Comments to the European Securities and Markets Authority on its Consultation on Distributed Ledger Technology Applied to Securities Markets

Peter Van Valkenburgh
September 2, 2016

Introduction

Coin Center is a non-profit research and advocacy center focused on the public policy issues facing open and decentralized blockchain technologies. Specifically, our focus encompases cryptocurrencies (*e.g.* Bitcoin), decentralized computing platforms (*e.g.* Ethereum) and inter-ledger systems and protocols (*e.g.* sidechains). Our mission is to build a better understanding of these technologies and to promote a regulatory climate that preserves the freedom to innovate using them. We do this by producing and publishing policy research from respected academics and experts, educating policymakers and the media about decentralized blockchain technology, and by engaging in advocacy for sound public policy. In that spirit, please find below our comments on the discussion paper entitled, "The Distributed Ledger Technology Applied to Securities Markets."

Our message is focused on ESMA's conclusion in the discussion paper that "open" or "permissionless" blockchains may be inappropriate for financial services:

> 3. Importantly, ESMA understands that the DLT that would be used for financial services would differ from the Blockchain designed for Bitcoins in a number of ways. In particular, while the Bitcoin Blockchain is an open system where all can contribute to the validation process ('permissionless' system), the DLT that is likely to be used in financial markets would be a permissioned-based system with authorised participants only. This difference is important to keep in mind because it has a number of consequences in terms of potential benefits and risks. [1]

[1] ESMA, Discussion Paper: The Distributed Ledger Technology Applied to Securities Markets (Feb. 6, 2016) https://www.esma.europa.eu/sites/default/files/library/2016-773_dp_dlt.pdf

The presumption that only permissioned-based systems are "likely" to be used in financial markets is premature. As of yet, all permissioned-based systems remain in the proof-of concept ("POC") stage of development, and none are, as-of-yet, securing high-value information. Permissionless systems have been running in public for almost ten years, and they are battle-tested. The largest, Bitcoin, secures a ledger that describes roughly $10 Billion worth of valuable assets. This is, in essence, a multi-billion-dollar bug-bounty challenge to break a permissionless blockchain's security model, a challenge that has yet to be met by anyone in the world. ESMA has rightly articulated some risks that come with these permissionless systems. However, until permissioned POCs begin to face similar scrutiny from security researchers and hackers, it is difficult to conclude that these systems better address risks or can be as robustly secured as Bitcoin has thus far proven to be.

> 36. We understand that the DLT that is likely to be applied to securities markets would be 'permission-based' in contrast to the 'permissionless' system that was originally designed for virtual currencies, e.g., Bitcoins, for a number of reasons, including efficiency, security and privacy purposes. In a permission-based system, only 'permissioned' participants can act as a node, i.e., validate transactions. Their identity is typically known to the rest of the network.[2]

We would like to address the implied superiority of permissioned systems when it comes to efficiency, security, privacy in turn.

Efficiency

It is true that present day permissionless blockchains generally have a slow settlement speed (for Bitcoin 10 to 60 minutes depending on one's desire for multiple block confirmations; and for Ethereum 15 seconds to 2 minutes), and also that these systems may have limited throughput, which for Bitcoin is ~5 transactions per second (TPS) globally. Nonetheless these limitations are not fundamental to the technology.

Bitcoin's maximum TPS exists because of an arbitrary value in a line of the core software code that could be changed. Ethereum has no such hard limit to TPS in the code, and stress tests to the early testnet chain have indicated that the network could handle as many as 20 to 30 TPS. Ethereum's faster block time and higher theoretical throughput are indicative of the growth of technology in the permissionless space. Similarly, several improvements to both Bitcoin and Ethereum are being researched, among them the Lightning Network, proof-of-stake, and consensus sharding. While still speculative, such technologies could potentially deliver settlement times and

[2] *Id.*

throughput that surpasses even highly sophisticated centralized payment or clearing intermediaries today.

Security

Generally speaking, both permissioned and permissionless systems are similar in that they are only trustworthy so long as a majority of the validators (either identified consortium members or open participants, respectively) are behaving honestly. Permissioned systems, however, exhibit two fundamental weaknesses in security as compared to open systems.

First, in a permissioned system we must also consider the integrity of the entity that identifies and then grants credentials to the consortium members. If this identifying member or authority is corrupted (whether by negligence, internal malfeasance, or hacking), it could potentially shift the balance of power on the permissioned network by granting more participatory rights to one or another consortium member than was assumed to be fair and agreed upon by the other members. The sanctity of the consensus mechanism, and thus the immutability of a ledger, is only upheld by trust in an identifying agent and the safekeeping of identity credentials by participants (who could also, of course, be compromised).[3] This is in contrast to permissionless networks where the consensus mechanism relies on a provable sacrifice of resources by participants—thus reducing the viability of an attack by adding true economic costs to the efforts of a hacker.

Second, the nature of an identified consortium may make it easier for some subset of the consensus members to find each other and collude to defraud the rest of the network.

Privacy

The information upon which a network reaches consensus cannot be truly private. For a network to reach agreement on the authoritative state of a ledger, some information about participant transactions must be public to the group of validators—whether an

[3] Identity and identity authentication is notoriously difficult online. The technology currently employed to provide secure and identified communications channels on the web is known as SSL (Secure Sockets Layer) and it relies on certificate authorities to issue digital certificates for identity and authentication purposes. Security expert Bruce Schneier writes of these systems, "As it is used, with the average user not bothering to verify the certificates exchanged and no revocation mechanism, SSL is just simply a (very slow) Diffie-Hellman key-exchange method. Digital certificates provide no actual security for electronic commerce; it's a complete sham." The harms generated by this sham, writes Schneier, are mitigated not via any form of technological security but rather by the willingness of credit card companies to reverse fraudulent transactions. Bruce Schneier, *Secrets and Lies*, Chapter 15, "Certificates and Credentials," section "PKIs On The Internet" (page 238).

open or closed set. Faced with an essential trade-off wherein verifiability requires transparency, but privacy requires that user-data remain opaque, there are essentially two design options:

- **Perimeter Security:** Leave all data relevant to the consensus transparent, but restrict the set of verifiers with whom one is comfortable sharing otherwise private data.
- **Data Minimization:** Only reveal data essential to group consensus if it is absolutely necessary to verification, and allow the group of verifiers to be open and global.

Perimeter security follows an older approach in network security generally: if there are things to be kept secret, we build a secure perimeter, restrict the flow of sensitive information to within that perimeter and only allow authorized parties into that perimeter.

Data minimization takes an alternative approach: there will be no secure perimeter, all information in the system can be presumed to be public and available, but the only information ever put into to the system is the minimum amount of information necessary to accomplish the goal.

It has been suggested that open consensus mechanisms are not suitable for enterprise or financial services applications because they are not sufficiently private, and Bitcoin presents us with an example of this weakness: bitcoin address pseudonyms are too easily identified and transaction histories of users are too vulnerable to public scrutiny. However, faced with this dilemma, there are a variety of solutions. A commonly cited solution is to build only closed, consortium-consensus networks for these sensitive use-cases. The only privacy gain inherent to this approach is the creation of perimeter security.

The Banking technology consortium R3, for example, has described it's Corda distributed ledger product as follows:

> The foundational object in our concept is a state object, which is a digital document which records the existence, content and current state of an agreement between two or more parties. It is intended to be shared only with those who have a legitimate reason to see it.[4]

[4] Corda Introductory Whitepaper (Aug. 24, 2016) http://static1.squarespace.com/static/55f73743e4b051cfcc0b02cf/t/57bda2fdebbd1acc9c0309b2/1472045822585/corda-introductory-whitepaper-final.pdf.

Privacy is thus ensured by only sharing the "state object" with one's trusted counterparties, and with those "who have a legitimate reason to see it." The agreement is made private by placing it behind a secure perimeter, not necessarily by limiting the contents of the agreement to data relevant to consensus over that agreement. If any of the "legitimate" parties are compromised, the contents of the agreement could become public. In this sense the consortium model on its own does little to change the state of information security beyond what we see from existing centralized financial intermediaries. Indeed, it may be on balance a more vulnerable system, because the secure perimeter now includes employees at other firms. Additionally, if the entire contents of the agreement are private to the relevant parties, independent validation of the data cannot occur in a fully trust-minimized manner (*i.e.* from an open and global network of impartial transaction validators performing proof-of-work or proof-of-stake); one only gets validation from the set of parties permitted by the users to enter the secure perimeter.

To R3's credit, it is investigating various other approaches to better enhance privacy as described in their near to mid-term roadmap:

> Privacy enhancements using technology such as address randomization, zero-knowledge proofs.[5]

These, however, are approaches that apply equally well in consortium as well as open consensus driven systems, and have been primarily pioneered in the Bitcoin and related cryptocurrency communities.

Address randomization is effectively the attempt to create more robust pseudonyms that fail to yield to forensic identification techniques. Most research into the development of these techniques is occurring in the Bitcoin space where, without robust address randomization, privacy is fairly poor as previously described. Notable pioneering advances in this approach are the Coin Join[6] and Coin Shuffle[7] protocols, which create decentralized communications channels to facilitate the shuffling of bitcoins between several addresses in a manner that makes it difficult to link a set of addresses to one particular user. Additionally, changes to the Bitcoin core protocol have been researched and proposed that would obscure the value of each transaction as it appears in the

---

[5] *Id.*

[6] Blockchain.info, SharedCoin and other CoinJoin implementations: Uses and Limitations (June 10, 2014) https://blog.blockchain.com/2014/06/10/sharedcoin-and-other-coinjoin-implementations-uses-and-limitations/

[7] Tim Ruffing, Pedro Moreno-Sanchez, and Aniket Kate, CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin https://crypsys.mmci.uni-saarland.de/projects/CoinShuffle/coinshuffle.pdf

ledger (thus making it more difficult to identify a pseudonym based on a pattern of transactions), a project referred to as Confidential Transactions.[8]

Zero-knowledge proofs are a cryptographic tool for proving some important fact (*e.g.* this transaction is valid, these bitcoins have never been spent by this sender before) without revealing any other information aside from the proof (*e.g.* this is the identity of the sender, the recipient, and the amount sent). Integrating zero-knowledge proofs into an open consensus blockchain could potentially allow a decentralized open set of transaction validators to prove that all recent transactions have been appropriately funded, signed, and not double-spent, without revealing any additional information about who sent how much to whom. Zero-Knowledge proofs are cutting edge science and few people in the world have the expertise to develop and implement these novel tools. The Zcash Electric Coin Company has been pioneering these technologies in the form of Zcash, an open proof-of-work-driven digital currency protocol based on Bitcoin.[9] Not only is Zcash testing the viability of a truly data minimized approach to privacy and consensus, the protocol also allows users to selectively disclose information about their transactions to whomever they choose.

> Zcash transactions automatically hide the sender, recipient and value of all transactions on the blockchain. Only those with the correct view key can see the contents. Users have complete control and can opt-in to provide others with their view key at their discretion.[10]

A system thus specified would in many ways be ideal: Trust in the scarcity of the underlying tokens, the validity of the ledger, and the non-repudiability of transactions is generated by an open set of impartial validators, rather than a consortium of identified but potentially corrupt or infiltrated parties. Privacy is guaranteed by neglecting to share any information about transactions with these validators except for the minimized amount of information necessary to prove scarcity, signature validity, and non-repudiation. And selective disclosure ensures that counterparties and third parties, such as regulators, can be given visibility into the details of any particular transaction whenever the initiator wishes to be transparent or is compelled to be transparent by law.

---

[8] The Elements Project, Confidential Transactions https://www.elementsproject.org/elements/confidential-transactions/

[9] Zcash Technology Preview, https://z.cash/

[10] Giulio Prisco, Zcash Creator on the Upcoming Zcash Launch, Privacy and the Unfinished Internet Revolution (Aug. 30, 2016) https://bitcoinmagazine.com/articles/zcash-creator-on-the-upcoming-zcash-launch-privacy-and-the-unfinished-internet-revolution-1472568389

This technology, pioneered in the *open, permissionless* blockchain space may be most likely to satisfy the desires of ESMA with respect to financial DLTs:

> Different levels of access to the network, depending on the exact nature and scope of the participant might also be needed.[11]

Conclusion

Great work is being done to develop secure, private, and efficient distributed consensus systems, both from those working on open as well as closed blockchain technologies. We urge ESMA to avoid writing-off open technologies. As we have explained, nothing fundamental to these technologies makes them inappropriate for use in financial markets.

[11] ESMA, Discussion Paper: The Distributed Ledger Technology Applied to Securities Markets (Feb. 6, 2016) https://www.esma.europa.eu/sites/default/files/library/2016-773_dp_dlt.pdf