When Trust Saves Energy: A Reference Framework for Proof of Trust (PoT) Blockchains*

Leila Bahri
Royal Institute of Technology - KTH, Stockholm
lbahri@kth.se

Sarunas Girdzijauskas
Royal Institute of Technology - KTH, Stockholm
sarunasg@kth.se

ABSTRACT

Blockchains are attracting the attention of many technical, financial, and industrial parties, as a promising infrastructure for achieving secure peer-to-peer (P2P) transactional systems. At the heart of blockchains is proof-of-work (PoW), a trustless leader election mechanism based on demonstration of computational power. PoW provides blockchain security in trusless P2P environments, but comes at the expense of wasting huge amounts of energy. In this research work, we question this energy expenditure of PoW under blockchain use cases where some form of trust exists between the peers. We propose a Proof-of-Trust (PoT) blockchain where peer trust is valuated in the network based on a trust graph that emerges in a decentralized fashion and that is encoded in and managed by the blockchain itself. This trust is then used as a waiver for the difficulty of PoW; that is, the more trust you prove in the network, the less work you do.

KEYWORDS

Blockchain, Proof of Work, Bitcoin, Distributed Ledger, PoW is expensive, Proof of Trust, PoW alternative

ACM Reference Format:
Leila Bahri and Sarunas Girdzijauskas. 2018. When Trust Saves Energy: A Reference Framework for Proof of Trust (PoT) Blockchains. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23–27, 2018, Lyon, France.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3184558.3191553

1 INTRODUCTION

Bitcoin has proved that trustless peers can create and exchange value over the internet without the intermediation of any central trusted authority. The main technology underneath Bitcoin is *Blockchain*, which is a public ledger that logs system transactions as they happen through time, and that is managed by the Bitcoin P2P network. Blockchains have been attracting the attention of many parties from the financial,[1] industrial,[2] as well as the academic (e.g., [10], [4]) domains, and are considerred, by many, as a revolution in P2P computing [16]. However, blockchain, as exemplified by Bitcoin, is also a beast that survives on consuming massive amounts of energy to achieve its core consensus protocol of Proof-of-Work (PoW). The current estimated annual electricity consumption of Bitcoin is of about 39.5 TWh, which is a bit above the annual consumption of whole countries, such as Qatar, and Bulgaria.[3]

In this work, we explore an alternative mechanism for blockchain operation, that provides the same P2P transaction capabilities without having to consume such huge amounts of energy. Our approach waives energy consumption by trust, introducing the concept of Proof-of-Trust (PoT) blockchains. We model PoT for use case scenarios where peers can express opinion about each other and declare trust links within the system. The trust can be declared based on explicit factors in the system (e.g., transactions happening between the peers, behavior observed in the network), or on other implicit elements, such as business relationships between peers or any other criteria relating to the underlying application supported by Blockchain. Such use cases are getting more pronounced with the raising interest in the conceptualization of private permissioned blockchains, or other non-currency blockchain based systems (e.g., identity management, access control, data provenance, consortium management, etc.). In PoT each peer gains a trust value in the network based on a dynamic trust graph that emerges in a decentralized fashion and that is encoded in and managed by the blockchain itself. This gained trust is used as a waiver to the amount of energy that has to be spent for traditional PoW.

We developed PoT framework, and analyzed its security and operability qualities. Our ongoing experimental exploration and security analysis under defined attack models show very promising results. For instance, initial experiments on two real world trust datasets, including the BitcoinOTC one [11], have shown that trust network dynamics always guarantee the emergence of trusted peers, who make, in average, less than 10% of the whole network; allowing PoT to achieve at least 20 fold energy improvements compared to not-waived PoW.

2 BLOCKCHAIN TECHNOLOGY

A blockchain, as in Bitcoin, is a distributed ledger available to anyone participating in the network, where all related transactions are announced publicly to all the participating peers. Peers commit valid transactions into blocks that are cryptographically locked to the previously committed block. These blocks of committed transactions form a chronological and sealed sequence of blocks; i.e., a blockchain. A blockchain can be either public or private. A public blockchain operates under a *permissionless* model by which peers

*A short Paper - This is the definitive version available from publisher with specified DOI

[1] e.g., https://hbr.org/2017/03/how-blockchain-is-changing-finance

[2] e.g., IBM's HyperLedger: https://www.hyperledger.org/

[3] https://digiconomist.net/bitcoin-energy-consumption

WWW '18 Companion, April 23–27, 2018, Lyon, France

ACM ISBN 978-1-4503-5640-4/18/04.
https://doi.org/10.1145/3184558.3191553

can join the network without any identity management being required. In a private blockchain, participating peers are subject to a *permissioned* model by which their identities need to be approved (by some third party authority that could be the blockchain's network itself) before they can participate in the network. Bitcoin's blockchain is permissionless, and our PoT framework can be fit to both models.

Consensus technique - Proof-of-Work (PoW). Blockchain management requires a consensus mechanism on who will propose the next block that will be extending the blockchain at each new round, all while considering the trustless P2P byzantine environment imposed in the system. This decentralized consensus is achieved by a cryptographic hash puzzle that allows both sealing blocks together, as well as leader election for proposing the next block. The cryptographic hash puzzle requires finding a cryptographic nonce that, together with the information from the previous block, hashes into a number that is smaller than a given agreed on value, known as the difficulty level. The use of cryptographic hashing provides both a tamper-proof quality to the blockchain, as well as a computationally intensive puzzle. This enables a solution to the byzantine *leader election* problem, where leaders are selected (and trusted) proportionally to their computational capacity. As an incentive to engage in the proof-of-work challenge, leaders get rewards in Bitcoin, and this is also the mechanism by which new Bitcoins are produced in the system (hence the term, Bitcoin mining).

PoW is very expensive. PoW is an energy exhaustive mechanism, where energy expenditure grows proportionally to the total computational capacity available in the system. Indeed, the PoW difficulty level is self-adjusted in the the network in order to maintain a predefined time in between consecutive rounds. Also known as the block time, this reflects the time duration between every two confirmed blocks in the blockchain and serves to synchronize communication in the network. The block time needs to be large enough to account for network latency and to ensure that most of the nodes in the P2P network have heard about a block before the next one is announced. More importantly, this parameter is positively correlated to the security of a PoW blockchain, as the larger the block time is, the more resilient to attacks the blockchain is [6]. Higher block times are also equivalent to slow transaction confirmation; that is, speed and security are usually traded off in PoW blockchains. For instance, the block time in Bitcoin protocol is conventionally set to 10 minutes and the transaction confirmation time is considered to be 60 minutes; i.e., a transaction is securely confirmed in Bitcoin when it is included in a block that is at least 6 blocks deep in the corresponding blockchain.[4] In order to maintain the predefined block time and its corresponding security guarantees, the difficulty of PoW increases as the computational capacity available in the system gets higher (CPUs are getting more powerful and/or more people join the network). For instance, in 2016, Bitcoin miners had been hashing at an average rate of 10^{18} hashes per second; whereas in 2017, this rate increased to a magnitude in the order of 10^{21}.[5] This clearly results in ever increasing energy expenditure rates for the same volume of transactions. Besides, high PoW difficulty levels make it very expensive for normal peers with commodity hardware to participate in the system. Indeed, Bitcoin is currently being run by a few mining-pools that are controlling its management; hence it is effectively centralized in the hands of a few players. [6]

These downsides of PoW have raised concerns about its scalability, as well as motivated the investigation of other more energy efficient alternatives. Most of the suggested alternatives[7] rely either on the ownership of physical resources (e.g. [17], [12]), or on the ownership of monetary value (i.e., stake) in the blockchain system (e.g., [9]). The premise under this line of thought is that demonstrating ownership of something that has value in real life could substitute the missing trust between peers. Effectively, it follows the assertion, "the more you own the more you are trusted", or "the more you own, the more legitimacy you have to leadership".

3 OUR PROPOSAL: PoT

We question the approach of solely relying on ownership of physical or monetary resources under use case scenarios where some form of trust can be expressed and/or captured within the system itself. Such use cases are indeed getting more pronounced with the raising interest in the conceptualization of private permissioned blockchains, or other non-currency blockchain based systems. Therefore, and starting from the assumption that trust is inherent to any collaborating system (as peers can form opinion about each other through time based on the transactions they share or the behavior they observe in the network), we investigate the usage of proof of trust (PoT) as a waiver for the amount of work (or physical richness) that peers need to demonstrate for leader election. The objective is to minimize the amount of energy spent on PoW as more trusted peers appear in the network. That is, install a mindset of "the more trusted you are, the less work you are allowed to perform". Our core idea relies on decreasing the amount of work that needs to be performed for PoW, by relying on a trust metric that is used as a work-waiver; hence achieving consensus with less energy consumption. As represented in Figure 1, whilst the amount of energy required for PoW is drastically increasing as the total computational power in the system grows,[8] it is expected to be decreasing in PoT as more trusted peers emerge in the network.

Designing PoT requires answering a number of design and technical challenges. First, it is required to design a model based on which trust will be managed in the network, and a mechanism by which every peer in the network is associated with a trust metric. This should be achieved with consideration to the decentralized property of blockchains. In addition to that, it is required to design a PoT protocol that should use the generated trust metrics as a calibration for PoW without subverting the security of the blockchain or affecting its performance. One of the issues here is in ensuring that the system will not be dominated by a few highest trusted peers, and that the network can be self-adaptive to avoid monopolized trust power.

[4] https://en.bitcoin.it/wiki/Confirmation
[5] https://blockchain.info/charts/hash-rate?timespan=2years
[6] https://www.technologyreview.com/s/610018/
[7] Most alternatives are discussed in forums and a few only have been formalized in research papers.
[8] Available computational power is indeed growing with time, and statistics on the Bitcoin network show an increase in the available hash rate from teta to zeta hashes per second within the year of 2017 only: https://blockchain.info/charts/hash-rate

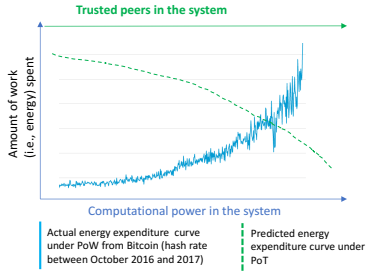

Figure 1: Exemplification of the core difference between PoW and PoT in terms of amount of energy (i.e., amount of work) that needs to be spent for leader election. The solid blue curve represents the increase in energy expenditure as the total computational power increases in the system, and is inferred from the actual change in PoW difficulty in Bitcoin between January 2016 and October 2017. The dotted green curve represents how energy is expected to decrease as more trusted peers appear in the network and is illustrative only.

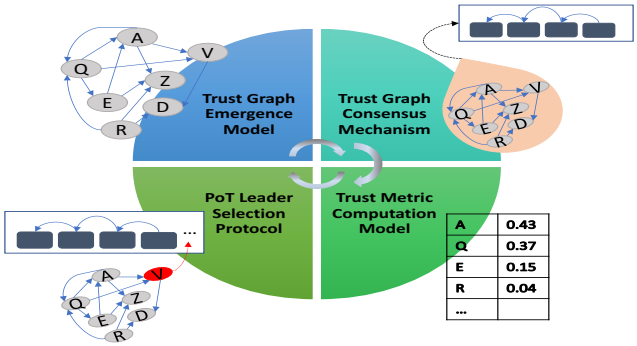

Figure 2: The PoT Reference Framework: four connected modules with graphical illustration.

Considering the above mentioned issues, we design a PoT reference framework that represents the main building modules that need to be addressed to achieve a PoT blockchain. Illustrated on Figure 2, our proposed PoT framework is comprised of 4 phases that are run consecutively following a decentralized and secure PoT protocol.

Trust graph emergence model. PoT implies a trust graph based on which peers can be annotated with a *trust metric*. Ideally, each peer should be able to declare, in a decentralized uncontrolled fashion, her trust towards any member of the system, thereby constructing a trust graph which later can be unambiguously encoded into the blockchain itself. Trust graphs could also emerge in a variety of other ways, such as being extracted from the peers social networks, inferred from the interactions between peers in the underlying system, etc. These options depend on the extent to which identity management is enforced (i.e., public with permissionless participation or private with permissioned membership), and should not in any case break the decentralized property of the environment.

Trust graph consensus mechanism. The trust graph is the basis for computing trust metrics that should be available to all the peers in the network for consensus. The most intuitive idea is to have the trust graph written in the blockchain itself. The limitation might be on the size of the graph, as compared to the block size limitation imposed on blockchains because of network throughput considerations. We propose that *only a secure digest of the trust graph be committed in the blockchain*, providing as such light-weight decentralized consensus on its status.

Trust metric computation model. Based on a trust graph that everyone agrees upon (i.e., encoded in the blockhain) there is need to extract a trust metric that would, on one hand give more trust to nodes that have more incoming trust links, while on the other hand would not allow any centralized superhubs to appear. The literature contains a plethora of algorithms that could be used to compute trust metrics (e.g., [5], [7], [8], [14]). The key element here is determinism, so as all peers can reach the same result independently of each other. In addition to that, the metric should be resilient to forged trust, such what can be achieved using spam farms. Indeed, this is a well known issue in graph-based trust (i.e., centrality) measures, as nodes can create forged identities with the aim of increasing their incoming links; hence boosting their importance in the network [3]. The literature has a number of proposals to address this problem, mostly focusing on establishing white-lists of trusted nodes, based on which suspect malicious nodes can be identified and quarantined in the system [15]. The need for white-lists comes from the fact that the connectivity patterns exhibited by highly central nodes look similar, regardless whether the node gained this centrality legitimately (really highly trusted) or illegitimately (backed up by a spam farm). This makes it hard to differentiate between truly highly trusted nodes and those that look like highly trusted but are in fact backed up by bogus identities. However, in our scenario, we may be interested in discarding the highest trusted nodes, regardless of how their trust has been obtained. On one hand, it is obvious that illegitimately highly trusted nodes should be discarded. On the other hand, top highly trusted nodes, even if legitimate, may introduce the risk of centralized power within their small circle (of highly trusted nodes). Indeed, *power-law behavior* is usually exhibited in natural graphs, where only a few nodes achieve considereably higher centrality (i.e., hubs) compared to the rest of the nodes in the graph [1]. In addition to that, the distribution of trust in natural trust graphs is also known to be slowly changing, with the phenomena of -the rich gets richer and the poor gets poorer. As such, the small subset of highly trusted nodes would easily be maintaining its position through time, making it also challenging for new nodes to join. Considering these elements, designing a PoT that naively selects the highly trusted nodes for each block introduces the risk of making the system quasi-centralized, with power in the hands of the few richest nodes. Therefore, it is needed to develop a trust metric that takes all these aspects into consideration to smartly and adaptively redistribute power among all trusted nodes, instead of having it centralized within the small circle of richest nodes only. Our initial studies on combining different centrality measures, such as pagerank [14] and Katz centrality [2], show promising results.

PoT leader selection. One of the most straightforward ideas is to elect peers based on a round-robin among the K top trusted peers where K is an agreed on number in the network. However, this will result in a centralized management of the blockchain in the hands of the K peers only. In order to avoid that, we develop our PoT protocol

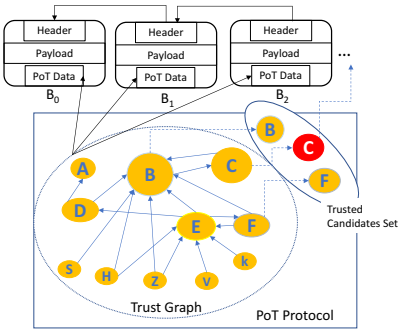

Figure 3: PoT overview: the trust graph is encoded in the blockchain, PoT protocol is run and produces, per every round, a randomly selected node proportional to its total trust in the underlying trust graph, and that should be the leader to propose the next block in the managed blockchain.

based on waiving PoW by trust. Under PoT, nodes continue running PoW with one key difference - the difficulty of the cryptographic puzzle for each node is adjusted inversely proportionally to its trust value extracted from the trust graph. In effect, the system waives the PoW effort for the high trust nodes and elects a leader proportionally to its total trust in the system. Low trusted nodes will have very minimal chances to be selected, as the more trusted ones are privileged to mine at much lower difficulty levels. Thus, it is expected that the nodes with low trust will not have incentives to even attempt PoW, thereby reducing the overall energy footprint of the system. However, this does not mean that the low trusted nodes will be discarded from the system, as they can always manage to increase their trust over the future rounds. Considering this, we represent on Figure 3 a practical suggestion based on limiting the selection of leaders only among nodes with trust higher than a predefined threshold depending on the blockchain application domain. At each round, a *Trusted Candidate Set (TCS)* is generated (by the leader of the previous round and based on the updated trust graph encoded in the blockchain), and a winning miner is selected using a low difficulty PoW. This difficulty level is tuned in the network based on the mining capacity available in *TCS* at each round. As explained earlier, this tuning is required to maintain a steady block time in the system. We note that, in PoT, trusted peers are allowed to mine at lower difficulty levels compared to pure PoW, resulting as such in much shorter block times. Since participation in PoT is intrinsically prohibitive by the achieved trust, shorter block times could be tolerated without having the same side effects on the system's security as in pure PoW (i.e., trusted nodes are assumed to be honest).

Security of PoT. The security of PoT relies on the premise that members of the *TCS* are honest (since they are trusted). The main issue is to be able to control the extent to which malicious peers can subvert the underlying trust graph so as to illegally increase their resulting trust metric. Besides, there is also the risk of falling into a quasi-centralized scenario where the blockchain management is dominated by a few highly trusted nodes. These could also form conglomerates for monopolized power to make it expensive (probably impossible) for new nodes to win the competition. Therefore, there is need for a smart trust metric, that should also be self-adaptive in the network, to prevent both malicious behavior from bogus identities, as well as monopolizing behavior from a few giants of trust. Our initial experimental analysis reveal promising results. For example, investigation on combining two different centrality measures, pagerank [14] and Katz [2], show that legitimate superhubs can, in 95% of cases, be differentiated from those backed by spam farms that follow the same structure as in [3].

Energy saving under PoT. PoT is expected to reduce the amount of energy required to maintain blockchain based systems, especially for use case scenarios that are not necessarily dealing with pure digital currency management. Trust is inherent to a plethora of application domains, where participating peers are usually not equally trusted. Using the concept of PoT is expected to inhibit low trusted peers from participating in the mining process, as their chances would be considerably low compared to the *privileged* more trusted nodes. By only achieving that, energy that would have been spent by the low trusted nodes would be naturally saved. In addition to that, the more trusted nodes would mine at lower difficulty levels, resulting as such in faster transaction confirmation rates; hence further reducing energy consumption per transaction. Our initial experiments on the BitcoinOTC dataset [11], a trust network of peers who transact on Bitcoin OTC platform with about 5k nodes and 11k edges,[9] show that peers are not equally trusted, with only 5.75% highly trusted nodes (i.e., trust value at least 25 times higher than all the low trusted nodes). With PoT waiving, the 94.25% least trusted nodes at every epoch will not have incentives to participate in PoW, making a saving of at least as much energy.

4 CONCLUSION

In this work, we propose the concept of PoT blockchains where trust is used as a waiver for PoW. The novelty of our work is a PoT framework that defines the critical modules needed to achieve PoT blockchains and formalizes the required properties to ensure their robustness and security. PoT saves the decentralized and security qualities of PoW blockchains, while providing potential to significantly decrease energy consumption. We are conducting an extensive study of the security of the system under formalized attack models, as well as working on making PoT self-adaptive to prevent dominance of highly trusted (i.e., richest) nodes, and to make the network resilient to illegitimate boosting of trust (e.g., resilient to spam farms phenomena).

REFERENCES

[1] Lada A Adamic, Rajan M Lukose, Amit R Puniyani, and Bernardo A Huberman. 2001. Search in power-law networks. *Physical review E* 64, 4 (2001), 046135.
[2] Stephen P Borgatti and Martin G Everett. 2006. A graph-theoretic perspective on centrality. *Social networks* 28, 4 (2006), 466–484.
[3] Ye Du, Yaoyun Shi, and Xin Zhao. 2007. Using spam farm to boost PageRank. In *Proceedings of the 3rd international workshop on Adversarial information retrieval on the web*. ACM, 29–36.
[4] Ittay Eyal, Adem Efe Gencer, Emin Gün Sirer, and Robbert Van Renesse. 2016. Bitcoin-NG: A Scalable Blockchain Protocol.. In *NSDI*. 45–59.
[5] Rino Falcone, Giovanni Pezzulo, and Cristiano Castelfranchi. 2002. A fuzzy approach to a belief-based trust computation. In *Workshop on Deception, Fraud and Trust in Agent Societies*. Springer, 73–86.
[6] Juan A Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The Bitcoin Backbone Protocol: Analysis and Applications.. In *EUROCRYPT (2)*. 281–310.

[9] https://www.bitcoin-otc.com/

[7] Ferry Hendrikx, Kris Bubendorfer, and Ryan Chard. 2015. Reputation systems: A survey and taxonomy. *J. Parallel and Distrib. Comput.* 75 (2015), 184–197.

[8] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision support systems* 43, 2 (2007), 618–644.

[9] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference*. Springer, 357–388.

[10] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. 2016. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 839–858.

[11] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. http://snap.stanford.edu/data. (June 2014).

[12] Mitar Milutinovic, Warren He, Howard Wu, and Maxinder Kanwal. 2016. Proof of Luck: An efficient Blockchain consensus protocol. In *Proceedings of the 1st Workshop on System Software for Trusted Execution*. ACM, 2.

[13] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).

[14] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web.* Technical Report. Stanford InfoLab.

[15] Nikita Spirin and Jiawei Han. 2012. Survey on web spam detection: principles and algorithms. *ACM SIGKDD Explorations Newsletter* 13, 2 (2012), 50–64.

[16] Don Tapscott and Alex Tapscott. 2016. *Blockchain Revolution: How the technology behind Bitcoin is changing money, business, and the world.* Penguin.

[17] Fan Zhang, Ittay Eyal, Robert Escriva, Ari Juels, and Robbert van Renesse. 2017. REM: Resource-Efficient Mining for Blockchains. *IACR Cryptology ePrint Archive* 2017 (2017), 179.

[18] Guy Zyskind, Oz Nathan, et al. 2015. Decentralizing privacy: Using blockchain to protect personal data. In *Security and Privacy Workshops (SPW), 2015 IEEE*. IEEE, 180–184.