The Evolution of Improvised Cryptocurrency – UVCoin

Jasmine Bedi Khurana
(Mentor)
Assistant Professor – 1
Galgotias College, G. Noida
Noida

Utkarsh Wadhwa
(Student)
4th year, B. Tech. (IT)
Galgotias College, G. Noida

Vivek Tomar
(Student)
4th year, B. Tech. (IT)
Galgotias College, G.

ABSTRACT
Cryptocurrencies which evolved with bitcoin has a decentralized structure based on the ledger which is handled via proof of work mechanism, indeed generating a monetary supply. We all agree that decentralization save us from the cruel national political system but has a limitation of computational cost involved and problem related to scalability. The idea is to introduce a new cryptocurrency named UV Coin which is a cryptocurrency framework having control of the central banks but involves distributed set of authorities to prevent double spending. This coin will maintain enough transparency. The proof of the benefits is partial centralization such as elimination of wasteful hashing and involves a scalable system to avoid double spending attack.

General Terms
Cryptocurrency, Altcoins, Litecoin, Ripple, Algorithm, Bitcoin

Keywords
Sybil attack, Proof of Work, Proof of stake, UVCoin, Double-Spending Attack, Cryptocurrency, Cryptographic Algorithm

1. INTRODUCTION
Bitcoin [1], introduced in 2009, and the many alternative cryptocurrencies it has inspired (e.g., Litecoin and Ripple), have achieved enormous success: financially, in November 2015, Bitcoin held a market capitalization of 4.8 billion USD and 30 cryptocurrencies held a market capitalization of over 1 million USD. In terms of visibility, cryptocurrencies have been accepted as a form of payment by an increasing number of international merchants, such as the 150,000 merchants using either Coinbase or Bitpay as a payment gateway provider.

Recently, major financial institutions such as JPMorgan Chase [2] and NASDAQ [3] have announced plans to develop blockchain technologies. The potential impacts of cryptocurrencies have now been acknowledged even by government institutions: the European Central Bank anticipates their "impact on monetary policy and price stability" [4]; the US Federal Reserve their ability to provide a "faster, more secure and more efficient payment system" [5]; and the UK Treasury vowed to "support innovation" [6] in this space. This is unsurprising, since the financial settlement systems currently in use by central banks (e.g., CHAPS, TARGET2, and Fed wire) remain relatively expensive and — at least behind the scenes — have high latency and are stagnant in terms of innovation.

Despite their success, existing cryptocurrencies suffer from a number of limitations. Arguably the most troubling one is their poor scalability: the Bitcoin network (currently by far the most heavily used) can handle at most 7 transactions per second, faces significant challenges in raising this rate much higher, whereas PayPal handles over 100 and Visa handles on average anywhere from 2,000 to 7,000. This lack of scalability is ultimately due to its reliance on broadcast and the need to expend significant computational energy in proofs- of-work — by some estimates [7], comparable to the power consumption of a large power plant — in order to manage the transaction ledger and make double-spending attacks prohibitively expensive. Alternative cryptocurrencies such as Litecoin try to distribute this cost, and Permacoin [8] tries to repurpose the computation, but ultimately neither of these solutions removes the costs. A second key limitation of current cryptocurrencies is the loss of control over monetary supply, providing little to no flexibility for macroeconomic policy and extreme volatility in their value as currencies.

Against this backdrop, we present UVCoin, a cryptocurrency framework that decouples the generation of the monetary supply from the maintenance of the transaction ledger. Our design decisions were largely motivated by the desire to create a more scalable cryptocurrency. Indeed, as Bitcoin becomes increasingly widespread, we expect that this will be a question of interest to many central banks around the world.

UVCoin's radical shift from traditional cryptocurrencies is to centralize the monetary supply. Every unit of a currency is created by a central bank, making cryptocurrencies based on UVCoin significantly more palatable to governments. Despite this centralization, UVCoin still provides the benefit over existing (non-crypto) currencies of a transparent transaction ledger, a distributed system for maintaining it, and a globally visible monetary supply. This makes monetary policy transparent, allows direct access to payments and value transfers, supports pseudonymity, and benefits from innovative uses of blockchain and digital money.

Centralization of the monetary authority also allows UVCoin to address some of the scalability issues of fully decentralized cryptocurrencies. Since mintettes are — unlike traditional cryptocurrency miners — known and may ultimately be held accountable for any misbehavior, UVCoin supports a simple and fast mechanism for double-spending detection. UV Coin adapts a variant

and performance scales linearly as we increase the number mintettes. Most transactions take less than one second to clear, as compared to many minutes in traditional cryptocurrency designs.

Beyond scalability, recent issues in the Bitcoin network have demonstrated that the incentives of miners may be misaligned, and recent research suggests that this problem — namely, that miners are incentivized to produce blocks without fully validating all the transactions they contain — is only exacerbated in other cryptocurrencies [9]. Our

hope is that this framework can lead to a more robust set of incentives. In a real deployment of UVCoin, we furthermore expect mintettes to be institutions with an existing relationship to the central bank, such as commercial banks, and thus to have some existing incentives to perform this service.

The ultimate goal for UVCoin is to achieve not only a scalable cryptocurrency that can be deployed and whose supply can be controlled by one central bank, but a framework that allows any central bank to deploy their own cryptocurrency. In fact, there is interest [10] to allow other entities to not only issue instruments that hold value (such as shares and derivative products), but to furthermore allow some visibility into transactions concerning them. With this in mind, we will discuss that what is needed to support some notion of interoperability between different deployments of UVCoin, how different currencies can be exchanged in a transparent and auditable way, and how various considerations — such as a pair of central banks that, for either security or geopolitical reasons, do not support each other — can be resolved without fragmenting the global monetary system.

2. AN OVERVIEW OF UVCOIN

At a high level, UVCoin introduces a degree of centralization into the two typically decentralized components of a blockchain-based ledger: the generation of the monetary supply and the constitution of the transaction ledger. In its simplest form, the UVCoin system assumes two structural entities: the central bank, a centralized entity that ultimately has complete control over the generation of the monetary supply, and a distributed set of mintettes [11] that are responsible for the maintenance of the transaction ledger. The interplay between these entities — and an overview of UVCoin as a whole — can be seen in Figure 1.

Briefly, mintettes collect transactions from users and collate them into blocks, much as is done with traditional cryptocurrencies. These mintettes differ from traditional cryptocurrency miners, however, in a crucial way: rather than performing some computationally difficult task, each mintette is simply authorized by the central bank to collect transactions. In UVCoin, this authorization is accomplished by a PKI-type functionality, meaning the central bank signs the public key of the mintette, and each lower-level block must contain one of these signatures in order to be considered valid. The time interval in which blocks are produced by mintettes are referred as an epoch, where the length of an epoch varies depending on the mintette.

Because these blocks are not ultimately incorporated into the main blockchain and are referred as lower-level blocks. Mintettes are collectively responsible for producing a consistent ledger, and thus to facilitate this process they communicate internally throughout the course of an epoch and ultimately reference not only their own previous blocks but also the previous blocks of each other. This means that these lower-level blocks form a cross-referenced chain. At the end of some longer pre-defined time interval called a period, the mintettes present their blocks to the central bank, which merges these lower-level blocks to form a consistent history in the form of a new block. This higher-level block is what is ultimately incorporated into the main blockchain, meaning a user of UVCoin need only keep track of higher-level blocks. (Special users wishing to audit the behavior of the mintettes and the central bank, however, may keep track of lower-level blocks.)

Interaction with UVCoin can thus be quite similar to interaction with existing cryptocurrencies, as the structure of its blockchain is nearly identical, and users can create new pseudonyms and transactions in the same way as before. In fact, we stress that

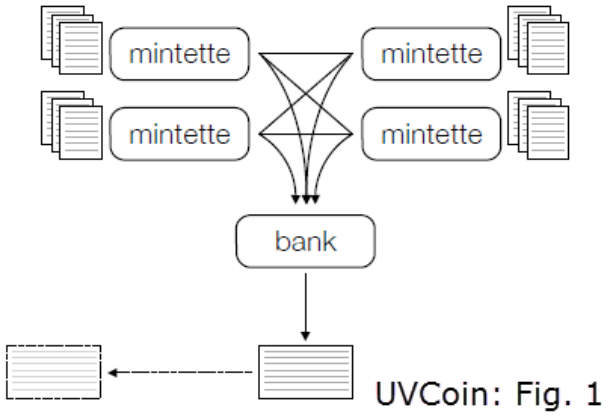

Fig. 1: The overall structure of UVCoin. Each mintettes maintains a set of lower-level blocks, and (possibly) communicates with other mintettes (either directly or indirectly). At some point, the mintettes send these blocks to the central bank, which produces a higher-level block. It is these higher-level blocks that form a chain and that are visible to external users.

UVCoin is intended as a framework rather than a stand-alone cryptocurrency, so one could imagine incorporated techniques from various existing cryptocurrencies to achieve various goals.

3. ACHIEVING CONSENSUS

In the previous, it has been described how mintettes send so-called “lower-level blocks” to the central bank at the end of a period. In this, a consensus protocol will be described by which these blocks can already be made consistent when they are sent to the central bank, thus ensuring that the overall system remains scalable by allowing the central bank to do the minimal work necessary.

As described in the introduction, one of the major benefits of centralization is that, although the generation of the transaction ledger is still distributed, consensus on valid transactions can be reached in a way that avoids the wasteful proofs-of-work required by existing cryptocurrencies. In traditional cryptocurrencies, the set of miners is neither known nor trusted, meaning one has no choice but to broadcast a transaction to the entire network and rely on proof-of-work to defend against Sybil attacks. Since our mintettes are in fact authorized by the central bank, and thus both known and — because of their accountability — trusted to some extent, we can avoid the heavyweight consensus requirement of more fully decentralized cryptocurrencies and instead use an adapted version of Two-Phase Commit (2PC), as presented in Figure 2.

Fig. 2

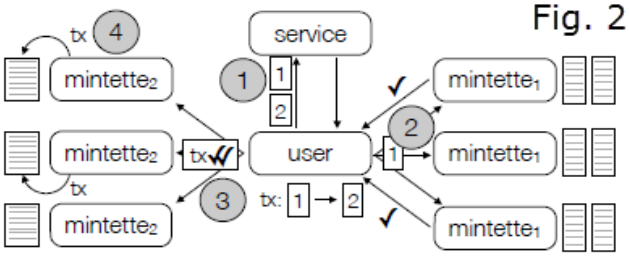

Fig. 2: The proposed protocol for validating transactions; each mintette mi is an owner of address i. In (1), a user learns the owners of each of the addresses in its transaction. In (2), the user collects approval from a majority of the owners of the input addresses. In (3), the user sends the transaction and these approvals to the owners of the transaction identifier. In (4), some subset of these mintettes add the transaction to their blocks.

A generic consensus protocol, ensuring total ordering of transactions, is not necessary for double-spending prevention; instead, a weaker property — namely that any transaction output features as a transaction input in at most one other transaction — is sufficient. UVCoin builds its consensus protocol for double-spending prevention based on this insight.

Below is described a threat model for the consensus protocol before going on to present a basic protocol that achieves consensus on transactions, an augmented protocol that allows for auditability of both the mintettes and the central bank, and a performance evaluation.

4. THREAT MODEL AND SECURITY

It is always assumed that the central bank is honest and that the underlying cryptography is secure; i.e., no parties may violate the standard properties offered by the hash function and digital signature. Honest mintettes follow the protocols as specified, whereas dishonest mintettes may behave arbitrarily; i.e., they may deviate from the prescribed protocols, and selectively or broadly ignore requests from users. Finally, honest users create only valid transactions (i.e., ones in which they own the input addresses and have not yet spent their contents), whereas dishonest users may try to double-spend or otherwise subvert the integrity of UVCoin.

Two threat models need to be proposed.

The first threat model assumes that each transaction is processed by a set of mintettes with an honest majority; this is different from assuming that a majority of all mintettes are honest.

The second threat model assumes that no mintette is honest, and that mintettes may further violate the integrity of UVCoin. This is a very hostile setting, but show that some security properties still hold for honest users. Additionally, it shows that mintettes that misbehave in certain ways can be detected and ultimately held accountable, which may serve as an incentive to follow the protocols correctly.

In the face of these different adversarial settings, below are some of the following key integrity properties:

- No double-spending: Each output address of a valid transaction will only ever be associated with the input of at most one other valid transaction.
- Non-repudiable sealing: The confirmation that a user receives from a mintette — which promises that a transaction will be included in the ledger — can be used to implicate that mintette if the transaction does not appear in the next block.
- Timed personal audits: A user can, given access to the lower-level blocks produced within a period, ensure that the implied behavior of a mintette matches the behavior observed at the time of any previous interactions with that mintette.
- Universal audits: Anyone with access to the lower-level blocks produced within a period can audit all transactions processed by all mintettes. In particular, mintettes cannot retroactively modify, omit, or insert transactions in the ledger.
- Exposed inactivity: Anyone with access to the lower-level blocks produced within a period can observe any mintette's substantial absence from participation in the 2PC protocol. (In particular, then, a mintette cannot retroactively act to claim transaction fees for services not provided in a timely manner.)

To see how to satisfy these security properties. We prove that at least some subset of these security properties can be captured in both our threat models, and that exposure may disincentive mintettes from violating those that we cannot capture directly.

5. THE UVCOIN SYSTEM

With the consensus protocol in place, the structure of UVCoin, focusing on the interaction between the mintettes and the central bank, and on the overall parameters and properties of the system. Firstly, the structure and usage of UVCoin and then address considerations that arise in how to allocate fees to mintettes overlay UVCoin on top of an existing cryptocurrency like Bitcoin incentivize mintettes to follow the consensus protocol and present a collectively consistent ledger to the central bank and set concrete choices for various system parameters. A lower-level-block produced by a mintette m within period i looks like b = (h,txset,σ,mset), where h is a hash, txset is a collection of transactions, and σ is a signature from the mintette that produced this block. The fourth component mset specifies the cross-chain property of lower-level blocks by identifying the hashes of the other previous blocks that are being referenced. Denote by pkbank the bank's public key and by DPKi the set of mintettes authorized by the bank in the previous higher-level block.

To form a lower-level block, a mintette uses the transaction set txset it has formed throughout the epoch and the hashes (h1, ... ,hn) that it has received from other mintettes and creates mset ← (h1,...,hn), other blocks ← h1k...khn.

5.1 Higher-level blocks

The higher-level block that marks the end of period i looks like B(i) bank = (h,txset,σ,DPKi+1), where these first three values are similar to their counterparts in lower-level blocks (i.e., a hash, a collection of transactions, and a signature), and the set DPKi+1 contains pairs (pkm,σ(m) bank); i.e., the public keys of the mintettes authorized for period i+1 and the bank's signatures on the keys.

To form a higher-level block, the bank must collate the inputs it is given by the mintettes. To create a consistent

transaction set txset, a vigilant bank might need to look through all the transaction sets it receives to detect double-spending, remove any conflicting transactions, and identify the mintette(s) responsible for including them. As this would require the bank to perform work proportional to the number of transactions (and thus somewhat obviate the reason for mintettes), we also consider an optimistic approach in which the bank relies on the consensus protocol and instead simply merges the individual transaction sets to form txset. The bank creates the set of authorized mintettes using a decision process.

1.) Coin generation and fee allocation: In addition to this basic structure, each higher-level block could also contain within txset a special coin generation transaction and an allocation of fees to the mintettes that earned them in the previous period. Semantically, the coin generation could take on the same structure as in Bitcoin; i.e., it could be a transaction tx($\emptyset$ n $\longrightarrow$ addrbank), where addrbank is an address owned by the bank, and fees could be allocated using a transaction tx(addrbank f $\longrightarrow$ addrm), where f represents the fees owed to them. The interesting question is thus not how central banks can allocate fees to mintettes, but how it decides which mintettes have earned these fees. In fact, the provided action logs allow the central bank to identify active and live mintettes and allocate fees to them appropriately.

This mechanism (roughly) works as follows. The central bank keeps a tally of the mintettes that were involved in certifying the validity of input addresses; i.e., those that replied in the first phase of the consensus protocol. The choice to reward input mintettes is deliberate: in addition to providing a direct incentive for mintettes to respond in the first phase of the protocol, it also provides an indirect incentive for mintettes to respond in the second phase, as only a transaction output that is marked as unspent can later be used as an input (for which the mintette can then earn fees). Thus, rewarding input mintettes provides incentive to handle a transaction throughout its lifetime.

The action logs also play a crucial role in fee allocation. The "exposed inactivity" security property from prevents an inactive mintette from becoming active at a later time and claiming that it contributed to previous transactions, as an examination of the action logs can falsify such claims. Additionally, if fee allocation is determined based on a known function of the action logs, anyone with access to the action logs can audit the actions of the central bank.

Finally, it is mentioned that although the logs are sent only to the central bank, the expectation is that the central bank will publish these logs to allow anyone to audit the system. As we assume the central bank is honest, this does not present a problem, but in a stronger threat model in which less trust were placed in the central bank, one might instead attempt to adopt a broadcast system for distributing logs (with the caveat that this approach introduces significantly higher latency). In such a setting, anyone with access to the logs could verify not only the actions of the mintettes, but could also replay these actions to compare the ledger agreed upon by the mintettes and the ledger published by the bank; this would allow an auditor to ensure that the bank was not engaging in misbehavior by, e.g., dropping transactions.

2.) A simplified block structure: The above description of higher-level blocks (and the previous description of lower-level blocks) contains a number of additional values that do not exist in the blocks of existing cryptocurrencies, making UVCoin somewhat incompatible with their semantics.

5.2 Incentivizing mintettes

One might naturally imagine that this structure, as currently described, places the significant burden on the central bank of having to merge the distinct blocks from each mintette into a consistent history. By providing appropriate incentives, however, we can create an environment in which the presented ledger is in fact consistent before the bank even sees it. If mintettes deviate from the expected behavior then, they can be held accountable and punished accordingly (e.g., not chosen for future periods or not given any fees they have earned).

One direct incentive for mintettes to collect transactions, which is fees. As described, mintettes are rewarded only for active participation, so that an authorized mintette needs to engage with the system in order to earn fees. Another direct incentive, which is the authorization of mintettes by the central bank. For semantic purposes, the value "X" used to authorize each mintette for the next period could be arbitrarily small. As an incentive, however, this value could be larger to directly compensate the mintettes for their services.

Finally, we expect that the central bank could be a national or international entity that has existing relationships with, e.g., commercial banks. There thus already exist strong business incentives and regulatory frameworks for such entities to act as honest mintettes.

5.3 Setting system parameters

As described, the system is parameterized by a number of variables, such as the length of epochs, the length of a period, and the number of mintettes. The length of an epoch for an individual mintette is entirely dependent on the rate at which it processes transactions. Mintettes that process more transactions will therefore have shorter epochs than ones that do so less frequently. There is no limit on how short an epoch can be, and the only upper limit is that an epoch cannot last longer than a period.

It might seem desirable for periods to be as short as possible, as ultimately a transaction is sealed into the official ledger only at the end of a period. To ease the burden on the bank, however, it is also desirable to have longer periods, so that central banks must intervene as infrequently as possible (so that central banks can potentially perform certain optimizations to reduce transaction bloat). The methods by which mintettes could "promise" (in an accountable way) to users that their transactions would be included, so that in practice near-instantaneous settlement can be achieved even with longer periods, so long as one trusts the mintette. Nevertheless, we do not expect periods to last longer than a day.

For the purposes of having a fair and competitive settlement process, it is desirable to have as many mintettes as possible; this is also desirable from a performance perspective, as the performance of the UVCoin system (measured in the rate of transactions processed) scales linearly with the number of mintettes. Adding more mintettes, however, also has the effect that they earn less in transaction fees, so these opposing concerns must be taken into account when settling on a

concrete number (to give a very rough idea, one number that has been suggested is 200).

6. OPTIMIZATION

What is presented here is a (relatively) minimal version of UVCoin, which allows to achieve the basic integrity and scalability properties that are crucial for any currency designed to be used on a global level. Here, we briefly sketch some extensions that could be adopted to strengthen either of these properties, and leave a more detailed analysis of these or other solutions as interesting future research.

A. Pruning intermediate transactions at the end of a period, the central bank publishes a higher level block containing the collection of transactions that have taken place in that time interval; it is only at this point that transactions are officially recorded in the ledger. Because mintettes provide evidence on a shorter time scale that a user's transaction is valid and will be included in the ledger, however, users might feel more comfortable moving currency multiple times within a period than in traditional cryptocurrencies (in which one must wait for one or several blocks to avoid possible double-spending). It therefore might be the case that at the end of a period, the central bank sees not just individual transactions, but potentially multiple "hops" or even whole "chains" of transactions. To limit transaction bloat, the bank could thus prune these intermediate transactions at the end of the period, so that ultimately only the start and end points of the transaction appear in the ledger, in a new transaction signed by the central bank. On its surface, this idea may seem to require a significant amount of trust in the central bank, as it could now actively modify the transaction history. The action logs, however, would reveal the changes that the bank had made and allow users to audit its behavior, but nevertheless the alterations that could be made would need be significantly restricted.

B. Further incentives for honest behavior in addition to the existing incentives for honest behavior outlined in, mintettes could adopt a sort of proof-of-stake mechanism, in which they escrow some units of currency with the central bank and are allowed to collate only a set of transactions whose collective value does not exceed the escrowed value. If any issue then arises with the transactions produced by the mintette (e.g., it has accepted double-spending transactions), the central bank can seize the escrowed value and remove the double-spending transactions, so the mintette ultimately pays for this misbehavior out of its own pocket (and maybe even pays additional fines). This mechanism as described is not fully robust (as in particular the mintette might accept many expenditures of the same unit of currency, not just two), but it does have an interesting effect on the length of periods. The length of earlier periods will necessarily be quite small, as mintettes will not have much capital to post. As mintettes accumulates to rest of currency, however, period scan grow longer. This is a natural process, as it also allows for a trial period in the beginning to ensure that authorized mintettes don't misbehave, and then for a more stable system as a set of trustworthy mintettes emerges.

6.1 Multiple banks and foreign exchange

In a global setting, one might imagine that each central bank could develop their own version of UVCoin; this would lead, however, to a landscape much the same as today's Bitcoin and the many Altcoins it has inspired, in which multiple implementations of a largely overlapping structure lead to an infrastructure fragmentation: bugs are replicated across codebases and compatibility across different Altcoins is artificially low. An attractive approach is for different central banks to instead use the same platform, to prevent this fragmentation and to allow users to seamlessly store value in many different currencies. While this allows, the currencies generated by different central banks to achieve some notion of interoperability, it is expected that different blockchains will be kept separate; i.e., a particular central bank does not—and should not—have to keep track of all transactions that are denominated in the currency of another central bank. (Mintettes, however, may choose to validate transactions for any number of central banks, depending on their business interests.) While every central bank does not necessarily need to be aware of transactions denominated in the currency of another central bank, this awareness may at times be desirable. For example, if a user would like to exchange some units of one currency into another belonging to a central bank that is relatively known to and trusted by the first (e.g., exchange GBP for USD), then this should be a relatively easy process. The traditional approach is to simply go to a third-party service that holds units of both currencies, and then perform one transaction to send units of the first currency to the service, which will show up in the ledger of the first currency, and another transaction to receive units of the second currency, which will show up in the ledger of the second currency.

Although this is the approach by far most commonly adopted in practice (both in fiat currency and cryptocurrency markets), it has a number of limitations, first and foremost of which is that it is completely opaque: even an outside observer who is able to observe both ledgers see two transactions that are not linked in any obvious way. One might naturally wonder, then, if a more transparent mechanism is possible, in which the currency exchange shows up as such in the ledger. Briefly, to achieve this fair exchange, the adoption of a protocol to achieve atomic cross-chain trading, which provides a Bitcoin compatible way for two users to fairly exchange units of one currency for some appropriate units of another currency; i.e., to exchange currency in a way that guarantees that either the exchange is successful or both users end up with nothing (so in particular it cannot be the case that one user reclaims currency and the other does not). If one is less concerned about compatibility with Bitcoin, then a slightly simpler approach such as "pay on reveal secret" [12] could be adopted. To fit the setting, in which central banks may want to maintain some control over which other currencies their currency is traded into and out of (and in what volume), the existing protocol has been modified to require a third party to sign both transactions only if they are denominated in currencies that are viewed as "exchangeable" by that party. This serves to not only signal the third party's blessing of the exchange, but also to bind the two transactions together across their respective blockchains.

The proposal of this protocol thus enables transparent exchanges that can be approved by a third party, but does not (and cannot) prevent exchanges from taking place without this approval. Importantly, however, an auditor can now—with access to both blockchains—observe the exchange.

7. CONCLUSIONS AND FUTURE WORK

In this paper, the first cryptocurrency framework UVCoin has been presented that provides the control over monetary policy that entities such as central banks expect to retain. By constructing a blockchain-based approach that makes relatively minimal alterations to the design of successful cryptocurrencies such as Bitcoin, it has been demonstrated that this centralization can be achieved while still maintaining the transparency guarantees that have made (fully) decentralized cryptocurrencies so attractive. A new consensus mechanism have been proposed based on 2PC and measured its performance, illustrating that centralization of some authority allows for a more scalable system to prevent double spending that completely avoids the wasteful hashing required in proof-of-work-based systems.

To ensure privacy for transactions, one could adapt existing cryptographic techniques such as those employed by Zerocoin [13], Zerocash [14], Pinocchio Coin [15], or Groth and Kohlweiss [16]. The goals of these cryptocurrencies are somewhat orthogonal to the goals of this paper so it will be interesting to see if privacy-enhancing and other techniques could be combined with UVCoin and that will be a great avenue for future work.

8. REFERENCES

[1] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008, bitcoin.org/bitcoin.pdf.

[2] D. O'Leary, V. D'Agostino, S. R. Re, J. Burney, andHoffman, "Method and system for processing Internet payments using the electronic funds transfer network," Nov. 2013. [Online]. Available: www.google.com/patents/US20130317984

[3] Nasdaq, "Nasdaq launches enterprise-wide blockchain technology initiative," May 2015, www.nasdaq.com/press-release/nasdaq-launches-enterprisewide-blockchain-technology-initiative-20150511-00485.

[4] European Central Bank, "Virtual currency schemes - a further analysis," Feb. 2015, www.ecb.europa.eu/pub/pdf/other/virtualcurrencyschemesen.pdf.

[5] B. Bernanke, Nov. 2013, qz.com/148399/ben-bernankebitcoin-may-hold-long-term-promise/.

[6] HM Treasury, "Digital currencies: response to the call for information," Mar. 2015, www.gov.uk/government/uploads/system/uploads/ attachment data/file/414040/digital currencies response to call for information final changes.pdf.

[7] A. Narayanan, J. Bonneau, E. Felten, A. Miller, and S. Goldfeder, Bitcoin and cryptocurrency technologies. [Online].Available:piazza.com/princeton/spring2015/btctech/resources

[8] A. Miller, A. Juels, E. Shi, B. Parno, and J. Katz, "Permacoin: Repurposing Bitcoin work for data preservation," in Proceedings of the IEEE Symposium on Security and Privacy, 2014.

[9] L. Luu, J. Teutsch, R. Kulkarni, and P. Saxena, "Demystifying incentives in the consensus computer," in proceedings of ACM CCS 2015, 2015, to appear.

[10] Bank of England, Private communication, 2015.

[11] B. Laurie, "An efficient distributed currency," 2011,www.links.org/files/distributedcurrency.pdf.

[12] T. Young, "Atomic cross-chain exchange," 2014,upcoder.com/11/atomic-cross-chain-exchange/.

[13] I. Miers, C. Garman, M. Green, and A. D. Rubin, "Zerocoin: Anonymous distributed e-cash from bitcoin," in 2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013. IEEE Computer Society, 2013, pp. 397–411. [Online].Available:http://dx.doi.org/10.1109/SP.2013.34

[14] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," in 2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014. IEEE Computer Society, 2014, pp. 459–474. [Online]. Available: http://dx.doi.org/10.1109/SP.2014.36

[15] G. Danezis, C. Fournet, M. Kohlweiss, and B. Parno, "Pinocchio coin: building zerocoin from a succinct pairing-based proof system," in PETShop'13, Proceedings of the 2013 ACM Workshop on Language Support for Privacy-Enhancing Technologies, Co-located with CCS 2013, November 4, 2013, Berlin, Germany, Franz, A. Holzer, R. Majumdar, B. Parno, and Veith, Eds. ACM, 2013, pp. 27–30. [Online]. Available: http://doi.acm.org/10.1145/2517872.2517878

[16] J. Groth and M. Kohlweiss, "One-out-of-many proofs: Or how to leak a secret and spend a coin," in Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II, ser. Lecture Notes in Computer Science, E. Oswald and M. Fischlin, Eds., vol. 9057. Springer, 2015, pp. 253–280. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-46803-6 9