## DIGITAL TWINS FOR DYNAMIC MANAGEMENT OF BLOCKCHAIN SYSTEMS

Georgios Diamantopoulos

University of Birmingham Edgbaston Birmingham, B15 2TT, UNITED KINGDOM Southern University of Science and Technology 1088 Xueyuan Blvd Shenzhen, 518055, CHINA Nikos Tziritas

Department of Informatics and Telecommunications University of Thessaly Papasiopoulou 2-4 Lamia, 35100, GREECE

Rami Bahsoon

School of Computer Science University of Birmingham Edgbaston Birmingham, B15 2TT, UNITED KINGDOM Georgios Theodoropoulos

Department of Computer Science and Engineering Research Institute for Trustworthy Autonomous Systems Southern University of Science and Technology 1088 Xueyuan Blvd Shenzhen, 518055, CHINA

# ABSTRACT

Blockchain systems are challenged by the so-called Trilemma tradeoff: decentralization, scalability, and security. Infrastructure and node configuration, choice of the Consensus Protocol, and complexity of the application transactions are cited among the factors that affect the tradeoff balance. Given that Blockchains are complex, dynamic systems, a dynamic approach to their management and reconfiguration. This paper introduces the utilization of DigitalTwins for this purpose. The novel contribution of the paper is the design of a framework and conceptual architecture of a Digital Twin that can assist in maintaining the Trilemma tradeoffs of time-critical systems. The proposed Digital Twin is illustrated via an innovative approach to the dynamic selection of Consensus Protocols. Simulation results show that the proposed framework can effectively support the dynamic adaptation and management of the Blockchain.

# **1 INTRODUCTION**

Blockchain has seen a huge leap in popularity since its inception as an immutable, decentralized ledger used by Bitcoin (Nakamoto 2008) and the plethora of other applications that soon followed. In Blockchain, entities that wish to transact with each other form a P2P network through which cryptographically signed transactions are batched into blocks, broadcasted, and stored in a chain of blocks by every entity individually. A transaction is defined as the transfer of a digital token which can be designed to model any functionality through a process called tokenization (Li et al. 2019). A major distinction between Blockchains is their type, with the two main categories being permissionless (or public) and permissioned (or private) (Helliar et al. 2020); with a consortium Blockchain (Li et al. 2017), being a hybrid type encompassing features of both main ones. In a permissionless Blockchain, the P2P network is public and everyone can participate anonymously. As a result, the network topology is unknown and no a-priori assumptions can be made

about nodes or the expected load of the system. In the permissioned case, the P2P network is private, and only verified nodes can participate, thus providing more knowledge about the state of the system. Finally, in a consortium Blockchain, the network is public and everyone can participate but only verified nodes can produce blocks.

Blockchain has been increasingly utilized in a wide range of applications including IoT, supply chain systems, e-government systems, medical databases and more recently metaverse type applications (Zheng et al. 2018; Al-Jaroodi et al. 2019; Maesa et al. 2020; Dai et al. 2019; Monrat et al. 2019; Gamage et al. 2020; Yang et al. 2022). The potential of Blockchain technology to support sustainable development is also increasingly being acknowledged, while tokenization is viewed as the key technology to promote and power ESG, impact investment, and sustainable finance (Uzsoki et al. 2019; Freire et al. 2021).

Despite the widely acknowledged potentials of Blockchain, several factors limit, if not prohibit its adoption in time-critical applications: low scalability, high latency, coupled with high power consumption, and an expanding carbon footprint are among the most cited factors (Yu et al. 2018). As an indicative example, Bitcoin can confirm an average of 4 transactions per second (TPS) and Ethereum's public implementation can confirm an average of 14 TPS (Graphs 2022) in comparison, VISA, a traditional transaction processing system, claims to process more than 24,000 TPS (Visa 2022). Henceforth, the designers of Blockchain-based systems are pressured by the need to develop secure, scalable, speedy, and sustainable solutions.

Aspiring to contribute to this endeavour, this paper presents an approach for the dynamic management and optimization of permissioned Blockchain systems utilizing Digital Twins. The novel contribution of this paper is the design of a framework and a conceptual architecture leveraging Digital Twin technology to assist application designers in maintaining the so-called Trilemma tradeoff in Blockchain-based systems (coined as suggested by Ethereum's Vitalik Buterin): decentralization, scalability, and security.

Our approach views the Digital Twin as a "combination of a computational model and a real-world system, designed to monitor, control and optimize its functionality". The objective of using Digital Twins is to dynamically assist in managing and optimizing the Trilemma tradeoffs in Blockchain-based systems. Our approach is fundamentally grounded on the premise that Digital Twins are essentially Dynamic Data-Driven Application Systems (DDDAS), wherein a real-time info-symbiotic feedback loop between the model and the real system allows data from an observed system to be absorbed into a simulation of the system to continually adapt the model to the reality, and if necessary, making changes to the assumptions on which it is based to gradually increase the reliability of its forecasts. Additionally, the predictions of the simulation can be fed back to the observed system to change or optimize its behaviour in real-time and direct the data collection and sampling (Darema et al. 2008).

Digital Twins and DDDAS have been utilized in a wide range of applications (Blasch et al. ; Jones et al. 2020; Minerva et al. 2020; dos Santos et al. 2021; Barricelli et al. 2019), including autonomic management of computational infrastructures (Liu et al. 2012; Onolaja et al. 2010; Faniyi et al. 2012; Abar et al. 2014). The last few years have witnessed several efforts to bring together Blockchain and Digital Twins, however, these have focused on utilizing the former to support the latter; a comprehensive survey is provided in (Suhail et al. 2022). Similarly, in the context of Dynamic Data-Driven Application Systems (or DDDAS), Blockchain technology has been utilized to support different aspects of DDDAS operations and components (Blasch et al. 2019; Xu et al. 2019; Xu et al. 2020). In contrast, this paper aims to address the reverse challenge namely how can the DDDAS paradigm and Digital Twin technology be utilized to support the dynamic management and optimization of blockchain systems.

The novel contributions of the paper are the following (1) It is the first to propose the utilization of Digital Twins to dynamically manage the Trilemma Tradeoffs in Blockchain systems. (2) It presents a generic reference architecture of Digital Twins for managing the Trilemma in Blockchain systems. (3) It demonstrates how the architecture can be instantiated to optimize for performance and to inform the dynamic selection and management of consensus in Blockchain-based systems. (4) It presents a quantitative analysis of dynamically adapting Consensus Protocols to optimize performance.

The rest of the paper is structured as follows: Section 2 discusses the factors affecting the performance of Blockchain systems, their dynamic management and the challenge of managing them. Section 3 presents a reference architecture of a Digital Twin for permissioned Blockchain systems, outlining its main components and illustrating an example instantiation of the architecture for the dynamic management of Consensus Protocols. Section 4 presents a quantitative analysis. Finally, section 5 concludes the paper and outlines paths for future research.

## 2 MANAGING BLOCKCHAIN DYNAMICS

The design of Blockchain-based systems is challenged by the well-known Trilemma tradeoff, coined by Vitalik Buterin, the co-founder of Ethereum: decentralization, scalability, and security. Factors that affect the behaviour of the Blockchain and change the balance between these three attributes relate to computational infrastructure and node configuration, the Consensus Protocol, and the complexity of the application transactions. Parameters such as network topology, bandwidth and latency, CPU and storage capacity, mining power utilization, number of nodes, distribution of mining power, block size, block interval, number of block producers, and orphaning/fork probability determine the transaction throughput and energy profile of the Blockchain system (Eklund et al. 2019; Odiljon et al. 2019; Hafid et al. 2020; Gencer et al. ; Xiong et al. 2018; Klarman et al. 2019).

The Consensus Protocol is at the core of influencing the Trilemma tradeoff of Blockchain-based systems. In the field of distributed systems, consensus algorithms have been thoroughly studied and optimized (Lamport 2001; Ongaro et al. 2015). Based on these well know and established algorithms, new variants have emerged for the blockchain. These variants have been generally effective in small-scale systems and can be best suited for permissioned or consortium Blockchain-based systems; their application to permissionless cases is not straightforward. This is attributed to the fact that permissioned/consortium Blockchain systems require a relatively smaller number of selected nodes to be in charge of producing blocks, where classical consensus algorithms can be effective. As the complexity of applications benefiting from Blockchain increases, several Consensus Protocols have been proposed to improve efficiency, scalability, transaction throughput, and convergence. However, providing solutions that maintain consistent performance over varying workloads, and in the face of changing environmental conditions and parameters remains a challenge (Giang-Truong et al. 2018). The challenge calls for dynamic and adaptive consensus to better address the Trilemma Tradeoffs in Blockchain-based systems. The concept of dynamic adaptation of consensus algorithms is further discussed in section 3.2.

With regard to the application transactions, smart contract systems are essentially complex systems with nonlinear profiles and emergent properties; their impact on the performance of the Blockchain system can not be determined a priori (Kim et al. 2017; Santos et al. 2017; Soloviev et al. 2019; dos Santos et al. 2018). The increasing complexity of transactions, partially attributed to smart contract logic-validation, has an observable impact on the performance of Blockchain-based systems.

A dynamic approach to the management and reconfiguration at runtime is deemed necessary to reflect on changes in the state of the infrastructure and application. Efforts in this direction have already commenced, looking at different aspects of Blockchain systems such as selection of neighbour nodes (Hamza et al. 2022) and optimization techniques for revenue maximization (Zhao et al. 2021).

In (Liu et al. 2019b) a framework in which a Reinforcement Learning (RL) agent is used to optimize a Blockchain system is proposed. The agent is tasked with solving a constraint optimization problem, that is, minimizing latency while not compromising on decentralization. This work provides a useful optimization exercise with some interesting insights into the ability of the agent to select the best algorithm for the state provided. However, as is typical of RL, the agent is trained on historical data and cannot provide a nonlinear extrapolation of future scenarios, which is essential when modelling complex systems (as is the case of smart contract systems).

A Digital Twin can overcome the deficiencies of solely relying on RL, as its simulation infrastructure can allow for what-if analysis and can act as a surrogate to explore alternative future scenarios (Theodoropoulos



Figure 1: A reference model for a Digital Twin managed Blockchain.

2015; Tolk 2015). Additionally, the Digital Twin can be used in conjunction with an RL agent enriching the training dataset with what-if generated scenarios and further increasing the performance of RL-based optimizers. It can also provide support for the dynamic off-chain simulation and evaluation of smart contract systems (Kim et al. 2017; Kampik et al. 2020; Hu et al. 2021; Kim and other 2021); the smart contract system can then be executed off-chain in the Digital Twin environment or uploaded to the Blockchain system thus supporting a hybrid on/off-chain execution model (Solaiman et al. 2021).

## **3 TWINNING A BLOCKCHAIN**

Blockchain is a distributed ledger technology which allows for trustless interactions between entities without a trusted middleman. Blockchain achieves the above by keeping a completely distributed and immutable ledger that stores ownership data of tokens representing physical or digital entities. A transaction in the Blockchain is defined as the change of ownership of an existing token or the generation and the assignment of ownership of a new token. In Blockchain, nodes connect with each other by forming a peer-to-peer (P2P) network and any node which wishes to send a token to another node, create a transaction and broadcasts it over the network. Asymmetric cryptography is used to prove the identity of nodes, by requiring every transaction and message sent in the Blockchain to be signed by a node's private key for identification.

When enough transactions are gathered, special nodes called block producers, batch the transactions into a block and broadcast it to the network. For a new block to be valid and accepted by the rest of the nodes it needs to have been agreed upon by the Consensus Protocol. The Consensus Protocol acts as a voting mechanism in which the block producers vote on candidate blocks to be added next to the Blockchain.

In this paper, we consider a generic permissioned Blockchain system with K nodes denoted as  $P = \{p_1, p_2, ..., p_K\}$  M of which are block producers denoted as  $B = \{b_1, b_2, ..., b_M\}$ ,  $B \subset P$  which take part in the Consensus Protocol and are responsible for producing the blocks. Figure 1 illustrates the described Blockchain system (labelled as Physical System). Each node  $p \in P$  holds a local copy of the Blockchain (BC). Additionally, block producers  $b \in B$  also hold a transaction pool (TP) which stores broadcasted transactions that took place in the system. When a node is ready to propose a new block, the oldest transactions from the TP are selected first to populate it.

#### 3.1 A Reference Architecture

A generic reference model for the proposed Digital Twin managed Blockchain is illustrated in figure 1. The model implements a typical MAPE-K approach (Kephart et al. 2003). Following the basic philosophy of

the DDDAS paradigm, data from the Blockchain system is fed to the Digital Twin at selected time intervals. The data deemed necessary to create and update the replica of the Blockchain include the following: (a) a list of transactions received in the interval (b) a list of all new blocks added since in the interval (c) the state of the block producers (d) the state of of the computational platform and workload state information as appropriate. A new block may contain the following: (a) the transactions included in the block (b) the list of block producers (c) the Consensus Protocol (CP) used to mine the block (d) the list of timestamped validator votes for a block (CP dependent) (e) votes to remove block producer rights from a node (f) votes proposing new block producers.

Due to the decentralized nature of the Blockchain, connecting the system with its digital representation is not a straightforward process. Unlike traditional centralized systems, with known and high-speed network topologies, Blockchain's P2P network infrastructure poses a challenge in data collection. Information about nodes cannot be easily requested and aggregated. Additionally, in most Blockchain applications, nodes are assumed to be byzantine and thus any non-validated information is assumed to be malicious which further complicates data collection. One approach, proposed in this paper, is to take advantage of the verifiable transaction data broadcasted to the network and the frequent communication between the block producers as part of the consensus process, assigning a single block producer as the data provider to the Digital Twin.

The Digital Twin part encompasses three main components: The Scenario generator, The Simulator, and the Optimiser. The scenario generation module can be viewed as a high-level model of the system nodes tasked with producing hypothetical workloads.

The scenarios will be fed to the Simulator which is at the heart of the Digital Twin. This may encapsulate different data-driven models to support a holistic, contextual analysis of the system, including: (a) a model of the Blockchain system and associated infrastructure (b) agent-based models of smart contract systems (c) models of the context, e.g. in the case of a Blockchain in the energy sector, this could be models of trading, models of the regulatory and compliance framework and a model of the energy supply chain (Andoni et al. 2019). The simulator executes faster than real-time multiple what-if scenarios for different views of the system, each view exploring an abstract aspect of the system to optimize for, for instance, an energy view, a trust view, a performance view, etc.

The final component is the optimizer, which is responsible for evaluating the simulation results and selecting the best strategy under the optimization goals. Pareto fronts and knee points analysis may be utilized to analyze the different tradeoffs involved (e.g. the cost of adaptation vs the sort and long-term benefits) and make a decision as to what is the best strategy to reconfigure the Blockchain. The results of the simulation can be used to enhance the training of an intelligent optimizer. In (Zhang et al. 2020; Zhang et al. 2022) we have discussed the design of intelligent Digital Twins and have presented an analysis of the tradeoffs for the adaptation of Digital Twins of agent-based systems.

The completion of the feedback loop, namely the communication and application of the optimizer output back to the Blockchain system presents an interesting challenge. One approach is to communicate the outcome to the entrusted controlled BP and allow this node to propagate it to the rest of the network. This may be achieved by piggybacking the information in the next block to be forwarded or through a broadcast to all other nodes.

### 3.2 An Instantiation for Dynamic Consensus Management

Given the centrality of the Consensus Protocol in the behaviour of Blockchain systems, as an illustrative example, this section considers the application of the reference Digital Twin model for the dynamic management of the Blockchains Consensus Protocol. As discussed in section 2, each of the existing protocols seem to work well under certain Blockchain configurations and workload conditions while none is able to deliver a consistently good general solution (Giang-Truong et al. 2018; dos Santos et al. 2018). Hybrid algorithms aim to exploit the comparative advantages of different protocols but they fail to reflect dynamic changes of the Blockchain and the associated workloads (Huang et al. 2021; Liu et al. 2019a).



Figure 2: A Digital Twin for dynamic Consensus Protocol selection.

It is therefore desirable for the consensus mechanism to adapt dynamically and switch to the appropriate protocol. Figure 2 illustrates an instantiation of generic Digital Twin architecture presented in the previous section, that specifically focuses on the dynamic selection of the consensus protocol. In this particular example, we look at transaction latency as a metric to ensure that the system is optimized for the Trilemma Tradeoff without compromising the Quality of Experience (QoE).

The optimization process begins with the new transactions and new blocks being fed into the Digital Twin. The number of offline nodes and the network state may be extracted from the blocks. Specifically, offline nodes can be inferred by the lack of block votes from a particular node while the network delay is calculated individually for every node as the average delay of their votes. Given the above, the simulation conducts a what-if exploration of different scenarios for the different Consensus Protocols to predict the average transaction latency for different configurations. The results of the simulation module are fed into the optimizer which makes the final decision as to which Consensus Protocol to be selected and this decision is communicated back to the system resulting in a dynamic switch of the Consensus Protocol used by the Blockchain nodes.

Transaction latency is defined as the time it takes from the moment a transaction is broadcasted to the system to the moment that transaction is packed into a block that gets accepted by the system. The system latency is primarily affected by the network state, the number of honest and malicious and/or faulty nodes, and the workload. Additionally, parameters such as the block-size and the block-interval can further be used to fine-tune the latency of the system, although they are not taken into consideration for this specific instantiation. Each of the different existing consensus protocols aims to offer optimal performance under specific system configurations.

## **4** EVALUATION

To demonstrate the suitability of the proposed approach, this section presents a quantitative analysis focusing on the optimization of transaction latency by dynamically switching the Consensus Protocol. The analysis is based on a prototype implementation of the conceptual model presented in section 3.2. The results obtained show that dynamically switching Consensus Protocols to reflect changes in the Blockchain system leads to better performance.

## 4.1 Consensus Algorithms

Two Consensus Protocols have been used for the experiments, the Istanbul Byzantine Fault Tolerance (IBFT) (Moniz 2020) and BigFoot (Saltini 2022), with IBFT having the ability to tolerate less stable

network conditions and node failures and BigFoot being very efficient under stable ones. Both algorithms require 3f+1 nodes for tolerating f faulty nodes and can achieve consensus with 2f+1 replies which is the theoretical optimal (Lamport 2003).

The two protocols are illustrated in Figure 3 (in our analysis, block producer and validator are used interchangeably). BigFoot has two phases, namely *Fast-path* and *Fallback-path*. In the Fast-path, BigFoot is efficient, requiring 2 message delays (namely Pre-prepare and Prepare) to reach consensus but 3f replies i.e. every node in the system must be online, in sync, and able to reply in time. To guarantee that BigFoot will eventually reach consensus even under less stable conditions, the Fast-path phase ends after a time-out period if less than 3f but more than 2f replies are received (since the proposing node also counts as a validator, 2f replies imply 2f+1 validator votes). In this case, an extra Fallback-path phase is initiated, achieving consensus in 1 extra message delay (commit) and 2f+1 replies. PBFT on the other hand always requires 3 message delays (pre-prepare, prepare, commit) and 2f+1 replies to reach consensus.

It is evident from the above, that both algorithms sacrifice performance under certain conditions to excel in others, a fact, that the proposed Digital Twin approach can take advantage of, and dynamically switch between the two.

# 4.2 Simulation

The conduct the experiments, the BlockSim (Alharby et al. 2019) simulator was used and extended to support the modelling of permissioned Blockchains and satisfy the requirements for the system. To the best of our knowledge, our extension of BlockSim is the only one supporting dynamically changing the Consensus Protocols (IBFT and BigFoot) during runtime and one of only two tools supporting permissioned Blockchain simulation, the other being Talaria (Xing et al. 2021). Specifically, the Node, block, and Consensus modules of the BlockSim simulator were re-implemented to model a permissioned system. The Events structure of the system was changed from supporting high-level events such as consensus and propagate block, to being able to model the lowest possible level events representing individual messages between the nodes to allow for more accurate modelling of the system. Finally, the Network module was augmented to model a unique network state for each node and allow for more complex network states to be modelled.

Using BlockSim, a prototype model of the system illustrated in figure 2 has been developed. The block producers produce blocks of size *BS* with a block interval of *BI*. The Consensus Protocol works in rounds. In each round, a block producer is selected to propose a block, and initiate the consensus process by broadcasting the proposed block. After a block is accepted by a node, that node automatically advances to the next round. Each round has a timeout period after which the nodes initiate the round change process to agree on which round to advance to next. This guarantees the system's liveliness under faulty or malicious block producers. Finally, the Blockchain node syncing protocol is modified to include the current Consensus Protocol along with the new blocks.

Data is provided to the Digital Twin periodically in time intervals (TI) of equal length. The system remains stable for a number of TI, and state changes occur every TS (TS being an integral multiple of TI). The parameters used in the system are the following: (a) no. of block producers, (b) state of block producers over time, (c) network state, (d) state of network over time, (e) avg. transactions per second, (f) transaction size, (g) max block size, (h) min block interval, and (i) round timeout.

In the simulation phase, the scenario simulates the two protocols for the next time interval  $TI_{n+1}$  and the results of the simulation i.e. the blocks produced, are fed to the optimizer. The optimizer selects the Consensus Protocol that yields the best average transaction latency.

## 4.3 Evaluation

**Experimental setup:** The physical system model was designed to represent the worst-case scenario, with faulty nodes periodically going offline, an unstable network, and a large number of transactions overloading

Diamantopoulos, Tziritas, Bahsoon, and Theodoropoulos



Figure 3: The two Consensus Protocols used in the analysis.

the system. The model consists of ten block producers with two of them being faulty. The system fluctuates between states every 100 time steps (TS = 100), with the updates occurring every 25 time steps (TI = 25).

For performance evaluation, three metrics were used: (a) average transaction latency (b) average inter-block time, and (c) throughput. The average transaction latency was measured for each new block B added as follows  $\frac{\sum_{i}^{T_B} Time_B - Time_{T_i}}{T_B}$  with  $T_B$  denoting the number of transactions in B,  $T_i$  the  $i_{th}$  transaction in B and  $Time_B$ ,  $Time_{T_i}$  the time the block B and transaction  $T_i$  were added to the system, respectively. The inter-block time was measured as follows  $\frac{\sum_{i}^{B_{BC}} Time_{B_i} - Time_{B_{i-1}}}{B_{BC}}$  with  $B_{BC}$  denoting the number of blocks. Finally, the throughput is defined as the number of transactions that the system is able to process per second and is defined as  $\frac{\sum_{i}^{B_{BC}} T_{B_i}}{T}$  with T denoting the total system runtime.

The Digital Twin was tasked with optimizing the above system by dynamically switching between the two consensus protocols (IBFT and BigFoot) and for comparison, two other identical simulations were executed, one using the IBFT protocol and the second the BigFoot without any protocol changes. Specifically, the parameter values for the evaluation are the following: (a) no. of block producers is set to 10 (BP=10), (b) 2 faulty nodes periodically going offline (f=2), (c) network state ranges from 0.7MB/s to 2MB/s over time (e) avg. transactions per second is set to 50T/s, (f) transaction size is set to 5KB (g) max block size is set to 1MB (h) min block interval is set to 0.1s, and (i) the round timeout is set to 10s.

**Results:** Figure 4 depicts the results of the experiments (note that black points in 4a and 4b denote the mean values), with Dynamic denoting the system optimized by the Digital Twin. It is evident that the dynamic protocol switching delivers the best performance. It achieves lower average transaction latency and inter-block times, as well as higher throughput. These results confirm that dynamic management of the blockchain utilizing a Digital Twin is a viable approach to optimizing a blockchain system by adapting the system parameters to reflect system and workload changes.



(a) Latency (BP=10, f=2)

(b) Inter Block Times (BP=10, f=2)

(c) Throughput (BP=10, f=2)

Figure 4: Performance results.

## 5 CONCLUSIONS

This paper has put forward the idea of utilizing Digital Twins to dynamically manage Blockchain systems in order to address the Trilemma tradeoff. It has proposed a generic reference Digital Twin architecture and has demonstrated how the architecture can be instantiated for the dynamic selection and management of consensus in Blockchain-based systems to optimize performance, as a core influencer of these tradeoffs. The experimental analysis has indicated that a Digital Twin can serve as a viable approach for dynamically managing the Trilemma tradeoff and help in improving performance.

Future work will further develop the architecture and refine it to dynamic analysis and management of a richer set of scenarios and complex time-varying tradeoffs. We will extend the architecture to incorporate multiple views, each abstracting finer aspects of the Trilemma. A more sophisticated optimizer will be developed that will utilize dynamic many optimization and machine learning techniques to optimize within and across views to assist in planning and what-if analysis. For the what-if analysis, we will utilize Distributed Simulation techniques to scale the analysis in the digital world and achieve richer and faster than real-time simulation. Other dimensions that can influence the Quality of Experience (QoE) and the Quality of Service (QoS) of the system will be considered.

In some smaller-scale systems, the performance gains delivered by the Digital Twin may be out weighted by the added computational overhead introduced by the Digital Twin. The introduction of the optimizer further increases the complexity of the system. A Digital Twin is a data-driven system and thus its performance heavily depends on the quality of the data used; this issue is exacerbated by the difficulty of collecting data from the blockchain given its distributed nature. We plan to investigate these issues as part of our future work.

#### Acknowledgements

This research was supported by: Shenzhen Science and Technology Program, China (No. GJHZ2021070514 1807022); SUSTech-University of Birmingham Collaborative PhD Programme; Guangdong Province Innovative and Entrepreneurial Team Programme, China (No. 2017ZT07X386); SUSTech Research Institute for Trustworthy Autonomous Systems, China. Georgios Theodoropoulos is a corresponding author.

#### REFERENCES

Abar et al. 2014. "Automated Dynamic Resource Provisioning and Monitoring in Virtualized Large-Scale Datacenter". In 2014 IEEE 28th International Conference on Advanced Information Networking and Applications, edited by T. E. F. X. Leonard Baroli, Kin Fun Li and M. Takizawa, 961–970.

Al-Jaroodi et al. 2019. "Blockchain in Industries: A Survey". IEEE Access 7:36500-36515.

- Alharby et al. 2019. "BlockSim: A Simulation Framework for Blockchain Systems". ACM SIGMETRICS Performance Evaluation Review 46(3):135–138.
- Andoni et al. 2019. "Blockchain Technology in the Energy Sector: A Systematic Review of Challenges and Opportunities". *Renewable and Sustainable Energy Reviews* 100:143–174.
- Barricelli et al. 2019. "A Survey on Digital Twin: Definitions, Characteristics, Applications, and Design Implications". *IEEE Access* 7:167653–167671.

Blasch et al. Handbook of Dynamic Data Driven Applications Systems. Switzerland.

- Blasch et al. 2019. "A Study of Lightweight DDDAS Architecture for Real-Time Public Safety Applications through Hybrid Simulation". In *Proceedings of the Winter Simulation Conference*, edited by S. L. M. M. R. C. S. P. H. N. Mustafee, K.-H.G. Bae and e. Y-J. Son, WSC '19, 762–773.
- Dai et al. 2019. "Blockchain for Internet of Things: A Survey". *IEEE Internet of Things Journal* 6(5):8076–8094.
- Darema et al. 2008. "Dynamic Data Driven Applications Systems (DDDAS) A Transformative Paradigm". In *Computational Science – ICCS 2008*, edited by M. Bubak, G. D. van Albada, J. Dongarra, and P. M. A. Sloot, 5–5. Berlin, Heidelberg: Springer Berlin Heidelberg.

- dos Santos et al. 2018. "PoW, PoS, & Hybrid protocols: A Matter of Complexity?". arXiv preprint arXiv:1805.08674. https://arxiv.org/abs/1805.08674, accessed 20.4.202.
- dos Santos et al. 2021. "Decision Support in Productive Processes Through DES and ABS in the Digital Twin Era: A Systematic Literature Review". *International Journal of Production Research*:1–20.
- Eklund et al. 2019. "Factors That Impact Blockchain Scalability". In *MEDES '19: Proceedings of the 11th International Conference on Management of Digital EcoSystems*. November 12<sup>th</sup>-14<sup>th</sup>, Limassol, Cyrpus, 126–133.
- Faniyi et al. 2012. "A Dynamic Data-Driven Simulation Approach for Preventing Service Level Agreement Violations in Cloud Federation". *Procedia Computer Science* 9:1167–1176.
- Freire et al. 2021. "Harnessing Blockchain For Sustainable Development: Prospects And Challenges". United Nations Conference On Trade And Development, Switzerland.
- Gamage et al. 2020. "A Survey on Blockchain Technology Concepts, Applications, and Issues". *SN Computer Science* 1(114).
- Gencer et al. "Decentralization in Bitcoin and Ethereum Networks". In *Financial Cryptography and Data Security*, edited by S. Meiklejohn and K. Sako, 439–457. Berlin, Heidelberg: Springer-Verlag.
- Giang-Truong et al. 2018. "A Survey about Consensus Algorithms Used in Blockchain". Journal of Information Processing Systems 14(1):101–128.
- Blockchain Graphs 2022. "Blockchain Graphs". https://www.blockchain.com/charts. accessed 20.4.2022.
- Hafid et al. 2020. "Scaling Blockchains: A Comprehensive Survey". IEEE Access 8:125244-125262.
- Hamza et al. 2022. "DONS: Dynamic Optimized Neighbor Selection for Smart Blockchain Networks". *Future Generation Computer Systems* 130:75–90.
- Helliar et al. 2020. "Permissionless and Permissioned Blockchain Diffusion". International Journal of Information Management 54:102–136.
- Hu et al. 2021. "A Comprehensive Survey on smart Contract Construction and Execution: Paradigms, Tools, and Systems". *Patterns* 2(2):100–179.
- Huang et al. 2021. "Incentive Assignment in Hybrid Consensus Blockchain Systems in Pervasive Edge Environments". *IEEE Transactions on Computers*.
- Jones et al. 2020. "Characterising the Digital Twin: A Systematic Literature Review". CIRP Journal of Manufacturing Science and Technology 29:36–52.
- Kampik et al. 2020. "Simulating, Off-Chain and On-Chain: Agent-Based Simulations in Cross-Organizational Business Processes". *Information* 11(1).
- Kephart et al. 2003. "The Vision of Autonomic Computing". Computer 36(1):41-50.
- Kim et al. 2017. "A Perspective on Blockchain Smart Contracts: Reducing Uncertainty and Complexity in Value Exchange". In *Proceedings of the IEEE Workshop on Workshop on Privacy, Security, Trust Blockchain Technologies*: http://dx.doi.org/10.2139/ssrn.2975770.
- Kim, and other. 2021, July. "An Off-The-Chain Execution Environment for Scalable Testing and Profiling of Smart Contracts". In 2021 USENIX Annual Technical Conference (USENIX ATC 21), edited by I. Calciu and G. Kuenning, 565–579: USENIX Association.
- Klarman et al. 2019. "bloXroute: A scalable trustless blockchain distribution network, Whitepaper v1.2". *Bloxroute Labs.*
- Lamport, L. 2001. "Paxos made simple". ACM SIGACT News Distributed Computing Column 32:51-58.
- Lamport, L. 2003. "Lower Bounds for Asynchronous Consensus". In *Future directions in distributed computing*, 22–23. Springer.
- Li et al. 2017. "Consortium Blockchain for Secure Energy Trading in Industrial Internet of Things". *IEEE transactions on industrial informatics* 14(8):3690–3700.
- Li et al. 2019. "Tokenization: Open Asset Protocol on Blockchain". In *IEEE International Conference on Information and Computer Technologies (ICICT)*, 204–209.

- Liu et al. 2012. "Towards an agent-based symbiotic architecture for autonomic management of virtualized data centers". In *Proceedings of the Winter Simulation Conference*, edited by R. P. O. R. Christoph Laroque, Jan Himmelspach and A. M. Uhrmacher, WSC '12, 1–13.
- Liu et al. 2019a. "Fork-free Hybrid Consensus with Flexible Proof-of-Activity". *Future Generation Computer Systems* 96:515–524.
- Liu et al. 2019b. "Performance Optimization for Blockchain-enabled Industrial Internet of Things (IIoT) Systems: A Deep Reinforcement Learning Approach". *IEEE Transactions on Industrial Informatics* 15(6):3559–3570.
- Maesa et al. 2020. "Blockchain 3.0 Applications Survey". Journal of Parallel and Distributed Computing 138:99–114.
- Minerva et al. 2020. "Digital Twin in the IoT Context: A Survey on Technical Features, Scenarios, and Architectural Models". *Proceedings of the IEEE* 108(10):1785–1824.
- Moniz, H. 2020. "The Istanbul BFT Consensus Algorithm". *arXiv preprint arXiv:2002.03613*. https://arxiv.org/abs/2002.03613, accessed 20.4.202.
- Monrat et al. 2019. "A Survey of Blockchain From the Perspectives of Applications, Challenges, and Opportunities". *IEEE Access* 7:117134–117151.
- Nakamoto, S. 2008. "Bitcoin: A Peer-to-Peer Electronic Cash System". Decentralized Business Review:21260.
- Odiljon et al. 2019. "Efficiency Issues and Solutions in Blockchain: A Survey". In *Smart Blockchain*, edited by B. S. G. W. Elisa Bertino, Wen Gao and M. Yung, 76–86. Cham: Springer International Publishing. Ongaro et al. 2015. "The Raft Consensus Algorithm".
- Onolaja et al. 2010. "Conceptual Framework for Dynamic Trust Monitoring and Prediction". Procedia
- Computer Science 1(1):1241–1250.
  Saltini, R. 2022. "BigFooT: A Robust Optimal-latency BFT Blockchain Consensus Protocol with Dynamic Validator Membership". Computer Networks 204:108–122.
- Santos, D. et al. 2017. "On the Philosophy of Bitcoin/Blockchain Technology: Is it a Chaotic, Complex System?". *Metaphilosophy* 48(5):620–633.
- Solaiman et al. 2021. "Implementation and Evaluation of Smart Contracts Using a Hybrid On- and Offblockchain Architecture". *Concurrency and Computation: Practice and Experience* 33(1).
- Soloviev et al. 2019. "Complex Systems Theory and Crashes of Cryptocurrency Market". In *Information and Communication Technologies in Education, Research, and Industrial Applications*, edited by V. Ermolayev, M. C. Suárez-Figueroa, V. Yakovyna, H. C. Mayr, M. Nikitchenko, and A. Spivakovsky, 276–297. Cham: Springer International Publishing.
- Suhail et al. 2022. "Blockchain-based digital twins: research trends, issues, and future challenges". ACM Computing Surveys 54(11):1–34.
- Theodoropoulos, G. 2015. "Simulation in the Era of Big Data: Trends and Challenges". In Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, edited by D. R. J. Rajive L Bagrodia, Damal K Arvind and J. Y. B. Lin, SIGSIM PADS '15, 1. New York, NY, USA: Association for Computing Machinery.
- Tolk, A. 2015. "The Next Generation of Modeling & Simulation: Integrating Big Data and Deep Learning". In Proceedings of the Conference on Summer Computer Simulation, edited by I.-C. M. Saurabh Mittal and E. Syriani, SummerSim '15, 1–8. San Diego, CA, USA: Society for Computer Simulation International.
- Uzsoki et al. 2019. "Impact Tokens A blockchain based solution for impact investing". International Institute for Sustainable Development, Switzerland.
- Visa 2022. "VisaFor Retail Information". https://usa.visa.com/run-your-business/small-business-tools/retail. html. accessed 20.4.2022.
- Xing et al. 2021. "Talaria: A Framework for Simulation of Permissioned Blockchains for Logistics and Beyond". *arXiv preprint arXiv:2103.02260*. https://arxiv.org/abs/2103.02260, accessed 20.4.2022.

- Xiong et al. 2018. "When Mobile Blockchain Meets Edge Computing". *IEEE Communications Magazine* 56(8):33–39.
- Xu et al. 2019. "Exploration of Blockchain-enabled Decentralized Capability-based Access Control Strategy for Space Situation Awareness". *Optical Engineering* 58(4).
- Xu et al. 2020. "Hybrid Blockchain-enabled Secure Microservices Fabric for Decentralized Multi-domain Avionics Systems". In *Proceedings of Sensors and Systems for Space Applications XIII*, edited by G. Chen and K. D. Pham, Volume 11422.
- Yang et al. 2022. "Fusing Blockchain and AI with Metaverse: A Survey". *arXiv preprint arXiv:2201.03201*. https://arxiv.org/abs/2201.03201, accessed 20.4.2022.
- Yu et al. 2018. "Virtualization for Distributed Ledger Rechnology (vDLT)". IEEE Access 6:25019–25028.
- Zhang et al. 2020. "Towards Engineering Cognitive Digital Twins with Self-Awareness". In 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC). October 11<sup>th</sup>-14<sup>th</sup>, Toronto, Canada, 3891-3891.
- Zhang et al. 2022. "Knowledge Equivalence in Digital Twins of Intelligent Systems". *arXiv* (10.48550/ARXIV.2204.07481). https://arxiv.org/abs/2204.07481, accessed 20.4.2022.
- Zhao et al. 2021. "A Stochastic Model and Social Optimization of a Blockchain System Based on a General Limited Batch Service Queue". *Journal of Industrial and Management Optimization* 17(4):1845–1861.
- Zheng et al. 2018. "Blockchain Challenges and Opportunities: A Survey". *International Journal of Web and Grid Services* 14(4):352–375.

# **AUTHOR BIOGRAPHY**

**GEORGIOS DIAMANTOPOULOS** received a B.Sc degree in Computer Science and Telecommunications from the University of Thessaly, Greece in 2021. He is a PhD candidate on a split-site programme between the University of Birmingham and the Southern University of Science and Technology. His current research interests include Blockchain, Digital Twins, Optimisation, Simulation, Reinforcement Learning and Deep Neural Networks. His e-mail address is gxd192@student.bham.ac.uk

**NIKOS TZIRITAS** received the PhD degree from the University of Thessaly, Greece, in 2011. He was an associate professor with the Chinese Academy of Sciences. In 2020, he joined the University of Thessaly as an assistant professor. He has authored or co-authored more than 60 publications. He was the recipient of the Award for Excellence for Early Career Researchers in Scalable Computing from IEEE Technical Committee in Scalable Computing in 2016. His e-mail address is nitzirit@uth.gr

**RAMI BAHSOON** is a Reader in Autonomous and at the School of Computer Science, University of Birmingham, UK. He conducts research in the fundamentals of self-adaptive and managed software architectures and its application to emerging paradigms such as cloud, microservices, IoT, Blockchain, Digital Twins and CPS etc. He co-edited four books on Software Architecture. He holds a PhD in Software Engineering from University College London(2006). He is a fellow of the Royal Society of Arts, Associate Editor of IEEE Software and Editor of Wiley Practice and Experience. His e-mail address is r.bahsoon@bham.ac.uk

**GEORGIOS THEODOROPOULOS** is currently a Chair Professor at the Department of Computer Science and Engineering at SUSTech in Shenzhen, China's Silicon Valley. He joined SUSTech from Durham University, UK where he was the inaugural Executive Director of the Institute of Advanced Research Computing, a Chair Professor in Computer Engineering and the Head of the Innovative Computing Group at the School of Engineering and Computing Sciences. He has been a Senior Research Scientist with IBM Research, held an Adjunct Chair at the Trinity College Dublin and senior faculty positions at the Nanyang Technological University, Singapore and the University of Birmingham, UK, where he was also founding Director of one of UK's National e-Science Centres of Excellence. He is a Chartered Engineer and holds a Ph.D. from the University of Manchester, UK. He is a Fellow of the World Academy of Art and Science. His e-mail address is georgios@sustc.edu.cn