Report from Dagstuhl Seminar 18461

Blockchain Security at Scale

Edited by

Rainer Böhme[1], Joseph Bonneau[2], and Ittay Eyal[3]

1 **Universität Innsbruck, AT,** rainer.boehme@uibk.ac.at
2 **New York University, US,** jbonneau@gmail.com
3 **Technion – Haifa, IL,** ittay@technion.ac.il

Abstract

38 researchers affiliated with over 25 different institutions in 7 countries met during Dagstuhl Seminar 18461 for discussing open problems regarding "Blockchain Security at Scale." The seminar was split into eight blocks of two presentations each. The mode for each talk was 15 minutes of blackboard-only presentation followed by 30 minutes of discussion. Discussions not fitting into this limit were resumed in smaller break-out groups. This report documents the scheduled talks as well as the improvised sessions for in-depth discussion.

Seminar November 11–16, 2018 – http://www.dagstuhl.de/18461
2012 ACM Subject Classification Networks → Network security, Computer systems organization → Peer-to-peer architectures
Keywords and phrases Blockchain, Consensus, Cryptography, Distributed Systems, Game Theory, Scaling
Digital Object Identifier 10.4230/DagRep.8.11.21
Edited in cooperation with Assimakis Kattis, Patrik Keller, Itay Tsabary

1 Executive Summary

Rainer Böhme (Universität Innsbruck, AT)

The security of blockchain-based systems has attracted great interest in the research community following the initial financial success of Bitcoin. Several security notions for blockchain-based systems have been proposed, varying in degree of formality and applicability to real-world systems. However, a major blind spot remains about the environment surrounding blockchain-based systems. This environment is typically assumed to be static (irresponsive to activities of the blockchain system). This is a sound starting point for security analysis while the stakes involved are small compared to the environment (i. e., the global economic and political system). However, if blockchain-based systems truly offer compelling advantages over legacy systems, they may eventually become the dominant form of organizing certain social choice problems. This "scale change" challenges the assumption that the blockchain-based system remains below the threshold of relevance for the parts of its environment that are vital for its security. One instance where this may already occur is the influence of mining puzzles on hardware design and electricity prices.

Blockchain Security at Scale, *Dagstuhl Reports*, Vol. 8, Issue 11, pp. 21–34
Editors: Rainer Böhme, Joseph Bonneau, Ittay Eyal
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The purpose of the seminar was to bring together researchers with expertise in various subfields of blockchain-based systems to jointly revisit security foundations. The primary goal was to incorporate explicit consideration of reciprocity effects between properties of cryptocurrency protocols and their environment.

The primary intended outcome of this seminar was proposing a new design principle, viewing security as a key scalability property to consider in addition to performance and efficiency. Second, the seminar aimed to converge on standard terminology for security notions that are robust to scale. Third, we applied this new methodology to Bitcoin specifically as a test case, producing a sort-of "break glass in case of rampant runaway growth" security plan.

Specific questions were:

1. **From micro-level to macro-level incentives** Bitcoin's ecosystem remains small relative to large multinational corporations. What happens to incentives when a cryptocurrency reaches a scale similar to large national economies?
2. **Cryptographic agility** How does the ability to upgrade cryptographic algorithms might change in the future as cryptocurrency protocols become widely embedded in hardware and/or codified in the law.
3. **Reciprocity effects on hardware design** How will the hardware industry be affected by the increasing importance of superior hardware for mining, and possibly trusted execution environments (TEE) in the future?
4. **Mining economics at scale** How will mining economics change in the future, in particular, dynamics between miners at large-scale power consumption levels, with mass availability of cheap commodity mining hardware (including TEE-based), and with different incentives, e.g., in a high-valued fee-only revenue model.
5. **Reconsidering non-monetary incentives** Can cryptocurrencies be resilient to disruptive nation-level attacks that are not due to monetary incentives?
6. **Governance at scale** To date, cryptocurrencies largely rely on informal leadership from a small group of influential software developers. Can this be translated into a more democratic model? What does democratic control mean for a cryptocurrency when the *demos* is not clearly defined?

2 Table of Contents

3 Overview of Talks

3.1 Using Differential Privacy to Analyze Cryptocurrencies Anonymity

Foteini Baldimtsi (George Mason University – Fairfax, US), transcription from abstract book

We investigate whether techniques inspired from the area of differential privacy can be used to construct anonymous cryptocurrencies offering an interesting set of trade-offs between the level of offered privacy, efficiency and underlying assumptions.

Our motivation rises from the fact that Monero, one of the most popular private cryptocurrencies has been recently analyzed to find that approximately 60% of transactions provide no or very limited privacy due to a very small anonymity set. We propose a protocol inspired by Monero (utilizing ring signatures) and formally analyze it while preserving differential privacy for users. Specifically we would like to claim that two neighboring transaction graphs are nearly equal to give rise to the same chain. In order to keep the size of each individual ring signature small while providing a large number of potential options for the real transaction we have users submit several ring signatures in a sequence of rounds that "pipe-in" an ever-increasing number of possible mix-ins.

3.2 STARKs for Blockchain Scalability

Eli Ben-Sasson (Technion – Haifa, IL)

An interactive proof system is defined to be a STARK if it satisfies the following conditions:

Scalability: for statements referring to computations of nondeterministic time T, proving time scales quasi-linearly in T and verification time scales poly-logarithmically in T.

Transparency: all verifier messages are public random coins ARgument of Knowledge: There exists a polynomial time extractor that, interacting with a valid prover, reconstructs a non-deterministic witness for the statement.

STARKs are unique in their scalability capabilities, when compared to the zk-SNARKs deployed in Zcash, the recursive SNARKs suggested for Coda and the BulletProofs system used by Monero. The talk described the essential properties of STARKs and compared them to SNARKs, recursive SNARKs and BulletProofs.

3.3 Proof of Work and Resource Hardness

Alex Biryukov (University of Luxembourg, LU), transcription from abstract book

We have defined the properties a resource hard proof of work should have. We have defined classed of R-hardness depending on prover-verifier capabilities (hard, easy with secret, publicly easy) and with regard to the specific resource R: Time (sequential or total computation),

memory, code size. In this setting PoWs in the Hard(R) for the prover, and publicly easy to verify. We recalled the scrypt construction and a new scheme we call Diodon – weakly MHF with easy verification with secrets. We have shown Equihash – a memory and computation hard PoW based on generalized birthday and shown its relation to VDFs and proofs of sequential work.

3.4 Trusted Execution Environments

Mic Bowman (Intel – Hillsboro, US), transcription from abstract book

Hardware-based trusted execution has the potential to dramatically improve the scale, efficiency and performance of decentralized applications. There are, however, two positions that are often taken with respect to trusted execution environments (TEE). The first is that the TEE works perfectly. If that were the case, then large portions of cryptographic research would become irrelevant. We know that TEEs (and any other system component) can be attacked so this idealistic assumption is unreasonable. At the other end we could assume that because it can be attacked, a TEE is useless and be ignored. That position is also because it ignores the difficulty in attacking a TEE. A more appropriate view is that the TEE is part of a larger security context. This approach allows for some performance and efficiency improvements and still preserves the overall system security objectives.

3.5 Asymmetric Trust

Christian Cachin (IBM Research – Zürich, CH)

Joint work of Björn Tackmann, Christian Cachin

The Ripple and Stellar blockchain consensus protocols aim at relaxing the strict assumptions in classical consensus and develop so-called federated consensus methods. They intend to stand between traditional BFT consensus (in the sense that the set of nodes is known and pre-agreed) and decentralized consensus (where participation is completely open to anyone). In practice this means that every node declares a list of other nodes which it "trusts." Consensus decisions can be federated in this way, from groups of participants potentially unknown to each other that may have different trust assumptions. For example, one might first reach consensus in small subsets and subsequently combine those partial results in later protocol steps, to reach consensus across the complete system.

Protocols that aim at this goal are not understood well, even though some are used in live blockchains. For example, a recent paper by authors from Ripple casts doubts on the claimed properties of the Ripple protocol. We sketch a model for asymmetric Byzantine quorum systems that explains such protocols precisely. It strictly generalizes existing Byzantine quorum systems. Well-known consensus protocols for Byzantine quorum systems can easily be extended to work in this new model with asymmetric trust.

3.6 PERUN – Virtual Payment and State Channels

Lisa Eckey (TU Darmstadt, DE)

Joint work of Stefan Dziembowski, Lisa Eckey, Sebastian Faust, Daniel Malinowski
Main reference Stefan Dziembowski, Lisa Eckey, Sebastian Faust, Daniel Malinowski: "PERUN: Virtual Payment Channels over Cryptographic Currencies", IACR Cryptology ePrint Archive, Vol. 2017, p. 635, 2017.
URL http://eprint.iacr.org/2017/635

One approach to securely scale blockchain protocols is to move some of the transaction load off-chain. Payment and State channels allow secure, optimistic, off-chain execution of multiple transactions and even contracts, while only relying on two on-chain interactions with a smart contract. As long as all connected parties agree to the current state of the channel, this method allows for fast and cheap off-chain execution of state changes. The overall security of each direct state channels is guaranteed through a single on-chain contract. Additionally, two existing state channels can be composed to form a new "virtual" state channel, that can be opened and closed off-chain and indirectly connects two parties through a network of state channels. In case of disputes, virtual state channels can be resolved in two ways, either the parties use the connecting intermediary or they directly go to the blockchain. The former solution adds a layer of protection for honest parties since they might not have to deal with on-chain disputes themselves, but it adds risks to the intermediary who might be forced to pay a high amount of transaction fees. Therefore the direct resolve offers a more fair way to resolve disagreement.

3.7 Proof of Personhood

Bryan Ford (EPFL Lausanne, CH), transcription from abstract book

Joint work of Maria Borge, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Bryan Ford
Main reference Maria Borge, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Linus Gasser, Nicolas Gailly, Bryan Ford: "Proof-of-Personhood: Redemocratizing Permissionless Cryptocurrencies", in Proc. of the 2017 IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2017, Paris, France, April 26-28, 2017, pp. 23–26, IEEE, 2017.
URL http://dx.doi.org/10.1109/EuroSPW.2017.46

Decentralization and membership foundations for permissionless blockchains such as proof-of-work, -stake, -storage, -elapsed -time, etc. are all "proof-of-investment": anyone who can and is willing to invest more gets more voting power and rewards. All proof-of-investment schemes are subject to re-centralization. ("rich get richer") due to economies of scale, and cannot provide a human-centric notion of fairness or equity. We propose Proof-of-Personhood, a democratic foundation for blockchain decentralization that gives each (real) person one equal "vote" or unit of "stake." Proof-of-Personhood (PoP) can be implemented in principle using government-issued identities, social trust networks, biometrics or physical presence tests such as pseudonym parties or PoP parties. We are developing tools and processes for PoP parties because they can be made privacy-preserving, low-cost, implementable anywhere including paperless (undocumented) refugees and migrants, and can potentially enforce strong and transparent "one-per-person" Sybil attack resistance. An "anytrust" group of organizers,

secures each party locally, while an inter-party federation or trust network with evidence-based transparency and cross-witnessing processes secures the system globally against corrupt parties. Applications include "one-per-person" accountable anonymous Web browsing or website login tokens, membership tokens for online reputation systems or social media forums, voting tokens for online deliberative forums such as liquid democracy, and one-per-person cryptocurrency minting tokens for decentralized implementations of "Universal Basic Income."

3.8 Manipulating Incentives

Aljosha Judmayer (Secure Business Austria Research, AT)

The theoretical possibility of bribing attacks on cryptocurrencies has been known since 2015/2016, with various techniques proposed since. The majority of these proposals focus on in-band bribing attacks, executed within the same cryptocurrency they are designed to attack.

This talk presents unpublished research on out-of-band bribing attacks, which are capable of facilitating double-spend collusion across different blockchains.

The bribing logic is thereby managed by a smart contract on a funding cryptocurrency, which leverages cross-chain state verification of the target cryptocurrency to determine the attack outcome and react accordingly.

Contrary to existing schemes, colluding miners are reimbursed independently of success or failure of the attack. This allows our bribing attack to become cheaper than comparable bribing attacks (i.e., the whale attack).

Finally, to hinder counter bribing measures and further reduce the costs of the attack, the notion of crowdfunded bribing attacks is introduced, where the interests of several attackers are aligned by the smart contract to execute multiple double-spending attacks concurrently.

3.9 Redesigning Bitcoin's Fee Market

Ron Lavi (Technion – Haifa, IL)

Joint work of Ron Lavi, Or Sattath, Aviv Zohar
Main reference Ron Lavi, Or Sattath, Aviv Zohar: "Redesigning Bitcoin's fee market", CoRR, Vol. abs/1709.08881, 2017.
URL https://arxiv.org/abs/1709.08881

Two of Bitcoin's challenges are (i) securing sufficient miner revenues as block rewards decrease, and (ii) alleviating the throughput limitation due to a small maximal block size cap. These issues are strongly related as increasing the maximal block size may decrease revenue due to Bitcoin's pay-your-bid approach. To decouple them we analyze the "monopolistic auction" [Goldberg et al. 2006], showing: (i) its revenue does not decrease as the maximal block size increases, (ii) it is resilient to an untrusted auctioneer (the miner), and (iii) simplicity for transaction issuers (bidders), as the average gain from strategic bid shading (relative to bidding one's true maximal willingness to pay) diminishes as the number of bids increases.

3.10 What Can Blockchains Do for You?

Ian Miers (Cornell Tech – New York, US), transcription from abstract book

Blockchains are a form of limited trusted third party. Unlike many proposed schemes, they are deployed and readily available. In addition to analyzing how blockchains can be improved by computer science, we should ask what they can do to solve issues in e.g. computer security and cryptography. This talk explored how they can be used to achieve fairness in multi party computation and get general secure computation with state keeping, proof and publications, and assured F/O.

3.11 Incentive-Compatibility – A Brief Tutorial

Tim Roughgarden (Stanford University, US), summarized by Patrik Keller

Tim Roughgarden gave an introduction to game theory and its application to security. He highlighted the importance of incentive-compatibility and how it can be achieved. Relevant tools for showing incentive-compatibility are dominant strategies and equilibria. Tim explained these concepts in the context of auctions. For auctions, the goal is that all participants truthfully state their willingness to pay. He showed that this is indeed a dominant strategy when participating in a Vickrey auction.

3.12 Biologically-Inspired Scaling for Cryptocurrencies

Marie Vasek (University of New Mexico, US), transcription from abstract book

All organisms scale in similar ways. For example, the metabolism rate scales at a rate approximately to the 3/4 power. However, information systems in organisms scale differently. We use the immune system, a partially decentralized network, as inspiration for how cryptocurrencies can scale. For instance, all T cells are approved initially by a centralized authority, the thymus. Afterwards, T cells respond relatively decentralized to infection. We outline how cryptocurrencies have many centralized checkpoints and discuss scaling them and the inherent issues therein.

3.13 Consensus without Cryptography

Roger Wattenhofer (ETH Zürich, CH)

Distributed protocols often employ some form of cryptography, in particular digital signatures. This talk shed some light on the role of cryptography in distributed systems, in particular byzantine agreement (consensus). It presented the simple Ben-Or voting framework, and then discussed various versions of how to implement the random choice: By a local coin, by an oracle, by a pre-determined bit-string. The bit-string must be hidden with Shamir's secret sharing, or even worst-case scheduling will break the framework. In the end, the talk discussed recent alternative methods, in particular the idea to use at least a quadratic number of random bits.

This brings us the interesting question to what degree cryptography is needed: What distributed problems can or cannot be solved without cryptography, and what is still unknown?

4 Working Groups

Each of the talks was followed by a short discussion. The discussions which had to be aborted in order to stay in schedule were later continued in smaller break-out groups. We had working groups to the following topics:

- Incentives – Micro to Macro
- Crypto Agility
- Network Layer
- Hardware
- Asymmetric Trust
- Proof-of-X
- Privacy
- Responsible Disclosure
- Game Theory
- Transaction Fees
- Governance
- STARKs

Some of these smaller session were very fruitful while others came to end early. We thus provide abstracts for only some of the working groups.

4.1 Crypto Agility

Patrik Keller (Universität Innsbruck, AT)

The Crypto Agility session was mainly about protocol migration strategies in case of new weaknesses in the deployed implementations of cryptographic primitives. We started with enumerating the different types of cryptographic breaks:

1. Signature schemes: existential forgery, universal forgery, key recovery
2. Hash functions: collision, preimage recovery, bias towards lower hashes
3. Zero knowledge: soundness failure, failures in zero knowledge

While some breakages seem to be fatal (key recovery, reproducible hash collisions), a carefully crafted protocol may be resilient to a relevant subset of failures. We discussed three strategies to defend against or recover from breaks:

1. Key updates: voluntary or mandatory, gradual update to a different signature scheme if existing scheme is imminently broken.
2. Hybrid accounts: multiple signature schemes in parallel, moving funds requires signatures in all schemes.
3. Backup keys: global activation of a previously set up fallback signature scheme might help in case of a rapidly developing signature scheme break.

4.2 Asymmetric Trust

Patrik Keller (Universität Innsbruck, AT)

This session followed up on Christian Cachin's talk on Asymmetric Trust. We discussed the following three points:

1. The presented results are based on a predisposed and bounded set of nodes. Can this be adapted for a potentially unlimited number of nodes? Can the results be reused for the permissionless setting, where nodes can join and leave the network at any time?
2. We clarified that the presentation considered only the safety properties of asymmetric trust.
3. We attempted to relate the new notion of Asymmetric Trust to existing work on DAG protocols. Can it be used to obtain improved security proofs?

4.3 Proof-of-X

Patrik Keller (Universität Innsbruck, AT)

There are two kind of proof-of-X schemes. On the one hand, there are proof-of-consumption schemes where a resource is consumed in order to obtain the right to participate in the protocol. The scarcity of the consumed resource implies a rate limit on participation. On the other hand, there are proof-of-stake schemes where the ownership of a resource is demonstrated instead. The fundamental difference between consumption and ownership makes the two kinds hard to compare. Proof-of-consumption and proof-of-stake must thus be treated separately.

On the proof-of-consumption side, we further recapitulated the properties necessary for proof-of-work constructions as presented by Alex Biryukov.

4.4 Responsible Disclosure

Rainer Böhme (Universität Innsbruck, AT)

Responsible disclosure of security vulnerabilities poses specific challenges in the domain of cryptocurrencies. The group tried to understand the differences between the conventional debate for proprietary software, which often encompasses national security interests and the specifics of the cryptocurrency space. The participants found similarities (e.g., the absence of a single point of contact also applies to many open source projects), differences in severity (e.g., competition between cryptocurrency projects and the fact that some bugs are easily monetizable), and differences in quality (e.g., some bugs are "unfixable"). The group also compiled a list of issues and plans to write up the lessons learned from case studies in a joint publication.

4.5 Transaction Fees

Patrik Keller (Universität Innsbruck, AT)

This session arose from the presentation of a new model for Bitcoin transaction fees by Ron Lavi. The presented scheme is based on the assumption of equally sized transactions. We discussed how the new pricing scheme can be adapted to variable transaction size. Additionally we added the constraint of a fixed block size. Unfortunately, the originally tractable optimization problem used for fixing the fee becomes intractable under the additional assumptions.

Apart from that, we observed that cryptocurrencies in contrary to earlier discussed systems (e.g. music download pricing) allow for asymmetric auctions where the buyer pays more than the seller receives. The difference could be burned. Whether this additional freedom in the design space allows for better auction schemes is to be discussed.

4.6 Governance

Assimakis Kattis (New York University, US)

Governance issues around distributed systems centered around the various potential desirable properties that governance models should have. The main areas of discussion revolved around:

1. Separation of powers between stakeholders
2. Responsiveness to emergencies
3. Transparency in the governance process
4. Accountability of actors to each other and the general ecosystem

Furthermore, the techniques to ensure decentralization and a reliable practical implementation of governance goals were identified as important areas for further study. Interaction with the current legal system was also discussed, along with the questions it poses for the design of governance models.

4.7 STARKs

Assimakis Kattis (New York University, US)

This session was based around Eli Ben-Sasson's talk on STARKs and their potential for scalability. We looked at the theoretical foundations of STARKs and investigated what the main barriers for scalable implementations are. Discussions followed around STARK trust assumptions, proof sizes and use cases. The separation in prover efficiency between non-interactive protocols and PCPs/IOPs was analyzed, along with efficient constructions of the latter.

Mathematical tools that allow for efficient STARK representations were also investigated. Usage of low-degree extensions, the relationship between error and soundness, and FFTs for interpolation/evaluation of low-degree polynomials were the main topics covered around the design techniques for STARK proofs. In a subsequent session, the construction and guarantees of getting low degree polynomials to verify were discussed, since they are linked to final proof sizes. We also looked at the potential for compression of constraint checks in STARKs, as well as at hash functions with efficient SNARK/STARK representations.

5 Open Problems

Joseph Bonneau (New York University, US)

We ended the seminar with a brief discussion of open problems. Many arose throughout the breakout groups, but there were a number of topics we did not have time to full explore. Some of the most interesting included:

- **Unifying proof-of-stake with proof-of-work** Efforts in the "Proof-of-X" session to identify a unified model for analyzing proof-of-stake protocols with proof-of-work protocols did not succeed. Can a unified model be found?
- **Auction-based consensus protocols** There were some efforts to outline a consensus protocol based on miners bidding for the right to create a block. The idea seems promising but we could not agree on an exact protocol.
- **On-chain governance** Most of the discussion of governance focused on higher-level decision making about protocol governance. It is an interesting question to explore what sorts of governance decisions can be made automatically by voting on the chain itself.
- **Incorporating market impacts** It is widely agreed that game-theoretic models of cryptocurrency should eventually incorporate the notion of market impact: executing attack may affect exchange rates and hence hurt the attacker despite a nominal gain in rewards. We lack a clear roadmap for incorporating this phenomenon into models in a tractable way.

Participants

- Svetlana Abramova
Universität Innsbruck, AT
- Sarah Azouvi
University College London, GB
- Foteini Baldimtsi
George Mason University – Fairfax, US
- Eli Ben-Sasson
Technion – Haifa, IL
- Alex Biryukov
University of Luxembourg, LU
- Rainer Böhme
Universität Innsbruck, AT
- Joseph Bonneau
New York University, US
- Mic Bowman
Intel – Hillsboro, US
- Dominic Breuker
solarisBank AG, DE
- Christian Cachin
IBM Research-Zurich, CH
- Nicolas Christin
Carnegie Mellon University – Pittsburgh, US
- Lisa Eckey
TU Darmstadt, DE
- Ittay Eyal
Technion – Haifa, IL
- Bryan Ford
EPFL Lausanne, CH
- Christina Garman
Purdue University – West Lafayette, US
- Arthur Gervais
Imperial College London, GB
- Philipp Jovanovic
EPFL Lausanne, CH
- Aljosha Judmayer
Secure Business Austria Research, AT
- Ghassan Karame
NEC Laboratories Europe – Heidelberg, DE
- Assimakis Agamemnon Kattis
New York, US
- Stefan Katzenbeisser
TU Darmstadt, DE
- Patrik Keller
Universität Innsbruck, AT
- Ron Lavi
Technion – Haifa, IL
- Patrick McCorry
King's College London, GB
- Ian Miers
Cornell Tech – New York, US
- Malte Möser
Princeton University, US
- Tyler W. Moore
University of Tulsa, US
- Neha Narula
MIT – Cambridge, US
- Tim Roughgarden
Stanford University, US
- Tim Ruffing
Universität des Saarlandes, DE
- Emin Gün Sirer
Cornell University – Ithaca, US
- Yonatan Sompolinsky
The Hebrew University of Jerusalem, IL
- Itay Tsabary
Technion – Haifa, IL
- Florian Tschorsch
TU Berlin, DE
- Marie Vasek
University of New Mexico, US
- Roger Wattenhofer
ETH Zürich, CH
- Edgar Weippl
Secure Business Austria Research, AT
- Aviv Zohar
The Hebrew University of Jerusalem, IL