Cross-Blockchain Technologies: Review, State of the Art, and Outlook

Michael Borkowski*, Philipp Frauenthaler*, Marten Sigwart*, Taneli Hukkinen‡, Oskar Hladký‡, Stefan Schulte*

* Distributed Systems Group
TU Wien, Vienna, Austria
{m.borkowski, p.frauenthaler, m.sigwart,
s.schulte}@infosys.tuwien.ac.at

‡ Pantos GmbH
Vienna, Austria
contact@pantos.io

***Abstract*—Interoperability between blockchains remains an open problem, with current approaches providing very limited means of cross-blockchain interaction, mostly in the form of atomic swaps. However, very little means of cross-blockchain data exchange, including cross-blockchain token transfers, are found in literature. To address this issue, within the TAST research project, we aim to create a platform for enabling cross-blockchain interoperability, in order to counter the fragmentation of the research and development field of blockchains.**

In this paper, we review the current state of the art in the field of cross-blockchain technologies, including our own work within the TAST research project, answer former open questions from today's perspective, and give an outlook of current challenges and possible future work.

I. Introduction

Since the presentation of Bitcoin [13], the first implementation of a blockchain protocol in widespread use, the utility and feasibility of decentralized ledgers has been demonstrated for various use cases and fields [7]. As discussed in our previous work [3, 1, 2], research activities related to blockchains cover, among others, the addition of new layers to Bitcoin itself [21], improvements to the Bitcoin codebase [11], and the development of entirely new blockchains [22]. The diversity and richness of this research field comes with an increasing number of technologies and implementations [9], causing structural problems within the blockchain community. The vast amount of blockchains and other projects in existence causes severe fragmentation of the research and development field. Interoperability is mostly not foreseen, with blockchains instead competing for users and developers [2].

Therefore, in the Token Atomic Swap Technology (TAST) research project[1], we aim to create a platform for cross-blockchain interoperability. The overarching goal is to connect the fragmented field of research and development by investigating possible means of interconnecting various blockchain-related projects. For instance, this can be done by developing currencies and tokens usable on more than just one blockchain (cross-blockchain tokens), investigating the transfer of data across blockchains (cross-blockchain data storage), or by enabling more complex interactions such as calling smart contract functions from different blockchains (cross-blockchain smart contract invocations).

Earlier in the TAST project, we reviewed the state of the art in the blockchain research field with a focus on cross-blockchain technologies [3]. During the course of the project, new research questions arose and the very fast-paced blockchain community came up with a number of novel solutions. Therefore, in the work at hand, we revisit our original review, investigate the progress made by the individual projects and the research field in general, including our own contributions, state the current open problems, and give an outlook on future research directions.

To this end, Section II discusses the progress and contributions throughout the TAST research project, and Section III discusses the current state of the art in the field of cross-blockchain technologies. Section IV summarizes open problems in this field and provides an outlook on open challenges and future work. Finally, Section V concludes the paper.

II. Progress Within TAST

The TAST research project has resulted in numerous contributions throughout its progress to date. Including the work at hand, four white papers have been published. In addition, a research prototype[2] has been developed, demonstrating the use of the concepts discussed in the publications using Solidity. Two research papers are currently under preparation for submission to according research conferences.

In the first TAST white paper [3], we lay the groundwork for cross-blockchain technologies by investigating fundamentals about cross-blockchain token transfers. We provide an overview of the stages and goals of the TAST research project, identify challenges, and discuss possible implementation strategies. In addition, we provide an extensive review of the state of the art in blockchain technologies, with a focus on cross-blockchain aspects, at the start of the TAST project. For this, we survey twenty of the most relevant blockchains at that time, discuss some of their technical aspects and properties, e.g., their consensus protocol, whether or not they support user-issued assets (UIAs), and the extent of smart contracts

White Paper, TU Wien; March 2019, version 1.0.

[1] http://www.infosys.tuwien.ac.at/tast/

[2] https://github.com/pantos-io/dextt-prototype

supported. In addition to these twenty blockchains, we discuss fourteen operational and forthcoming cryptocurrency systems in the area of cross-blockchain technologies.

The second white paper [1] provides additional fundamental research. We formally define key terms used throughout blockchain literature, provide a formal model of transaction consensus and arbitration consensus, and from this, define what constitutes the *main chain* in a given blockchain context. Subsequently, we show the XPP by deriving the *lemma of rooted blockchains*, which implies that it is not possible on a blockchain A to verify the existence of a certain data block (transaction, event, contract call, etc.) on a blockchain B with practical effort. Specifically, verifying this existence would entail access by A to the lineage of the block of B in which said data is located. In addition, A would require a sufficiently powerful transaction consensus (e.g., smart contract instruction set) to mimic the transaction consensus of B. In practice, the latter requirement is feasible. However, the former requirement, without loss of generality, requires all data of A constituting the lineage of the data block to be accessible from B, which is not realizable with practical means.

Furthermore, in [1], we present a novel approach to cross-blockchain asset transfers, called *claim-first* transfers. Traditionally, *spend-first* transfers ensure that the transferred assets (tokens or native currencies) are marked as spent first. Only then can the asset be claimed by the receiver. In contrast, the proposed claim-first transfers allow the reversal of this order. We define a special CLAIM transaction, which can only be used if signatures from both the sender and the receiver are provided. These signatures verify that the sender indeed intends to transfer assets (in our case, tokens) to the receiver. Therefore, we call the information provided in the CLAIM transaction *Proof of Intent* (PoI). The CLAIM transaction is used on the destination blockchain by the receiver, and provides the receiver with the transferred tokens (temporarily allowing double ownership of the tokens by both the sender on the source blockchain, and the receiver on the destination blockchain). We then define another transaction, called DESTROY. The DESTROY transaction must be used in conjunction with a valid PoI on the source blockchain, and can be submitted by anyone (since the PoI is made public in the CLAIM transaction). The effect of the DESTROY transaction is that the sender loses the transferred tokens (as intended). This, together with the CLAIM transaction used by the receiver—potentially before the DESTROY transaction is used—concludes the transfer.

The concept of claim-first transfers presented in [1] works by rewarding witnesses for actually submitting the DESTROY transaction, thus ensuring eventual consistency. These witness rewards themselves pose an additional challenge. Paying these rewards can be done using native assets (e.g., Ether for Ethereum), using the transferred tokens themselves, or using an additional, auxiliary reward asset type. We discuss benefits and drawbacks for all three of these types in [1].

In the third white paper [2], we address the challenge of witness rewards by proposing the novel concept of *deterministic witnesses*. Previously, we proposed a witness contest on a *first-come, first-serve* basis, assigning a reward to the first witness submitting a DESTROY transaction. However, the question of which address receives the reward cannot be answered reliably, which—according to the aforementioned XPP—poses the problem of transferring this information (reliably) to other blockchains. In contrast, the concept of deterministic witnesses introduces a contest which leverages on the determinism prevalent in blockchain technologies. In this contest, not the timing of a witness decides on the reward decision, but a hash value generated from the witness address together with the PoI. This means that the contest winner is determined from the PoI data together with a pool of potential witnesses, which results in the outcome (the answer to the question of who receives the witness reward) to be deterministic, and therefore identical across all blockchains.

In addition, our research prototype serves as a reference implementation of the concepts discussed in our publications. The prototype uses claim-first transactions as well as deterministic witnesses, and is implemented in Solidity. The reference implementation has been used to verify the functionality of the concepts presented in the TAST publications (i.e., claim-first transfers and deterministic witnesses). Additionally, an extensive evaluation was performed using this prototype, the results of which are currently under peer review. The evaluation was performed in a multi-blockchain ecosystem testbed using `geth` in Proof-of-Authority (PoA) mode.

III. State of the Art Revisited

In previous work [3], we have provided a review of the state of the art in blockchains, with a focus on cross-blockchain technologies. In this section, we revisit this review with reagrds to blockchains and technologies, and revisit the previously stated open questions.

A. Blockchains and Technologies

First, we review the currently most relevant blockchains and asset types. In our previous review, we chose a market capitalization of two billion dollars as the threshold to determine relevance. Due to the fact that overall market capitalization in cryptocurrencies has decreased in the meantime, we now use a minimum of one billion, since otherwise, only Bitcoin, Ripple, and Ethereum would satisfy the threshold. The market capitalization values of 2019-02-12 are used in this work. While the five most relevant blockchains have remained in our list, namely, Bitcoin, Ripple, Ethereum, Bitcoin Cash, and Litecoin, three new assets are in our selection now: EOS [6], Tether [16], and Tron [17].

EOS.IO, operated by block.one, is a blockchain protocol, providing the native cryptocurrency EOS. It offers a smart contract platform for decentralized applications and is intended for the deployment of large-scale applications using an infrastructure with virtualized hardware like CPUs, RAM, and storage. As we have discussed in other previous work [2], certain capabilities of blockchains (and more specifically, features like smart contracts) are required for creating cross-blockchain asset transfers. As EOS.IO provides a feature-rich smart contract platform and toolset, it is a promising candidate for cross-blockchain asset transfers, as well as other kinds of interoperability. EOS.IO uses Delegated Proof of Stake (DPoS) as its consensus algorithm.

Tether, operated by Tether Limited, is a token on the Bitcoin blockchain using the OmniLayer [21] protocol. Tether Limited guarantees that each token of their USDT cryptographic asset is backed by one US dollar and that this backing assets can be claimed using USDT tokens. In light of cross-blockchain technologies, USDT, using OmniLayer, represents what we have described in our previous review [3] as *piggybacking* on top of the Bitcoin blockchain: OmniLayer transactions are regular Bitcoin transactions, on top of which additional transaction data is transported. Nodes unaware of the OmniLayer transaction details simply ignore this additional data. This technique is sometimes also called colored coins [15].

Tron, like EOS.IO, uses the DPoS consensus algorithm. Furthermore, Tron also provides smart contracts. Tron was cloned from Ethereum, and the Tron Virtual Machine (TVM) is based on the Ethereum Virtual Machine (EVM) with only minor changes introduced [17]. Therefore, there are no fundamental differences in applicability of cross-blockchain technologies between Ethereum and Tron.

With regards to technologies, in our previous review, we identified Metronome [12] as the project closest to our goal. Since the time of writing, to the best of our knowledge, no further technical details were published about how the Metronome project realizes cross-blockchain token transfers in detail. While the authors name a *Proof of Exit*, which can be used to claim tokens on the destination blockchain, no further technical details about this process are discussed, and it remains unclear how Metronome tackles challenges like the cross-blockchain proof problem (XPP) [1].

B. Questions and Challenges

In our first review [3], we identified certain questions to describe the general goals for TAST. In the following, we review these questions from today's perspective.

How are the tokens issued on the blockchains? Is a fixed pool of issued tokens used, or are they re-issued on a regular basis?

In the approaches we developed, cross-blockchain tokens do not need to have a specific issuing (minting) scheme. Our prototype uses fixed supply (predetermined by one specific minting account), however, none of the functionalities we introduce is tied to the minting process. Therefore, in principle, any minting scheme currently used by tokens (fixed or variable supply) can be used with cross-blockchain tokens [2].

How are tokens disabled as they are leaving the blockchain? Are tokens destroyed, locked, or stored in a wallet or contract?

Currently, we destroy tokens on the source blockchain, and the balance of the source wallet is reduced. In order to ensure eventual consistency across blockchains despite the XPP [1], we employ a concept called claim-first transfers as described in our previous work [2].

Are tokens re-balanced across blockchains to maintain liquidity, and if so, how often and by which entity?

This question remains mostly open, albeit in a more economic than technical context. In fact, when this question was formulated, the idea of claim-first transfers had not yet been developed, and a (semi-)centralized balancing entity seemed to be a viable option. Now, entirely decentralized token transfers are possible, and therefore, re-balancing is no longer required. The presence of tokens on a blockchain is simply determined by the market, i.e., by the users' desire to hold tokens on that blockchain.

Which blockchains are suitable for cross-blockchain token transfers?

This question has been answered by both our previous research, identifying Ethereum and Ethereum Classic as suitable blockchains, and in Section III-A of the paper at hand, where we discuss blockchains which emerged since our last review, and name EOS.IO and Tron. These new blockchains constitute candidates for further work. Note that this list is not exhaustive and reflects the current technological state of blockchains, which can change rapidly.

Which features (e.g., native user-issued assets, smart contracts, Turing-completeness) are required from a blockchain to support token transfers as proposed by TAST?

Currently, smart contract functionality is required to a certain degree in order to realize cross-blockchain transfers. Our current approach uses Solidity. This is not a conceptual requirement, i.e., in general, other languages could be applied as well, However, care must be taken when examining the suitability of different smart contract platforms. For instance, our current approach makes extensive use of the signature creation and verification features of Solidity, and such features must be supported by other smart contract platforms to facilitate cross-blockchain transfers as currently defined in TAST.

Can cross-chain transfers be realized despite lack of Turing-complete smart contracts?

Strictly speaking, the smart contract platform provided by a blockchain does not need to be Turing-complete in order to be used with the proposed cross-blockchain transfers. For instance, no loops are required in our approach. Instead, only a certain subset of operations are required. Nevertheless, some of these operations are relatively complex (e.g., creation and verification of signatures), so simple languages like Bitcoin Script are not supported.

We discuss in Section IV how a concept of cross-blockchain token transfers without Turing-complete smart contracts could be approached.

IV. Outlook and Open Questions

While significant progress has already been made within the contributions presented in Section II, there are still numerous open questions and challenges in the domain of cross-blockchain technologies.

A. Requirement of Smart Contracts

As stated in Section III-B, we currently require smart contract support with substantial functionality, not covered by simple contract languages like Bitcoin Script. Further analysis is necessary on the concrete requirements with respect to smart contract functionality posed to blockchains in order to be able to process the type of cross-blockchain transactions proposed in our work.

Our prototype is written in Solidity, a Turing-complete language. Therefore, due to the Church-Turing thesis, all Turing-complete languages can be used to realize the approach used in the prototype [18]. However, it is an open question which subset of functionality can be formulated as a minimal required set of features. Minimizing this set will extend applicability of the approach proposed within TAST.

B. Rewards and Incentive Analysis

We currently reward nodes performing the token destruction using the DESTROY transaction (once the PoI has been published in the CLAIM transaction) with a witness reward. This reward is comparable to the mining reward in Proof-of-Work (PoW) blockchains like Bitcoin or Ethereum. However, as alternative models of ensuring economic incentive for nodes in blockchains such as Proof of Stake (PoS) are emerging, we also deem alternative reward models for cross-blockchain transactions possible. Based on popular peer-to-peer networks for file sharing, we claim that reward-less systems are indeed possible, if mutual interest (participating in a network) exists.

C. Cross-Blockchain Smart Contract Execution

In addition to transferring tokens across blockchains, we envision the execution of smart contract calls from one blockchain to another. Currently, the authors of a smart contract develop, test, and deploy their contracts for one specific blockchain. Offering contract functionality on other blockchains usually requires severe effort, as there is no current standard for porting smart contract code comparable to the POSIX standard [20] used to ensure portability of a program across operating systems.

We currently envision two ways of realizing such cross-blockchain contract executions. First, similarly to Remote Procedure Calls (RPCs), a stub contract can be used on the calling side, which then uses novel cross-blockchain technologies to forward the call to the target blockchain. On the target side, a surrogate contract imitates the call. If necessary, upon finishing, the result is then transmitted back to the calling blockchain. The same principle can be used to transfer events across blockchains. In this scenario, a cross-blockchain publish/subscribe (pub/sub) pattern is envisioned. As there already exist concepts for decentralized pub/sub messaging [4], adapting these technologies for cross-blockchain communication is a promising research direction.

Another conceptual approach is to create a cross-blockchain virtual machine. Instead of adapting to the features and specifics of existing smart contract capabilities, we propose to create a virtual machine, similar to existing virtual machines like the EVM or the EOS Virtual Machine. However, in contrast to these blockchain-bound virtual machines, we propose to consider an ecosystem of multiple blockchains as an execution environment for smart contracts, creating a virtual machine spanning across multiple blockchains. This is comparable to how the EVM is a virtual machine spanning across individual nodes. Special care must be taken with regards to scalability, as this is already a main challenge of existing blockchains [8]. We refer to existing literature for a study on how the scalability of blockchains can be improved [19].

The former approach allows for an easier interoperability with blockchain-native smart contracts. However, cross-blockchain communication has the potential to require more complexity in the design. The latter approach has an increased complexity in interacting with native smart contracts, but promises to result in a more unified, homogeneous ecosystem.

D. Cross-Blockchain Data Exchange

As described before, our current contributions entail a token which can be transferred across blockchains. Abstracting from this idea, a cross-blockchain data exchange can be developed. Instead of merely storing and transferring data about token balances, we envision the storage of arbitrary data, including state and memory data of smart contracts. This can be used to store various kinds of data beyond token balances, including provenance information [10], business process runtime information [14], or IoT data [5].

Again, we propose two fundamentally different approaches. First, existing blockchain-native data (such as Ethereum's storage) can be transferred across blockchains on an on-demand basis. Second, similar to the cross-blockchain virtual machine proposed in Section IV-C, we propose a storage inherently shared between blockchains, i.e., a cross-blockchain data store.

In this field of future work, it is crucial to analyze the implications of such cross-blockchain data exchange. Currently, data stored on one blockchain does not suffer from concurrency problems (not taking into account short-term forks in the blockchain itself). Introducing cross-blockchain data exchange adds to blockchains what multi-threading and preemptive process scheduling adds to processors. Concurrency, locks, and synchronization mechanisms will be required to ensure data integrity.

V. Conclusion

In this paper, we have presented an outline of the work done within the TAST research project. We have revisited our previous review regarding the state of the art in cross-blockchain token transfers, providing a review of the current state of the art in comparison and discussed current open questions and an outlook on future work.

Disclaimer

Information provided in this paper is the result of research, partly based on publicly available resources of varying quality. Popular use of cryptocurrencies includes investment and speculation on price developments of currencies and assets. The goal of this paper is to describe technical aspects relevant for the TAST research project. Economic considerations or future price developments are therefore not discussed. Technologies are described from a purely technical point of view. Therefore, the information in this paper is provided for general information purposes only and is not intended to provide advice, information, predictions, or recommendations for any investment. We do not accept any responsibility and expressly disclaim liability with respect to reliance on information or opinions published in this paper and from actions taken or not taken on the basis of its contents.

Acknowledgment

The work presented in this paper has received funding from Pantos GmbH within the TAST research project.

References

[1] M. Borkowski et al. *Caught in Chains: Claim-First Transactions for Cross-Blockchain Asset Transfers*. 2018. URL: http://dsg.tuwien.ac.at/staff/mborkowski/pub/tast/tast-white-paper-2.pdf. White Paper, Technische Universität Wien. Version 1.1. Accessed 2019-02-14.

[2] M. Borkowski et al. *Deterministic Witnesses for Claim-First Transactions*. 2018. URL: http://dsg.tuwien.ac.at/staff/mborkowski/pub/tast/tast-white-paper-3.pdf. White Paper, Technische Universität Wien. Version 1.0. Accessed 2019-02-14.

[3] M. Borkowski et al. *Towards Atomic Cross-Chain Token Transfers: State of the Art and Open Questions within TAST*. 2018. URL: http://dsg.tuwien.ac.at/staff/mborkowski/pub/tast/tast-white-paper-1.pdf. White Paper, Technische Universität Wien. Version 1.2. Accessed 2019-02-14.

[4] G. Chockler et al. "Spidercast: a scalable interest-aware overlay for topic-based pub/sub communication". In: *Inaugural International Conference on Distributed Event-Based Systems (DEBS 2007)*. ACM. 2007, pp. 14–25.

[5] A. Dorri et al. "Towards an optimized blockchain for IoT". In: *2nd International Conference on Internet-of-Things Design and Implementation*. ACM. 2017, pp. 173–178.

[6] *EOS.IO Technical White Paper v2*. URL: https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md. White Paper. Version 2018-03-16 (Git: 2018-04-28). Accessed 2019-02-04.

[7] T. M. Fernández-Caramés et al. "A Review on the Use of Blockchain for the Internet of Things". In: *IEEE Access* 6 (2018), pp. 32979–33001.

[8] J. Herrera-Joancomartí et al. "Privacy in Bitcoin Transactions: New Challenges from Blockchain Scalability Solutions". In: *Modeling Decisions for Artificial Intelligence*. Ed. by V. Torra et al. Cham: Springer, 2016, pp. 26–44.

[9] I. Konstantinidis et al. "Blockchain for Business Applications: A Systematic Literature Review". In: *Business Information Systems*. Ed. by W. Abramowicz et al. Springer, 2018, pp. 384–399.

[10] X. Liang et al. "Provchain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability". In: *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2017)*. IEEE. 2017, pp. 468–477.

[11] *Litecoin*. URL: https://litecoin.org/. Website. Accessed 2019-02-14.

[12] *Metronome: Owner's Manual*. URL: https://www.metronome.io/pdf/owners_manual.pdf. White Paper. Version 0.967, 2018-04-17. Accessed 2019-02-14.

[13] S. Nakamoto. *Bitcoin: A peer-to-peer electronic cash system*. 2008. White Paper.

[14] C. Prybila et al. "Runtime verification for business processes utilizing the Bitcoin blockchain". In: *Future Generation Computer Systems* (2019).

[15] M. Rosenfeld. *Overview of Colored Coins*. 2012. URL: https://bitcoil.co.il/BitcoinX.pdf. White Paper. Version 2012-12-04. Accessed 2019-02-14.

[16] *Tether: Fiat currencies on the Bitcoin blockchain*. URL: https://tether.to/wp-content/uploads/2016/06/TetherWhitePaper.pdf. White Paper. Version 2018-03. URL not accessible, mirror at https://whitepaperdatabase.com/tether-usdt-whitepaper/. Accessed 2019-02-14.

[17] *TRON: Advanced Decentralized Blockchain Platform*. URL: https://tron.network/static/doc/white_paper_v_2_0.pdf. White Paper. Version 2.0, 2018-12-10. Accessed 2019-02-14.

[18] A. M. Turing. "Systems of logic based on ordinals". In: *Proceedings of the London mathematical society* 2.1 (1939), pp. 161–228.

[19] M. Vukolić. "The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication". In: *International Workshop on Open Problems in Network Security*. Springer. 2015, pp. 112–125.

[20] S. R. Walli. "The POSIX family of standards". In: *StandardView* 3.1 (1995), pp. 11–17.

[21] J. Willett et al. *Omni Protocol Specification*. 2017. URL: https://github.com/OmniLayer/spec. Version 0.5 (Git: 2017-01-23). Accessed 2019-02-14.

[22] G. Wood. *Ethereum: A Secure Decentralised Generalised Transaction Ledger*. 2018. URL: https://ethereum.github.io/yellowpaper/paper.pdf. White Paper. Version 69351d5, 2018-12-10. Accessed 2019-02-14.