General

S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2008. [Online].

This is the foundational Bitcoin paper--published online, pseudonymously, and with no peer review. Nakamoto proposes a system for storing and validating transactions of a virtual currency without a central authority. In particular, the paper focuses on solving the double-spending problem, in which one unit of currency can be spent more than once, violating the integrity of the system. Nakamoto describes a proof-of-work system, a high-level network architecture, and addresses some potential issues like storage space, fast payment verification, privacy, and the 51% attack. While the paper is important because it introduced the concept of a blockchain, it omits many details and does not fully address security and privacy concerns.

J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, “SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies,” in *2015 IEEE Symposium on Security and Privacy*, San Jose, CA, 2015, pp. 104–121.

This paper was published with the label Systematization of Knowledge (SoK), a type of paper solicited by this IEEE conference. As the title implies, the authors present a thorough overview of Bitcoin, synthesizing a wide variety of sources. They provide an extended history of Bitcoin, tracing its origins as far back as 1983. Their technical description of Bitcoin is extensive, covering details of the mining, consensus, and peer-to-peer network protocols. They examine the game-theoretic properties of the system, outlining known properties and cataloguing attacks, including those involving mining pools and wallets. They describe proposed and implemented modifications to Bitcoin, as well as alternatives to proof-of-work. In addition, they address the issue of privacy and outline potential solutions for an anonymous cryptocurrency. This paper provides a remarkably complete overview of Bitcoin and its extensions. It also serves as a valuable source of additional specialized literature. However, having been published in 2015, this paper does not include advances in the field from the last four years.

A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, “On the Security and Performance of Proof of Work Blockchains,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, New York, NY, USA, 2016, pp. 3–16.

This paper explores the differences in security of various Proof of Work (PoW) blockchains such as Litecoin, Bitcoin, Dogecoin, and Ethereum. Using a simulator, the authors assessed the effects of stale block rate, the interval, the size, and the throughput of each cryptocurrency’s

blockchain. They looked at both selfish mining and attempts at double spending. Their conclusions were that selfish mining is not always a “rational” strategy, that all other tested coins need many more confirmations to match Bitcoin’s security, that larger rewards protect better against double spending, and that the throughput of Bitcoin can be increased without worrying about security.

Applications

S. T. Ali, P. McCorry, P. H.-J. Lee, and F. Hao, “ZombieCoin: Powering Next-Generation Botnets with Bitcoin,” in Financial Cryptography Workshops, 2015.

This paper from 2015 proposes that blockchains are a potential new avenue for the propagation of botnets, networks of computers running coordinated malicious code. This is because there is no need for command-and-control infrastructure, there is some degree of anonymity, and finally because using blockchain means that there is no need for bots to communicate with each other. The researchers involved with this paper did a botnet simulation with 14 nodes and were able to run it cheaply and reasonably successfully, both of which are also important qualities that point towards the use of the blockchain as a possible means of controlling a botnet. The conclusion proposes that the threat of a botnet using blockchain is a real but currently unaddressed issue. This is an interesting proposition in terms of blockchain security.

A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, “MedRec: Using Blockchain for Medical Data Access and Permission Management,” in *2016 2nd International Conference on Open and Big Data (OBD)*, 2016, pp. 25–30.

This article proposes a solution to the issue of scattered medical records. It uses a blockchain to link different points holding a person’s medical records. The idea behind this is to allow better access to disparate records while also attempting to maintain certain permissions and confidentialities. While this article does consider privacy concerns related to DRM and HIPAA, they do not thoroughly address informed consent. They note growing acceptance of access to anonymous records and that there may be a “growing interest among patients, care providers...to responsibly share more data.” While wider access to medical records might be helpful in studies, it is unclear whether they have really considered the ethics of the situation. This is a technical paper, so it may be a stretch to hold it to these standards, but reading the paper raised several questions in regards to the ethics of blockchain based medical records.

I. Bentov, A. Gabizon, and A. Mizrahi, “Cryptocurrencies without Proof of Work,” arXiv:1406.5694 [cs], Jun. 2014.

As in many other non-Proof of Work papers, this work begins by citing the inherent issues with Proof of Work as well as those that arise in particular applications. This paper promotes the use

of Proof of Stake, a system in which those who have more coins have more power to add to the chain. This is based on the assumption that those with more stake (coins) will have more of an incentive to keep the currency valuable, incentivizing them not collude to attack the system. This paper does address the security issues of proof of stake - like Proof of Work, if a majority of people decide to be malicious, the system will be compromised. Peer Coin (called PPCoin in this paper) is used to explain many of their concepts, which include Chains of Activity (CoA) and Dense-CoA pure Proof of Stake systems. Much of the paper is spent discussing possible security issues and their resolutions. This paper was published in the 2016 Financial Cryptography Workshops, though it was written earlier.

S. Bragagnolo, H. Rocha, M. Denker, and S. Ducasse, “SmartInspect: solidity smart contract inspector,” in *2018 International Workshop on Blockchain Oriented Software Engineering (IWBOSE)*, 2018, pp. 9–18.

This paper discusses the difficulties in editing issues in smart contracts once they have been added to the blockchain. It looks specifically at Ethereum’s smart contract methods and the Solidarity language used to create them. The difficulty with smart contract debugging is that once the code has been accepted into the chain, there is no way to change any bugs that may arise, given the immutable nature of blockchains. This can have disastrous effects should a malicious party find a vulnerability in one of these smart contracts. The authors propose SmartInspect, a tool that would allow smart contract writers to understand how their code will work prior to deploying it, with the hope that this will mitigate issues with bugs in smart contracts.

K. Christidis and M. Devetsikiotis, “Blockchains and Smart Contracts for the Internet of Things,” *IEEE Access*, vol. 4, pp. 2292–2303, May 2016.

This paper provides a high-level abstract overview of blockchains and smart contracts. The authors divide blockchains into two categories: those that track the ownership and movement of assets (like Bitcoin), and those that execute arbitrary interactions between parties (like smart contracts). The authors focus on the applications of smart contracts to the Internet of Things, including firmware updates, smart locks for managing home rentals, electricity generation, and supply chain tracking. The supply chain tracking example is particularly well-explained. The authors present potential pitfalls for a blockchain approach to IoT, including low transaction throughput, privacy, transaction censoring, legal enforceability, and the autonomy of potentially rogue smart contracts. While this paper does not dive into the details of any particular application, is does provide a solid overview of smart contracts, especially as they might be used for IoT devices.

G. Destefanis, M. Marchesi, M. Ortu, R. Tonelli, A. Bracciali, and R. Hierons, “Smart contracts vulnerabilities: a call for blockchain software engineering?,” in *2018 International Workshop on Blockchain Oriented Software Engineering (IWBOSE)*, 2018, pp. 19–25.

This article generally deals with the emergence of smart contracts and the need to make such smart contracts more secure. They use both Bitcoin and Ethereum as examples, though they discuss Ethereum in more detail. In particular, they examine the issues with Parity, an Ethereum wallet that was attacked in 2017, freezing 500K ether. They begin with a basic explanation of blockchains and then go on to describe how Ethereum transactions and smart contracts function. They then analyze Parity, going through the code of the application in order to describe the failure of this wallet, which occurred because of a bug in the initWallet function. They then address possible ways of mitigating smart contract bugs, which include allowing for the correction of smart contracts, the ability to reference new libraries, and better testing. Their conclusion is that there is a need for what they call Blockchain Oriented Software Engineering to avoid such issues in the future.

S. Dziembowski, S. Faust, V. Kolmogorov, and K. Pietrzak, “Proofs of Space,” in *Annual Cryptology Conference*, pp. 585-605. Springer, 2015.

Proof of Work is the typical approach for making blocks challenging to publish, as this was the method proposed in the Bitcoin white paper. Noting that Proof of Work as a concept had existed since the 90s, this paper introduces another method, Proof of Space. The authors argue that Proof of Space addresses both energy and specialized hardware concerns raised by Proof of Work. Like Proof of Work, Proof of Space requires the delegation of a resource in order to make it difficult to modify blocks, though in this case, the resource is disk space. In particular, the authors propose using hard-to-pebble graphs as a Proof of Space challenge. The paper also addresses some security concerns about general Proof of Space systems. This paper was published in 2013 and therefore is not up to date on current applications, though the coin system (in this paper called Spacecoin) is furthered in a later paper: “Spacemint: A Cryptocurrency Based on Proofs of Space.” However, this paper is continuing to be cited even through 2018, and therefore has laid important groundwork.

I. Eyal, A. E. Gencer, E. G. Sirer, and R. van Renesse, “Bitcoin-NG: A Scalable Blockchain Protocol,” in *Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation* (NSDI ’16), pp. 45-69, 2016.

This paper proposes a variant of the Bitcoin consensus protocol that improves Bitcoin’s scalability. In particular, Bitcoin-NG (Next Generation) allows higher transaction throughput by using mining to elect a temporary leader. This leader is allowed to freely add transactions to the chain with no proof-of-work until the next miner solves a proof-of-work puzzle and becomes the new leader. In this way, new transactions are constantly being added to the chain by the current

leader, allowing many more transactions to be processed with lower latency. The paper takes care to address the vulnerabilities of the system, for instance to selfish mining, denial of service (DoS) attacks, and fraud by the leader. The authors derive reward values that make the system incentive-compatible. In addition, they simulate a Bitcoin-NG network to demonstrate its improved performance compared to Bitcoin. The paper provides a solid argument for the superiority of Bitcoin-NG over vanilla Bitcoin. We note that Eyal and Sirer were the authors of the selfish mining paper, so it is not surprising that they pay special attention to this particular vulnerability of Bitcoin-NG.

T. Hanke, M. Movahedi, and D.Williams, “DFINITY Technology Overview Series Consensus System,” arXiv:1805.04548, 2018.

Dfinity is an is an upcoming blockchain-based platform that has innovated a consensus protocol that relies on randomness to operate effectively. This paper focuses on the problem of selecting the next successful miner in a trustless manner, without relying on a proof of work competition. Dfinity achieves this through a Verifiable Random Function (VRF) in a process called threshold relay. It uses BLS (Boneh, Lynn, and Shacham) cryptography instead of RSA, as it allows for threshold signatures. A threshold signature requires that in a given group of participants, some minimum (threshold) number of them have to contribute a share of the signature in order to produce a single overall valid signature. In the context of a blockchain, the VRF selects a group of registered nodes to verify a block. Only those nodes will be able to verify that block. If a node in that chosen block wants to verify that block, they sign a message. Once the number of nodes who have chosen to verify that block passes a certain set threshold number, the VRF randomly selects the next group nodes who can verify the next block. This group is called the random beacon, and regardless of who signed the previous block, the VRF will produce the same random beacon. Once a block has been verified it is proposed to the network and has to notarized before the next block can be built on top of it. Only notarized blocks can be built upon, and only timely published block can be notarized. If the network receives a proposed block, but not the block notarization after a period, it can know for sure that the block proposal’s chain is dead. Since without notarization, it can no longer be built upon. This incentives a nodes to share their verified block quickly. A consensus point is reached when there is only a single active chain. Consensus is reached if precisely one notarized block is produced. Producing multiple notarized blocks is intentionally tolerated, so the consensus can be reached over time, rather than everyone having to reach consensus before moving on. Overall, without having to do proof of work, a block can be verified and notarized within a matter of seconds. Additionally, signatures consume less power than proof of work, and therefore this is more eco friendly and can scale better.

Security and Privacy

A, Biryukov, D. Khovratovich, and I. Pustogarov, “Deanonymisation of Clients in Bitcoin P2P Network,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pp. 15-29, Nov. 2014.

This paper describes a method for discovering the links between Bitcoin addresses (which the authors call pseudonyms) and IP addresses (which they call addresses). The attack consists of several phases. First, clients are prevented from using the Tor network by abusing a DoS-protection blacklist built into the Bitcoin client. By sending malformed messages to Bitcoin nodes through every Tor exit node, an attacker can cause all Tor exit nodes to be blacklisted. Then, an attacker can take advantage of the fact that every Bitcoin peer maintains 8 outgoing connections. By sending and receiving specific ADDR and GETADDR messages, an attacker can discover the set of outgoing connections maintained by any address they wish to deanonymize. Then, they merely have to pay attention to the set of nodes that relay transactions. By comparing these nodes to the set of outgoing connections maintained by peers, an attacker can then identify which transactions originate from which IP addresses. Since transactions contain pseudonyms, this completes the attack. This attack relies heavily on the details of the Bitcoin P2P protocol. The authors describe some countermeasures that could prevent this type of attack from taking place. Rather than uncovering an insurmountable obstacle for the privacy blockchains, this paper reveals that seemingly innocuous aspects of the network protocol can have profound impacts on the privacy of a blockchain system.

M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan, “On the Instability of Bitcoin Without the Block Reward,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS’16, Vienna, Austria, 2016, pp. 154–167.

Most discussions of miner incentives assume that block rewards and transaction fees are equivalent and entirely interchangeable. This paper challenges that assumption, introducing several new strategies that can exploit random fluctuations in transaction fees. One such strategy involves deliberately creating forks when high-fee blocks are mined. If the forker leaves half of the available fees out of their block, they can incentivize other miners to choose their branch and reap disproportionate profits. The authors also introduce an improved selfish mining strategy that also exploits variance in transaction fees. The basic idea behind this strategy is to publish high-value blocks immediately, but use low-value blocks to start selfish mining. This scheme guarantees better payouts than either the default strategy or vanilla selfish mining. These exploits cast doubt on whether cryptocurrencies (most notably Bitcoin) can function stably using only transaction fees. The authors use both analytical methods and simulation to demonstrate the effectiveness of their strategies. Their simulations agree remarkably well with

their theoretical results and provide compelling visualizations. Like the original selfish mining paper, this work deals a near-fatal blow to the notion that blockchains are incentive-compatible.

M. Conti, S. Kumar, C. Lal, and S. Ruj, "A Survey on Security and Privacy Issues of Bitcoin," *IEEE Communications Surveys & Tutorials*, May 2018. Early access.

This paper provides an overview of known Bitcoin vulnerabilities in four major categories: double spending attacks, mining pool attacks, client-side attacks, and network attacks. The authors collect and summarize the existing literature in each of these areas and describe possible countermeasures. In addition, they address privacy concerns surrounding methods for de-anonymization of Bitcoin transactions. This paper is particularly useful as a resource for finding more detailed accounts of each type of attack. The attacks are described in an even-handed way, without downplaying or exaggerating the effect of an attack on the viability of Bitcoin. Because this is an early access version of the paper, it contains many grammatical errors that make reading it challenging. It is also unnecessarily repetitive in sections. Despite these flaws, the paper's summary of attacks and thorough collection of references to other papers makes it a valuable starting point for understanding Bitcoin security.

A. Dorri, M. Steger, M, S.S. Kanhere, and R. Jurdak, "BlockChain: A Distributed Solution to Automotive Security and Privacy," *IEEE Communications Magazine, 55*, 2017.

This article discusses the possibility of using blockchains to make software in smart vehicles more secure. They list three main issues with smart vehicles, namely centralization, a lack of privacy, and general safety threats. The paper then describes blockchains and Lightweight Scalable Blockchains, which are optimized specifically for the Internet of Things, and therefore would be useful in smart vehicles as well. Once they explain their implementation, they then go on to explain the smart vehicle system that would use blockchains for software updates and security measures against people trying to take control of the car. They suggest that smart vehicles will be much better protected against threats if they use a blockchain-based system.

I. Eyal and E. G. Sirer, "Majority Is Not Enough: Bitcoin Mining Is Vulnerable," *Communications of the ACM*, vol. 61, no. 7, pp. 95-102, Jul. 2018.

The authors present a "selfish mining" strategy that allows a miner's share of the total revenue to exceed their share of the total mining power. This vulnerability is present in Bitcoin and, in principle, any other blockchain system with similar structure. Under the selfish mining strategy, a miner does not reveal a mined block but instead keeps it private, releasing it only when another block is found. This creates a fork, and other nodes will divide their work across the two branches while the selfish miner continues mining their own chain. If instead the selfish miner finds another block on their secret branch, they keep it secret until the network catches up within one block, at which point the selfish miner publishes their branch. Under all circumstances, this

strategy yields more than the miner's fair share of revenue if they control at least 33% of the hash power. The exact threshold depends on the fraction of nodes that choose to mine on the selfish miner's branch. However, the selfish miner can add fake nodes to the network (a sybil attack) to reduce this threshold to near 0. Miners are incentivized to join a selfish pool once one exists and selfish pools are also incentivized to gather more members, creating an instability in the protocol. The authors suggest a fix to Bitcoin node behavior to ensure that the threshold for selfish mining is no lower than 25%. Critics of this paper argue that selfish mining causes a decrease in profit, since even though a selfish miner receives more than their fair share, they also cause the total revenue to decrease. However, this criticism ignores the fact that the pool incentivizes new members to join--a selfish miner may endure a short term profit decrease to gain power in the network. It may be the case that selfish mining is too easy to detect to be a likely threat to the network. However, this paper establishes that the Bitcoin protocol is not incentive-compatible. A previous version of this paper was made available in 2013 on arXiv.

S. Goldfeder, H.A. Kalodner, D. Reisman, and A. Narayanan, "When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies." *CoRR, abs/1708.04748*. 2017.

In this article, the authors studied how third party tracking on the web is risky for cryptocurrency users in that their pseudonymity may be compromised, especially when it comes to payment on various sites. A large number of merchant websites intentionally give out payment data to third parties, which if intercepted or collected by malicious parties can make it easier to link transactions in bitcoin and other cryptocurrencies to specific identities. In their conclusion they discuss ways to mitigate this leak and possible compromise of identities - for example, using https on all pages, not outsourcing bitcoin payment, not leaking payment details, and avoiding unintentional leaks. They also suggest the use of various protections on the client side, including browser extensions and mixing of transactions. The authors point out a common issue with cryptocurrencies: while blockchain transactions themselves are fairly private, the user side can be compromised easily.

R. Henry, A. Herzberg, and A. Kate, "Blockchain Access Privacy: Challenges and Directions." *IEEE Security & Privacy* vol 16, pp. 38-45, Jul/Aug 2018.

The authors claim that despite the intentions of early adopters, Bitcoin and similar cryptocurrencies are less private than established banks, in some part due to lack of regulation. They discuss the issue that identities are not hidden in a blockchain, meaning that many users turn to external methods, especially Tor, which pose security concerns. They do acknowledge that there are many aspects of the blockchain system that are relatively secure, but note that there has been oversight in terms of network-level security, such as securing - or not - the IP addresses that are being used by people trying to conceal their identities. They discuss existing protocols that make transactions more private, noting Zcash and Monero's cryptographic distinctions from Bitcoin. However, the authors spend a great deal of time exploring the issues that arise when users employ Tor to add another layer of security to their transactions. Their

general conclusion is that there needs to be a great deal of rethinking done about how private transactions are.

A. Judmayer, N. Stifter, P. Schindler, and E. Weippl, “Pitchforks in Cryptocurrencies: Enforcing rule changes through offensive forking and consensus techniques,” in *Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology*, CBT '18, 2018.

This paper covers the steps in a Pitchfork attack on a blockchain, in which attackers fork a blockchain, creating a branch that misuses merged mining (Auxiliary Proof-of-Work) to reduce the utility of the original branch, causing the majority of miners to switch branches. Attackers start of creating a hard fork of a blockchain. This results in two branches, the original parent blockchain, and the new auxiliary blockchain. Attackers then update the protocol of the auxiliary branch so it supports merged mining (Auxiliary proof of work), which allows a parent block and an auxiliary block to be mined, by doing proof of work just on the parent block. Miners of the auxiliary branch then create an auxiliary block which contains recent transactions on that branch.They then create an empty parent block (just has the coinbase transaction), and then insert a hash of the auxiliary block in an unused section in the parent block. They mine the empty parent block. The auxiliary blockchain sees that the mined parent block contains a hash of an auxiliary block, and accepts that as proof of work, paying that miner. The parent blockchain also accepts that as a valid proof of work and pays that miner with that block’s coinbase transaction. The goal of this attack is to mine empty blocks, reduce the utility of the parent branch, and cause the majority of miners to switch to the auxiliary branch, which will kill off the parent branch. Once that occurs, miners go back to just mining the auxiliary block for a reward, instead of the parent block. Even if they fail, miners still get paid in the auxiliary branch for each block they mine, and also by the parent branch. Also, no unspent transactions are involved in the parent block, so it hard to track these attackers

A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou, “Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts,” in *2016 IEEE Symposium on Security and Privacy (SP)*, 2016, pp. 839–858.

This paper proposes a smart contract system that does not make the transactions themself vieweable from the blockchain, thus making them more private. This is being proposed because pseudonymity still allows for tracking balances and transactions by watching a certain public key and this can be used to the detriment of the public key holder. The authors claim that this aspect of the smart contract system is a major issue when it comes to making smart contracts more widely accepted.

M. Vasek, J. Bonneau, R. Castellucci, C. Keith, and T. Moore, “The Bitcoin Brain Drain: A Short Paper on the Use and Abuse of Bitcoin Brain Wallets,” *Financial Cryptography and Data Security*, *Lecture Notes in Computer Science*, Springer, 2016.

This paper discusses brute force attacks on a specific kind of wallet: a so-called “brain wallet” which generates private keys deterministically based on a user’s password. This renders wallets vulnerable to offline attacks, as the attackers can check passwords by generating private keys from them and searching the blockchain for those keys. The authors took a similar approach, checking 300 billion passwords against the Bitcoin blockchain and discovering 884 of these wallets, 98% of which had been drained at least once. Although they confirm that some of these events are associated with the owners of the wallets, they also confirm at least 14 drainers targeting multiple wallets.This supports their conclusion that although brain wallets are attractive for the ease of use, these wallets cannot be considered a secure way to use cryptocurrency. Ultimately, this is a single type of wallet, and not one in widespread use, so the paper does not have wide-ranging implications. However, it does support the point that much of the crime and attacks on cryptocurrency do not involve the cryptocurrency itself, but rather users’ insecure ways of interacting with it. If cryptocurrency and other applications of blockchain are to become more widespread, an equal amount of attention must be given to these issues as to the more esoteric attacks on the inner workings of blockchain.

G. Zyskind, O. Nathan, and A. Pentland, “Decentralizing Privacy: Using Blockchain to Protect Personal Data,” in *2015 IEEE Security and Privacy Workshops*, pp. 180-184, 2015.

This paper states that as users begin to care more and more about the privacy of their ever growing data presence, blockchains may prove to be a secure solution. The authors set out to address the lack of control over personal data held by a third party, and argue that blockchains are a good solution. Their approach is to combine a blockchain mechanism with a non-blockchain storage method in order to make personal data only accessible to the holder and those the holder would want to have access. This paper does make some unwarranted assumptions, for instance that users would be smart about the way they use their keys. They also make the assumption that blockchains are “tamper-free,” so the system would have to be large enough not to be vulnerable to a 51% attack.