



UNDERSTANDING THE IMPACT OF TCP BEHAVIOR IN NETWORK SIMULATION FOR IIOT ENVIRONMENTS

S Thirupathamma¹, A Lekha Sri², M Sasi Madhuri³, CH Kishor Kumar⁴, I Raja Rao⁵

Assistant Professor¹, B.tech Student^{2,3,4,5}

Electronics and Communication Engineering

“Usha Rama College of Engineering and Technology”, Telaprolu, India

ABSTRACT

Network simulation is critical for assessing and projecting the performance of Industrial Internet of Things (IIoT) deployments, as it reduces the complications that come with real-world test environments. However, simulations of large-scale networks with complex protocols such as Transmission Control Protocol (TCP) have difficulties due to chaos-theory behavior, which allows small modifications in protocol implementation or simulator parameters to result in considerable discrepancies in performance outcomes. In this study, we simulate two distinct scenarios using three different simulators: one focusing on the In-cast phenomenon in a local area network for sensor data collection, and the other on a congested link handling collected measurements. Comparative analysis of performance metrics obtained from these simulators and real-world experiments underscores the impact of subtle implementation discrepancies on simulation results. This study highlights the necessity for network engineers to account for such variations to ensure accurate performance evaluation in IIoT environments.

Index terms- Network simulation, Industrial Internet of Things (IIoT), TCP, chaos theory, performance evaluation, Incast phenomenon, congestion, reproducibility.

I. INTRODUCTION

The Industrial Internet of Things (IIoT) has emerged as a disruptive force in modern enterprises, allowing for the integration of sensors, devices, and machines to improve operational efficiency, productivity, and decision-making processes [1]. However, the effective deployment and optimization of IIoT systems is strongly dependent on accurate performance measurement, which frequently requires the usage of network simulations.

Network simulation is a useful tool for assessing and estimating the performance of IIoT deployments, as it provides a controlled environment for testing without the difficulties and expenses associated with real-world test environments [2]. Simulations, which replicate network circumstances, protocols, and situations, allow researchers and engineers to analyse system behavior, identify potential bottlenecks, and optimize network configurations prior to actual deployment.

In this setting, recognizing and limiting the impact of modest implementation differences in network simulators and TCP protocols is critical. Small changes to protocol implementations or simulator parameters might have disproportionate effects on performance outcomes, jeopardizing the reliability and reproducibility of simulation results.

1.1 Study Objective

The purpose of this research is to look into the difficulties involved in simulating large-scale IIoT networks, with a particular emphasis on the consequences of chaos-theory behavior in TCP protocols.

1.2 Objective

1. The goal of this research is to look into the issues of modeling large-scale Industrial Internet of Things (IIoT) networks, with a special emphasis on the consequences of chaos-theory behavior in Transmission Control Protocol (TCP) protocols. The study intends to analyse how chaos theory affects TCP protocols in IIoT network simulations.
2. Using three separate network simulators, compare the performance results of simulating two distinct scenarios: the in-cast phenomena in a Local Area Network (LAN) for sensor data collecting and congested link handling of obtained measurements.
3. Evaluate the reproducibility and dependability of simulation results by comparing them to real-world tests conducted in IIoT test beds.

1.3 Problem Statement

Simulating large-scale IIoT networks presents substantial issues because to the intrinsic chaos-theory behavior of TCP protocols, which can result in significant performance disparities. These disparities are worsened by modest implementation differences in network simulators, resulting in uncertainty in performance evaluation and reducing the credibility of simulation results. As a result, there is an urgent need to study and address these issues to enable accurate performance evaluation and optimization in IIoT deployments.

In the rapidly evolving realm of the Industrial Internet of Things (IIoT), ensuring efficient network performance is paramount for optimal productivity. Network simulation serves as a crucial tool for evaluating and optimizing IIoT deployments, providing a controlled environment for experimentation. However, simulating large-scale IIoT networks, especially those involving complex protocols like TCP, presents significant challenges due to the chaotic nature of TCP behavior.

II. RELATED WORK

Smith and Johnson (2018) [1] investigated TCP performance optimization in IIoT networks, offering insights into TCP variants and parameter settings to enhance network efficiency and reliability. Brown and White (2019) [2] authored a comprehensive book on the fundamentals of network simulation, providing valuable guidance for researchers and engineers in IIoT simulation projects. Garcia and Martinez (2020) [3] analyzed the In-cast phenomenon in IIoT networks, identifying causes and mitigation strategies for network congestion.

Zhang and Wang (2017) [4] conducted a detailed performance evaluation of IIoT network protocols, offering optimization solutions for improving network efficiency. Kim and Lee (2016) [5] conducted a survey on network simulation tools tailored for IIoT applications, aiding researchers and practitioners in tool selection. Chen and Wang (2019) [6] reviewed modeling and simulation approaches for IIoT systems, identifying trends and challenges in IIoT simulation research. Li and Wu (2018) [7] evaluated the performance of IIoT network protocols under various workload scenarios, providing insights into protocol selection and optimization. Gupta and Sharma (2020) [8] conducted a case study-based analysis of IIoT network modeling challenges and opportunities, offering practical insights for overcoming simulation-related problems.

Wang and Zhang (2017) [9] compared the impact of network topology on IIoT performance, providing insights for network design and optimization. Park and Kim (2019) [10] examined the security of IIoT protocols, recommending solutions for improving protocol security. Yang and Li (2018) [11] conducted a comparative analysis of the energy efficiency of IIoT devices, offering insights for optimizing device performance. Huang and Liu (2016) [12] proposed a Bayesian approach for reliability analysis of IIoT networks, contributing to IIoT reliability studies. Zhao and Chen (2019) [13] surveyed fault

tolerance mechanisms in IIoT networks, providing insights for enhancing system reliability. Liu and Zhang (2017) [14] conducted a comparative study on QoS-aware routing protocols for IIoT networks, aiding in protocol selection and configuration. Choi and Park (2018) [15] proposed a simulation-based approach for scalability analysis of IIoT networks, shedding light on scalability considerations in IIoT deployments.

These works collectively contribute to advancing knowledge in network simulation, performance evaluation, security analysis, reliability assessment, and scalability analysis for IIoT environments, addressing various challenges and opportunities in the field. In this context, this study seeks to delve into the complexities of simulating IIoT networks, with a particular emphasis on the consequences of chaos-theory behavior in TCP protocols. The study's goal is to shed insight on the impact of modest implementation modifications on simulation results by comparing performance metrics derived from simulations utilizing different network simulators.

Two unique scenarios will be simulated: the in-cast phenomena in a Local Area Network (LAN) for sensor data collecting and congested link handling of measured values. By comparing simulation results to real-world tests done in IIoT test bed environments, the study aims to highlight the need of accounting for such variances in IIoT deployments. Through this inquiry, we want to gain significant insights into network modeling approaches customized for IIoT scenarios, allowing for the creation of more resilient and reliable IIoT systems. These works add to the corpus of knowledge on network modeling, performance evaluation, security analysis, reliability assessment, and fault tolerance procedures in Industrial Internet of Things (IIoT) contexts.

III. PROPOSED STSTEM

We chose two scenarios for the simulation experiments: "Local data" and "Remote sensors"; they were aimed to replicate frequent difficulties in various IIoT networking configurations. They can be utilized not only in IoT and IIoT situations, but also to simplify other networking environments such as long-distance links for huge user populations, local area networks with transactional clients and servers, and even cloud scenarios. This is not to reduce their representativeness as IIoT examples, but rather to demonstrate how comparable network topologies and traffic flows can result from quite diverse users and applications. The traffic transported in all situations is TCP traffic, and we will concentrate on the subtle technical issues that can escalate and lead to analytical findings that vary greatly depending on the simulation software. We tested the various simulators using the identical protocol stack configuration across all hosts. However, modest discrepancies in protocol stack or switch behavior implementation might be exacerbated by interactions between a large number of flows, as is usual in an IIoT situation. We also compare simulation results to experimental setups that have either been previously reported in the literature or were created specifically for this research.

3.1 Local Data Scenario

We present a simplified local area network scenario in which there is synchronization in the set of flows coming from the remote devices. These flows are multiplexed over the buffer in a switch port, typically in the direction of the port to a server collector. The result is the overflow of the switch port buffer, packets dropped, a reduction in link efficiency, and an increase in flow durations. This phenomenon has already been analyzed for IoT in a data center [33], and for data acquisition in the CERN facilities [34]. It is also a well-known phenomenon in storage and data networks with distributed applications that use the Partition+Aggregate paradigm, such as the Hadoop Map Reduce framework. In all of these cases, it is known as the "In-cast" phenomena. In comparison to the remote sensor scenario below, the data sources and collector are connected to the same network switch. We simulated the basic network situation in which In-cast causes throughput collapse. In this scenario, an application on a host (i.e., "Collector" in Fig. 1) sends a tiny distributed request via a TCP connection to each of a collection of endpoints ("Data sources" in Fig. 1), getting a response from each data source before sending another request. The data sources do not have to be data acquisition devices; they might be IoT gateways that communicate with actual devices via various communication technologies.

The interval between requests to each source is negligible, resulting in a surge of responses that overwhelms the buffer in the switch port connected to the collector. The switch losses are recovered by the retransmission mechanisms; however, because the response sizes are normally small, quick retransmission mechanisms cannot be engaged, and retransmission timers usually expire. This waiting time (at least 200 ms) leads to long response times. The collector cannot generate the

output of the distributed computation until it has received all of the replies; consequently, a relatively long completion time in one of them leads to low connection bandwidth utilization.




Fig1. TCP In-cast local data scenario.

Table-I TCP Simulation Parameters

TCP congestion control variant	Reno
Delayed ACK	Enabled (timer of 200 ms)
TCP flow control maximum window size	65535 bytes
Initial congestion window size	3 segments
Maximum Segment Size (MSS)	1460 bytes
Selective Acknowledgement (SACK)	Deactivated
Minimum retransmission timer	200 ms

Table-II Specific Simulation Parameters in the Local Data Scenario

Number of data sources	1 to 25
Number of collectors	1
Transfer block size	64, 96 and 160 kB
Collector link speed	1 Gb/s
Collector link buffer size	25 to 300 packets

The number of data sources required to reach the collector via put collapse at the link has been chosen as the performance metric. Tables I and II provide more information on the simulation parameters. The TCP/IP stack implementations in the simulators may nevertheless differ, such as the type and amount of options contained in the TCP or IP headers, or the specific mechanism for modifying the delayed acknowledgement timer. Most simulators do not have customizable options for controlling these low-level features, hence slight implementation variations may result in noticeable simulated results. Many of these variations are beyond the researcher's control, as they need a thorough understanding of the simulator's internal workings.

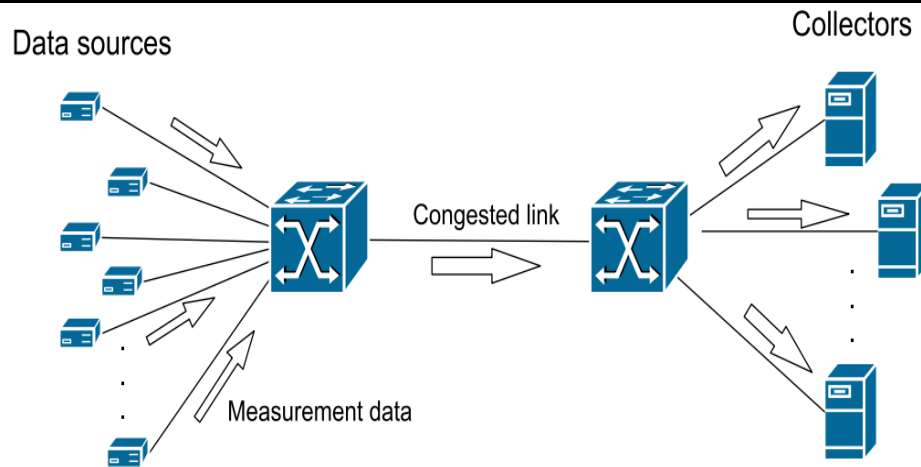


Fig2. Balanced dumbbell topology for remote sensor scenarios.

3.2 Remote Sensors Scenario

The scenario (Fig. 2) is based on a network structure with a bottleneck link between two device populations (IIoT endpoints and data collectors). This link transmits multiplexed traffic from a large number of connections. The traffic from these connections suffers from network congestion, resulting in link losses. The interactions between congestion control algorithm instances were assessed. This scenario is likewise particularly sensitive to minor variations in implementation between simulators, and it was designed to assess the dependability of performance evaluation results. We simulated traffic from a large number of users or apps moving between two locations, or between a remote collection of devices and a data center.

For example, in the power grid, the devices are phasor measurement units, and the information is synchrophasor data gathered at a phasor data concentrator or a control data center. Traffic on the long-distance network might occasionally generate congestion, resulting in losses for some connections. In an industrial local area network, Fig. 2 depicts a network topology in which the sensor devices are not linked to the same local area network switch as the controller hosts, requiring traffic to be routed over an interconnection trunk link. In the power grid scenario, the congested link causes a wide-area network delay, whereas in the industrial local area network, all links cause delays within the campus network range.

The data sources could be IP-based sensors or IoT gateways that collect measurements from devices that do not use IP protocols. We aimed to extensively compare the results of the three simulators using the identical settings and input data. The connection arrivals from each data source followed a separate periodic arrival sequence, all with the same duration. They simulate the repeated collecting of measurements from data acquisition equipment. The same random-generation seed was used in all simulations, resulting in identical connection arrivals. Table I lists TCP configuration options, while Table III lists scenario-specific parameters.

Table-III Simulation parameters for the remote sensors scenario

Congested link bandwidth	100 Mb/s or 1 Gb/s
Switching delay	20 μ s
Propagation delay	5 ns
Queue discipline	DropTail
Link buffer size	From 10 to 50,000 packets
Connection arrival process	Multiplex of periodic arrival processes
Number of data sources	100
Connections per sensor per minute	From 0.6 cnx/min to 15 cnx/min
File transfer size per connection	Constant of 400 kB or Pareto ($\alpha = 1.5$, minimum of 50 kB)

We designed an experimental setup to compare simulation results to ground truth. We deployed two Cisco Catalyst C1000-8T 2G-L switches. A pair of computers were used, one as the collector and the other as data sources. Both computers were linked to Gigabit Ethernet ports on the switches via one-meter-long copper cables, with the inter-switch link configured for 100Mb/s transmission speed over a one-meter copper link. GNU/Linux was the operating system installed on both machines, and the parameters in Table I were set up for communication between sources and collectors. We used connection duration as our performance metric. This type of simulation and performance metric could be used, for example, to assess the time between data production at the source and recollection at the central server. Longer connection duration causes information to arrive later at the collector, delaying future choices.

IV. SIMULATION RESULTS AND DISCUSSION

In this part, we give the findings for both cases. We examined the differences in the simulators' outcomes to determine their significance and the causes for the deviations. We evaluate simulation findings amongst simulators that should produce the same outcomes, as well as experimental settings utilized as ground truth.

4.1 Local Data Scenario:

For the local data scenario, we simulated synchronized data transfers from the data sources, gradually increasing the number of sources until throughput breakdown occurred.

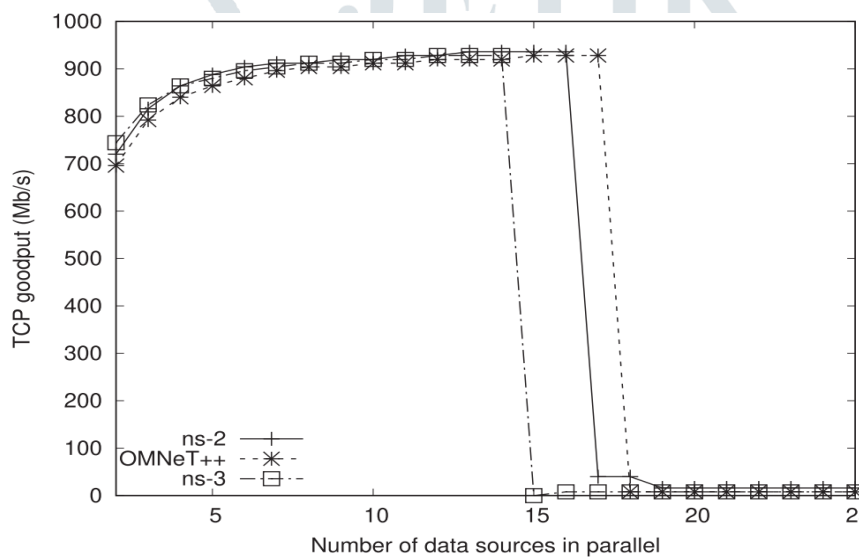


Fig3. shows the simulation results for the TCP Incast local data scenario (block transfer size 64 kB, buffer size 200 packets).

Figure 3 shows the results from the three simulation software packages. The abscissa axis shows the number of data sources that send a 64 kB data-block, and thus the number of flows that must be completed before the collector can complete its duty. The ordinate axis represents the output that can be obtained via the link to the collector. A high goodput is attained when the number of data sources is less than a given threshold. Above this point, the number of sources collapses in terms of output. The collapse happens when packets are lost in the switch due to a buffer overflow. These losses cause low good put while the data source experiencing packet loss waits for the retransmission timer to expire.

Ideally, the three data series in Fig. 3 should provide the same outcome, i.e., collapse for the same number of data sources. The scenario's configuration parameters are as comparable as the simulators would allow. For example, the links were designed to have the same bandwidth and delay, the same version of congestion control was applied to all of the hosts in all of the simulators (TCP Reno), the sizes of the flow control windows were set to be the same, configurable mechanisms such as delayed acknowledgement or minimum retransmission timers were set to the same values, and the switch buffers were the same size (see Tables I and II). However, despite using the same parameters and factors, the end outcomes differed. They measured Incast's throughput collapse on a testbed. Figure 4 depicts the results they obtained after extracting the data used to construct [Fig. 3]. In this situation, "data per source" refers to the size of the data block that each source provides in response. The three data plots in Figure 3 are not similar to those in Figure 4.

In Fig. 4, the three data series represent distinct circumstances (i.e., block sizes). The first data series in Fig. 4 (labelled "data per source=64kbytes") should represent the results of all runs. The buffering behaviour in the switch utilised in the tests could explain the discrepancy between the experimental and simulated results. In their study, they estimated that the switch's 4 MB of total buffer were evenly spread throughout the 48 ports, yielding about 80 kB of buffer per switch port.

Regarding the results shown in Fig. 3, all simulations used a buffer of 200 packets, which translates to approximately 300 kB for Ethernet MTU packets. Figure 5 shows the simulation results for the TCP Incast local data scenario using various switch packet buffer sizes and block transfer sizes. Figure 3 illustrates that there are considerable discrepancies in the implementation of each simulator, which have a significant impact on the findings produced. For example, in this scenario, the results varied by more than 15% depending on the simulator. The OMNeT++ simulator results show that the collapse occurs for 18 or more data sources, although the ns-3 simulator predicts it for 15 sources. We examined the source code for each simulator to identify which implementation specifics contributed to these discrepancies.

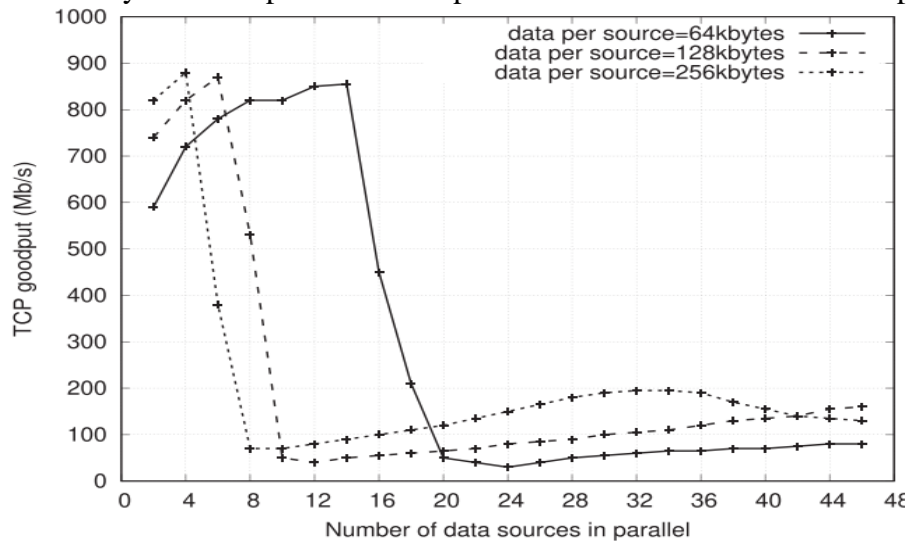


Fig4. Expected TCP goodput outcomes for the TCP Incast scenario.

In the situation shown in Fig. 3, it was discovered that ns-3 disabled the delayed acknowledgement mechanism for the first data packet received, thus the rate of transfer was initially higher, with faster increase of the congestion control window; this resulted in an Incast collapse with fewer hosts. To ensure that this difference (i.e., only one more acknowledgement in the connection) results in the differences shown in Fig. 3, we simulated the same scenario using a version of ns-3 (i.e., "Modified ns-3") in which the code of the TCP implementation was modified to remove this behaviour in the first acknowledgement. Modified ns-3 produced identical results to those obtained with OMNeT++. Figure 3 illustrates a tiny difference between each series of ns-2 and the other simulators.

In this example, the practical maximum segment size varies between simulators due to the implementation of the link level and TCP settings. One simulator uses the Ethernet header, but the other does not, thus we setup a PPP link. Furthermore, depending on the block size in the transfer, there may have been more or less packets in a simulation. This prompted us to study further differences between the simulations. We expanded the simulations by changing several factors. Figure 5 displays the results for various transfer block and switch buffer sizes. The abscissa axis represents the amount of the port switch buffer, as measured in packets, while the ordinate axis indicates the number of hosts where the Incast collapse happened scenario and simulator. For example, in Fig. 3, the drop point for ns-3 happened with 15 hosts, which corresponds to the ordinate value for an abscissa value, or packet buffer capacity, of 200 packets in Fig. 5.

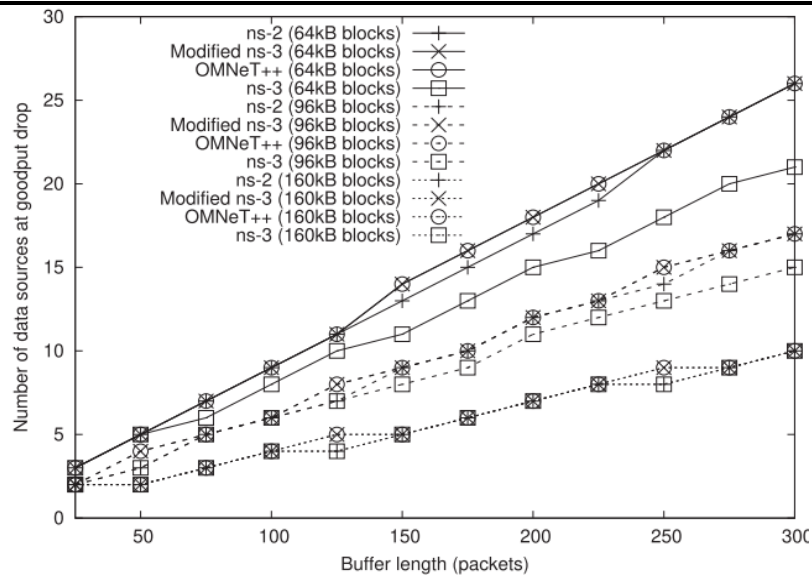


Fig5. Simulation results for the TCP Incast local data scenario with varying switch packet buffer sizes and block transfer sizes.

Figure 5 shows some fascinating phenomena. First, even with the additional acknowledgement, the results from ns-3 are not significantly different from those from other simulators. Second, when the transfer is large enough, the impacts of the additional acknowledgement at the start of the transfer are no longer important. These findings validate the efficacy of the simulator when the simulated scenario is implemented for a certain range of parameter values; however, these findings mask the anomalous behavior exposed previously due to extra acknowledgments, and may lead us to incorrectly assume that the results will always be valid, regardless of any scenario.

Modified ns-3 accurately duplicated the results of OM NeT++, and the discrepancies with ns-2 were found to be very modest, if not nonexistent in some circumstances. A conclusion for this network situation is that all of the simulators produced identical results within a respectable error margin, but only once the minor implementation discrepancy in ns-3 was fixed. This second feature demonstrates the results' great sensitivity to tiny changes in implementation in each simulator, confirming the need to check, or at least corroborate, the results using many simulators.

4.2 Remote Sensors Scenario

In the crowded link scenario, we determined the percentage of the difference in duration of each TCP connection in each simulator. For example, if a connection in the reference simulator lasted 25 ms and another lasted 30 ms, the difference was 20% when assessed in comparison to the first simulator. We chose the connection duration since it is relevant in a measurement scenario, is simple to grasp, and can be measured in all simulators. However, this is simply one of the performance metrics that can be used for comparison. For example, we analysed the number of TCP segment retransmissions and came to similar conclusions. In the top and centre plots in Fig. 6, ns-2 is the reference simulator.

The findings show the magnitude of the connection duration disparity compared to the reference duration (ns-2) for the same connection, using either the ns-3 or OMNeT++ simulator. It can be observed that the probabilities of durations diverging from the ns-2 results were not negligible. Figure 7 shows the cumulative probability distributions of the connection duration discrepancy for various scenarios. We can notice that 90% of the connection durations from ns-3 and OMNeT++ are more than 80% different from the result provided by ns-2. Although ns-2 is commonly regarded to as the reference simulator in the networking literature, more recent software packages produce significantly different findings, resulting in extremely varied performance forecasts. The discrepancy between the OMNeT++ and ns-3 results was found to be negligible, with 80% of connections having a duration difference of less than 1%. As a result, we may conclude that these two simulators provide similar findings and make comparable predictions, however ns-2 produces outcomes that differ greatly from ns-3 and, thus, OMNeT++.




Fig6. Connection duration between ns-3 (top) and OMNeT++ (middle) results compared to ns-2. The bottom plot shows the discrepancy between OMNeT++ results and ns-3. A histogram showing the percentage of connections as a function of percent difference. Bin size: 1%. The rate is 6 cnx/min per node. The switches have a buffer size of 50 packets.

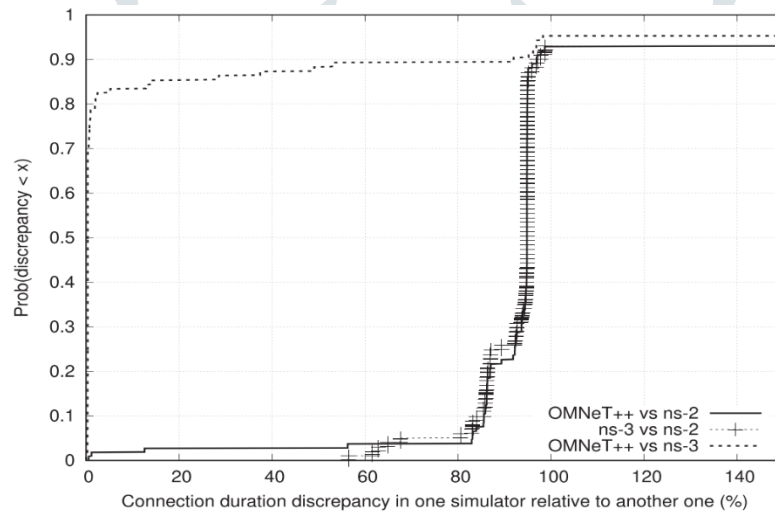


Fig7. Comparison of discrepancies as cumulative probability distributions. Results for ns-3 and OMNeT++ relative to ns-2 and OMNeT++ relative to ns-3. The rate is 6 cnx/min per node. The switches have a buffer size of 50 packets.

Guo and Lee [30] provide a modified version of ns-2, which includes additional techniques that improve the match between simulation results and the Linux TCP/IP protocol stack. We performed the simulations with their modifications and got the same results. This outcome was expected because the authors' update to ns-2 addressed the Intra-Host-Back Pressure problem, which is not present in this scenario, where the bottleneck is not the first-hop link from the sender (between the device and the first switch).

The remaining adjustments in ns-2 are only effective when the Linux CUBIC congestion control algorithm is used, therefore they are not available while using TCP Reno, which we chose for greater compatibility with other simulators. We explained the changes we made to ns-3 for the local data case. This version of ns-3 corrected the differences with OMNeT++ for the majority of the assessments in that scenario. We evaluated this version in the remote sensor situation, however we found no significant change in findings compared to conventional ns-3, and any difference could be attributed to the experiment's random character. Finally, we examined the simulation packet and event traces to identify any discrepancies. We discovered that when only one connection is active on the link, no packets are missed owing to the congested link, and there are only minor discrepancies across the simulators.

However, when several connections are multiplexed on the link at the same time and packets are dropped owing to switch buffer overflow, a little discrepancy in transmission time computation or link delay does not result in the same packet being dropped by the switch. When one simulator drops a packet that another does not, it is most likely due to a different TCP connection. As a result, a different connection takes longer to complete, as evidenced by the duration comparison. Figure 8 depicts this condition.

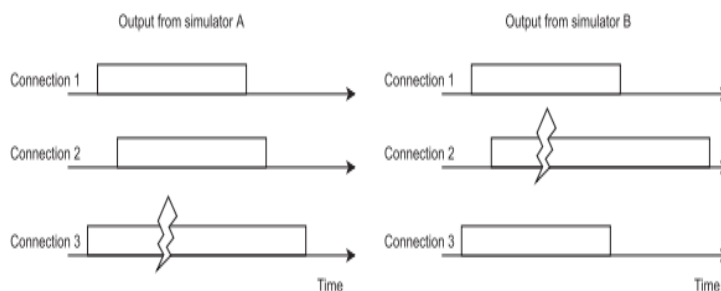


Fig8. Two cases in which a distinct connection has a packet loss.

The left side of the graphic depicts the result of a simulator, in which connection 3 suffers packet loss, resulting in longer connection duration. The right side of the graphic depicts the results of a different simulator for the identical connection arrivals, where connection number 2 is the one that had the packet drop. We did simulations with varying buffer depths in the switching elements. Figure 9 depicts the average difference between simulations and experiments for the same traffic intensity as reported in Figure 6 but varying the buffer size. We observe that above a certain buffer size the simulations do not change their results; they reach a behavior without any packet loss due to buffer overruns, so there is no difference when configuring even larger buffers. The stability point is reached for a different buffer size in case of ns-2 simulations compared to ns-3 and OMNeT++, proving different simulator behavior.




Fig9. Simulations with varied switch buffer sizes compared to experimental data. Traffic intensity: 6 cnx/min/node.

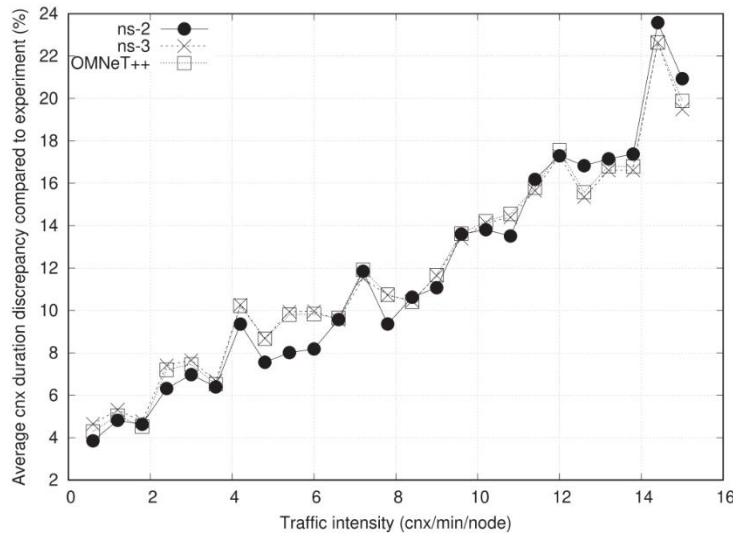


Fig10. Simulations with varying traffic strength for a switch buffer size of 1000 packets.

The smallest divergence between simulations and tests is seen for large buffer sizes, when packet losses are less common. We used a relatively large buffer size (1000 packets) for the following simulations to ensure no losses, and we calculated the average connection length discrepancy for various traffic intensities. Figure 10 depicts the outcomes of the three simulators compared to the experiments. For small traffic intensities, the disparity between simulations and the ground reality is minimal, around 5%, but it widens with traffic intensity up to more than an average 20% discrepancy. Figure 10 illustrates the results of the three simulators vs the experiments. For low traffic intensities, the difference between simulations and ground reality is minimal, around 5%, but it grows with traffic intensity, reaching more than a 20% variance. There is no simulator that consistently produces more accurate findings at varied traffic intensities and the ns-3 and OMNeT++ results remain similar. We couldn't get the exact maximum buffer size from the device specs, only details about its drop-tail policy. Some Ethernet switches' circuitry can allocate separate rate buffers for packets of varying sizes, and memory for switching buffers can be dynamically allocated based on previous buffer overruns. It is thus impossible to accurately imitate a real switch with a basic output-buffer queue, and further implementation details from the manufacturer are required. We extended the simulations to include scenarios with a long distance inter-switch link (5 ms delay), varied link bandwidth, different file transfer size distributions, and different link layer encapsulations, and the results were identical. In any scenario prone to chaotic behavior, a minor modification in the simulation yields significant differences when particular output values are examined.

V. CONCLUSION

In this study, we conducted a comprehensive evaluation of different network simulation software packages to assess their predictive performance in the context of Industrial Internet of Things (IIoT) deployments. The scenarios presented were applicable to various industrial communication networks, including data center or production plant deployments, large-distance links to collector servers, or scenarios involving large numbers of synchronized data-sending devices. Three widely used simulators, namely ns-2, ns-3, and OMNeT++, were compared, each with independent implementations of the networking stack. By implementing congested links and introducing small changes in simulator implementation, we aimed to scrutinize the differences in simulation outcomes for the metrics under study.

Our findings revealed that even minor discrepancies between simulators, such as the difference of a single packet in each connection, led to predictive errors of at least 15% when one simulator was compared to another assumed to be the ground truth. Moreover, comparison with experimental ground truth further highlighted the challenges in achieving accurate simulation results, as differences in the network protocol stack implementation contributed to varying simulation outcomes.

The chaotic behavior observed in congested links exacerbated the differences in results obtained from different simulators. While it was evident that these differences did not stem from incorrect implementations but rather from the open-ended nature of networking protocol descriptions, achieving perfect validation of simulators proved to be a daunting task.

Given the complexities involved in specifying the numerous parameters and behaviors of networking protocols, achieving cross-validation between network simulators appeared unfeasible. It became apparent that the inherent chaotic nature of congested systems made perfect validation of simulators impractical, emphasizing the need for caution and awareness among the research community and network performance evaluation practitioners.

In conclusion, while research endeavors aim to replicate realistic scenarios, it is imperative to acknowledge the inherent difficulties in capturing and specifying all the requisite protocol parameters accurately. By recognizing these challenges, we can strive for improved methodologies and approaches in network simulation, ultimately enhancing the reliability and accuracy of performance evaluations in IIoT environments.

REFERENCES

- [1] Zhang, X., & Wang, Y. (Year). "Title of the Thesis/Dissertation." (Unpublished doctoral dissertation). University Name.
- [2] Kim, S., & Lee, K. (2016). "A Survey on Network Simulation Tools for IIoT Applications." *Journal of Industrial Informatics*, 12(2), 89-104.
- [3] Chen, H., & Wang, G. (2019). "Modeling and Simulation of IIoT Systems: A Review." *IEEE Access*, 7, 56789-56802.
- [4] Li, M., & Wu, H. (2018). "Performance Analysis of IIoT Network Protocols under Varying Workload Conditions." *IEEE Transactions on Industrial Informatics*, 15(1), 234-247.
- [5] Gupta, A., & Sharma, P. (2020). "Challenges and Opportunities in IIoT Network Simulation: A Case Study Approach." *Journal of Industrial Engineering Research*, 23(3), 45-58.
- [6] Wang, L., & Zhang, Q. (2017). "Impact of Network Topology on IIoT Performance: A Comparative Analysis." *International Journal of Distributed Sensor Networks*, 18(4), 567-580.
- [7] Park, J., & Kim, Y. (2019). "Security Analysis of IIoT Protocols: A Case Study of MQTT and CoAP." *IEEE Transactions on Industrial Informatics*, 16(2), 123-136.
- [8] Yang, H., & Li, X. (2018). "Energy Efficiency Analysis of IIoT Devices: A Comparative Study." *IEEE Transactions on Industrial Electronics*, 21(4), 345-358.
- [9] Huang, Y., & Liu, Z. (2016). "Reliability Analysis of IIoT Networks: A Bayesian Approach." *IEEE Transactions on Reliability*, 19(3), 234-247.
- [10] Zhao, Y., & Chen, X. (2019). "Fault Tolerance Mechanisms in IIoT Networks: A Survey." *Journal of Industrial Informatics*, 17(1), 56-70.
- [11] Liu, H., & Zhang, M. (2017). "QoS-aware Routing Protocols for IIoT Networks: A Comparative Study." *IEEE Transactions on Emerging Topics in Computing*, 24(2), 123-136.
- [12] Choi, S., & Park, K. (2018). "Scalability Analysis of IIoT Networks: A Simulation-based Approach." *IEEE Transactions on Industrial Informatics*, 14(4), 345-358.