Blockchain: Cryptography meets Economics

Co-leads: Elaine Shi (CMU), Andrew Miller (UIUC)
Scribes: Saba Eskandarian (UNC), Fan Zhang (Duke), Yupeng Zhang (TAMU)

Blockchains and cryptocurrency have continued to grow on the world stage. It has become a trillion dollar industry, including for example 50 billion USD value locked in “Decentralized Finance” projects alone. This development has been very exciting from a tech transition viewpoint. The cryptocurrency industry’s ethos of “permissionless innovation” has made this space a proving ground where implementations of previously academic technology – especially zero knowledge proofs (ZKP) and multiparty computation (MPC) are rapidly developed and tested.

This observation of real-world adoption gives us the opportunity to revisit previously accepted assumptions around . Several themes stand out as remaining technical challenges. First, there are many remaining open problems centered around the incentive design of such systems, understanding how these shape the behavior of participants. Second, ZKP and MPC remain difficult-to-break bottlenecks when applying them at scale. Addressing these design questions can lead to not only improving blockchains and digital ledgers to reach their potential, but also can be utilized by private consortium blockchains, industry applications such as supply chains, and central bank digital currencies (CBDCs).

As with many security technologies, blockchains have dual uses, and the explosion of popularity has led to many clearly negative outcomes - ransomware, scams, and susceptibility to hacks and financial risk. Security researchers can play a role in defining directions that seek to better understand and mitigate the harmful uses of these as well.

1. Incorporating incentives into distributed systems research

Many traditional distributed systems are developed assuming a strong Byzantine/malicious adversary that can misbehave arbitrarily. However, as blockchains evolve, it is becoming apparent that most malicious activities are not arbitrary and motivated by incentives. It is important to understand incentive structures in the blockchain application scenarios and develop tailored solutions/mechanisms/design patterns that utilize these incentive structures to develop towards protocol with better resilience as well as computation and communication efficiency.

One of the early works that connect rationality and BFT is the BAR Incentive model: https://www.cs.cornell.edu/lorenzo/papers/sosp05.pdf. In blockchain literature, incentive-compatibility issues were first studied in the context of consensus protocols, and more recently in layer-2 or application layer protocols. In the consensus layer, there are several works in identifying bad Byzantine behavior in staking protocols such as Casper Ethereum, BFT

Protocol Forensics. There are also many efforts in identifying incentive-based attacks and mitigation including selfish mining attacks, fruitchain, colordag, undercutting. Another important line of research is on fee mechanism designs including transaction Fee Mechanism (TFM) design: EIP-1559, Game theoretic analysis, Bitcoin's fee mechanism. In layer 2, several bribery resistant protocols are introduced, including MAD-Incentives related to HTLC, Ponyta, He-HTLC specification: MAD-HTLC, He-HTLC, Ponyta. Finally, several game-theoretic oracle designs have been proposed, for example, Chainlink 2.0.

Transaction manipulation attacks: Frontrunning is a concern in blockchain systems and it occurs due to the inherent incentive structure in the blockchain. While frontrunning may not always be bad, it is unclear in which contexts frontrunning becomes unethical. New research may identify applications in which seemingly innocuous frontrunning can lead to critical impact. Can a model-driven approach help detect frontrunning? How do we prove a certain behavior as frontrunning? What can we do to recuperate from frontrunning after the fact?

Order fairness: One recent direction to address frontrunning attacks is to enforce fairness at the consensus level (e.g., CRYPTO 20). Many questions remain open. For example, various fairness notions have been proposed, but what are some metrics to evaluate different fairness notions? How to enforce fairness efficiently, especially when different applications might want different fairness guarantees.

Identify new incentive attacks and defenses: What are new bribery attacks that look innocuous but may turn into big profits? For example, eclipse or partitioning attacks can be performed by bribing nodes to change their connections to certain nodes in the network. Changing nodes' connection seems harmless, however, by rewiring the connections, the adversary can effectively eclipse a particular targeted node. This can further lead to front-running, censorship, or double spending attacks. What consensus rules and/or incentives need to be added to mitigate these attacks?

Transaction fee mechanism design: In public blockchains, an important question is how to incentivize participation, i.e., how to properly compensate the participants who must spend resources to support the system. Prominent blockchains employ the gas model, i.e., transaction fees that are calculated based on the computation consumption and the utilization of the network. Is gas the right model? E.g., Solana has a flat fee model. New research can compare different models for TFM.

Incentives and consensus: In general, most existing works in the design of consensus protocols assume the existence of some honest (altruistic) parties and some arbitrarily malicious parties. However, based on the design and use of blockchains, there are other incentives that may influence a party's decision making. For instance, depending on the earning a party can have through frontrunning or transaction fee mechanism design, parties may not be incentivized to participate as specified in the protocol. Thus, taking such external aspects into consideration in the design of consensus protocols is an important challenge.

Incentive attacks in applications beyond finance: As blockchains are getting considered to applications (e.g., decentralized games, identities, name services, digital art) beyond payments and finance, it is important to understand incentive structures in those.

Cost vs Decentralization tradeoffs: With many variants of blockchain protocols emerging, it is very important for developers and practitioners to understand the relative pros and cons and tradeoffs associated with each variant of blockchain protocol. How can we measure the degree of decentralization of a blockchain protocol? Several optimization techniques in blockchain systems are aimed at optimization and scalability and at times deviate from the original Bitcoin protocol. Such variants of blockchain may offer reduced levels of decentralization and democratization. New research may study and inform what kind of tradeoffs are acceptable and at what threshold point, a blockchain protocol is considered to provide inadequate decentralization and democratization.

New applications of blockchains

New positive uses of blockchains can also help shape the directions of this industry, and lead to new market-places and eco-systems. For example, financial inclusion is an important direction --- how can we use blockchains to provide banking and security trading to users in developing countries who may not have access to normal banking services, and perhaps users who may only have sporadic internet access? What about other promised applications of blockchains, such as decentralized ride-sharing and airbnb-type applications? What are the barriers towards deploying such applications and how can we overcome them?

Blockchains are a promising solution to problems involving groups of *mutually distrusting entities* who need to share information and to make collective decisions, but whose incentives are inherently in opposition. One such use case for blockchain is dynamic spectrum allocation (i.e. decentralized Spectrum access System (SAS). In this case multiple SAS admistors will have to collectively make decisions for their customers about the use of spectrum shared among all users (customers of multiple SAS administrators). Blockchain in this case provides a platform to address the “trust” and fairness issues among multiple SAS administrators.

2. Scalability of blockchains through the adoption of ZKP and MPC technologies.

Blockchains have been an unprecedented catalyst for moving cryptography (especially ZKP and MPC) from theory to practice. Now that we can observe how the industry has adopted these, we can revisit some of the traditional security assumptions and performance goals. For example, it is now clear that we need to consider incentives and diversity of faults in larger networks like 50-100 or more nodes, as opposed to the traditional setting of accepting for example a majority honest assumption among small networks.

We have identified the following specific research challenges:

1. Cross-chain payments and swaps via zero-knowledge proofs (ZKP) and secure-multiparty computations (MPC). Most existing formal modeling focuses on one blockchain system in isolation, while in industry it is common to use multi-blockchain and cross-blockchain protocols, many of which are complicated and have created high-profile exploits (such as the Axie blockchain hack and the Wormhole bridge attack). We are currently lacking theoretical abstractions and formal security definitions for tasks like cross-chain exchange. The development of formal modeling and security proofs will help prevent attacks on cross-chain payments which have caused 100-million-dollar losses.

2. Centralization and lack of privacy in layer 2 solutions. In one of the major approaches of Layer 2 solutions (rollups), the transactions are off-chain and hosted by the centralized service provider of the rollup. The storage of the transactions becomes centralized again, which violates the original goal of decentralization. Existing layer 2 solutions usually don't consider privacy of transactions.

3. Scalability of ZKP and MPC techniques on space efficiency. While ZKP and MPC performance have largely focused on computation and communication costs, the adoption of these in practice by the blockchain industry (along with developers' continued improvements and optimizations) has revealed storage and memory has become the new bottleneck. Designing space-efficient protocols that preserve the computational efficiency of existing protocols is an interesting and important topic.

4. The application gap of MPC. Threshold signatures for crypto wallets have become the first major use–case for MPC in blockchain applications, which is deployed by platforms such as Coinbase, as well as Partisia, Unbound. In these applications, communication and the number of rounds of interactions are key concerns and noninteractive protocols are preferred. Beyond threshold signatures, other applications such as dark pools with MPC auctions are not widely used in practice because of the high overhead of the MPC protocols to achieve privacy.

5. Micropayments and fair exchange via zero-knowledge contingent payments.

Micropayments for daily use transactions (groceries, gas, …) brings efficiency and scalability issues in terms of the huge number of transactions to be verified in a short time to provide the daily services. Possible solutions may include off chain transaction verification and submitting proof to the chain. A major challenge is double spending during the off chain verification. Contingent payments are a possible solution.

3. Mitigating the downsides and negative externalities of blockchains

Energy consumption: Large-scale proof-of-stake consensus is now relatively well-understood. However, there are still practical impediments and economical concerns for existing blockchains to transition from proof-of-work to proof-of-stake. Therefore, the interesting open questions include:

- How to securely transition from PoW to PoS? Techniques for securely making the transition can expedite the transition to PoS.
- Can we do better than having a limited number of major investors having all the stake when a PoS chain starts? How can we avoid the centralization of stake in Proof-of-Stake systems?

Measurement testbeds: Measurement and monitoring of blockchains can help us understand the attacks, fraud, crimes, and other usages that lead to negative externality. Measurement of negative externalities and problems on real-world blockchains would provide much needed visibility that could give insight into resolving these problems. However, measurement of blockchains is an extremely difficult problem. Although on-chain data is permanently logged and publicly available, currently it is difficult to get access to offchain data and transient network behaviors. Such data is massive in the amount, and expensive to store. Moreover, although various blockchain monitoring companies may have access to such data, they are not incentivized to share them. For example, today, we don't even have data to answer very rudimentary questions, such as, "does selfish mining attack actually happen in real-world blockchains?"

It would be especially useful to have measurement that not only shows that problems exist as expected, but also helps point toward ways to resolve the problems. We believe that it would be useful for funding agencies such as NSF to fund projects that create a global measurement infrastructure to provide public measurement data.

User education: It would be helpful to involve the HCI/usable security community to design educational tools to help people identify blockchain scams, risks, as well as the difference between legitimate blockchain risk and scams. The existence of edge cases shouldn't deter us from pointing out the clear scams.

Understanding scams and systemic risks: It would be important to fund research projects that bring together the economists, regulators, lawmakers and the CS community. For example, we would like to understand the difference between scams and systemic risk. This can help us mitigate/avoid cases such as the recent collapse of Lunar, a stablecoin that collapsed and lost billions of dollars due to the high leverage that led to instability of the ecosystem.

There are questions for regulators and lawmakers here too. It may be a great opportunity to reach out to collaborators and to expand our toolkits. Collaboration between the disparate

communities may be challenging due to the difference in our respective languages. Perhaps an effective way is for NSF to start by funding “seedling” projects aimed at bringing these communities together to first conduct the “meta-research” that defines the joint research agenda.

Mechanism Design meets cryptography
(This may overlap with findings of the “Incentive” group)

One way to protect the users is to provide safe mechanisms that have provably secure properties which we can articulate to the users. Traditional mechanism designs, however, often fail completely for the decentralized environment. For example, for the transaction fee mechanism, even the auctioneer (i.e., the miner) can be a strategic player and this is not captured by the models of traditional mechanism design. Another exciting challenge is the interaction of cryptography and mechanism design. This requires new models and new game theoretic notions since we are now considering games which are interactive protocols involving computational agents --- which departs from traditional modeling techniques in the game theory literature.

Since miners and users can easily collude through side contracts in a blockchain environment, it is also important to develop new theories and game theoretic models, leading to new mechanisms that are fair or game theoretically secure, even in the presence of arbitrary (but possibly bounded) external incentives.

EDU related aspects: There is a lot of excitement among middle-school, high-school and college going students on the use of blockchains. For instance, multiple universities have student-led informal blockchain clubs. Using blockchains as a conduit, this presents an important opportunity for us as computer science researchers as well as NSF, to attract students to learn computer science in general.

On the other hand, blockchains are also touted as a hammer to solve many potentially unrelated problems. As researchers working in this space, it is our responsibility to educate when blockchains should and should not be used.