

A Graph Database Approach to Wireless IIoT Workcell Performance Evaluation

Richard Candell*, Mohamed Kashef*, Yongkang Liu*, Karl Montgomery*, Sebti Foufou†

*National Institute of Standards and Technology, Gaithersburg, Maryland, USA

Email: {richard.candell, mohamed.kashef, yongkang.liu, karl.montgomery}@nist.gov

† University of Burgundy, Dijon, France, Email: sfoufou@u-bourgogne.fr

Abstract—The workcell is considered a main building block of various industrial settings. Hence, it is examined as a primary testing environment for studying wireless communication techniques in factory automation processes. A new testbed was recently designed and developed to facilitate such studies in workcells by replicating various data flows in an emulated production environment. In this paper, an approach to storing and analyzing network performance data from a manufacturing factory workcell is introduced. A robotic testbed was constructed using two collaborative grade robot arms, machine emulators, and wireless communication devices. A graph database approach was implemented to capture network and operational event data among the components within the testbed. A schema is proposed, developed, and elaborated; a database is then populated with events from the testbed, and the resulting graph is presented. Query commands are then presented as a means to examine and analyze network performance and relationships within the components of the network. Additionally, we demonstrate how to extract correlations between receive signal power and network delay within the testbed using the graph database query language. Finally, using the inherently interconnected nature of the graph database, we discuss applying the graph database approach toward examining more complex relationships between the wireless communications network and the operational system.

Index Terms—industrial wireless, factory automation, testbed, measurement, instrumentation, graph database.

I. INTRODUCTION

Wireless communication is a key enabling technology for the modernization of factory workcells. Modern factory workcells are highly-interconnected networked control systems in which various devices interact and collaborate to accomplish complex and adaptive production orders. The workcell often contains mobile robots that collaborate with other robots or human beings. Requirements of the workcell include the incorporation of mobile collaborative robotics with real-time coordination of motion and tool actuation. As such, the communication of sensor and control information must be ultra-reliable and of low-latency to assure trustworthy operation [1]. Due to an increased demand for ease of installation, reduced costs of deployment and maintenance, and flexibility, wired networks are being gradually replaced with wireless networks. This presents a real challenge for networks and control systems. Compared with wired connections, wireless links have their unique advantages in connecting field sensors and actuators with reduced cabling costs and natural support of mobility [2]; however, most current communications systems

lack the latency and reliability supports [3] mandated by factory owners [4], [5]. New wireless protocols are being designed to address reliability and latency concerns of real-time systems such as those in manufacturing automation. These new protocols include advancements, such as the Institute of Electrical and Electronics Engineers (IEEE) 802.11ax standard and 5G cellular evolution as defined by the International Mobile Telecommunications-2020 (IMT-2020) Standard. Both of these two standards employ improved diversity techniques for multiple access of devices as well as 1 ms latency and greatly improved reliability [6].

Evaluation of such systems used in manufacturing environments requires not only rigorously analyzing network performance, but also studying the impacts of networks on manufacturing systems. Additionally, industry lacks effective and easy-to-use strategies for test and evaluation of such systems in a way that correlates network performance with operational performance. Furthermore, since factory operators desire the ability to control operations within the workcell using wireless technology, we present a novel method to simultaneously capture network and operational event information using a graph database (GDB). The use of a GDB allows for more intuitive inferences to be made through the stored relationships and graph theoretic models [7]. In this paper, we present a GDB approach to the capture and analysis of factory workcell performance utilizing the Neo4j database platform. We present a proposed schema of the database, the process for capturing both network and operational events, and examples of querying the database for cyber-physical performance evaluation of the workcell for our collaborative two-robot machine-tending workcell [8].

II. GRAPH DATABASES

A GDB is the type of databases that uses nodes, edges, and properties to store and present data. A GDB is a part of a family of databases known as NoSQL databases that are used often to represent complex interrelated structures of data and their relationships. This can be very difficult with traditional relational databases. The GDB places a high priority on the relationships (edges) between units of information. In addition, the GDB does not enforce any particular schema or structure, and, therefore, provides greater flexibility in storing and representing information in which the parts of information

may vary among units. The relationships within a GDB are efficiently queried because they are persistently stored within the database. In a GDB, queries can be made based on relationships. This, in particular, presents an advantage when storing information regarding systems with correlations that are apparent but difficult to visualize or quantify. The GDB approach was chosen for this specific purpose—to quantify and visualize relationships between non-ideal communications within a workcell and its impact on the physical system.

III. RELATED WORK

Multiple surveys about GDBs have been presented to describe the associated models, tools, and their features in [7], [9], [10]. Also, examples of applications and implementations of GDBs are presented in [11] to show their use on enterprise data, social networks, and determining security and access rights. It was found that GDBs provide the much needed structure for storing data and incorporating a dynamic schema. On the other hand, query languages are used to extract data including traversing the database, comparing nodes properties, and subgraph matching [12]. The performance of different GDB tools and methodologies is analyzed and compared in [13], [14]. Multiple comparisons in these articles have shown improved performance of Neo4j in the general features for data storing and querying, and data modeling features such as data structures, query languages and integrity constraints.

Furthermore, industrial data analytics play an essential role in achieving the smart factory vision and improving decision-making in various industrial applications. Five main industrial data methodologies are studied including highly distributed data ingestion, data repository, large-scale data management, data analytics, and data governance [15]. Industrial data processing offers valuable information about various sections of industrial applications including inefficiencies in industrial processes, costly failures and down-times, and effective maintenance decisions [16]. In [17], a platform for performing industrial big data analysis is presented where the performance requirements are introduced to achieve a cost-effective operation. Various other frameworks for industrial data analysis can be found in [18], [19], where the importance of using data analysis in decision making is emphasized.

Due to its advantages including scalability, efficiency, and flexibility, NoSQL databases are a popular alternative to relational databases in the case of large amounts of data in various applications [20]. The GDB is a kind of NoSQL database approaches that additionally handles complex relationships [21]. GDBs are widely adopted in various industry-related applications and use cases such as network operations, fraud detection, and asset and data management [22]. Relationships in social networks have been modeled using a GDB for structural information mining and marketing [23]. On the other hand, business solutions for scenarios with multiple large data sources require distributed processing in decision making for various problems such as fraud detection, trend prediction, and product recommendation [24].

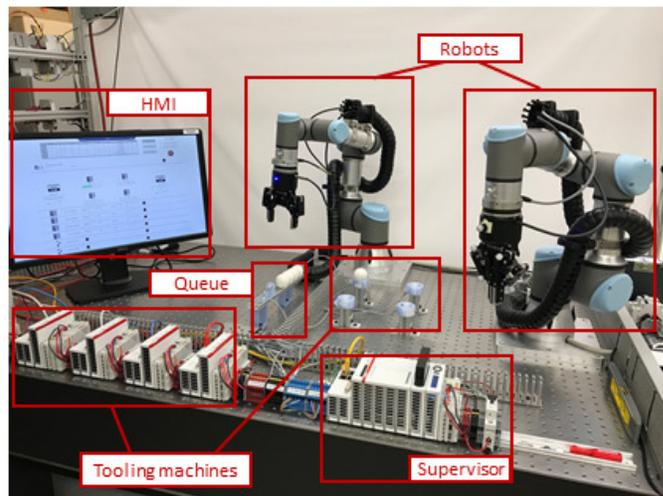


Fig. 1: Collaborative work cell testbed

IV. CASE STUDY: ROBOTIC MACHINE-TENDING

In this section, we present a two-robot machine tending workcell case study. We first present the design of the workcell followed by an elaboration of the database design called a schema. The information work-flow, i.e., the process for collecting, processing, and analyzing the network event data is presented. We then provide examples of the resulting graph and results of targeted queries that demonstrate the purpose of the database.

A. Workcell Design

To facilitate wireless network research and showcase the power of wireless technologies in industrial practices, a testbed shown in Fig. 1 has been developed at the National Institute of Standards and Technology (NIST) as described in [8]. The testbed is composed of two collaboration-grade robots, a supervisory programmable logic controller (PLC) used for the control of the workcell, four smaller PLCs serving as computer numerical control (CNC) machine emulators, and a human-machine interface (HMI). Each robot is equipped with a six-degrees-of-freedom (DOF) force-torque (FT) sensor and a two-finger gripper. A Modbus/TCP server is included within the supervisor PLC and is used for communication between the supervisor and the robots. The PLCs themselves communicate to each other using the Beckhoff Automation Device Specification (ADS) protocol. All elements within the workcell are synchronized to a stable and accurate grand-master clock. Therefore, as described in [8], the operational, network, and measurement elements are all synchronized to the grand-master clock through a precision time protocol (PTP)-capable switch.

Work-orders for the workcell are submitted through the supervisor. Each work-order consists of a work-plan for a part, and the work-plan determines how each part moves through the workcell until it is completed. The inspection of each part is conducted at each machining station, and after the final inspection, the part is placed back into the input queue. Under

normal operating conditions, the work continues until all work-orders have been processed. This continuous form of operation provides ample opportunity to collect statistically significant metrics of both the network and the operation of the workcell.

B. Database Schema

GDBs are NoSQL databases such that the database does not contain any predefined structure or rules to enforce such structure. This is a major difference between relational databases and GDBs. Nevertheless, it was necessary to sketch a pseudo-schema to capture the intended nodes and relationships that would be stored within the database (the terms pseudo-schema and schema will be used interchangeably). Before describing the schema itself, it is necessary to first explain the requirements of the schema. Therefore, the requirements for our schema are as follows:

Time Any manufacturing automation system is indeed a time-varying control system with network and operational events. The database schema must necessarily support time-based queries and, specifically, time-windowed queries.

Operational Events The schema must represent operational events such as the movement of a robot arm or the movement of a part.

Network Events The schema must represent network events such as the transmission of packets.

Message Grouping The schema must support grouping of logically related events such that those events can be correlated to a specific occurrence within the testbed.

Wireless Support The schema must support the capture of both wired and wireless network traffic without special provisions within queries for either.

QoS Support The schema must allow for the capture of quality of service (QoS) data when available.

Spectrum Monitoring The schema must support the capture and association of network events with observations from a spectrum monitoring system (SMS) if that information is available.

1) *Node Design*: Given the fore-mentioned requirements, a sample pseudo-schema is shown in Fig. 2, which represents the intended structure of the information within the GDB. It is important to remember that since GDB schemas are not really schemas, such as those found in relational databases, but representations of intent, the schema represented here should be considered a notional example of the final product. Within the graphs, nodes represent logical elements, and edges represent the relationships between those elements. Both nodes and edges may contain properties providing more description and labels that define categories or classes of the said nodes or edges. Our schema is designed such that the data within the graph is intuitive to understand and allows for time-based queries to occur. The facilitation of time-based queries was an essential requirement of our database design. Our schema is represented using the following node labels:

Actor A physical component within the factory workcell such as a robot, PLC, or other networked item.

NtwkID A network address item for an actor such as an Internet Protocol (IP) address.

Transaction A complete information exchange between two or more actors (multiple actors may participate in a transaction).

Message A network transmission event that occurs between two actors (messages are essentially packet transmissions captured at the transport layer; multiple messages support a transaction).

Physical Action A physical occurrence within a factory workcell associated with actors through multiple time based relationships.

SMS An SMS observes and records significant spectral events within the workcell and may report those events to actors within the workcell.

Sniffer Measurement device that records all transmissions conducted over the wireless medium and includes wireless header information for each wireless transmission detected.

Adapter Device that serves to connect an actor to a network (adapters are divided into sub-categories depending on the type of interface to a network).

Adapter:Ethernet A subcategory of adapter representing an Ethernet interface.

Adapter:Wireless A subcategory of adapter representing a wireless interface.

Adapter:Wireless:AP A subcategory of adapter representing a wireless access point interface.

Adapter:Wireless:UE A subcategory of adapter representing a wireless user equipment interface.

QoS Report Quality of service report of a message (not all messages will have a QoS report).

It is important to note that most network infrastructure components are not captured within the graph, but, instead, only basic interfaces between actors and the network are captured. Our intent when designing the graph was to make the graph network and protocol agnostic, such that the network is viewed as a black-box. Accordingly, the captured events of the physical system and the network are considered useful for the analysis of performance.

2) *Relationships (Graph Edges)*: Relationships are edges within the graph that capture the informational interactions between nodes. Relationships, like nodes, can contain labels and properties. As shown in Fig. 2, nodes are connected through defined relationships. A subset of the relationships are defined as follows:

PARTICIPATED_IN Actors will participate in transactions. A transaction exists for each logical set of messages between actors such as the setting of a Modbus register or the sending of a command to a robot. Therefore, actors will participate in many transactions, and multiple actors may participate in a single transaction.

SUPPORTED Messages (i.e., packets between actors) are associated with transactions through the SUPPORTED relationship. Depending on the protocol and the quality

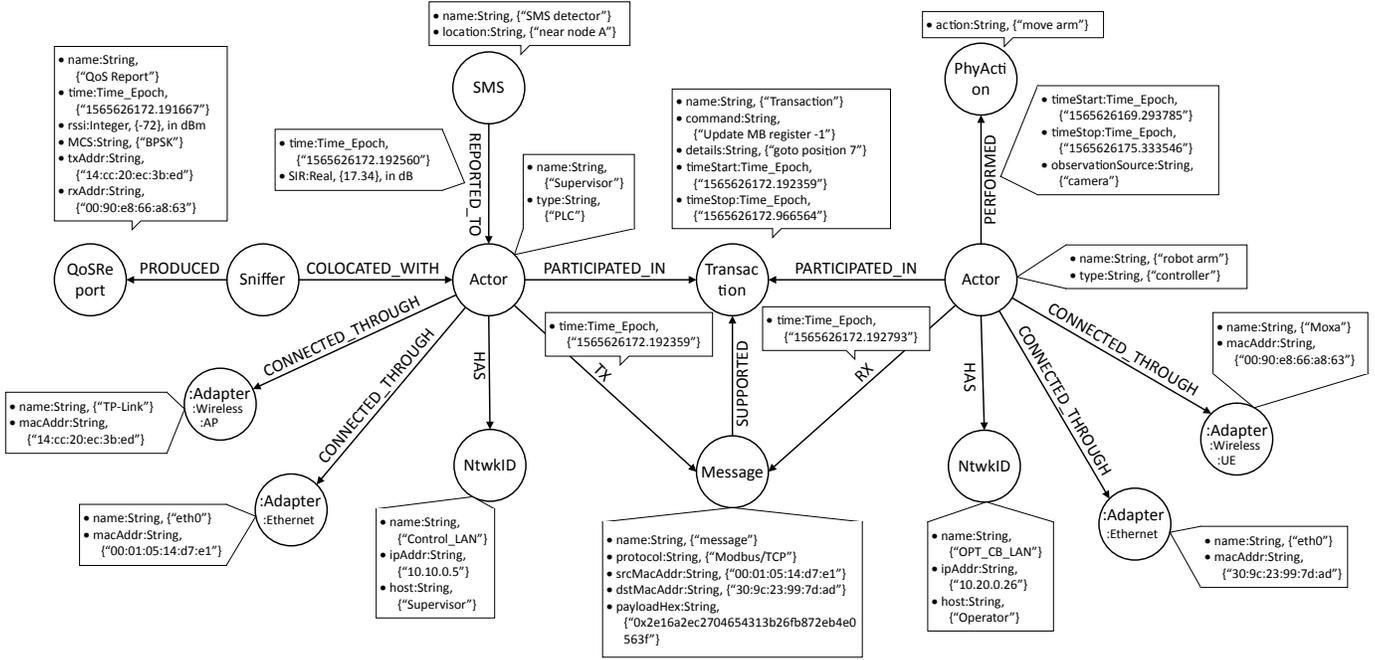


Fig. 2: The intended schema (i.e., pseudo-schema) of the graph database used for each operational run of the NIST wireless factory testbed. The graph is organized into nodes and edges, where the edges signify relationships among network elements and physical operational elements.

of the channel, a single transaction could have one or many messages connected through this relationship.

TX/RX An actor may either transmit (TX) or receive (RX) a message. Both the TX and RX relationships contain a timestamp in the format of an epoch time which is a floating point number in seconds since January 1, 1970, with a resolution of microseconds.

PERFORMED When an actor performs a physical action, a relationship is created between the actor and the physical action node. This relationship contains start and stop time properties as well as the source of the observation such as a networked camera.

REPORTED_TO An SMS may be a passive or active listener within a workcell. When an SMS operates as an active listener, spectral reports from the SMS may be sent to an actor such that the actor can respond intelligently to the spectral event. Reports from an SMS to an actor are captured within this relationship.

Other relationships shown in the schema of Fig. 2 but not explained above are considered self-explanatory.

3) *Closer Examination*: Examining the sample schema more closely, two actor nodes are represented. In this case, Actor A is the supervisory controller, and Actor B is a robot arm. Both nodes participate in a transaction, which, in this example, is a Modbus/TCP exchange. The transaction itself is associated with one or more messages (i.e., packets). Each message associated with a transaction manifests itself as a node in the graph. Multiple message nodes will exist for each transaction. Additionally, QoS reports may be associated with each actor node through a collocated sniffer node. By keeping

QoS reports separate, we have the flexibility of supporting different wired and wireless protocols within the same graph. Recall, that a graph database has no enforceable structure and thus affords this type of flexibility. Each actor node may have network identities associated with it through the use of network identifier nodes. Each network identifier node may contain address information such as a hardware address or a network address; however, this is dependent on the protocols and addressing schemes being used, and a node could have many different identification nodes.

4) *Physical Actions*: Finally, each actor in the graph may associate with physical actions. These actions exist in the database as automations such that every time a new action occurs, a new edge would be added between the actor and the physical action. Timestamps within the graph represent “measurement time” denoting that all timestamps are accurately synchronized to the grand-master clock. The method of synchronization is outside the scope of this paper and is explained fully in [8]. This paper describes the database structure and the process for preparing and inserting the data into the database, which is described in the following subsection.

C. Information Workflow

The workflow for collecting, processing, and inserting measurement data into the graph database is multi-pronged, as shown in Fig. 3. The workcell, which in our case study is a two-robot workcell with four CNC machines, is instrumented with network probes that capture transmitted and received packets as well as probes to capture operational movements

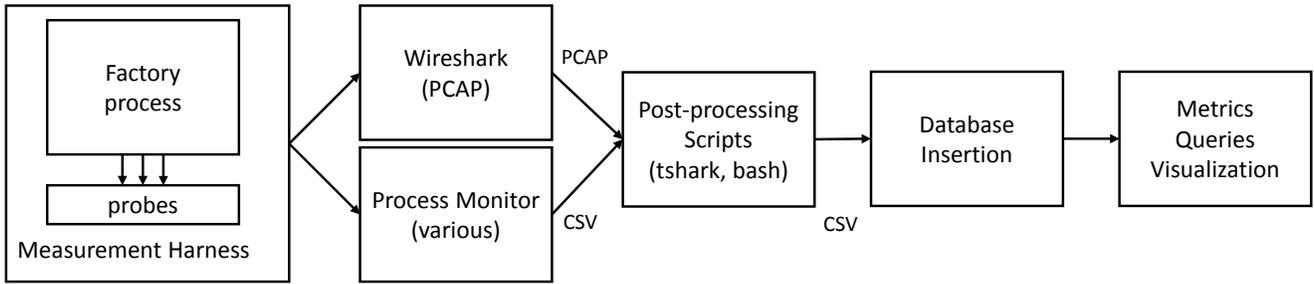


Fig. 3: Data processing flow from factory workcell to database

such as arm robot position and the state of the supervisor PLC. Network data is stored in packet capture (PCAP) files, while operational data is stored in comma separated value (CSV) files. We developed bash scripts to extract relevant information and prepare the data for insertion into the database. The scripts also contain rules for grouping packets together as transactions based on the application protocol and time. The scripts produce CSV files that are ready for insertion into the Neo4j database. Once the data resides within the database, we apply queries to extract information for the evaluation of workcell performance and visualization of network and operational events within the workcell. By tracking paths through the relationships within the graph, discerning how a network event such as interference is related to physical events such as position uncertainty or part throughput is possible. Various impairments may be introduced as a part of workcell operation. Examples of such impairments include competing wireless traffic, radio interference, and reflections and diffraction due to the multi-path environment [25]. We have shown that it is feasible to implement such impairments and measure the resulting physical performance manifestation [8]. This is accomplished through the use of a radio channel emulator as demonstrated in [26].

D. Sample of a Resulting Graph

In the following, we show results from a trial run of the NIST industrial wireless testbed. In this trial run, a single physical wireless link is used between a wireless adapter connected directly to the supervisor and a wireless access point connected to all the other actors in the testbed. The wireless nodes represent IEEE 802.11b/g/n devices and are connected through a variable radio frequency (RF) attenuator that allows us to vary the channel quality. During the trial run, the production of 10 parts was emulated, which resulted in 10 minutes of network activity.

After populating the database with data captured from the trial run, the resulting realized schema is shown in Fig. 4. The schema visualization is produced by invoking the command

```
call db.schema.visualization()
```

in Neo4j. It is important to note that a realized schema shows only one representation of each node and relationship whereas the intended pseudo-schema shown in Fig. 2 was developed to exemplify the relationships between types of nodes, labels, and

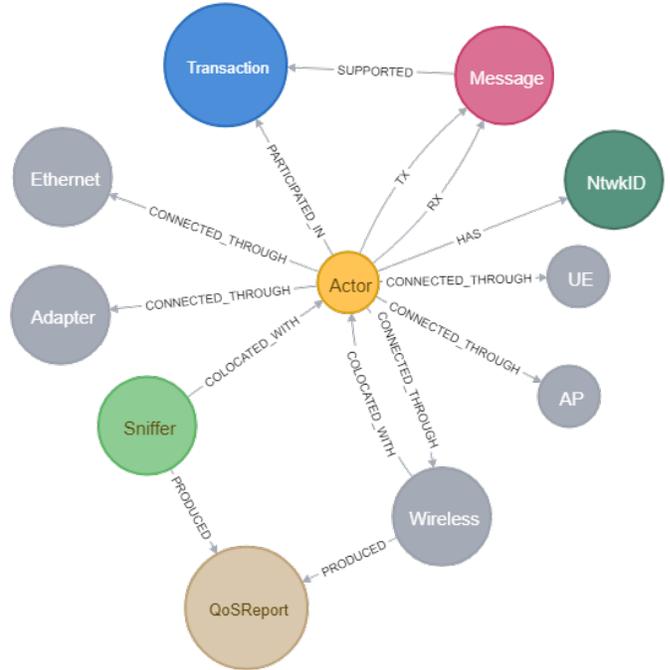


Fig. 4: Realized schema of the graph database fully populated after capturing network and operational data from the NIST industrial wireless testbed.

relationships. Where label inheritance is employed, such as the case for different adapter types, relationships are reproduced; however, this is a result of the visualization tool rather than the schema itself. Fig. 4 serves, therefore, to validate that the intended schema was indeed realized by the insertion of event data from the testbed. In the realized schema, inherited labels are shown as separate nodes.

As described in Section IV-B, the database includes every network transaction that occurs during the operation of the testbed. This includes any logical transaction nodes inserted into the database and any associated packets that happened to traverse the network. Therefore, for a short duration of time depending on packet transmission rates, the amount of data stored in the database can grow quickly. This presents a visualization challenge that graphs are designed to handle. A sample graph is shown in Fig. 5, which represents only 1 second of wired and wireless network data captured from

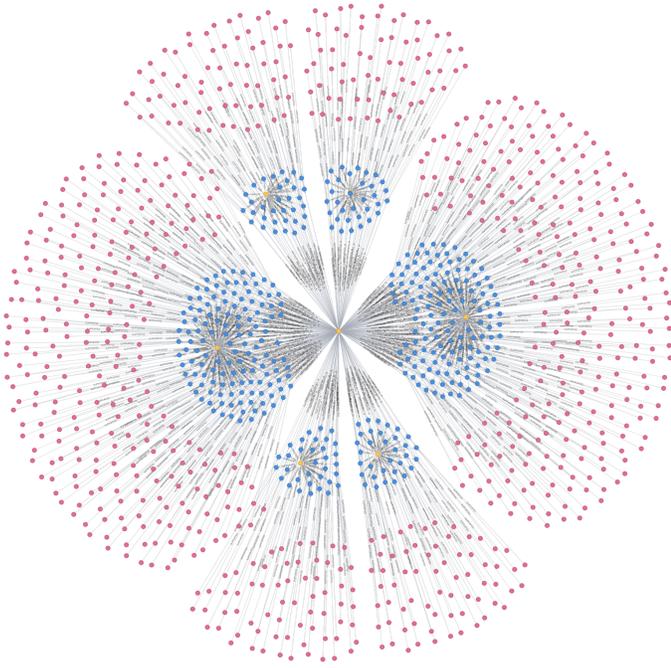


Fig. 5: Visualization of a node graph resulting from a testbed experiment.

the NIST two-robot pick-and-place wireless testbed described in [8].

This visualization is produced by calling the following query in the Neo4j.

```
MATCH p=(a:Actor {name:'Supervisor'})--(t:
  ↳ Transaction)--(b:Actor) , p2=(m:
  ↳ Message)-->(t) WHERE t.timeStart>T
  ↳ AND t.timeStop<T+1
RETURN p,p2
```

The colors of the resulting nodes follow the realized schema in Fig. 4 while only the actors, transactions, and messages are visualized in 5. The relationships between the messages and transactions are shown where a single transaction is connected to at least two messages to represent the communications between any two actors. The variable T in the query represents an arbitrary time variable in seconds within the trial run to capture the data within 1 second only.

We then show a more detailed visualization for all the nodes and relationships corresponding to a single transaction. This visualization is produced by calling the following query where Ts is an arbitrary time variable to specify the timeStart value for the single transaction.

```
MATCH p=(a:Actor {name:'Supervisor'})-->(t
  ↳ :Transaction {timeStart:Ts})<--(b:
  ↳ Actor {name:"CNC-1"})
WITH p, a, b, t
MATCH p1=(a)-->(m:Message)<--(b) , p2=(m)
```

```
↳ -->(t) , p3=(a)-[:HAS]-() , p4=(a)<--[:
  ↳ COLOCATED_WITH]-()-[:PRODUCED]->(q:
  ↳ QoSReport) , p5=(a)-[:
  ↳ CONNECTED_THROUGH]-() , p6=(b)-[:HAS
  ↳ ]-() , p7=(b)-[:CONNECTED_THROUGH]-()
WHERE q.time>t.timeStart AND q.time<t.
  ↳ timeStop
RETURN p,p1,p2,p3,p4,p5,p6,p7
```

In Fig. 6, the actor nodes are labeled by their names *CNC-1* and *Supervisor*, the transaction is labeled by its type *ADS*, the messages are labeled by their transmission role *Request* and *Response*, the NtwkIDs are labeled by their IP addresses, the Ethernet adapters are labeled by their names *eth0*, the wireless adapters are labeled by their names *Moxa* and *TP-Link*, the sniffer is labeled by its name *WLS1*, and the QoS reports are labeled by the received signal strength indicator (RSSI) value in dBm. This visualization includes only the QoS reports generated within the duration of the corresponding transaction.

E. Calculation of Metrics

We now present an example of deploying the proposed GDB approach in industrial wireless analysis. During the trial run time, we varied an RF attenuator value in the single wireless link. The attenuation can take the values {0, 10, 20, 30, 40, 50, 60} dB. We evaluate the message latency as the difference between the receive and transmit times of a message including transmission, processing, and retransmission times. Each message has been coupled to a QoS report that is the one reported at the closest time instant before the message transmit time. One of the parameters in the QoS report is the RSSI value captured by the Sniffer colocated with the supervisor's wireless adapter.

In Fig. 7, we present a scatter plot showing the correlation between the message latency in seconds to the measured RSSI values by the sniffer in dBm. The figure shows that the latency is higher at the lower RSSI values due to the increased number of retransmissions. Generally, an IEEE 802.11 transmission can occur using different IEEE 802.11 mode and modulation and coding scheme (MCS) based on the channel quality. At the RSSI of -90 dBm, the receiver should operate at the most robust communications mode and the lowest MCS index. However, retransmissions still occur due to having the received power near the sensitivity of the supervisor's wireless adapter. As shown in Fig. 7, retransmissions may occur at other RSSI values as well when the transmitter selects a higher mode of transmission and a higher MCS index. In this case, we assert that the receiver is operating close to the marginal sensitivity for a given mode (e.g., 802.11n) and MCS index. We observed that the transmitter switches to a lower MCS index for the retransmitted messages, which supports our supposition. Therefore, latency can in fact be high for better RSSI values. This effect is illustrated in Fig. 7 where certain messages can have latency values at -43 dBm, which are comparable to those in the case of -90 dBm.

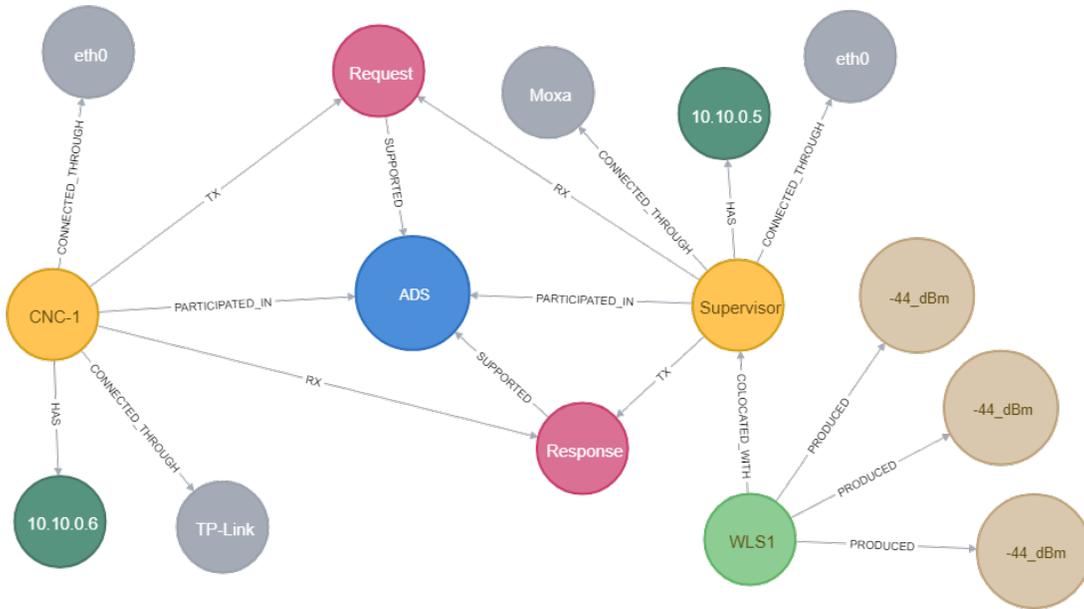


Fig. 6: A detailed visualization resulting from a testbed experiment for all the nodes and relationships corresponding to a single transaction.

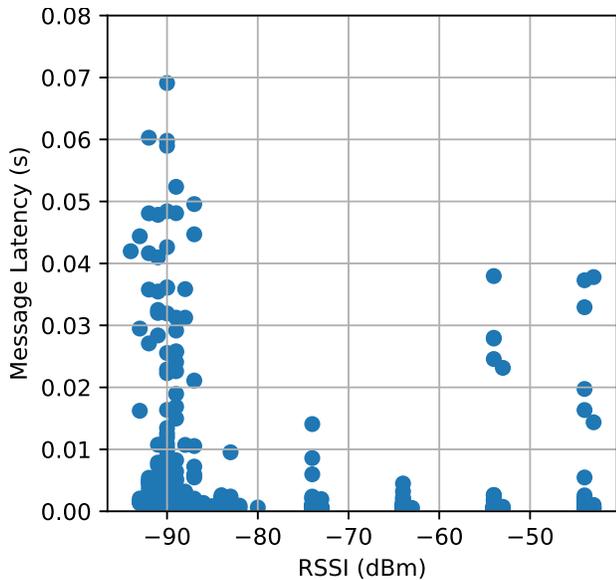


Fig. 7: Correlation between the message latency and the RSSI values by the sniffer.

V. FUTURE DIRECTION

We have presented the general architecture of a performance evaluation database. Future work will include developing query algorithms for the calculation of performance metrics as well as correlation of network events to the physical performance of the manufacturing system. Algorithms will be developed to indicate locations in time of lost information making more direct performance correlations possible. Beyond performance evaluation of the manufacturing system, we believe that the graph database approach enables

anomaly detection because relationships within the data are intrinsically stored and thus efficiently queried. The structure of the information is important within a graph database, and defining the relationships such that the paths through nodes may be discovered is essential to efficient queries and discovery of hidden correlations. Therefore, more research will be done in the modification of the presented schema. Machine learning will be applied for the detection of anomalies, performance degradation, signal quality, and network performance enhancement through more efficient resource allocation. Our plans also include online database insertions to enable the demonstration of online operation of the database. We also intend to examine the development of better control system strategies that mitigate network losses and delays. Since this approach is indifferent to the communications protocols used, we intend to extend our database approach to include various protocols, radio bands such as millimeter wave bands for 5G communications, and different use cases to include mobile robotic platforms and safety critical systems.

VI. CONCLUSIONS

We have presented in this paper a novel approach to capturing network and operational event information from a factory workcell with the purposes of 1) capturing and storing of network and operational events, 2) calculating performance metrics of the network, and 3) discovering performance dependencies between the network and the physical assembly of the workcell. Using a graph database, we have demonstrated that it is possible to construct such a database, compute network performance metrics and discover correlations. We have also developed the capability of examining the correlation between network events and the performance of physical actions. This will be a source of further research.

Future progress and measurement data will be deposited in the NIST public domain repository as a reference for industrial traffic modeling efforts and comparative studies on industrial wireless technologies [27].

DISCLAIMER

Certain commercial equipment, instruments, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

REFERENCES

- [1] L. L. Bello, J. Akerberg, M. Gidlund, and E. Uhlemann, "Guest editorial special section on new perspectives on wireless communications in automation: From industrial monitoring and control to cyber-physical systems," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 3, pp. 1393–1397, June 2017.
- [2] V. K. L. Huang, Z. Pang, C. A. Chen, and K. F. Tsang, "New trends in the practical deployment of industrial wireless: From noncritical to critical use cases," *IEEE Industrial Electronics Magazine*, vol. 12, no. 2, pp. 50–58, June 2018.
- [3] (ETSI), "Etsi tr 103 588 reconfigurable radio systems (rrs); feasibility study on temporary spectrum access for local high-quality wireless networks," Sophia-Antipolis, 2018.
- [4] H. Kagermann, W. Wahlster, and J. Helbig, "Recommendations for implementing the strategic initiative industrie 4.0, industrie 4.0 working group," 2013.
- [5] A. Barnard Feeney, S. Frechette, and V. Srinivasan, *Cyber-Physical Systems Engineering for Manufacturing*. Cham: Springer International Publishing, 2017, pp. 81–110. [Online]. Available: https://doi.org/10.1007/978-3-319-42559-74_4
- [6] Wireless Broadband Alliance (WBA), "Enhanced Wi-Fi - 802.11ax Decoded: Overview, Features, Use Cases and 5G Context." [Online]. Available: <https://www.wballiance.com/wp-content/uploads/2018/09/Enhanced-802.11ax-Overview-Use-Cases-Features-5G-Context-v1.asd-Final-Clean.pdf>
- [7] R. Angles and C. Gutierrez, "Survey of graph database models," *ACM Comput. Surv.*, vol. 40, no. 1, pp. 1:1–1:39, Feb. 2008. [Online]. Available: <http://doi.acm.org/10.1145/1322432.1322433>
- [8] Y. Liu, R. Candell, M. Kashef, and K. Montgomery, "A collaborative work cell testbed for industrial wireless communications the baseline design," in *2019 IEEE 28th International Symposium on Industrial Electronics (ISIE)*, June 2019, pp. 1315–1321.
- [9] R. Kumar Kaliyar, "Graph databases: A survey," in *International Conference on Computing, Communication Automation*, May 2015, pp. 785–790.
- [10] H. R. Vyawahare and P. P. Karde, "An overview on graph database model," *International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCCE)*, vol. 3, pp. 7454–7457, August 2015.
- [11] K. N. Satone, "Modern graph databases models," *International Journal of Engineering Research and Applications (IJERA)*, pp. 19–24, 2014.
- [12] P. T. Wood, "Query languages for graph databases," *SIGMOD Record*, vol. 41, pp. 50–60, 2012.
- [13] P. S. Jadhav and R. K. Oberoi, "Comparative analysis of graph database models using classification and clustering by using weka tool," 2015.
- [14] P. Macko, D. Margo, and M. Seltzer, "Performance introspection of graph databases," in *Proceedings of the 6th International Systems and Storage Conference*, ser. SYSTOR '13. New York, NY, USA: ACM, 2013, pp. 18:1–18:10. [Online]. Available: <http://doi.acm.org/10.1145/2485732.2485750>
- [15] J. Wang, W. Zhang, Y. Shi, S. Duan, and J. Liu, "Industrial big data analytics: Challenges, methodologies, and applications," *CoRR*, vol. abs/1807.01016, 2018. [Online]. Available: <http://arxiv.org/abs/1807.01016>
- [16] J. Lee, *Industrial Big Data (Mechanical Industry Press, China)*, 07 2015.
- [17] (GE), "Unlocking machine data to turn insights into powerful outcomes." [Online]. Available: <https://www.ge.com/digital/>
- [18] Brian Courtney, "Industrial big data analytics: The present and future." [Online]. Available: <https://www.isa.org/intech/20140801/>
- [19] (ABB), "Big Data and decision-making in industrial plants." [Online]. Available: <https://new.abb.com/cpm/production-optimization/big-data-analytics-decision-making>
- [20] J. Pokorny, "NoSQL databases: a step to database scalability in web environment," *International Journal of Web Information Systems*, vol. 9, no. 1, pp. 69–82, 2013. [Online]. Available: <https://doi.org/10.1108/17440081311316398>
- [21] Z. J. Zhang, "Graph databases for knowledge management," *IT Professional*, vol. 19, no. 6, pp. 26–32, November 2017.
- [22] J. Webber and I. Robinson, "The top 5 use cases of graph databases," *white paper, Neo4j*, 2015.
- [23] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," *ACM Trans. Knowl. Discov. Data*, vol. 5, no. 4, pp. 21:1–21:37, Feb. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2086737.2086741>
- [24] S. Skhiri and S. Jouili, *Large Graph Mining: Recent Developments, Challenges and Potential Solutions*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 103–124. [Online]. Available: https://doi.org/10.1007/978-3-642-36318-4_5
- [25] R. Candell, C. Remley, J. Quimby, D. Novotny, A. Curtin, P. Papazian, G. Koepke, J. Diener, and M. Kashef, "Industrial wireless systems: Radio propagation measurements," National Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep., 2017. [Online]. Available: <http://nvlpubs.nist.gov/nistpubs/TechnicalNotes/NIST.TN.1951.pdf>
- [26] R. Candell, K. Montgomery, M. Kashef, Y. Liu, and S. Foufou, "Wireless interference estimation using machine learning in a robotic force-seeking scenario," in *2019 IEEE 28th International Symposium on Industrial Electronics (ISIE)*, June 2019, pp. 1334–1341.
- [27] R. Candell, M. Hany, Y. Liu, and K. Montgomery, "Reliable, High Performance Wireless Systems for Factory Automation," 2019. [Online]. Available: <https://www.nist.gov/programs-projects/reliable-high-performance-wireless-systems-factory-automation>