Workshop on Cryptographic Tools for Blockchains 2024 Program

Session 1

Invited Talk: Consensus in blockchains: Overview and recent results

Speaker: Christian Cachin

Reaching consensus despite faulty or corrupted nodes is a central question in distributed computing; it has received renewed attention over the last few years because of its importance for cryptocurrencies and blockchain networks. Modern consensus protocols in this space have relied on a number of different methods for the nodes to influence protocol decisions. Such assumptions include (1) traditional voting, where each node has one vote, (2) weighted voting, where voting power is proportional to stake in an underlying asset, and (3) proof-of-X, which demonstrates a cryptographically verifiable investment of a resource X, such as storage space, time waited, or computational work. This talk will give an overview of blockchain consensus methods and then highlight recent work on constructing new consensus protocols and analyzing existing ones.

Christian Badertscher, Mahdi Sedaghat and Hendrik Waldner. Unlinkable Policy-Compliant Signatures for Compliant and Decentralized Anonymous Payments

Speaker: Christian Badertscher

Privacy-preserving payment systems face the difficult task of balancing privacy and accountability: on one hand, users should be able to transact privately and anonymously, on the other hand, no illegal activities should be tolerated. The challenging question of finding the right balance lies at the core of the research on accountable privacy that stipulates the use of cryptographic techniques for policy enforcement.

In this talk, we present unlinkable Policy-Compliant Signatures (ul-PCS), an enhanced cryptographic primitive extending the work of Badertscher et al. (TCC 21). We give rigorous definitions, formally proven constructions, and benchmarks using our prototype developed using CharmCrypto. Unlinkable PCS has the following unique combination of features:

1. It is an enhanced signature scheme where the public key encodes in a privacy-preserving way the user's verifiable credentials (obtained from a credential authority).

2. Signatures can be created (and later publicly verified) by additionally specifying a recipient's public key aside of the to-be-signed message. A valid signature can only ever be created if the attributes x of the signer and the attributes y of the receiver fulfill some global policy F(x,y).

3. The signature can be created by the signer just knowing the recipient's public key; there is no further interaction needed no attributes are leaked (beyond the validity of the policy).

4. Once credentials are obtained, a user can generate fresh public keys without interacting with the credential authority.

By merging the act of signing a transaction with the act of providing an assurance about the involved participants being compliant with complex policies, yet retain that participants are able to change addresses without the involvement of an authority, we showcase how ul-PCS constitutes a crucial step towards achieving a technology that improves regulatory compliance of privacy coins such as Monero or Zcash.

Elizabeth Crites, Aggelos Kiayias and Amirreza Sarencheh. SyRA: Sybil-Resilient Anonymous Signatures with Applications to Decentralized Identity

Speaker: Amirreza Sarencheh

We introduce a new cryptographic primitive called Sybil-Resilient Anonymous (SyRA) signature, which enables users to generate, on demand, unlinkable pseudonyms tied to any given context and issue digital signatures on their behalf. Concretely, given a personhood relation, an issuer (who may be a distributed entity) enables users to prove their personhood and extract an associated long-term key, which can then be used to issue signatures for any given context and message. Sybil-resilient anonymous signatures achieve two key security properties: 1) Sybil resilience, ensuring that every user is entitled to at most one pseudonym per context, and 2) anonymity, requiring that no information about the user is leaked through their various pseudonyms or the signatures they issue on their pseudonyms' behalf.

We conceptualize SyRA signatures as an ideal functionality in the Universal Composition (UC) setting and realize the functionality via an efficient, pairing-based construction that utilizes two levels of verifiable random functions (VRFs) and may be of independent interest. One of the key features of this approach is the statelessness of the issuer: we achieve the core properties of Sybil resilience and anonymity without requiring the issuer to retain any information about past user interactions.

SyRA signatures have various applications in multiparty systems such as e-voting (e.g., for decentralized governance), privacy-preserving regulatory compliance (e.g., AML/CFT checks), and cryptocurrency airdrops, making them an attractive option for deployment in decentralized identity (DID) systems.

Session 2

Noemi Glaeser, István András Seres, Michael Zhu and Joseph Bonneau. Cicada: A framework for private non-interactive on-chain auctions and voting

Speaker: Noemi Glaeser

Auction and voting schemes play a crucial role in the Web3 ecosystem. Yet currently deployed implementations either lack privacy or require at least two rounds, hindering usability and security. We introduce Cicada, a general framework for using linearly homomorphic time-lock puzzles (HTLPs) to enable provably secure, non-interactive private auction and voting protocols. We instantiate our framework with an efficient new HTLP construction and novel packing techniques that enable succinct ballot correctness proofs independent of the number of candidates. We demonstrate the practicality of our approach by implementing our protocols for the Ethereum Virtual Machine (EVM).

Hao Chung, Elisaweta Masserova, Elaine Shi and Sri Aravinda Krishnan Thyagarajan. Rapidash: Atomic Swaps Secure under User-Miner Collusion

Speaker: Elisaweta Masserova

User-to-user cross-chain trading is a fundamental primitive in the blockchain space and Decentralized Finance (DeFi). One way to achieve cross-chain trading in a truly decentralized manner, i.e., without using trusted third parties, is by using atomic swaps. However, recent works revealed that Hashed Time-Lock Contract, a crucial building block of the existing atomic swap schemes, is entirely insecure in the presence of user-miner collusion. Specifically, a user can bribe the miners of the blockchain to help it cheat — a phenomenon known as Miner Extractable Value.

In this work, we provide the first rigorous formal treatment of fair trading on blockchains, where users and miners may enter arbitrary binding contracts on the side. We propose a new atomic swap protocol called Rapidash and prove that it is incentive-compatible in the presence of user-miner collusion. In particular, we show that Rapidash satisfies a coalition-resistant Nash equilibrium absent external incentives. Finally, to showcase the practicality and compatibility of Rapidash with a wide range of blockchain systems, we present instantiations of Rapidash that are compatible with Bitcoin and Ethereum, and incur only a minimal overhead in terms of costs for the users.

Yue Guo, Harish Karthikeyan and Antigoni Polychroniadou. PriDe CT: Towards Public Consensus, Private Transactions, and Forward Secrecy in Decentralized Payments

Speaker: Harish Karthikeyan

Anonymous Zether, proposed by Bunz et al. (FC 2020) and subsequently improved by Diamond (IEEE S&P 2021), is an account-based confidential payment mechanism that works by using a smart contract to achieve privacy (i.e., the identity of receivers to transactions and payloads is hidden). In this work, we look at simplifying the existing protocol while also achieving batching of transactions for multiple receivers, ensuring consensus and forward secrecy. To the best of our knowledge, this work is the first to formally study the notion of forward secrecy in the setting of blockchain, borrowing a very popular and useful idea from the world of secure messaging. Specifically, we introduce:

- FUL-Zether, a forward-secure version of Zether (Bunz et al., FC 2020).
- PRIvate DEcentralized Confidential Transactions (PriDe CT), a much-simplified version of Anonymous Zether that achieves competitive performance and enables batching of transactions for multiple receivers.
- PRIvate DEcentralized Forward-secure Until Last update Confidential Transactions (PriDeFUL CT), a forward-secure version of PriDe CT.

We also present an open-source, Ethereum-based implementation of our system. PriDe CT uses linear homomorphic encryption like Anonymous Zether but with simpler zero-knowledge proofs. PriDeFUL CT uses an updatable public key encryption scheme to achieve forward secrecy by introducing a new DDH-based construction in the standard model.

In terms of transaction sizes, Quisquis (Asiacrypt 2019), which is the only cryptocurrency that supports batchability (albeit in the UTXO model), has 15 times more group elements than PriDe CT. Meanwhile, for a ring of N receivers, Anonymous Zether requires 6 log N more terms even without accounting for the ability to batch in PriDe CT. Further, our implementation indicates that, for N=32, even if there were 7 intended receivers, PriDe CT outperforms Anonymous Zether in proving time and gas consumption.

David Arroyo, Sergio Chica, Jesús Díaz and Andrés Marín-López. 10 years of implementation of an usable group signature library: contributing to the design of decentralized identity management systems

Speaker: David Arroyo

This communication summarizes our efforts in designing, implementing, and applying group signatures to concrete practical scenarios where privacy and security must be conveniently harmonized. We will provide an overview of the definition of a "utility, privacy, and then utility again" paradigm, and some insights about the way the provision of usable privacy-enhancing tools can help the implementation of such a paradigm. We will focus on explaining our Libgroupsig software library, which is a rare example of a usable and easy-to-deploy library for the development of privacy-respectful identity management using group signatures.

The library, which is publicly available, is expected to help the community verify the good implementation of group signatures. In addition, we will discuss the potential of leveraging the different types of group signatures in the library to achieve different types of "utility-privacy trade-offs" and will showcase some practical applications implementing the notion of "privacy by design," while maintaining as much compatibility and interoperability with available infrastructures as possible.

Session 3

Invited Talk

Speaker: Jens Groth

Zero-knowledge proofs are powering an increasing number of applications in the blockchain space such as rollups, bridges, and L1s. The Ethereum Foundation is talking about a zero-knowledge singularity, where most of the activity takes place in L2s and the main job of the Ethereum chain is to order and publish zero-knowledge proofs of transaction batches. A main driver of adoption is that proofs can be succinct, giving low-cost verifiability without having to recompute.

Designing proofs directly for applications is cumbersome and error-prone. Zero-knowledge virtual machines (zkVMs) attempt to make it easy for developers to express the statements they want to prove. When using a zkVM, you compile a program written in a high-level language, e.g., Rust or Solidity, to a VM program. The zkVM then executes the VM program and attaches a succinct proof to the VM output that it has been correctly computed.

Istvan Andras Seres, Noemi Glaeser and Joseph Bonneau. Naysayer proofs

Speaker: Noemi Glaeser

This work introduces the notion of naysayer proofs. We observe that in numerous (zero-knowledge) proof systems, it is significantly more efficient for the verifier to be convinced by a so-called naysayer that a false proof is invalid than it is to check that a genuine proof is valid. We show that every NP language has logarithmic-size and constant-time naysayer proofs. We also show practical constructions for several example proof systems, including FRI polynomial commitments, post-quantum secure digital signatures, and verifiable shuffles. Naysayer proofs enable an interesting new optimistic verification mode potentially suitable for resource-constrained verifiers, such as smart contracts.

Andrea Cerulli, Aisling Connolly, Gregory Neven, Franz-Stefan Preiss and Victor Shoup. vetKeys: How a Blockchain Can Keep Many Secrets

Speaker: Aisling Connolly

We propose a new cryptographic primitive called verifiably encrypted threshold key derivation (vetKD) that extends identity-based encryption with a decentralized way of deriving decryption keys. We show how vetKD can be leveraged on modern blockchains to build scalable decentralized applications (dapps) for a variety of purposes, including preventing front-running attacks on decentralized finance (DeFi) platforms, end-to-end encryption for decentralized messaging and social networks (SocialFi), cross-chain bridges, as well as advanced cryptographic primitives such as witness encryption and one-time programs that previously could only be built from secure hardware or using a trusted third party. All of that is achieved by secret-sharing just a single secret key.

Session 4

Chen-Da Liu-Zhang, Elisaweta Masserova, João Ribeiro, Pratik Soni and Sri Aravindakrishnan Thyagarajan. Improved YOSO Randomness Generation with Worst-Case Corruptions

Speaker: Elisaweta Masserova

We study the problem of generating public unbiased randomness in a distributed manner within the recent You Only Speak Once (YOSO) framework for stateless multiparty computation, introduced by Gentry et al. in CRYPTO 2021. Such protocols are resilient to adaptive denial-of-service attacks and are, by their stateless nature, especially attractive in permissionless environments. While most works in the YOSO setting focus on independent random corruptions, we consider YOSO protocols with worst-case corruptions, a model introduced by Nielsen et al. in CRYPTO 2022.

Prior work on YOSO public randomness generation with worst-case corruptions designed information-theoretic protocols for t corruptions with either n = 6t + 1 or n = 5t roles, depending on the adversarial network model. However, a major drawback of these protocols is that their communication and computational complexities scale exponentially with t. In this work, we complement prior inefficient results by presenting and analyzing simple and efficient protocols for YOSO public randomness generation secure against worst-case corruptions in the computational setting. Our first protocol is based on publicly verifiable secret sharing and uses n = 3t+2 roles. Since this first protocol requires setup and somewhat heavy cryptographic machinery, we also provide a second lighter protocol based on ElGamal commitments and verifiable secret sharing which uses n = 5t + 4 or n = 4t + 4 roles depending on the underlying network model. We demonstrate the practicality of our second protocol by showing experimental evaluations, significantly improving over prior proposed solutions for worst-case corruptions, especially in terms of transmitted data size.

Peter Gaži, Aggelos Kiayias and Alexander Russell. Fait Accompli Committee Selection: Improving the Size-Security Tradeoff of Stake-Based Committees

Speaker: Peter Gaži

We study the problem of committee selection in the context of proof-of-stake consensus mechanisms or distributed ledgers. These settings determine a family of participating parties, each of which has been assigned a non-negative "stake," and are subject to an adversary that may corrupt a subset of the parties. The challenge is to select a committee of participants that accurately reflects the proportion of corrupt and honest parties, as measured by stake, in the full population. The trade-off between committee size and the probability of electing a committee that over-represents corrupt parties is a fundamental factor in security and efficiency considerations for proof-of-stake consensus, as well as committee-run layer-two protocols.

We propose several new committee selection schemes that improve upon existing techniques by adopting low-variance assignment of certain committee members who hold significant stakes.

These schemes provide notable improvements to the size-security trade-off arising from the stake distributions of many deployed ledgers.

Chen-Da Liu-Zhang, Christian Matt, Ueli Maurer, Guilherme Rito and Søren Eller Thomsen. Practical Provably Secure Flooding for Blockchains

Speaker: Søren Eller Thomsen

In recent years, permissionless blockchains have received a lot of attention from both industry and academia, where substantial effort has been spent to develop consensus protocols that are secure under the assumption that less than half (or a third) of a given resource (e.g., stake or computing power) is controlled by corrupted parties. The security proofs of these consensus protocols usually assume the availability of a network functionality guaranteeing that a block sent by an honest party is received by all honest parties within some bounded time. To obtain an overall protocol that is secure under the same corruption assumption, it is therefore necessary to combine the consensus protocol with a network protocol that achieves this property under that assumption. In practice, however, the underlying network is typically implemented by flooding protocols that are not proven to be secure in the setting where a fraction of the considered total weight can be corrupted. This has led to many so-called eclipse attacks on existing protocols and tailor-made fixes against specific attacks. To close this apparent gap, we present the first practical flooding protocol that provably delivers sent messages to all honest parties after a logarithmic number of steps. We prove security in the setting where all parties are publicly assigned a positive weight and the adversary can corrupt parties accumulating up to a constant fraction of the total weight. This can directly be used in the proof-of-stake setting, but is not limited to it. To prove the security of our protocol, we combine known results about the diameter of Erdős–Rényi graphs with reductions between different types of random graphs. We further show that the efficiency of our protocol is asymptotically optimal. The practicality of our protocol is supported by extensive simulations for different numbers of parties, weight distributions, and corruption strategies. The simulations confirm our theoretical results and show that messages are delivered quickly regardless of the weight distribution, whereas protocols that are oblivious to the parties' weights completely fail if the weights are unevenly distributed. Furthermore, the average message complexity per party of our protocol is within a small constant factor of such a protocol.

Bernardo David, Felix Engelmann, Tore Frederiksen, Markulf Kohlweiss, Elena Pagnin and Mikhail Volkhov. Updatable Privacy-Preserving Blueprints

Speaker: Mikhail Volkhov

Privacy-preserving blueprints enable users to create escrows using the auditor's public key. An escrow encrypts the evaluation of a function P(t,x), where t is a secret input used to generate the auditor's key and x is the user's private input to escrow generation. Nothing but P(t,x) is revealed even to a fully corrupted auditor. The original definition and construction (Kohlweiss et al., EUROCRYPT'23) only support the evaluation of functions on an input x provided by a single user.

We address this limitation by introducing updatable privacy-preserving blueprint schemes (UPPB), which enhance the original notion with the ability for multiple parties to non-interactively update the private value in a blueprint. Moreover, a UPPB scheme allows for verifying that a blueprint is the result of a sequence of valid updates while revealing nothing else. We present uBlu, an efficient instantiation of UPPB for computing a comparison between private user values and a private threshold t set by the auditor, where the current value x is the cumulative sum of private inputs, which enables applications such as privacy-preserving anti-money laundering and location tracking. Additionally, we show the feasibility of the notion generically for all value update functions and (binary) predicates from FHE and NIZKs.

Our main technical contribution is a technique to keep the size of primary blueprint components independent of the number of updates and reasonable for practical applications. This is achieved by elegantly extending an algebraic NIZK by Couteau and Hartmann (CRYPTO'20) with an update function and making it compatible with our additive updates. This result is of independent interest and may find additional applications thanks to the concise size of our proofs.