Due to GDPR and other privacy laws, it is not possible for banks to directly share meta-information about accounts or its owner without their consent. Furthermore, if a bank finds a flagged account it believes is engaging in illegal activities then when informing authorities, it must be able to *explain* to said authorities how they came to this conclusion. Hence the bank's judgements must be auditable by a third party. If the conclusion depends on data received from other financial institutions, the bank must be able to point to this data and the third party must trust it as well. While data from banks from within the same legal framework (the EU, USA, etc.) is usually considered as valid, data from international banks, in particular those from countries with a history of corruption, has less trustworthiness.

Implementing sufficient and efficient AML techniques is also difficult due to the quantities of information involved. AML technologies should ideally be scalable to include all transactions and accounts. At the same time, even a limited AML technology which only covers cross-country transfers or an arbitrary subset of accounts, could still make a substantial dent into the suspected large amount of money laundering currently going unnoticed.

Cryptography and AML. The conjunction of AML and MPC is new and the main bodies of work on the topic are by Zand *et al.* [115] and Egmond *et al.* [110]. Zand *et al.* show how computation on secret data can be used to notify an auditor of suspicious behavior. Egmond *et al.* show, in collaboration with multiple banks, how to use additively homomorphic encryption to obliviously update risk scores, and eventually (with consent from collaborating banks) decrypt the risk scores and flag accounts and customers appropriately.

However, related to this is the area of auditability of confidential transactions. As for example discussed by Tomescu *et al.* [108] where users are given a limited monthly “anonymity” budget. This budget is a certain amount of currency they are able to transfer anonymously per month. However, transfers surpassing this amount, is subject to deanonymization and clearance by a trusted auditor.

⁸ In practice it turns out that about 95% of automatically flagged accounts are false-positives.

Finally, we note that a survey of real world concepts using PETs to combat financial crime has been conducted by the Future of Financial Intelligence Sharing consortium [81]. Unfortunately, many of their mentioned solutions require a trusted party to be involved.

As mentioned above, AML in centralized banking is challenging as the transaction graph is hidden due to e.g. privacy regulations. However in the decentralized finance space, such transaction graphs are usually visible. This is why most popular cryptocurrencies, such as Bitcoin, Ethereum or Cardano, are only pseudonymous and not anonymous⁹. Cryptocurrency exchanges such as Coinbase and Binance allow to turn large amount of cryptocurrency into Fiat currencies. These exchanges are required by law to enforce know-your-customer (KYC) rules. Through the help of transaction graph analysis firms such as Chainalysis, it has become hard to launder money using pseudonymous cryptocurrencies.

Researchers have also proposed mechanisms to enforce AML even if transactions are kept private. This includes using an escrow system where anonymity and privacy can be broken in case suspicious activities occur, such as transfers to or from an account known to be used by criminals [97, 86, 8, 48]. Such escrow mechanisms does not necessarily imply the usage of a trusted third party, as the data for escrow activities can e.g. be shared using Threshold Secret Sharing. Another approach is to specify a small budget per client which they can use every month for anonymous payments. After the client has made more transactions than covered by this budget, any future transactions can be traced [113, 108]. Although seemingly a good compromise between privacy and security, this does still pose a risk to smurfing.

4 Markets & Transaction Settlement

In this section, we systematize PETs in market and settlement applications.

In financial markets, there is a need for auctions and markets with *fairness* guarantees, as rational actors are incentivized to collude and front-run honest parties, if the true valuation or trade-intent of the latter is revealed. Here, we first consider the *traditional finance* setting (§4.1), where the settlement of transactions is handled by traditional, asynchronous settlement processes. In the presence of a *public ledger* (§4.2), settlements occur *synchronously* and *immediately* after a transaction is completed. Such a mechanism also permits the “netting” of inter-bank payments (§4.3) to minimize the liquidity requirements on participating banks; this must done with PET approaches, since the public ledger would otherwise expose all individual payment orders, a clear breach of consumer privacy.

Finally, we highlight approaches to achieve bidder privacy in *demand-response* electricity markets (Appendix C in [10]), which coordinate the remote scheduling of power consuming devices to match forecast production from sustainable production sources; the submission of granular device-level information to an auction in the clear can reveal the activity and presence of customers at home, violating their privacy.

4.1 Markets in Traditional Finance

The first setting reflects an idealized view of *traditional finance*, where accounts and balances are generally maintained by financial institutions and considered private. Here, the settlement of auctions, or exchange transactions, occur asynchronously; whilst the *counterparty risk* from

⁹ We note that there exist privacy-focused blockchains like ZCash, Monero, or Dash that hide the transaction graph. Moreover, one can build private transactions on top of non-privacy focused blockchains using e.g. Zether [22] that leverages encryption and ZK proofs. Finally, mixers such as Tornado [92] take transfers from many users and put them into a holding account, from which they can later be transferred to the intended recipient.

■ **Table 3** Auctions & Markets: no benchmarks (○), preliminary benchmarks (◐), benchmarks achieving real-world performance with traditional market parameters (●).

Setting	Applications		Privacy	Benchmarks	PET	Works
Markets in Traditional Finance (§4.1)	Distributed sealed-bid auctions	Single-sided	Bid privacy	○	MPC	[56], [70], [25], [85]
		Double-sided	Bid privacy	◐	MPC	[19], [18]
		Public verifiability	-	○	HE+ZK	[89]
	Distributed Dark Pools	Continuous matching	Order privacy	◐	MPC	[33]
		Periodic matching	Order privacy	●	MPC	[34], [44]
			Order privacy	◐	FHE	[7]
		<i>with many assets</i>	Order privacy	●	MPC+HE	[34]
		<i>with many servers</i>	Order privacy	●	MPC	[44]
Public verifiability	-	○	HE+ZK	[107]		
Markets on Public Ledgers (§4.2)	Decentralized sealed-bid auctions	Single-sided	Order privacy	○	MPC+ZK	[6], [49], [60]
	Privacy-preserving Decentralized Exchanges	Futures	Net position privacy	◐	MPC+ZK	[80]
		Periodic matching	Balance privacy	◐	MPC+ZK	[68]
			Partial order privacy	Order privacy (Balance privacy)	◐	MPC+ZK
	Intent-based order matching	Balance privacy	Order privacy	◐	ZK	[20], [114]
		Balance privacy	Order privacy	●	WKA	[87]
	Settlement on Public Ledgers (§4.3)	Liquidity preserving inter-bank netting	-	Payment privacy	◐	ZK
-			Payment privacy & Robustness	◐	MPC+ZK	[50]
Demand-Response Markets (Appendix C in [10])	Distributed auctions for demand flexibility	Double-sided (Single buyer)	Device power constraint privacy	◐	MPC	[1], [119], [59]

defaulting on obligations implied by pending transactions is real, we consider it an orthogonal challenge addressed in the public ledger setting (§4.2, §4.3). Clearing prices and executed volumes are considered public information as this information is forwarded to institutions executing the settlement. This first setting intends to achieve resilience against dishonest venue operators and participants attempting to obtain a financial gain from unwarranted information flow. Communication between parties generally assumes direct, authenticated channels, implying the knowledge of identities and a public key infrastructure.

Distributed Sealed-bid Auctions. One-off, sealed-bid auctions are frequently performed in the sale of frequency-spectrum rights, government contracts, real-estate and other private items, such as art. In open-cry, single-sided auctions, bids are broadcast publicly until no additional bids are made. However, the leakage of bids or orders can be exploited by the adversary for financial gain. In Vickrey auctions, where the winner pays the price submitted by the second-highest bid, the auction operator collecting the submitted bids is incentivized to collude with other bidders to increase the second-highest bid price and maximize auction fees. The auction operator must also be *trusted* to not reveal anything about submitted bids, in order for all bidders to submit their *true valuation*. The advent of the public, commercial internet coincides with the first protocol proposals which permit the execution of one-time, sealed-bid auctions by distributing the role of the auction operator, thereby removing the need for a trusted auction venue.

Franklin et al. [56] propose a one-sided auction protocol, where the auction venue is distributed amongst multiple servers; bidders submit their signed bids as *verifiable secret-shares* (VSS) to participating servers during the bidding phase. Subsequently, bids and signatures are jointly reconstructed by all servers, upon which all bid information becomes public. Verifiable secret-sharing ensures that bidders submit well-formed bids. As long as a single server is honest, the reconstruction of bids will not occur before the end of the bidding phase. However, it is often important to protect the privacy of bids, even if they are not successful; the valuation may reveal a bidding strategy for another, related auction.

In the work of Harkavy *et al.* [70], MPC is deployed to maintain the privacy of all submitted bids; only the winning bid is made public. Naor *et al.* [85] propose a variant of MPC with garbled circuits, to reduce the rounds of communication, thereby improving performance. Cachin [25] builds a purpose-built, privacy-preserving protocol, which permits the comparison ($>$) of *prices* between two parties with the help of an untrusted third party. An auction determining the highest bidder is then constructed from this primitive.

The first work to demonstrate the feasibility of privacy-preserving (one-time) *double-sided, sealed-bid auctions* was proposed by Bogetoft *et al.* [19]. Later secure auctions were used in practice [18], to facilitate the auctioning of sugar beet delivery contracts in Denmark. Concretely, farmers producing sugar beets hold contracts which represent an obligation and right to deliver beets to the the (single) Danish sugar beet processor Danisco. The trading of such contracts permits the reallocation of contracts to the most efficient producers, but such an exchange run by Danisco would permit it to learn information about the economic circumstances of producers, potentially compromising sugar beet farmers during contract negotiations. The matching and determination of a price computation from 1229 buy and sell orders was achieved in approximately half an hour by an MPC committee of 3 servers; a throughput volume sufficient for a one-time auction, but unacceptable for traditional electronic security exchanges. Notice, however, that cryptographic techniques have improved drastically since this work. Such an auction would run much more efficiently today.

Publicly verifiable auction operators are proposed in the work of Parkes *et al.* [89], a weaker alternative to implementing the auction operator with MPC; instead, the dark pool venue is still operated by single entity, but provides cryptographic proofs that the auction algorithm is performed correctly by the venue operator. Whilst this prevents the auction venue from *manipulating* the correct evaluation of auction bids, it does not prevent the auction operator from leaking bid information to malicious participants.

Distributed Dark Pools. Dark pools have gained adoption in traditional finance as venues where submitted orders are not publicly accessible, thus minimizing the potential price impact caused by signalling *trade intent* to front-running market participants. As is the case with auction venues, the dark pool operator must be trusted to not *share* the order flow information with malicious participants, a trust assumption that is frequently violated in practice (Table 1 in [34]), motivating the need for distributing the role of the venue operator.

Whilst the matching of orders in a *secret* order book is a natural domain of MPC, we observe that current proposals illustrate specific configurations and architectures for distributed dark pools that may or may not match throughput observed in traditional, centralized dark pool markets. In particular, the choice of auction clearing algorithm remains a deciding feasibility factor. Whilst secret-sharing based MPC schemes permit secret computation with n participating servers, the runtime bottleneck generally lies in the amount of *communication* that is required between servers. This is because MPC has a specific model of defining computations, and certain auction clearing algorithms can be realized with less communication overhead in MPC than others.


In the case of auction clearing algorithms, which must compute a *clearing price* from current bids and sell orders, the *sorting* thereof by *price limit* induces many *comparisons* ($>$, $<$, $=$) between secret-shared values, which in turn imply sub-protocols generating the majority of communication cost. Alternatively, order matching based on volume only (where prices are determined by third party price feeds) can greatly accelerate throughput, as the expensive clearing price evaluation is not required. Furthermore, whether orders are processed *continuously* or *periodically* also greatly affects the real-world applicability of the following distributed dark pool protocols.

Continuous Double Auctions with MPC: A recent line of work by Smart *et al.* [33, 34, 44, 45] implements and examines real-world, double-side auction algorithms. In the initial work [33], continuous double auctions (CDA) are implemented in a distributed fashion across servers running an MPC. Continuous double auctions maintain a limit order book (LOB) where buy and sell orders are ordered by ascending and descending price respectively; *each incoming order* is matched against one or more LOB orders if it crosses the “spread” between best buy (or sell) prices; its remaining volume is then inserted into the LOB. It is also the most *expensive* exchange algorithm since (1) each single order must be matched against m other fulfilled orders and (2) its remaining trade volume must be inserted into a (potentially large) order book of N size. Benchmarking such an algorithm requires specifying the *expected* state of the order book, given the sensitivity of CDA run-time on (1) the average number of matched orders m ; and (2) the expected order book length N . In the dark CDA implementation of Cartlidge *et al.* [33], run with 3 servers and Shamir-sharing based MPC, a worst-case throughput of 34 – 43 orders per second for LOB parameters $m = 3$ and $N \approx 30$ is achieved. This work demonstrates that distributed CDA with MPC cannot yet match the throughput volumes of traditional CDA venues¹⁰. In contrast, *periodic* order matching greatly improves the performance of distributed dark pools.

Periodic Double Auctions with MPC: Periodic auctions in the dark pool setting implemented with MPC promise throughput that match those of traditional dark pool markets, as shown in these works by Cartlidge *et al.* [33, 34]. In periodic auctions, limit orders are submitted during an open auction period after which a clearing price is computed during the clearing phase, which maximizes the volume of matched orders (unmatched orders are carried over to the next round). In contrast to CDA algorithms, where orders are processed individually against a potentially large order book, periodic auctions only need to compute a single clearing price for the entire batch in a given period. In fact, real-world order execution throughput has been achieved with a realistic number of asset pairs.

In [34], the London Stock Exchange Group’s *Turquoise Plato Uncross*, a widely-used traditional dark pool supporting thousands of assets, is implemented with promising results. Based on volume observed on the real-world Turquoise Plato Uncross venue, it is assumed that order book clearing occurs at most every 5 seconds, where at most 2000 newly input orders must be processed across an asset universe of 4000 financial instruments. This throughput was successfully handled by smaller MPC committee sizes of 2 (dishonest majority) and 3 (honest-majority), but required multiple MPC instances, each handling orders trading a small subset of all assets (Figure 1). For example, ~ 280 MPC committee instances are each randomly assigned 16 assets in each round by a *gateway* engine. This gateway periodically

¹⁰In their work, the runtime of MPC pre-processing is neglected, which represents a non-trivial “hidden” computational cost that can be performed during “offline” hours or outsourced to dedicated pre-processing workers; pre-processing generally does not limit the maximum, sustainable throughput of MPC.




■ **Figure 1** In [34], a gateway MPC (A) distributes inbound received orders across multiple MPC committees (B,C,D) to improve order clearing throughput. MPC servers never learn the asset pairs its committee is assigned in each round.

reassigns asset subsets to new MPC instances in order to break potential linkages between orders across time periods. We note that auction algorithms are not entirely *oblivious*. Indeed, the direction of orders are leaked in [33, 34], whilst volumes and order limits remain private.

Complementary follow-up work by Da Gama *et al.* [44, 45] both focus on (single-asset) privacy-preserving *volume matching*, which enables further performance gains since the clearing prices are determined by an external reference price; [44] introduces an improved MPC volume matching algorithm which permits dummy orders and hides the trade direction. [45] scales volume matching up to MPC instances consisting of ~ 100 servers and shows the economic costs associated with operating such a single server in a MPC instance of up to 100 servers to be below ~ 0.10 USD and ~ 0.025 USD for computation and network communication respectively in each auction round; the negligible cost demonstrates the feasibility of *market participants* contributing to the distributed operation of dark pools.

Periodic Double Auctions with FHE: JP Morgan has demonstrated initial results in work by Balch *et al.* [7] to realize dark pool venues where the venue operator is not distributed, but computes the periodic volume matching over data encrypted under a jointly controlled public key; here, the secret encryption key material used in the *threshold fully homomorphic encryption* is held in secret-shared form by all participants (Figure 2). While the computation can be done by one party on the ciphertexts (not knowing their plain values), the participants then later take part in a so-called distributed decryption protocol which reconstructs the outcomes to the venue operator. Whilst [7] benchmark periodic volume matching implemented with threshold FHE, the omission of the *partial decryption sub-protocol* complicates the evaluation of its performance. Still, FHE offers an alternative approach to secret sharing-based MPC which promises competitive performance; fewer communication rounds are required, since computation is performed locally by the dedicated venue operator over encrypted data, although local computation (over encrypted data) is more costly.

Publicly, verifiable dark pool operator: Similar to verifiable one-sided auctions [89], a weaker notion of order privacy for dark pools is proposed in the following work by Thorpe *et al.* [107], where the venue operator only reveals a homomorphically encrypted order book to traders; the operator itself, however, maintains the encryption key and can thus compute over the order book plaintext. Each update to the public, encrypted order book triggered by a submitted order is accompanied with a zero-knowledge range proof generated by the operator; any public party can locally re-compute the claimed update over encrypted values and verify zero-knowledge range proofs that guarantee that the plaintext values lies within certain ranges, thereby enabling verification of comparison statements between encrypted values. This system ensures the correctness of each order book update without revealing the order details themselves. [107] implements the CDA algorithm in a publicly verifiable manner; as in [89], this approach does not prevent any misuse of the order information held by the operator.




■ **Figure 2** In [7], market auctions are implemented with fully homomorphic encryption (FHE); here, the encryption key is jointly generated by key servers (P1-P3), but the FHE evaluation is solely performed by the evaluator, who never learns the plaintext of the inputs or intermediary results. Decryption of the FHE output requires interaction with all key servers. In contrast, private computation with MPC in [33] requires interaction amongst MPC servers (P4-P6) for each (multiplicative) operation.

4.2 Markets on Public Ledgers

The advent of public ledger protocols [61, 73, 64] resulting from permissionless participation of servers across the public internet promises a truly “server-less” system of transaction settlements, no longer dependent on any single trusted intermediary. The state of the ledger is public and its integrity is publicly verifiable (by any online party) through local verification of all previously finalized transactions sequenced in form of an append-only list, or *blockchain*. Observe that access to such a global transaction history also implies a Turing-complete *state machine*. Hence it can be used to realize smart contracts, which represent user-deployed programs run by the “servers” that run the blockchain protocols. Smart contracts can be used to construct custom ledgers [ERC20, ERC721], but also to facilitate more advanced functionalities such as decentralized auctions or decentralized exchanges (DEX), which forgo the need for trusted venue operators. In contrast to traditional finance, market applications in the public ledger setting offer instant settlement; any market application implemented with smart contracts instances permits the simultaneous evaluation and settlement between participants. Despite scalability challenges arising from the vast number of participants running the blockchain backbone protocol, the promise of instant settlement would allow the mitigation of counter-party risk, a real cost to transactions conducted in traditional finance today.

However, the public verifiability of a public ledger also introduces novel challenges for financial applications; account balances are public by default and leak information about submitted bids, trades or margin positions; the latter must be backed by valid balances. In decentralized finance (DeFi) [112], front-running is indeed rampant in decentralized exchanges (DEX) [109], since pending transactions leak trade intent to the adversary which can precisely order and inject transactions to execute optimal front-running strategies. Thus, proposals have been made to implement private balances on public ledgers with publicly verifiable, non-interactive zero-knowledge [99, 23]. However, a privacy-preserving ledger (even with standard smart contract support) is generally *not sufficient* for privacy-preserving financial applications such as exchanges [13].

Privacy-preserving ledgers generally complicate the realization of *smart contracts*, since these must verify and update account balances known only to its *owners* according to an agreed-upon transition logic. For decentralized exchanges implemented in the privacy-preserving ledger setting, this requires the presence of a secure multiparty computation instance, to which users can privately input their trade orders and private balances; the



■ **Figure 3** We sketch the architecture of privacy-preserving smart contract applications in MPC with instant settlement on a (confidential) ledger; clients provide input parameters to the MPC instance, and forward financial deposits to a smart contract in a confidential manner. The MPC privately returns computation output to clients, but also authorizes a new financial distribution which is paid out to the clients by the smart contract functionality.

MPC then computes an updated DEX state and private balances, which are then updated on the ledger (Figure 3). Enforcing consistency between the *secret*, internal MPC state and *private* account balances on the ledger requires protocol design advances illustrated in the subsequent paragraphs. We emphasize that counter to popular belief, zero-knowledge is not sufficient to realize universally expressive, privacy-preserving smart contracts, as the witness (or secret state) for decentralized privacy-preserving applications are partially held by separate, distrusting parties; instead, function evaluation over private inputs from separate parties and secret-shared data is the natural domain of secure multiparty computation.

We note there are privacy-preserving smart contract proposals which shield *private data* [104, 103] held by individual users or *private contract logic* [20], but such techniques are generally limited in their expressiveness. The work of Bove *et al.*[20] only supports two parties, and is not widely used to realize privacy-preserving financial applications.

Sealed-bid Auctions (with Instant Settlement). The first work by Bag *et al.* [6] to realize sealed-bid auctions specifically in the setting of public ledgers focuses on using the blockchain as a *communication medium* instead of a settlement layer; as a permissionless protocol, any party can anonymously post an arbitrary message to the bulletin board, visible to all other parties. For protocols with low communication rounds, this is a practical solution; in particular, the simplicity of evaluating single-sided sealed-bid auctions permits task-specific secure multiparty protocols which only require public message broadcasts. The SEAL [6] protocol proposes the use of a *anonymous veto protocol* [69] requiring only two communication rounds, that is then repeated once by auction bidders for *each bit* of their bid price, thereby removing the necessity an auctioneer role entirely. In particular, the veto protocol of [69] receives the private input $b_i \in \{0, 1\}$ for party $i \in [n]$, for each of the m bits representing the permissible price range; the parties learn the highest bid, bit by bit. Each execution of the veto protocol will thus publicly output 1 if one of the users submits a veto; thus, by repeating the veto protocol for each bit position, all participants receive the highest bid without revealing the prices of failed bids. [6] assumes participants to behave according to the protocol (semi-honesty).

This mechanism was later adopted and hardened by FAST [49] to be secure against malicious participants not adhering to the protocol. Furthermore, [49] introduces *guaranteed settlement* on the public ledger featuring privacy-preserving deposits. Participants are thus committed to execute the payment for their bids if these are successful during the auctions. Cheating participants are penalized by having their deposits slashed and reimbursed to other


parties; non-interactive zero-knowledge proofs from all parties ensure that parties only submit a veto if it is consistent with their initial bid (in commitment form on the ledger); still, despite the privacy-preserving aspects of the anonymous veto protocol, privacy leakage occurs when the highest bidder learns when he overtakes the second highest bidder. The work of Ganesh *et al.* [60] adopts a similar construction which is proven to be game-theoretically secure; it is rational to pursue the honest protocol despite any strategy chosen by other players. The work of Chin *et al.* [40] does not employ zero-knowledge proofs to shield committed funds for a single-sided, sealed-bid auction; deposits are sent to committed, yet undeployed contracts and are thus indistinguishable from normal Ethereum transactions, a similar technique used in Breidenbach *et al.* [21] to commit inputs to smart contracts without revealing them to the front-running adversary. This approach only guarantees k -anonymity and relies on the presence of other, unrelated transactions.

Privacy-preserving Decentralized Exchanges. We note a number of recent proposals for privacy-preserving DEX applications in recent years; intent-based privacy-preserving DEX applications mirror the functionality of over-the-counter (OTC) venues (in traditional finance) and only require a public ledger, but do not scale well and are not widely deployed. Privacy-preserving and front-running secure DEX protocols generally involve private ledger deposits and perform the order matching in an MPC instance, as is the case in distributed Dark Pool proposals previously described in Section 4.1, but offer instant settlement following each DEX round (Figure 3).

Intent-based, privacy-preserving DEX. In the works of [20] and [87], a simpler model of a decentralized exchange is implemented; a bulletin board functionality provided by a public ledger permits a “maker” to broadcast their trade intent. An interested counter-party or “taker” then directly opens an authenticated communication channel with the maker to jointly perform a privacy-preserving atomic swap on the public ledger [20]. Intent-based DEX protocols resemble over-the-counter models in traditional finance. [87] introduces a “witness key agreement” (WKA) construction which preserves the privacy of the maker’s offer; the WKA allows a taker to establish a shared secret key with a maker which has posted its order in commitment form to the ledger. The key agreement protocol succeeds if the committed, private order fulfills a relation determined by the taker. This key permits subsequent anonymous communication with the maker to finalize the transaction.

Privacy-preserving Futures DEX. An interesting example of decentralized exchanges is illustrated by Massacci *et al.* [80], which realizes a futures exchanges modelled closely after the Chicago Mercantile Exchange; here, the *future obligation* (or contract) to buy or sell a commodity is traded. The net position of a market participant is the sum of both current liquidity balances and future obligations; importantly, a party holding a future to “sell” a given commodity, must always hold sufficient liquidity to acquire the respective commodity, as it otherwise would default on its contractual obligation. Thus, a net position that falls below zero must be liquidated to protect the counter-party of any futures contract held by the liquidated party. Achieving this in a privacy-preserving manner without *revealing the net position* of a party is the goal of the work of [80].

If the net position of a participant is revealed, price manipulations could be conducted with the explicit intention of forcing the liquidation of otherwise valid positions. Thus, [80] proposes a similar scheme to [99], where the net position of each account is committed in a cryptographic accumulator. The validity of each update to the account is proven in



■ **Figure 4** We adopt a netting example from [111]; processing of individual payment orders may fail due to a lack of liquidity (left), as balances must remain positive following execution of each individual payment. Netting relaxes this constraint; balances need only to be positive following execution of all payments orders (right).

zero-knowledge, whilst the trading venue is executed in a MPC instance, similarly to the Dark Pool proposals in Section 4.1. [80] requires parties participate in the protocol for each account update, even if this means the liquidation of their own account. We note that the subsequent privacy-preserving smart contract framework instantiated with MPC and a confidential ledger [9] achieves the privacy guarantees of [80] without permitting users to block application liveness.

Front-running Secure DEX. A general motivation for privacy in decentralized exchanges is the front-running of DEX applications in Decentralized Finance due to public transactions and accounts in the default ledger setting; Despite offering instant settlement of trades and transactions, pending user input authorizations generally broadcast a users trade intent before their finalization. To this end, P2DEX [11] proposes the first privacy-preserving decentralized exchange, which can operate and settle transactions across multiple ledger instances; clients submit orders to an MPC committee which computes the order matching and subsequently settles these on the respective public ledgers; since the trade inputs are private, front-running is mitigated. Follow-up work [9] generalizes this model to a setting with confidential accounts; here, all zero-knowledge proofs are moved *outside* the MPC computation, as computing such proofs *inside* the MPC remains generally unfeasible for real-world application. The work of Govindarajan *et al.* [68] realizes a privacy-preserving DEX in a similar manner; here, however, the actual order matching is computed in the clear of a smart contract over anonymized trade lists to accelerate the determination of a clearing price.

4.3 Inter-bank Netting on Public Ledgers

Inter-bank payment requests are currently submitted to the real-time gross settlement (RTGS) system managed by the central bank to update the accounts of sending and receiving financial institutions. In times of low liquidity, a bank may fail to honor *individual payment instructions*, as the liquidity requirement may exceed its balance and credit line granted by the central bank; a *gridlock* occurs, when a failed payment settlement prevents further payment instructions from being processed. Given the large payment volumes processed each day, liquidity saving mechanisms are implemented which settle payment instructions on a *netting basis* (Figure 4).

Recent work has proposed *distributing* the role of the RTGS operator with a public ledger protocol, whilst implementing efficient netting protocols with smart contracts [111, 84], thereby increasing system resiliency as the operational liability burden on the central bank

operator today is very high. Whilst the aforementioned works implement inter-bank netting of queued payments, the nature of public ledgers means that payment instructions are revealed to parties participating in the underlying blockchain backbone protocol. Instead, [32] proposes payment instructions to be posted to the ledger in *commitment form* accompanied with non-interactive zero-knowledge proofs attesting their well-formedness. Here, local netting solutions are computed *by each participating bank* and verified by a coordinating smart contract, which verifies correctness of all submitted, local netting solutions without revealing amounts and the identity of institutions. Since parties must compute partial netting solutions, the protocol of [32] is not *robust* against cheating participants, who can stall or abort the netting process by posting invalid partial netting proposals. In contrast, [50] computes the netting solution inside an MPC instance, thereby achieving *fault tolerance* against dishonest participants. Despite initial implementation benchmarks provided by works above, it remains an open question in what configuration such systems can scale to real-world payment settlement volume and what netting frequency is required in practice.

5 Future applications

We will now outline which other PET use cases could be of interest in the financial sector in the foreseeable future. While many of the use cases previously described in this work may also not yet be production-ready, we want to highlight areas in this section which we think deserve more attention by researchers and practitioners. This necessarily is of speculative nature, so the reader may see this as food for thought.

Voting. Voting is a standard mechanism in deciding on future policies. While in many cases it is sufficient to make the whole voting process public, this is not always possible. For example, a voter may fear repercussions or embarrassment if his or her vote becomes public. Hence, to ensure *honest digital voting*, cryptographic voting algorithms have to be used. These ensure that election outcomes can be computed while individual votes cannot be attributed to participants. While such cryptographic voting can be realized using MPC or FHE, a dedicated line of work started by Chaum [38] presents highly efficient dedicated voting protocols. Cryptographic voting mechanisms find interesting applications in the DeFi space, e.g. for privacy-preserving Decentralized Anonymous Organizations (DAOs). In particular, a treasury system for DAOs based on electronic voting has been proposed in [116] and a board room voting scheme based on smart contracts (and this amenable to the DAO scenario) has been proposed in [82]. We believe that these techniques may also be useful for coordination among classical banking institutions and other financial operations (*e.g.* shareholder meetings).

Fraud detection. Both insurance and gambling are known as industries where companies in the sector exchange information on their customers in order to detect fraud or exploitative customers¹¹. This information sharing may be problematic for privacy reasons, and it also leaks information about suspected but ultimately honest customers if done in plain. PETs such as PSI might be an interesting tool to construct a trusted intermediary. This intermediary can obtain information from participating companies and alerts them if e.g. more than 3 of them share the same customer. Here, PSI can ensure that only those customers are revealed that appear often enough.

¹¹ For example the infamous “Griffin Book”.

Better and fairer pattern recognition. In section 3 we have outlined how AML does benefit from recognition of suspicious patterns. Such patterns, if one wants to keep up in the digital age, must be learned from a large dataset and must be updated frequently. Moreover, many companies in an industry have an interest in pooling their data with other institutions for the purpose of learning these patterns. At the same time, they may not want to share raw customer or transaction data. Another, related area is assessing the credit risk of potential customers. Here, the risk becomes more accurate the more participants can contribute information or models. At the same time, input providers have an interest to keep their data private (for data protection or to protect intellectual property).

Both applications fall into the area of privacy-preserving Machine Learning [78, 54, 83] or confidential benchmarking [47] which are subfields of MPC. While these areas have received much attention recently¹², optimized applications to finance seem to be lacking.

Another important aspect is that (automatically generated) models should not be biased against certain groups. While fair machine learning itself is a rapidly developing field, its application to finance [51] may deserve more attention.

Privacy preserving mitigation of systemic risk. Audits of financial institutions guarantee that their balances plus credit cover outstanding obligations. This reduces counter-party risk and means that the overall system can rely less on biasable methods such as ratings and reputation. At the same time, an audited company may not want to open its books fully to the public, or it might not be guaranteed that these books are correct. [80] have shown how audits can be realized using ZK proofs, although limited to the futures market. We believe that this concept may be generalized to the wider financial system to permit privacy-preserving audits.

References

- 1 Aysajan Abidin, Abdelrahman Aly, Sara Cleemput, and Mustafa A Mustafa. An mpc-based privacy-preserving protocol for a local electricity trading market. In *Cryptology and Network Security: 15th International Conference, CANS 2016, Milan, Italy, November 14-16, 2016, Proceedings 15*, pages 615–625. Springer, 2016.
- 2 Shashank Agrawal, Peihan Miao, Payman Mohassel, and Pratyay Mukherjee. PASTA: PASsword-based threshold authentication. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 2042–2059. ACM Press, October 2018. doi:10.1145/3243734.3243839.
- 3 Ghada Almashaqbeh and Ravital Solomon. Sok: Privacy-preserving computing in the blockchain era. In *2022 IEEE 7th European Symposium on Security and Privacy (EuroSP)*, pages 124–139, 2022. doi:10.1109/EuroSP53844.2022.00016.
- 4 Nasser Alsalami and Bingsheng Zhang. Sok: A systematic study of anonymity in cryptocurrencies. In *2019 IEEE Conference on Dependable and Secure Computing (DSC)*, pages 1–9, 2019. doi:10.1109/DSC47296.2019.8937681.
- 5 Elli Androulaki, Jan Camenisch, Angelo De Caro, Maria Dubovitskaya, Kaoutar Elkhyaoui, and Björn Tackmann. Privacy-preserving auditable token payments in a permissioned blockchain system. In *AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY, USA, October 21-23, 2020*, pages 255–267. ACM, 2020. doi:10.1145/3419614.3423259.
- 6 Samiran Bag, Feng Hao, Siamak F Shahandashti, and Indranil Ghosh Ray. SEAL: Sealed-bid auction without auctioneers. *IEEE Transactions on Information Forensics and Security*, 15:2042–2052, 2019.

¹²Privacy-preserving Machine Learning opens up interesting use cases, but it does not come without its own problems. See [52] for a good overview.

- 7 Tucker Balch, Benjamin E Diamond, and Antigoni Polychroniadou. SecretMatch: inventory matching from fully homomorphic encryption. In *Proceedings of the First ACM International Conference on AI in Finance*, pages 1–7, 2020.
- 8 Amira Barki and Aline Gouget. Achieving privacy and accountability in traceable digital currency. Cryptology ePrint Archive, Report 2020/1565, 2020. URL: <https://eprint.iacr.org/2020/1565>.
- 9 Carsten Baum, James Hsin-yu Chiang, Bernardo David, and Tore Kasper Frederiksen. Eagle: Efficient Privacy Preserving Smart Contracts. *Cryptology ePrint Archive (To appear in Financial Cryptography and Data Security 2023)*, 2022. URL: <https://eprint.iacr.org/2022/1435>.
- 10 Carsten Baum, James Hsin-yu Chiang, Bernardo David, and Tore Kasper Frederiksen. SoK: Privacy-Enhancing Technologies in Finance, 2023. Full version. URL: <https://eprint.iacr.org/2023/122>.
- 11 Carsten Baum, Bernardo David, and Tore Kasper Frederiksen. P2DEX: privacy-preserving decentralized cryptocurrency exchange. In *International Conference on Applied Cryptography and Network Security*, pages 163–194. Springer, 2021.
- 12 Carsten Baum, Tore Kasper Frederiksen, Julia Hesse, Anja Lehmann, and Avishay Yanai. PESTO: proactively secure distributed single sign-on, or how to trust a hacked server. In *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*, pages 587–606. IEEE, 2020. doi:10.1109/EuroSP48549.2020.00044.
- 13 Carsten Baum, James Hsin yu Chiang, Bernardo David, Tore Kasper Frederiksen, and Lorenzo Gentile. Sok: Mitigation of front-running in decentralized finance. Cryptology ePrint Archive, Paper 2021/1628, 2021. To appear on the Proceedings of the The 2nd Workshop on Decentralized Finance (DeFi) in Association with Financial Cryptography 2022. URL: <https://eprint.iacr.org/2021/1628>.
- 14 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. ACM Press, May 1988. doi:10.1145/62212.62213.
- 15 Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society Press, May 2014. doi:10.1109/SP.2014.36.
- 16 Abhishek Bhowmick, Dan Boneh, Steve Myers, Kunal Talwar, and Karl Tarbe. The apple PSI system, 2021. Accessed on 02/08/2023. URL: https://www.apple.com/child-safety/pdf/Apple_PSI_System_Security_Protocol_and_Analysis.pdf.
- 17 Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012. doi:10.1145/2090236.2090263.
- 18 Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael I. Schwartzbach, and Tomas Toft. Secure multiparty computation goes live. In Roger Dingledine and Philippe Golle, editors, *FC 2009*, volume 5628 of *LNCS*, pages 325–343. Springer, Heidelberg, February 2009.
- 19 Peter Bogetoft, Ivan Damgård, Thomas Jakobsen, Kurt Nielsen, Jakob Pagter, and Tomas Toft. A practical implementation of secure auctions based on multiparty integer computation. In Giovanni Di Crescenzo and Avi Rubin, editors, *FC 2006*, volume 4107 of *LNCS*, pages 142–147. Springer, Heidelberg, February / March 2006.
- 20 Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 947–964. IEEE, 2020.

- 21 Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. Enter the hydra: Towards principled bug bounties and {Exploit-Resistant} smart contracts. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1335–1352, 2018.
- 22 Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In Joseph Bonneau and Nadia Heninger, editors, *FC 2020*, volume 12059 of *LNCS*, pages 423–443. Springer, Heidelberg, February 2020. doi:10.1007/978-3-030-51280-4_23.
- 23 Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In *International Conference on Financial Cryptography and Data Security*, pages 423–443. Springer, 2020.
- 24 Joseph Burleson, Michele Korver, and Dan Boneh. Privacy-protecting regulatory solutions using zero-knowledge proofs: Full paper, 2020. URL: <https://a16zcrypto.com/privacy-protecting-regulatory-solutions-using-zero-knowledge-proofs-full-paper/>.
- 25 Christian Cachin. Efficient private bidding and auctions with an oblivious third party. In *Proceedings of the 6th ACM Conference on Computer and Communications Security*, pages 120–127, 1999.
- 26 Jan Camenisch, Manu Drijvers, and Maria Dubovitskaya. Practical UC-secure delegatable credentials with attributes and their application to blockchain. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 683–699. ACM Press, October / November 2017. doi:10.1145/3133956.3134025.
- 27 Jan Camenisch and Thomas Groß. Efficient attributes for anonymous credentials. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 2008*, pages 345–356. ACM Press, October 2008. doi:10.1145/1455770.1455814.
- 28 Jan Camenisch, Anja Lehmann, Anna Lysyanskaya, and Gregory Neven. Memento: How to reconstruct your secrets from a single password in a hostile environment. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 256–275. Springer, Heidelberg, August 2014. doi:10.1007/978-3-662-44381-1_15.
- 29 Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 93–118. Springer, Heidelberg, May 2001. doi:10.1007/3-540-44987-6_7.
- 30 Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Heidelberg, August 2004. doi:10.1007/978-3-540-28628-8_4.
- 31 Jan Camenisch and Els Van Herreweghen. Design and implementation of the idemix anonymous credential system. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 21–30. ACM Press, November 2002. doi:10.1145/586110.586114.
- 32 Shengjiao Cao, Yuan Yuan, Angelo De Caro, Karthik Nandakumar, Kaoutar Elkhyaoui, and Yanyan Hu. Decentralized privacy-preserving netting protocol on blockchain for payment systems. In Joseph Bonneau and Nadia Heninger, editors, *Financial Cryptography and Data Security*, pages 137–155, Cham, 2020. Springer International Publishing.
- 33 John Cartlidge, Nigel P Smart, and Younes Talibi Alaoui. MPC joins the dark side. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*, pages 148–159, 2019.
- 34 John Cartlidge, Nigel P Smart, and Younes Talibi Alaoui. Multi-party computation mechanism for anonymous equity block trading: A secure implementation of turquoise plato uncross. *Intelligent Systems in Accounting, Finance and Management*, 28(4):239–267, 2021.
- 35 Benjamin Case, Richa Jain, Alex Koshelev, Andy Leiserson, Daniel Masny, Thurston Sandberg, Ben Savage, Erik Taubeneck, Martin Thomson, and Taiki Yamaguchi. Interoperable private attribution: A distributed attribution and aggregation protocol. Cryptology ePrint Archive, Paper 2023/437, 2023. URL: <https://eprint.iacr.org/2023/437>.

- 36 David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (abstract) (informal contribution). In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, page 462. Springer, Heidelberg, August 1988. doi:10.1007/3-540-48184-2_43.
- 37 David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *20th ACM STOC*, pages 11–19. ACM Press, May 1988. doi:10.1145/62212.62214.
- 38 David L Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.
- 39 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. ACM Press, October / November 2017. doi:10.1145/3133956.3134061.
- 40 Kota Chin, Keita Emura, Kazumasa Omote, and Shingo Sato. A Sealed-bid Auction with Fund Binding: Preventing Maximum Bidding Price Leakage. In *2022 IEEE International Conference on Blockchain (Blockchain)*, pages 398–405. IEEE, 2022.
- 41 Benny Chor, Shafi Goldwasser, Silvio Micali, and Baruch Awerbuch. Verifiable secret sharing and achieving simultaneity in the presence of faults (extended abstract). In *26th FOCS*, pages 383–395. IEEE Computer Society Press, October 1985. doi:10.1109/SFCS.1985.64.
- 42 Coinbase. Coinbase to acquire leading cryptographic security company, Unbound Security, November 2021. URL: <https://www.coinbase.com/blog/coinbase-to-acquire-leading-cryptographic-security-company-unbound-security>.
- 43 dlonyslus. Zero-knowledge rollups, 2023. Accessed on 02/08/2023. URL: <https://ethereum.org/en/developers/docs/scaling/zk-rollups/>.
- 44 Mariana Botelho da Gama, John Cartlidge, Antigoni Polychroniadou, Nigel P Smart, and Younes Talibi Alaoui. Kicking-the-bucket: Fast privacy-preserving trading using buckets. In *International Conference on Financial Cryptography and Data Security*, pages 20–37. Springer, 2022.
- 45 Mariana Botelho da Gama, John Cartlidge, Nigel P Smart, and Younes Talibi Alaoui. All for one and one for all: Fully decentralised privacy-preserving dark pool trading using multi-party computation. *Cryptology ePrint Archive*, 2022. URL: <https://eprint.iacr.org/2022/923>.
- 46 Mariana Botelho da Gama, John Cartlidge, Nigel P Smart, and Younes Talibi Alaoui. Privacy-preserving dark pools. *Cryptology ePrint Archive*, 2022. URL: <https://eprint.iacr.org/2022/923>.
- 47 Ivan Damgård, Kasper Damgård, Kurt Nielsen, Peter Sebastian Nordholt, and Tomas Toft. Confidential benchmarking based on multiparty computation. In Jens Grossklags and Bart Preneel, editors, *FC 2016*, volume 9603 of *LNCS*, pages 169–187. Springer, Heidelberg, February 2016.
- 48 Ivan Damgård, Chaya Ganesh, Hamidreza Khoshakhlagh, Claudio Orlandi, and Luisa Siniscalchi. Balancing privacy and accountability in blockchain identity management. In Kenneth G. Paterson, editor, *CT-RSA 2021*, volume 12704 of *LNCS*, pages 552–576. Springer, Heidelberg, May 2021. doi:10.1007/978-3-030-75539-3_23.
- 49 Bernardo David, Lorenzo Gentile, and Mohsen Pourpouneh. FAST: fair auctions via secret transactions. In *International Conference on Applied Cryptography and Network Security*, pages 727–747. Springer, 2022.
- 50 Angelo De Caro, Andrew Miller, and Amit Agarwal. Privacy-Preserving Decentralized Multi-Party Netting, September 29 2022. US Patent App. 17/216,644, <https://patents.google.com/patent/US20220309492A1/en>.
- 51 Leo de Castro, Jiahao Chen, and Antigoni Polychroniadou. Cryptocredit: securely training fair models. In *Proceedings of the First ACM International Conference on AI in Finance*, pages 1–8, 2020.
- 52 Emiliano De Cristofaro. A critical overview of privacy in machine learning. *IEEE Security & Privacy*, 19(4):19–27, 2021. doi:10.1109/MSEC.2021.3076443.

- 53 Apple Differential Privacy Team. Learning with privacy at scale. Accessed on 02/08/2023. URL: <https://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf>.
- 54 Wenliang Du, Mikhail J Atallah, et al. Privacy-preserving cooperative scientific computations. In *csfw*, volume 1, page 273, 2001.
- 55 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 265–284. Springer, Heidelberg, March 2006. doi:10.1007/11681878_14.
- 56 Matthew K Franklin and Michael K Reiter. The design and implementation of a secure auction service. *IEEE Transactions on Software Engineering*, 22(5):302–312, 1996.
- 57 Tore Kasper Frederiksen. A holistic approach to enhanced security and privacy in digital health passports. In Delphine Reinhardt and Tilo Müller, editors, *ARES 2021: The 16th International Conference on Availability, Reliability and Security, Vienna, Austria, August 17-20, 2021*, pages 133:1–133:10. ACM, 2021. doi:10.1145/3465481.3469212.
- 58 Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 1–19. Springer, Heidelberg, May 2004. doi:10.1007/978-3-540-24676-3_1.
- 59 Mariana Gama, Fairouz Zobiri, and Svetla Nikova. Multi-party computation auction mechanisms for a p2p electricity market with geographical prioritization, 2022. URL: <https://www.esat.kuleuven.be/cosic/publications/article-3526.pdf>.
- 60 Chaya Ganesh, Bhavana Kanukurthi, and Girisha Shankar. Secure Auctions in the Presence of Rational Adversaries. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 1173–1186, 2022.
- 61 Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 281–310. Springer, 2015.
- 62 Christina Garman, Matthew Green, and Ian Miers. Decentralized anonymous credentials. In *NDSS 2014*. The Internet Society, February 2014.
- 63 Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. doi:10.1145/1536414.1536440.
- 64 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th symposium on operating systems principles*, pages 51–68, 2017.
- 65 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, October 1986. doi:10.1109/SFCS.1986.47.
- 66 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. ACM Press, May 1987. doi:10.1145/28395.28420.
- 67 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985. doi:10.1145/22145.22178.
- 68 Kavya Govindarajan, Dhinakaran Vinayagamurthy, Praveen Jayachandran, and Chester Rebeiro. Privacy-preserving decentralized exchange marketplaces. In *2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9. IEEE, 2022.
- 69 Feng Hao and Piotr Zielinski. A 2-round anonymous veto protocol. In *International Workshop on Security Protocols*, pages 202–211. Springer, 2009.
- 70 Michael Harkavy, J Doug Tygar, and Hiroaki Kikuchi. Electronic auctions with private bids. In *USENIX Workshop on Electronic Commerce*, 1998.

- 71 Space systems — Definition of the Technology Readiness Levels (TRLs) and their criteria of assessment. Standard, International Organization for Standardization, Geneva, CH, November 2013.
- 72 Stanislaw Jarecki, Aggelos Kiayias, Hugo Krawczyk, and Jiayu Xu. TOPPSS: Cost-minimal password-protected secret sharing based on threshold OPRF. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki Kikuchi, editors, *ACNS 17*, volume 10355 of *LNCS*, pages 39–58. Springer, Heidelberg, July 2017. doi:10.1007/978-3-319-61204-1_3.
- 73 Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual international cryptology conference*, pages 357–388. Springer, 2017.
- 74 Markulf Kohlweiss, Anna Lysyanskaya, and An Nguyen. Privacy-preserving blueprints. *IACR Cryptol. ePrint Arch.*, page 1536, 2022. URL: <https://eprint.iacr.org/2022/1536>.
- 75 Team KZen. Bitcoin wallet powered by two-party ECDSA extended abstract. URL: <https://github.com/ZenGo-X/gotham-city/blob/master/white-paper/white-paper.pdf>.
- 76 Protocol Labs. Filecoin: A decentralized storage network, 2017. URL: <https://filecoin.io/filecoin.pdf>.
- 77 Yehuda Lindell and Ariel Nof. Fast secure multiparty ecDSA with practical distributed key generation and applications to cryptocurrency custody. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1837–1854, 2018.
- 78 Yehuda Lindell and Benny Pinkas. Privacy preserving data mining. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 36–54. Springer, Heidelberg, August 2000. doi:10.1007/3-540-44598-6_3.
- 79 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. IEEE Computer Society Press, May 2021. doi:10.1109/SP40001.2021.00038.
- 80 Fabio Massacci, Chan Nam Ngo, Jing Nie, Daniele Venturi, and Julian Williams. FuturesMEX: secure, distributed futures market exchange. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 335–353. IEEE, 2018.
- 81 Nick Maxwell. Case studies of the use of privacy preserving analysis to tackle financial crime, January 2021. URL: https://www.future-fis.com/uploads/3/7/9/4/3794525/ffis_innovation_and_discussion_paper_-_case_studies_of_the_use_of_privacy_preserving_analysis_-_v.1.3.pdf.
- 82 Patrick McCorry, Siamak F. Shahandashti, and Feng Hao. A smart contract for boardroom voting with maximum voter privacy. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 357–375. Springer, Heidelberg, April 2017.
- 83 Payman Mohassel and Yupeng Zhang. SecureML: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy*, pages 19–38. IEEE Computer Society Press, May 2017. doi:10.1109/SP.2017.12.
- 84 Ken Naganuma, Masayuki Yoshino, Hisayoshi Sato, Nishio Yamada, Takayuki Suzuki, and Noboru Kunihiro. Decentralized netting protocol over consortium blockchain. In *2018 International Symposium on Information Theory and Its Applications (ISITA)*, pages 174–177. IEEE, 2018.
- 85 Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 129–139, 1999.
- 86 Neha Narula, Willy Vasquez, and Madars Virza. zkLedger: Privacy-Preserving Auditing for Distributed Ledgers. In Sujata Banerjee and Srinivasan Seshan, editors, *15th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2018, Renton, WA, USA, April 9-11, 2018*, pages 65–80. USENIX Association, 2018. URL: <https://www.usenix.org/conference/nsdi18/presentation/narula>.

- 87 Chan Nam Ngo, Fabio Massacci, Florian Kerschbaum, and Julian Williams. Practical witness-key-agreement for blockchain-based dark pools financial trading. In *International Conference on Financial Cryptography and Data Security*, pages 579–598. Springer, 2021.
- 88 Christian Paquin. U-Prove Technology Overview V1.1. Tech report, Microsoft Corporation, April 2013. URL: <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/U-Prove20Technology200verview20V1.120Revision202.pdf>.
- 89 David C Parkes, Michael O Rabin, Stuart M Shieber, and Christopher Thorpe. Practical secrecy-preserving, verifiably correct and trustworthy auctions. *Electronic Commerce Research and Applications*, 7(3):294–312, 2008.
- 90 THE EUROPEAN PARLIAMENT and THE COUNCIL OF THE EUROPEAN UNION. Directive (EU) 2015/2366 of the european parliament and of the council, November 2015. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32015L2366&from=EN>.
- 91 Juha Partala, Tri Hong Nguyen, and Susanna Pirttikangas. Non-interactive zero-knowledge for blockchain: A survey. *IEEE Access*, 8:227945–227961, 2020. doi:10.1109/ACCESS.2020.3046025.
- 92 Alexey Pertsev, Roman Semenov, and Roman Storm. Tornado Cash Privacy Solution, version 1.4, December 2019. URL: https://web.archive.org/web/20211026053443/https://tornado.cash/audits/TornadoCash_whitepaper_v1.4.pdf.
- 93 Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *21st ACM STOC*, pages 73–85. ACM Press, May 1989. doi:10.1145/73007.73014.
- 94 Peter Reuter and Edwin M. Truman. *Chasing Dirty Money: The Fight Against Money Laundering*. Peterson Institute for International Economics, 2004.
- 95 Ronald L Rivest, Len Adleman, Michael L Dertouzos, et al. On data banks and privacy homomorphisms. *Foundations of secure computation*, 4(11):169–180, 1978.
- 96 Umut Sahin. How to get started with concrete - zama’s fully homomorphic encryption compiler, 2023. Accessed on 02/08/2023. URL: <https://www.zama.ai/post/how-to-started-with-concrete-zama-fully-homomorphic-encryption-compiler>.
- 97 Tomas Sander and Amnon Ta-Shma. Flow control: A new approach for anonymity control in electronic cash systems. In Matthew Franklin, editor, *FC’99*, volume 1648 of *LNCS*, pages 46–61. Springer, Heidelberg, February 1999.
- 98 Olivier Sanders. Efficient redactable signature and application to anonymous credentials. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 628–656. Springer, Heidelberg, May 2020. doi:10.1007/978-3-030-45388-6_22.
- 99 Eli Ben Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *2014 IEEE symposium on security and privacy*, pages 459–474. IEEE, 2014.
- 100 Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- 101 Manu Sporny, Dave Longley, and David Chadwick. Verifiable credentials data mode, 2022. URL: <https://www.w3.org/TR/vc-data-model/>.
- 102 Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Ori Steele, and Christopher Allen. Decentralized identifiers (DIDs), 2022. URL: <https://www.w3.org/TR/did-core>.
- 103 Samuel Steffen, Benjamin Bichsel, Roger Baumgartner, and Martin Vechev. ZeeStar: Private Smart Contracts by Homomorphic Encryption and Zero-knowledge Proofs. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1543–1543. IEEE Computer Society, 2022.
- 104 Samuel Steffen, Benjamin Bichsel, Mario Gersbach, Noa Melchior, Petar Tsankov, and Martin Vechev. zkay: Specifying and enforcing data privacy in smart contracts. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 1759–1776, 2019.

- 105 Zhichuang Sun, Ruimin Sun, Long Lu, and Alan Mislove. Mind your weight(s): A large-scale study on insufficient machine learning model protection in mobile apps. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 1955–1972. USENIX Association, August 2021.
- 106 Shahroz Tariq, Sowon Jeon, and Simon S. Woo. Am I a real or fake celebrity? evaluating face recognition and verification apis under deepfake impersonation attack. In Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini, editors, *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, pages 512–523. ACM, 2022. doi:10.1145/3485447.3512212.
- 107 Christopher Thorpe and David C Parkes. Cryptographic securities exchanges. In *International Conference on Financial Cryptography and Data Security*, pages 163–178. Springer, 2007.
- 108 Alin Tomescu, Adithya Bhat, Benny Applebaum, Ittai Abraham, Guy Gueta, Benny Pinkas, and Avishay Yanai. UTT: Decentralized ecash with accountable privacy. *Cryptology ePrint Archive*, Report 2022/452, 2022. URL: <https://eprint.iacr.org/2022/452>.
- 109 Christof Ferreira Torres, Ramiro Camino, and Radu State. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In Michael Bailey and Rachel Greenstadt, editors, *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 1343–1359. USENIX Association, 2021. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/torres>.
- 110 Marie Beth van Egmond, Thomas Rooijakkers, and Alex Sangers. Privacy-Preserving Collaborative Money Laundering Detection. *ERCIM News*, 2021(126), 2021. URL: <https://ercim-news.ercim.eu/en126/special/privacy-preserving-collaborative-money-laundering-detection>.
- 111 Xin Wang, Xiaomin Xu, Lance Feagan, Sheng Huang, Limei Jiao, and Wei Zhao. Inter-bank payment system on enterprise blockchain platform. In *2018 IEEE 11th international conference on cloud computing (CLOUD)*, pages 614–621. IEEE, 2018.
- 112 Sam M Werner, Daniel Perez, Lewis Gudgeon, Ariaah Klages-Mundt, Dominik Harz, and William J Knottenbelt. Sok: Decentralized finance (defi). *arXiv preprint arXiv:2101.08778*, 2021. URL: <https://arxiv.org/abs/2101.08778>.
- 113 Karl Wüst, Kari Kostiaainen, Vedran Capkun, and Srdjan Capkun. PRCash: Fast, private and regulated transactions for digital currencies. In Ian Goldberg and Tyler Moore, editors, *FC 2019*, volume 11598 of *LNCS*, pages 158–178. Springer, Heidelberg, February 2019. doi:10.1007/978-3-030-32101-7_11.
- 114 Alex Luoyuan Xiong, Binyi Chen, Zhenfei Zhang, Benedikt Bünz, Ben Fisch, Fernando Krell, and Philippe Camacho. Veri-zexe: Decentralized private computation with universal setup. *Cryptology ePrint Archive*, 2022.
- 115 Arman Zand, James Orwell, and Eckhard Pfluegel. A Secure Framework for Anti-Money-Laundering using Machine Learning and Secret Sharing. In *2020 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, pages 1–7, 2020. doi:10.1109/CyberSecurity49315.2020.9138889.
- 116 Bingsheng Zhang, Roman Oliynykov, and Hamed Balogun. A treasury system for cryptocurrencies: Enabling better collaborative intelligence. In *NDSS 2019*. The Internet Society, February 2019.
- 117 Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. Town crier: An authenticated data feed for smart contracts. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 270–282. ACM Press, October 2016. doi:10.1145/2976749.2978326.
- 118 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. ACM Press, November 2020. doi:10.1145/3372297.3417239.

12:30 SoK: Privacy-Enhancing Technologies in Finance

- 119 Fairouz Zobiri, Mariana Gama, Svetla Nikova, and Geert Deconinck. A Privacy-Preserving Three-Step Demand Response Market Using Multi-Party Computation. In *13th Int. Conf. Innov. Smart Grid Technol.(ISGT North Am. 2022), Washingt. DC (to Appear)*, 2022. URL: <https://www.esat.kuleuven.be/cosic/publications/article-3451.pdf>.