

2020

A Blockchain Simulator for Evaluating Consensus Algorithms in Diverse Networking Environments

Peter Foytik

Old Dominion University, pfoytik@odu.edu

Sachin Shetty

Old Dominion University, sshetty@odu.edu

Sarada Prasad Gochhayat

Old Dominion University, sgochhay@odu.edu

Eranga Herath

Old Dominion University, cmedawer@odu.edu

Deepak Tosh

See next page for additional authors

Follow this and additional works at: https://digitalcommons.odu.edu/vmasc_pubs



Part of the [OS and Networks Commons](#), and the [Systems and Communications Commons](#)

Original Publication Citation

Foytik, P., Shetty, S., Gochhayat, S. P., Herath, E., Tosh, D., & Njilla, L. (2020). A blockchain simulator for evaluating consensus algorithms in diverse networking environments. 2020 Spring Simulation Conference (SpringSim), Fairfax, Virginia (Virtual Event), May 18-21, 2020.

This Conference Paper is brought to you for free and open access by the Virginia Modeling, Analysis & Simulation Center at ODU Digital Commons. It has been accepted for inclusion in VMASC Publications by an authorized administrator of ODU Digital Commons. For more information, please contact digitalcommons@odu.edu.

Authors

Peter Foytik, Sachin Shetty, Sarada Prasad Gochhayat, Eranga Herath, Deepak Tosh, and Laurent Njilla

A BLOCKCHAIN SIMULATOR FOR EVALUATING CONSENSUS ALGORITHMS IN DIVERSE NETWORKING ENVIRONMENTS

Peter Foytik
Sachin Shetty
Sarada Prasad Gochhayat
Eranga Herath

Deepak Tosh
University of Texas at El Paso
El Paso, Texas USA
dktosh@utep.edu

Virginia Modeling Analysis and Simulation Center
Old Dominion University
1030 University Blvd.
Suffolk, VA, USA
{pfoytik, sshetty,sgochhay,cmedawer}@odu.com

Laurent Njilla
Air Force Research Lab
Rome, NY USA
laurent.njilla@us.af.mil

ABSTRACT

The massive scale, heterogeneity and distributed nature of Internet-of-Things (IoT) presents challenges in realizing a practical and effective security solution. Blockchain empowered platforms and technologies have been proposed to address aspects of this challenge. In order to realize a practical Blockchain deployment for IoT, there is a need for a testing and evaluation platform to evaluate performance and security of Blockchain applications and systems. In this paper, we present a Blockchain simulator that evaluates the consensus algorithms in a realistic and configurable network environment. Though, there are several Blockchain evaluation platforms, they are either wedded to a specific consensus protocol and do not allow evaluation in a configurable and realistic network environment. In our proposed simulator, we provide the ability to evaluate the impact of the consensus and network layer that will inform practitioners on the appropriate choice of consensus algorithms and the impact of network layer events in congested or contested scenarios in IoT. To accomplish this a generalized representation for consensus methods is proposed. The Blockchain simulator uses a discrete event simulation engine for fidelity and increased scalability. We evaluate the performance of the simulator by varying the number of peer nodes and number of messages required to find consensus.

Keywords: Blockchain, Internet-of-Things, Consensus, NS3, Simulator

1 INTRODUCTION

Blockchain is a distributed digital ledger for keeping track of transactions, contracts and is characterized by its inherent transparency, resilience and immutable properties. These properties are desirable to realize a secure Internet-of-Things (IoT) architecture capable of ensuring auditability and accountability. Blockchain has the potential to track authenticity of data transactions among millions of devices. However, deployment of IoT devices in congested and contested environments can cause performance challenges. In addition to the properties that are already baked into the Blockchain technology, there are several claims about high transaction throughput and faster response time. Recent incarnations of the Blockchain technology has also manifested into a highly complex distributed system like IoT and it is unclear if the claimed advantages are due to the strengths of the fundamental components of Blockchain or the additional system implementation. There is a need to provide a quantitative analysis to validate the claimed performance advantages of a practical Blockchain platform prior to deploying the technology in a production environment.

Researchers and practitioners have proposed several Blockchain development platforms, simulation environments, and testbeds to evaluate Blockchains and the applications for various sectors (Dinh, Wang, Chen, Liu, Ooi, and Tan 2017, Technologies 2017). These environments analyze security and performance (Gervais, Karame, Wüst, Glykantzis, Ritzdorf, and Capkun 2016), scalability based on parameter tuning (Goswami 2017), transaction latency for PoW (Wang, Chen, Yao, Liu, Xu, and Zhu 2018) and quality of the infrastructure (Yasaweerasinghelage, Staples, and Weber 2017). These platforms and testbeds range from an open deployment on publicly accessible machines to permissioned or privately owned peer-to-peer networks and simulated architecture environments to support a wide range of Blockchain-based applications. However, as researchers and practitioners have pointed out the main performance benefits of Blockchain are closely tied to the choice of consensus plane and the network plane (Zhong and Cole. 2018). The aforementioned environments offer a static network topology and do not provide the ability to configure the network plane, as needed, to evaluate the feasibility in resource constrained environments, such as, IoT. In addition, the consensus plane is not configurable, which therefore limits the ability to deploy hybrid consensus protocols in an on-demand basis. In addition, Blockchain technology's resilience property can only be put to test in the presence of a configurable network plane. This capability would prove the effectiveness for applications that would like to measure the resilience of the Blockchain technology against various network attacks and fragility of network.

Blockchain platforms, including the Hyperledger Suite (Androulaki, Barger, Bortnikov, Cachin, Christidis, De Caro, Enyeart, Ferris, Laventman, Manevich, Muralidharan, Murthy, Nguyen, Sethi, Singh, Smith, Sorniotti, Stathakopoulou, Vukolić, Cocco, and Yellick 2018), Corda (Richard Gendal Brown 2016), and Tendermint (Buchman 2016), etc., typically abstract the network plane for easy development of application specific use cases, which makes the life of application developer easy and simple. However, this abstraction can have major bottlenecks when such Blockchain platforms are incorporated in a resource constrained environments, such as, IoT. For instance, while evaluating Hyperledger Sawtooth platform in a simulated environment, with 6 validator nodes sending transactions at different rates, we observe that the average number of packets sent from the nodes grow exponentially, as seen in Fig. 1. For instance, if each node sends 10 transaction at every 2 seconds, the total number of packets sent under a completely connected network is more than 400K. The performance is highly dependent on the considered network topology and the underlying consensus scheme. Although, this evaluation is conducted for only the Hyperledger Sawtooth based on the Proof-of-Elapsed-Time (PoET) consensus model, the network plane metrics for other consensus models need to be investigated so we can better understand the feasibility of Blockchain integration in various resource constrained and diverse network connectivity environments.

The impact of the vagaries in the network layer on the Blockchain performance, especially, in IoT environments, has motivated us to develop a Blockchain simulator that allows to enable a configurable network layer to evaluate the performance of selected consensus protocols. In this paper, we propose the develop-




Figure 1: Number of packet transmitted in Hyperledger Sawtooth.

ment of a Blockchain simulator that utilizes a generalized representation of consensus protocols that would provide insights into the performance of the consensus protocols under various networking conditions. The simulated network could be an IOT network or a separate network that interfaces with an IOT network. The simulation would provide users the ability to swap in and out various consensus protocols, configure the network settings/topology, and assess the performance of the consensus protocols. The paper will present an approach to generalize consensus protocols then discuss simulating a Blockchain-based system using the generalized consensus protocol. The implementation of the simulator will be described and finally the results will be presented. The paper shows that the generalized and simulated system was able to calibrate to the observed consensus protocol, and a sensitivity analysis is shown to provide some insight into the input effects to the simulator.

The paper is organized as follows. Section 2 presents the overview of the methodology followed to design the proposed simulator. In Section 3, the implementation details and data collection techniques are described. Results related to Raft consensus model implemented in the simulator is presented in Section 4. Finally, Section 5 concludes the paper along with future plans on the proposed simulator.

2 METHODOLOGY

2.1 Overview

A set of abstraction layers, such as (1) Network; (2) Consensus; (3) Storage; and (4) View and Side, have been proposed (Croman, Decker, Eyal, Gencer, Juels, Kosba, Miller, Saxena, Shi, Gün Sirer, Song, and Wattenhofer 2016) in past to provide a hierarchical design structure for the decentralized Blockchain ecosystem. Testing performance of specific consensus based algorithms is studied in the literature (Segent 1997)(Gervais, Karame, Wüst, Glykantzis, Ritzdorf, and Capkun 2016), but a general model for network layer evaluation of Blockchains is yet to be proposed. It is important to understand the complexity and performance metrics at each layer prior to deployment of Blockchain-based solutions. Therefore, we propose the development of a simulator that implements these layers and provides statistical insights on performance and security metrics via realization of the Network and Consensus layer.

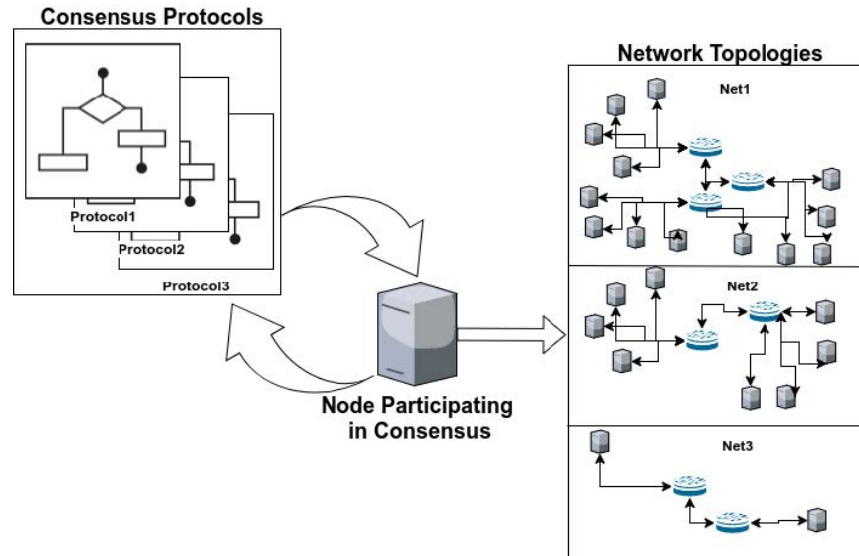


Figure 2: Consensus and network layers in blockchain simulator.

Our choice of focusing on consensus and network layer is driven by the need and potential of Blockchain deployment in IoT environments where nodes are heterogeneous and have stringent constraints on networking, computation, and communication capabilities. Given the bottom-up design (Croman, Decker, Eyal, Gencer, Juels, Kosba, Miller, Saxena, Shi, Gün Sirer, Song, and Wattenhofer 2016), the consensus and network layer is the core of Blockchain platforms and its cost/complexity needs must be analyzed prior to integration of the additional layers. This approach will provide inner insights for the community of Blockchain adopters about the performance and security metrics of the underlying consensus protocols under various network conditions. In Fig. 2, the architecture overview is presented which incorporates different network topology for establishment of a P2P network and implements various distributed consensus protocols for evaluation of performance metrics. We present an overview of the proposed architecture wherein consensus protocols and the network architectures are interchangeable.

2.2 Establishing Discrete Event Simulator

We propose a discrete event simulator to implement the consensus and network layers wherein the simulation time step is considered to be system events rather than regular time steps to facilitate quicker computation. The design goals considered while developing the simulator are:

1. consensus protocols need to be generalized,
2. underlying network needs to be re-configurable,
3. simulator must be able to handle large networks so that tests of scale can be performed.

The interchangeability of consensus protocols will ensure that the key aspect to the Blockchain systems performance is considered and evaluated instead of efficiency of the whole platform itself. The simulator will contribute to the benchmarking of consensus protocols for different network topologies. Using the simulator, developers will be able to make informed decisions on which consensus protocol would be preferred in certain types of network environments where the Blockchain is considered for by understanding the performance metrics of each protocol, which is discussed in the following subsection.




Figure 3: Demonstration of the data collection from the ideal system (Blue indicating data gathered from individual nodes and red indicating data gathered from total system).

2.2.1 Performance Metrics

The key metrics considered in the simulator are discussed briefly in the following.

- *Throughput*: Number of successful transactions/second.
- *Latency*: Response time/transaction.
- *Fault Tolerance*: Ability to find consensus and complete transactions with sub-optimal network topology.
- *Heterogeneity*: Meet above three requirements under heterogeneous network conditions and device characteristics.

The performance and resilience evaluations of consensus protocols on the considered Blockchain-integrated network are mainly conducted at the node level. Modeling the system at the node level enables better representation of the specific protocols' reliance on number of nodes and how they operate with varying network topologies. Therefore, we first observe and record the key aspects from an ideal system that has minimum influence from external factors, which is discussed more in the following subsection 2.2.2. Then, we use the observed data to model various components of the simulator for performance evaluation. Figure 3 shows the collection and use of the data from both the ideal system and the simulated system. Finally, we fine tune the developed model and validate to provide a meaningful measure on its effectiveness to represent the ideal system.

2.2.2 Ideal System

An ideal system is a set up of the consensus protocols where the external engagements and influences are minimum. The purpose of the ideal system is to provide a platform that can be easily modeled which will provide results for the baseline scenarios. In this context, the ideal system has several nodes where every node may have either transient or permanent network connectivity with each other in a mesh topology. A good candidate for implementing this is to use a virtual network and containerized technology on a single machine using docker containers (Docker 2018). The docker nodes in an ideal system act as validator nodes and run the consensus protocol. The number of networked nodes is kept variable to identify how performance metrics change as number of nodes scale. Logging is also enabled on the individual nodes to record performance aspects of the consensus protocol. In the next subsection, we discuss different kinds of gathered data and collection mechanism to better understand the state and performance of the system.

2.3 Data Collection

The ideal system provides the ground truth for the proposed simulator and will be run with each consensus protocol modeled. We collect various features in order to represent the key performance metrics of the modeled system as discussed in subsection 2.2.1. The details of data points to be recorded are discussed below.

- *nMessage*: Number of exchanged messages to achieve consensus
- *nodeDelay*: Processing delay at node level required for achieving consensus.
- *nodeConsen*: Number of participating nodes required to achieve consensus
- *consenDelay*: Network delay for the consensus protocol to achieve consensus
- *consenPauseState*: Probability that the consensus state is paused for a special case
- *pauseStateTime*: Time that the consensus system remains in the paused state

2.4 Consensus Layer Modeling

To fully represent the consensus protocols' effects on different numbers of participating nodes, limitations on network, computing resources, and network topology, several simulation runs with samples of the input data is needed. Hence, to complete these runs in a reasonable amount of time, discrete event simulation is used, where a networked message is considered as an event. The discrete event simulation engine schedules each message as the simulation runs based on how each modeled component affects the system. The component affects are modeled as statistical representations of observed data from the ideal system or distributions.

Each participating node runs a validation instance that models the actual workload of a peer in the consensus protocol. The validation process includes the necessary inputs required to represent the underlying consensus protocol steps. All of the input data for each state of the consensus protocol is stored in the form of distributed libraries or variables in the application. As the simulation runs, the participating peer either begins a new transaction or participates in the consensus of an existing transaction. The participating peer continues this operation until either a certain number of transactions have been processed or a period of time has elapsed.

Transaction processing starts when one of the current peers proposes a transaction to the system. It causes initialization of the transaction module, which then sends out messages to the participating peer nodes. The transacting node schedules the message to be delivered to every peer node that it can communicate with on the simulated network through utilizing the existing routing libraries. Based on the simulated arrival time of

the message to each node, the simulation engine orders the events accordingly where the events are handled in the order they are scheduled in the event list.

As participating peer nodes receive messages from the consensus process, they internally record the number of messages processed during the current consensus iteration. If a threshold number messages have been processed by a peer node as determined by $nMessage$, then the node's state is assigned to "completed". Peer nodes with status "completed" send a "completed" message to the other participating nodes stating that it has completed its consensus round. If the peer node has not processed enough messages, it will schedule another process message with the simulation engine for a delay sampled value based on the $nodeDelay$ distribution. The node continues to process messages until the appropriate number of messages have been processed and upon completion, its state changes to "completed". The leader peer that started the transaction monitors the system until $nodeConsen$ number of participating nodes reach a state of "completed". After enough participating nodes have reached a "completed" state, another random peer can start a transaction.

The general consensus system is modeled to change leaders periodically based on the consensus protocol it is modeling. If the system changes states, the transactions will be halted for period of time ($stateLength$). The next event would be scheduled with a new transaction and the start time is set after the delay $stateLength$ completes.

3 IMPLEMENTATION

In this Section, we discuss the steps taken to implement the proposed Blockchain simulator in details. The simulation paradigm is considered to be discrete event simulation. The general consensus protocol described in the methodology is built on Network Simulator-3 (NS3).

3.1 Simulation Data Gathering

For this initial test, Raft (Ongaro and Ousterhout 2013) is chosen as the consensus protocol because of its simplicity and the many available packages that exist to execute just the consensus protocol without any additional features. We set up Raft to run on a single machine with three to five participating nodes. For each of these settings one of the nodes acts as a leader node throughout the test. The nodes run as docker containers on that machine and are connected through a virtual network to provide as ideal connectivity as possible so that the time to process each message does not get effected by network delays. As discussed in Section 2.3, consensus data is collected for the Raft protocol to test the methodology. We record the time to find consensus for one thousand iterations and the time for each node to process a single message in each of the thousand iterations. We subsequently fit the gathered data to a distribution. Figure 4 shows the distribution of a node's processing data represented as a histogram and the result of the curve fitting (using R stats) for Weibull, Lognormal, and Normal distributions. It is observed that node's processing data for the Raft protocol follows the Weibull distribution as it best fits the ideal Raft consensus with a scale parameter of 1.35 and a shape parameter of 249857.5.

3.2 Network layer Implementation

The network simulator handles the flow messages among the participating peers. Researchers have implemented a proof-of-work consensus protocol with the Bitcoin-Simulator using NS3 (Gervais, Karame, Wüst, Glykantzis, Ritzdorf, and Capkun 2016). It has the ability to generate large scale networks and communicate using various networking methods such as P2P and gossip-relay types of network. For this paper, we modified the main controller module, as well as the Bitcoin node and built a new application extending the Bitcoin node for the general consensus node.




Figure 4: Histogram and curve fitting of node message process times in milliseconds, where the histogram shows the observed data of the node messaging and with the three lines representing the fit for weibull (red), lognormal (blue), and normal (green) distributions.

The network layer is built to represent the ideal system that is used to gather data for the simulation. To best represent this, the simulator uses a mesh P2P network using the point-to-point NS3 module. The links use data rate in the order of gigabits per second, and the link delay is set to zero with the intention of modifying this for calibration. Although the Bitcoin simulator is capable to generate more complex network topologies spanning multiple regions of the world, in this paper we consider that all nodes are connected through a mesh P2P network in the same region.

3.3 Consensus layer Implementation

As described in the methodology section, each participating node has certain variables and distributions to represent the behavior of the consensus protocol it is modeling. A leader node is specified at the beginning and will broadcast a message to the participating peers indicating the start of consensus. The message is sent using the NS3 socket class and the send ability. NS3 handles the message as a set of discrete events traversing the network. The peer nodes then proceed to process consensus messages as specified in the methodology section and consensus is found when *nodeCompleted* nodes have processed *nMessages*.

3.4 Calibration and Simulation Sensitivity

With the model developed and functioning based on the observed data from the participating nodes, the total system is calibrated to observe how close it performs with respect to the ideal system. Calibration is used to tune the model so that it performs good enough compared to the ideal system. Tuning the model and simulation is done by increasing network background traffic or adding delay to network to create more delay in the message processing. Calibration is done by tuning the distributions used to represent the message

processing. This is done using constants as a scaling factor for selected input data sets as needed. In our proposed simulator, calibration is done by: 1) increasing network background traffic or adding delay to network to create more delay in the message processing, and 2) tuning the distributions using constants as a scaling factor for selected input datasets as needed. It is an iterative process and can use optimization methods to minimize the error in time to find consensus between the model system and the observed data from the ideal system.

4 RESULTS AND DISCUSSION

In this section, we provide the results obtained both at network layer and consensus layer with respect to the time taken to find consensus in the system.

4.1 Network layer observation

The network delay is modeled using a lognormal distribution. The lognormal mean and variance are found to be good at 12.85 and 0.57 respectively. These values are fitted with an iterative process, running the simulation comparing results, then we use these results to adjust the parameters and feedback this change to the model. This process continues for several iterations until a close fit is found. The calibration is done for the ideal system, with three peers running the Raft implementation. The results of the calibrated model are shown in Fig. 5 with two scatter plots. The scatter plot on the left is calibration of the simulation to the ideal system. The x-axis represents time in seconds for a consensus transaction by the simulator. The y-axis represents time in seconds for a consensus transaction by the ideal system. The red straight line is the computed regression line of the two data sets. The left scatter plot shows that there is a tight grouping and positive correlation with the simulation and the observed ideal system. The scatter plot also shows that for the few extreme values the two systems differ some. This is likely because the curve fitting is mostly matching the common results in the observed data. The distribution could be modified and calibration could be pursued further to better match the extremem values.

The scatter plot to the right in Fig 5 is the results of the calibrated model and simulation of the ideal system but run with 5-nodes instead of 3. The simulated results are compared with observed results for the same system. In this plot, the simulator is not calibrated for 5-peers, rather we wanted to observe if the behavior is good enough to represent the system’s state using the calibration from the 3-peers scenario. This comparison in Fig 5 was done as a verification process to see if the behavior of the calibrated simulation is transferable to other scenarios.

Peers	Average (ms)	Standard Deviation	Minimum	Maximum
3	1318.48	27.87	-3.96%	4.81%
4	1505.0	26.44	-2.69%	5.72%
5	1651.89	30.73	-3.68%	3.30%
10	2168.09	40.53	-3.96%	4.06%
16	2575.94	45.39	-4.01%	4.24%

Table 1: Consensus times (in ms) for 1-Message scenario

4.2 Consensus layer observation

In order to get insight into how much the simulation of the general consensus protocol changes and understand the effect of the number of peers and the number of messages on the time to find consensus consensus, we conduct sensitivity analysis. In this analysis, we modify the number of peers and the number of messages required for the general consensus protocol, and observe their impact.

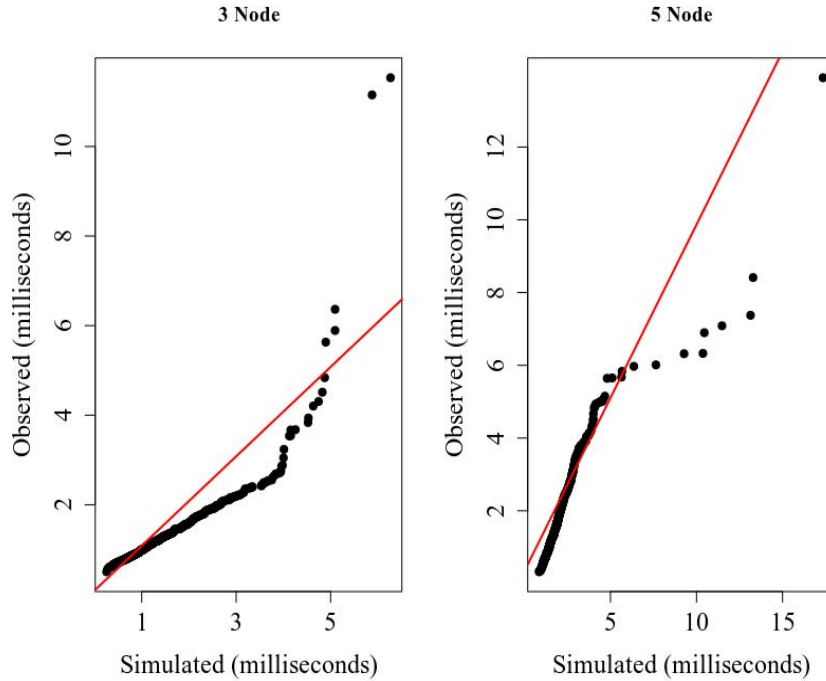


Figure 5: Calibration results of simulation consensus times versus observed consensus times in milliseconds for the 3 peer node and 5 peer node scenarios.

For this test, 30 simulation runs are conducted for each scenario to handle the stochastic nature of the sampled distributions used. The first scenario consisted of three required participating peers and the requirement for each one to process one message. The simulation processed 1000 consensus transactions and then recorded the time it took to complete the transactions. Table 4.1 shows the results for the test while varying the number of peers. In this table, the average time to process 1000 consensus transactions is shown for each number of peers tested. We can observe that the time to reach consensus increases as the number of peers increases. The table also includes the standard deviation and the minimum and maximum values as a percentage of the average to provide information on the variance over 30 simulation runs. This gives an idea on the amount of variability the simulator is producing in its results for each number of peers used. Generally speaking the standard deviation increases as the number of peers increases indicating more variance in the results. This is expected since the system becomes more complex with more participating peers.

In order to further observe the impact of the number of messages on the time to find consensus, we ran the above described simulation while increasing the number of messages. Fig 6 shows the same data for the 1 message test as shown in Table 4.1, but compared to similar tests with 4 message consensus and 10 message consensus. With this figure we observe increments in time to process 1000 consensus transactions between the different scenarios. Also, the time difference between different number of messages decreases when the number of peers increase. This finding provides some insight into the factors that cause the most delay. In this case it is shown that the number of peers has a greater effect on the delay of the system than the number of messages.

5 CONCLUSIONS AND FUTURE WORK

We present a Blockchain simulator for testing and evaluating consensus algorithms in a realistic and configurable network layer. This paper introduces and shows a generalized consensus protocol method that can be used to model various consensus protocols. By observing the ability for the system to operate in various




Figure 6: Average time (in ms) to confirm 1000 transactions with different message requirements.

configurable network topology’s we begin to study under what conditions the system fails to operate. This can occur when portions of the network do not allow the peer nodes to quickly message in participation of transactions. The unique characteristics of the simulator are a general consensus algorithm operating in a realistic and configurable network environment. The discrete event simulation engine allows to specify the consensus algorithm operations at faster than real-time fidelity without loss of scalability. In future work, we plan to implement and calibrate additional consensus algorithms to accommodate bootstrapping scenarios and evaluate performance in diverse network configurations. Future work with this model and simulation will include the effects of the consensus special state using the *consenPauseState* variable, which will be modeled to get a better comparative test between different consensus protocols.

REFERENCES

- Androulaki, E., A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, and J. Yellick. 2018. “Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains”. In *Proceedings of the Thirteenth EuroSys Conference*, EuroSys ’18, pp. 30:1–30:15. New York, NY, USA, ACM.
- Ethan Buchman 2016. “Tendermint: Byzantine Fault Tolerance in the Age of Blockchains”. Last accessed 29 April 2019.
- Croman, K., C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. Gün Sirer, D. Song, and R. Wattenhofer. 2016. “On Scaling Decentralized Blockchains”. In *Financial Cryptography and Data Security*, edited by J. Clark, S. Meiklejohn, P. Y. Ryan, D. Wallach, M. Brenner, and K. Rohloff, pp. 106–125. Berlin, Heidelberg, Springer Berlin Heidelberg.
- Dinh, T. T. A., J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan. 2017. “BLOCKBENCH: A Framework for Analyzing Private Blockchains”. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD ’17, pp. 1085–1100. New York, NY, USA, ACM.
- Docker 2018, Aug. “Docker Documentation”. <https://docs.docker.com/>.

- Gervais, A., G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun. 2016. “On the security and performance of proof of work blockchains”. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 3–16. ACM.
- Sneha Goswami 2017. “Scalability analysis of blockchains through blockchain”. Last accessed 29 April 2019.
- Ongaro, Diego and Ousterhout, John 2013. “In search of an understandable consensus algorithm (extended version)”.
- Richard Gendal Brown, James Carlyle, Ian Grigg, Mike Hearn 2016. “Corda: An Introduction”. Last accessed 29 April 2019.
- Segent, N. 1997. “Performance evaluation of a consensus algorithm with Petri nets”. In *Proceedings of the Seventh International Workshop on Petri Nets and Performance Models*, pp. 143–152. IEEE.
- Huawei Technologies 2017. “Caliper: A Blockchain Benchmark framework”. Last accessed 29 April 2019.
- Wang, B., S. Chen, L. Yao, B. Liu, X. Xu, and L. Zhu. 2018. “A Simulation Approach for Studying Behavior and Quality of Blockchain Networks”. In *Blockchain – ICBC 2018*, edited by S. Chen, H. Wang, and L.-J. Zhang, pp. 18–31. Cham, Springer International Publishing.
- Yasaweerasinghelage, R., M. Staples, and I. Weber. 2017, April. “Predicting Latency of Blockchain-Based Systems Using Architectural Modelling and Simulation”. In *2017 IEEE International Conference on Software Architecture (ICSA)*, pp. 253–256.
- Qi Trey Zhong and Zak Cole. 2018. “Analyzing the Effects of Network Latency on Blockchain Performance and Security Using the Whiteblock Testing Platform.”. Last accessed 29 April 2019.

AUTHOR BIOGRAPHIES

PETER FOYTIK is a Senior Project Scientist at Old Dominion University’s Virginia Modeling Analysis and Simulation Center (VMASC). His research interests include modeling and simulation. His email address is pfoytik@odu.edu.

DEEPAK TOSH is an Assistant Professor in the Department of Computer Science at the University of Texas at El Paso. His research interests include cyber security and applied game theory. His email address is dktosht@utep.edu.

SACHIN SHETTY is an Associate Professor in Old Dominion University. His research interests include cyber security, machine learning and the cyber physical security. His email address is sshetty@odu.edu.

SARADA PRASAD GOCHHAYAT is an Adjunct Assistant Professor at Old Dominion University and a Postdoctoral Research Fellow at the Virginia Modeling Analysis and Simulation Center (VMASC). His research interest include formal specification of decentralized distributed systems and cyber security research. His email address is sgochhay@odu.edu

ERANGA HERATH is a Project Scientist at Old Dominion University’s Virginia Modeling Analysis and Simulation Center (VMASC). His research interest include application development of web based applications and decentralized distributed systems. His email is cmjedawer@odu.edu

LAURENT NJILLA is a research electronics engineer at the Cyber Assurance Branch of the U.S. Air Force Research Laboratory in Rome, New York. His email is laurent.njilla@us.af.mil