NISTIR 8227

Manufacturing Profile Implementation Methodology for a Robotic Workcell

Timothy A. Zimmerman

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8227

National Institute of
Standards and Technology
U.S. Department of Commerce

NISTIR 8227

Manufacturing Profile Implementation Methodology for a Robotic Workcell

Timothy A. Zimmerman
Intelligent Systems Division
Engineering Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8227

May 2019

U.S. Department of Commerce
Wilbur L. Ross, Jr., Secretary

National Institute of Standards and Technology
Walter Copan, NIST Director and Undersecretary of Commerce for Standards and Technology

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept 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 entities, materials, or equipment are necessarily the best available for the purpose.

National Institute of Standards and Technology Interagency or Internal Report 8227
Natl. Inst. Stand. Technol. Interag. Intern. Rep. 8227, 41 pages (May 2019)

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8227

Abstract

The National Institute of Standards and Technology has constructed a testbed to measure the performance impact of cybersecurity technologies on Industrial Control Systems (ICS). The testbed was chosen to support the implementation of the Cybersecurity Framework Manufacturing Profile: a voluntary, risk-based approach for managing cybersecurity activities and reducing cyber risk to manufacturing systems. This report focuses on the Collaborative Robotics System, one of the two manufacturing systems within the testbed. A methodology for implementation of technical solutions to meet the Profile language is described, as well as a comprehensive review of the testbed measurement systems, and the comparative analysis procedures used for identifying performance impacts. Finally, an example comparative analysis is performed and the characterization of the workcell is discussed.

Key words

cybersecurity; robotics; manufacturing; performance; industrial control system; PLC; measurement methodology; testbed; workcell.

Disclaimer
Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by NIST, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Table of Contents

List of Tables

List of Figures

1. Introduction

The National Institute of Standards and Technology (NIST) has developed a voluntary, risk-based approach for managing cybersecurity activities and reducing cyber risk to manufacturing systems, called the Cybersecurity Framework Manufacturing Profile [1]. Concurrently, NIST constructed the Cybersecurity for Smart Manufacturing Systems (CSMS) testbed to measure the performance impact of cybersecurity technologies on industrial control systems (ICS) [2]. The Manufacturing Profile will be implemented into the testbed to validate the outlined approach, and to measure performance impacts realized after cybersecurity tools and technologies are installed in the testbed's manufacturing systems.

A major goal of this research is to determine the role each component of the Manufacturing Profile implementation contributes to the overall system performance. To meet the research goals, a robust methodology for experimentation was developed, including a characterization procedure and comparative analysis methods for identifying and quantifying performance impacts to the manufacturing system.

2. Cybersecurity for Smart Manufacturing Systems Testbed

The Manufacturing Profile implementation will be performed on the Cybersecurity for Smart Manufacturing Systems testbed [2], located at the NIST campus in Gaithersburg, Maryland. The testbed contains two manufacturing systems: the Collaborative Robotics System (CRS) [3], and the Process Control System (PCS) [4]. The CRS is a discrete manufacturing process, while the PCS is a continuous manufacturing process. This report focuses specifically on the Collaborative Robotics System.

The CRS workcell, shown in Figure 1, contains two robotic arms that perform a material handling process called *machine tending* [3]. Robotic machine tending utilizes robots to interact with machinery, performing physical operations a human operator would normally perform (e.g., loading and unloading of parts in a machine, opening and closing of machine doors, interacting with the operator control panel, etc.). A human operator interfaces with the workcell through a human-machine interface (HMI) and a control panel external to the work area.

The workcell was designed and constructed to be reconfigurable, allowing numerous types of operational methodologies, network topologies, and industrial networking protocols to be investigated. The two robots collaborate to transport parts through the manufacturing process, as a single robot cannot physically reach all four stations, and having two robots increases workcell efficiency.

Parts are moved by the robot arms through four simulated machining operations, known as *stations*. Each station is comprised of: a fixture for holding the part, an infrared proximity sensor for detecting the part, a single board computer simulating the actions and communications of a typical machining center, and a liquid crystal display (LCD) for displaying the operational status of the station. The stations communicate with the supervisory programmable logic controller (PLC) over the workcell's local area network (LAN).

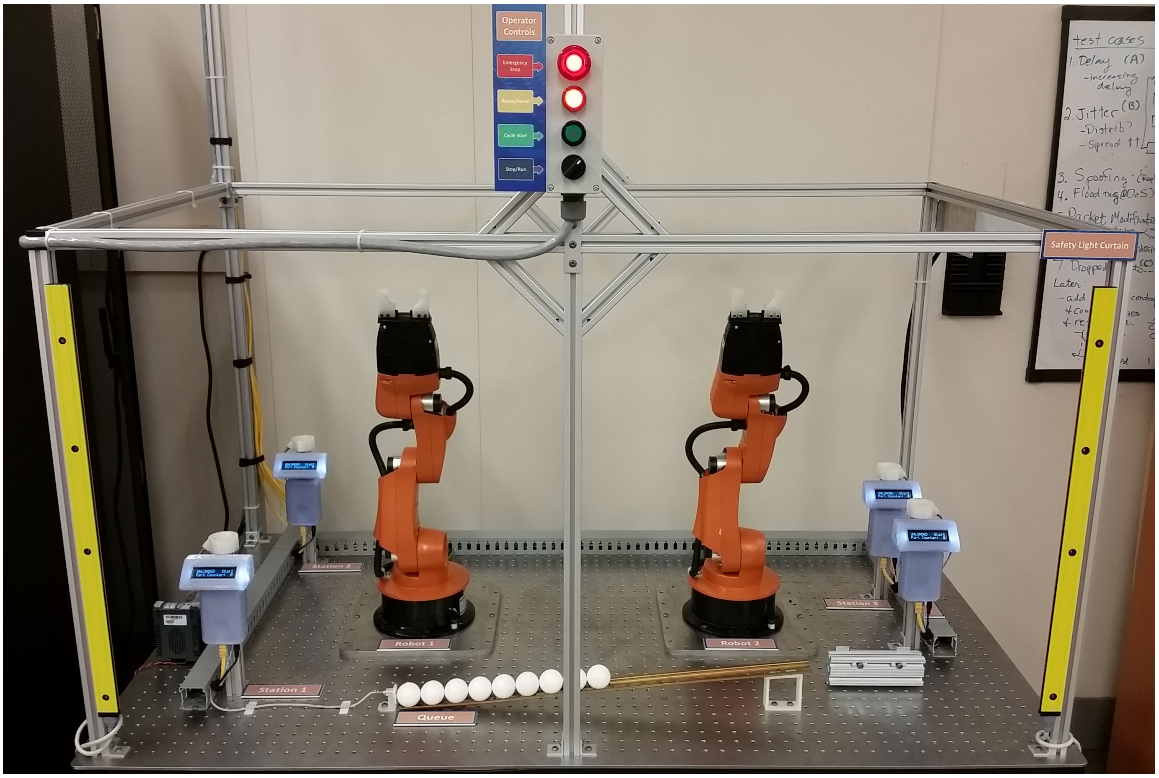

Fig. 1. The CRS workcell in standby, waiting for the operator to initiate the manufacturing process. The operator control panel is visible at the top of the figure.

The supervisory PLC monitors and controls all aspects of the manufacturing process. Manufacturing data from the four machining stations are used by the PLC to determine which operations (known as *jobs*) the robots must perform to keep the parts moving through the sequential manufacturing process. The PLC also communicates with the HMI for operator visibility and control.

The workcell is supported by a shared infrastructure of networked servers, measurement tools, and other technologies. The infrastructure is used for many research functions including: testing, deployment, and hosting of cybersecurity tools; measurement systems for network traffic; creation and manipulation of network traffic for inducing anomalous network activity; and archival storage of experiment data. A virtualized server infrastructure was installed to support the numerous cybersecurity technologies and tools required for the implementation.

3. Manufacturing Profile

A major milestone of the Cybersecurity for Smart Manufacturing Systems project at NIST was the creation of the Manufacturing Profile [1]. The Manufacturing Profile provides a roadmap and actionable guidance for implementing the Cybersecurity Framework within a manufacturing environment. The prescribed methodology can be used to identify gaps in cybersecurity posture (when paired with a "current profile"), and addresses each gap with actionable language derived from industry-specific standards, guidelines, and best practices.

As with many documents relating to the cybersecurity of manufacturing systems, the Manufacturing Profile does not attempt to identify potential performance impacts that may be realized once a cybersecurity solution is installed. Implementing prescribed cybersecurity technologies and tools without affecting ICS performance is a challenge. Impacts to the manufacturing system can be difficult to predict, and differ between manufacturing processes due to a wide spectrum of manufacturing processes, operating procedures, control schemes, industrial hardware, software, industrial protocols, and network topologies.

Meeting the Manufacturing Profile language will involve sourcing and procuring cybersecurity tools, reconfiguration of existing workcell systems, and instituting organizational policies and procedures. Detailed information describing the installation, configuration, operation, and financial cost of each tool will be recorded during the implementation.

3.1 Security Levels

The Manufacturing Profile identifies three Security Levels that categorize the potential impact a manufacturer may realize if the manufacturing system is compromised by a cybersecurity incident: low, moderate, and high. Examples of how each security level aligns with business objectives are shown in Table 1.

The implementation of each security level within the workcell is performed independently, each with a specific manufacturing use case that is representative of the chosen security level. A use case describes a hypothetical manufacturer and all of the operational details required to inform the Manufacturing Profile implementation process. Use cases are further discussed in Section 4.2.

The security levels will be implemented in sequential order from the low security level to the high security level to measure the performance impact each has on the system. As the security level increases, some of the cybersecurity tools previously implemented for the lower-security level may be reconfigured to meet the language of the higher level.

Table 1. Manufacturing system impact levels from the Manufacturing Profile.

Impact Category	Low Impact	Moderate Impact	High Impact
Injury	Cuts, bruises requiring first aid	Requires hospitalization	Loss of life or limb
Financial Loss ($)	Tens of thousands	Hundreds of thousands	Millions
Environmental release	Temporary damage	Lasting damage	Permanent damage, off-site damage
Interruption of production	Minutes	Days	Weeks
Public image	Temporary damage	Lasting damage	Permanent damage

4. Implementation of the Manufacturing Profile

Meeting the guidance language for each security level can be achieved by implementing technical solutions within the CRS (i.e., cybersecurity tools and technologies). For this research, the language of each subcategory is reviewed by the researchers to determine if a technical solution can be used to meet the subcategory intent. Commercially available technical solutions are then reviewed by the researchers, procured (if the solution meets the language and fits the use case), and finally installed within the workcell as part of the implementation.

Each of the three Manufacturing Profile security levels (low, moderate, and high) will be implemented and tested on the workcell, starting with the low security level. The technical solutions identified to satisfy each subcategory will be installed in no specific order. The subcategories for the low security level implementation identified as requiring technical solutions are shown in Table 2.

Performance measurements will occur after each technical solution is installed. Technical solutions that have multiple modes of operation will be tested for each mode that meets the Manufacturing Profile language, as each mode may affect the manufacturing process differently. All modes of a technical solution selected for performance measurements will be aligned with the requirements of the current security level and its applicability validated against the use case.

After the measurements for a technical solution with multiple modes of operation have been completed, the mode of operation that best meets the Manufacturing Profile language for the current security level, and that has the least impact to the manufacturing process, will be used for the remainder of the security level implementation. A flow chart describing the the implementation methodology is shown below in Figure 2, which includes the three measurement types to be performed, as described in Section 4.1.

Table 2. Low security level Cybersecurity Framework subcategories identified as requiring technical solutions.

Identify (ID)	Protect (PR)	Detect (DE)	Respond (RS)
Asset Management	*Identity Management and Access Control*	*Anomalies and Events*	*Analysis*
ID.AM-1	PR.AC-1	DE.AE-1	RS.AN-3
ID.AM-2	PR.AC-2	DE.AE-2	*Mitigation*
ID.AM-3	PR.AC-5	DE.AE-3	RS.MI-2
ID.AM-4	*Data Security*	*Security Continuous Monitoring*	RS.MI-3
Risk Assessment	PR.DS-3		
ID.RA-1	PR.DS-5	DE.CM-1	
	Information Protection Processes and Procedures	DE.CM-2	
	PR.IP-1	DE.CM-3	
	PR.IP-2	DE.CM-4	
	PR.IP-3	DE.CM-6	
	PR.IP-4	DE.CM-7	
	PR.IP-6	DE.CM-8	
	Maintenance	*Detection Processes*	
	PR.MA-1	DE.DP-3	
	PR.MA-2		
	Protective Technology		
	PR.PT-1		
	PR.PT-2		
	PR.PT-3		
	PR.PT-4		

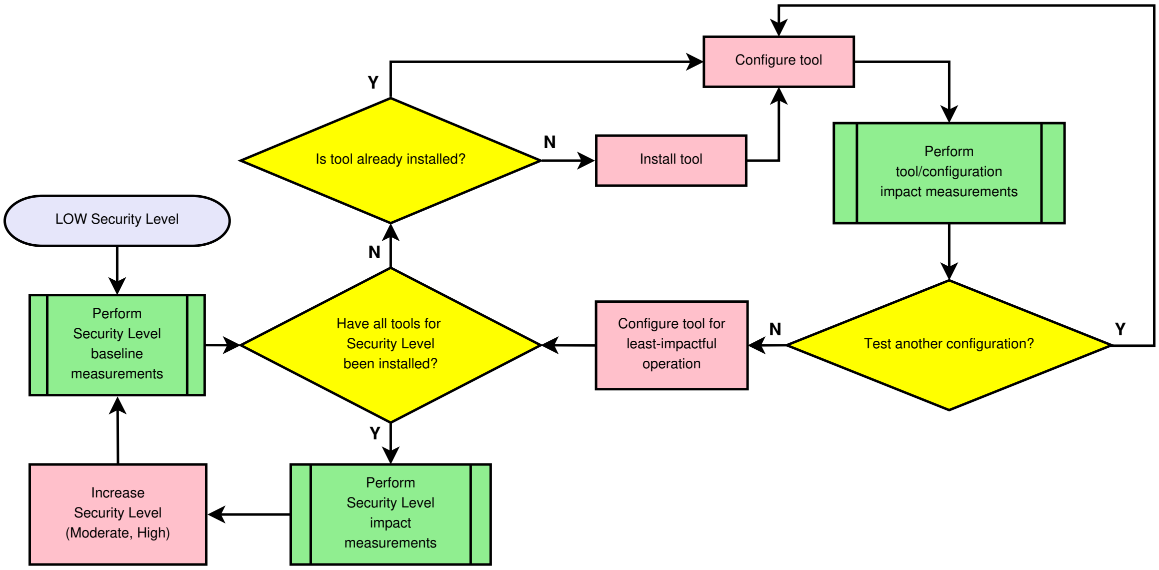
Configure tool
Y
Is tool already installed?
N
Install tool
Perform tool/configuration impact measurements
LOW Security Level
N
Perform Security Level baseline measurements
Have all tools for Security Level been installed?
Configure tool for least-impactful operation
N
Test another configuration?
Y
Y
Increase Security Level (Moderate, High)
Perform Security Level impact measurements

Fig. 2. Flow chart describing the implementation process for all Manufacturing Profile security levels.

4.1 Performance Measurements

There are three types of performance measurements that will be performed during the implementation of all three security levels: baseline measurements of the initial workcell performance, impact of individual technologies or configurations, and impact of the completed security level implementation. Each of these measurements are represented in Figure 2 as green boxes. The process of sequentially implementing and measuring enables the detection of performance-impacting interactions that may occur between the technical solutions.

- **Security level baseline** - Before any changes are made to the workcell, baseline measurements must be captured. Since all experiments are meant to be comparative, a baseline reference of system performance must be obtained to determine if the manufacturing process or its sub-systems have been impacted after a technical solution is installed or reconfigured. Interim baselines may also be required to provide a ground truth after workcell operational modifications unrelated to the Manufacturing Profile implementation (e.g., enabling or disabling services, control system software modifications).

- **Technology/configuration implementation impact** - These measurements are performed after each technical solution is installed and configured to meet the security level requirements. Some technical solutions may provide multiple modes of operation that meet the security level requirements and have the potential to affect the manufacturing process differently. Measurements are performed for each unique configuration to compare its impact to previous configurations. The first technical solution implemented for a given security level is compared to the security level baseline.

- **Security level implementation impact** - These measurements are performed after all technical solutions have been installed and configured for a given security level. These measurements are compared with the Low security level baseline to determine the total impact to the manufacturing process, and compared with other security level implementation impact measurements to determine the relative performance impact between the security levels.

4.2 Manufacturer Use Cases

To properly implement the Manufacturing Profile within the CRS, use cases must be created for each of the security levels. A use case describes a hypothetical manufacturer and all of the operational details required to inform the Manufacturing Profile implementation process (e.g., how many employees does the company have, how large is the facility, what are the business mission objectives, and what company data must be protected). A summary of operational details provided by a use case is shown in Table 3.

The details provided for each use case are designed to be representative of a manufacturer classified to the chosen security level, and to inform the selection for appropriate cybersecurity tools, configuration, and operating modes.

Table 3. Summary of manufacturer operational details provided in the security level use cases.

Name	Description
Employees	Titles and roles for all employees.
External Personnel	Facility operations outsourced to external entities.
Supporting Services	Services provided by external entities that support the manufacturing system.
Supply Chain	Upstream and downstream supply chain details.
Critical Infrastructure	Industry categorization as a Department of Homeland Security Critical Manufacturing sector [5].
Manufacturing Process	Parts that are manufactured, processes and technologies utilized.
Critical Systems	Systems that are critical for the manufacturing system to operate properly and safely.
Data	Types of data used, stored, or created by the manufacturer.
Mission Objectives	Manufacturing Profile-defined mission objectives of importance to the manufacturer, and how those objectives apply to the manufacturing operations.

5. Testbed Measurement Systems

Many of the systems within the workcell perform measurements natively while the manufacturing process is operating. Each of these systems are listed in Table 4, and are further described in the following subsections. As noted in [3], the high-level production measurements give an indication if the manufacturing process performance is impacted, but these measurements alone are not be able to identify the underlying cause. To solve this problem, numerous other low-level measurements are captured beyond the manufacturing process performance (e.g., raw network traffic, time-stamped events from machining stations).

Figure 3 shows the flow of data between workcell subsystems required for manufacturing operations. Because of the tight interaction between the many workcell subsystems, it is critical that these data flows continue to operate effectively and without error. An introduction to the problem of performance impact propagation through the workcell systems is described in detail in [6].

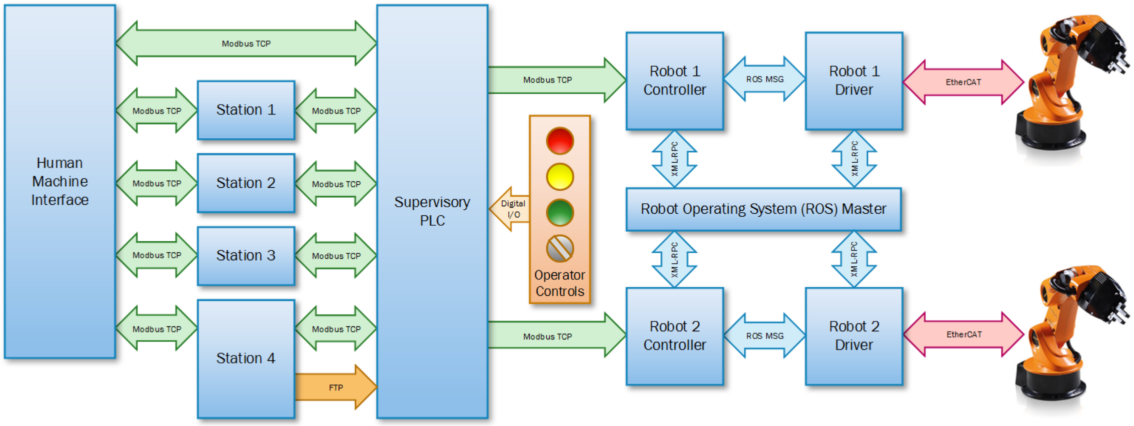

Fig. 3. Data flows between workcell subsystems and their respective network protocols, as required for normal manufacturing operations.

Due to the large amount of infrastructure required (e.g., network taps, cabling, additional servers) and intrusiveness of the measurement systems, manufacturers should not attempt to reproduce the described measurement systems, especially on production systems. Guidance for implementing measurement systems to baseline existing manufacturing systems will be produced in future work. Most of the measurements performed for this research are used to understand *how* performance impacts propagate through networked ICS.

For example, assume an increase in network bandwidth is observed by a manufacturer after a technical solution was installed. Although there is an observable and measurable performance impact at the network level (a portion of the available network bandwidth is being consumed by the technical solution), the manufacturer cannot immediately conclude manufacturing system performance will be affected by this change.

It is hypothesized that performance impacts to the manufacturing system can be detected through existing measurement systems used by many manufacturers (e.g., data historians), and production data obtained from these measurement systems are likely the most important to monitor, since a measurable deviation in key performance indicators at the process level indicates the manufacturing system has failed to maintain production goals.

5.1 Key Performance Indicators

Key performance indicators are computed from the measurement data after an experiment completes. The KPI provide quantifiable and communicable indicators of performance impacts from the Profile implementation [3]. If any performance impacts are detected, further analysis is performed to assist with locating the source of the performance degradation. The measurement data used for computation of the KPI are listed in Table 4, and the resulting KPI files are listed in Table 5.

5.2 Supervisory PLC

The Supervisory PLC is responsible for monitoring the workcell safety systems, supervising the four machining stations, and disseminating jobs to the robots based on the stateful data received from each station. For measurement purposes, the PLC has direct connection to the proximity sensors located on each machining station (for detecting when a part is present), enabling the PLC to independently capture timestamps of each individual part entering and exiting a station. The PLC measurement data is the primary source of all manufacturing production metrics.

Each of the files that are generated from PLC measurement data are described in the following subsections.

5.2.1 Experiment Metadata

This file contains general information about the experiment, as recorded by the PLC. The most important data logged by this file is the configuration of the workcell, which is used to validate that it was properly configured.

5.2.2 PLC Part Data

This file contains the timestamps each part generates as it triggers events within the workcell, as recorded by the PLC. Each numeric timestamp is relative to the start of the experiment, and is measured in tens of milliseconds (e.g., a timestamp value of 2313 is equal to 23.13 seconds). The file also contains the serial number of each part, and the inspection result. An inspection can result in a value of `NOT INSPECTED`, `PASS`, or `FAIL`, encoded respectively as an integer `0`, `1`, or `2`.

Table 4. Measurement files created by the testbed components, and their resulting archive folder location.

System Name	Archive Folder	Measurement Files
Packet Capture Files	./pcaps/	capture.pcap
		hmi.pcap
		plc.pcap
		station1.pcap
PLC	./plc/	ExperimentMetadata.dat
		PLCPartData.csv
		PLCPerformanceData.csv
Robots	./robots/	Robot1_Performance.csv
		Robot2_Performance.csv
Servers	./servers/	vController1_ServerPerformance.csv
		vController2_ServerPerformance.csv
Machining Stations	./stations/	Station1_Events.csv
		Station2_Events.csv
		Station3_Events.csv
		Station4_Events.csv
Time Synchronization	./timesync/	Mintaka.loopstats
		Polaris.loopstats
		vController1.loopstats
		vController2.loopstats
		Station1.loopstats
		Station2.loopstats
		Station3.loopstats
		Station4.loopstats

Table 5. KPI files calculated from measurement data, and their resulting archive folder location.

System Name	Archive Folder	KPI Files
Network	./network/	6.1_<HOST>to<HOST>_ppd.csv
		6.2_<HOST>to<HOST>_ipd.csv
		6.3_rtt.csv
		6.4_ir_calcs.csv
		6.4_ir_hostinout.csv
		6.4_ir_hosttohost.csv
		6.5_6.6_rates_hostinout.csv
		6.5_6.6_rates_hosttohost.csv
		6.8_protocols.txt
		6.8_protocols_calcs.txt
PLC	./plc/	KPI_ProductionTimes_CycleTimes.csv
	./reports/	PLCReport_SingleRun.pdf
		ProductionDataReport_SingleRun.pdf
Robots	./robots/	KPI_RobotJobActuationTimes.csv
		KPI_RobotJobExecutionTimes.csv
		KPI_RobotPoseTravelTimes.csv
		KPI_RobotPower.csv
		Report_RobotJobActuationTimes.txt
		Report_RobotJobExecutionTimes.txt
		Report_RobotPoseTravelTimes.txt
		Report_RobotPower.txt
Servers	./servers/	KPI_vController1.csv
		KPI_vController2.csv
	./reports/	PerformanceReport_vController1.pdf
		PerformanceReport_vController2.pdf

5.2.3 PLC Performance Data

This file contains the performance of the PLC internal operating system, and the logic program (i.e., task) written specifically to supervise the workcell. The PLC is configured to log its performance every one second while an experiment is underway. A timestamp (seconds since the start of the experiment) is recorded, along with the processor usage ratio, average amount of time the PLC task required to execute (measured in microseconds) since the previous measurement, and the maximum amount of time the PLC task required to execute during since the last measurement (measured in microseconds).

5.3 Robot Controllers

The robot controllers are drivers that utilize the Robot Operating System (ROS) [7] framework and Python [8]. The controller facilitates communications between the robots and PLC, and executes the actions defined by each job script. Two robot controllers (one for each robot) are required while the manufacturing process is operating.

The robot controllers are pre-programmed with a series of discrete instructions organized into scripts that are enumerated by a job number (e.g., Job 1: move to Station 1, close gripper, move to station 2, open gripper). Which job a robot is to perform at any given time is sent by the PLC to a robot controller using the industrial network protocol Modbus TCP. When a robot is idle, it continues to query the PLC until a new job has been published. After the robot controller receives a job, it executes the script instructions before returning to the idle state and querying the PLC for the next job.

A performance logger for robot telemetry is implemented natively in each robot controller. The logger combines three different sources of performance data into a single file: robot joint states, robot position commands, and robot controller events. Joint state messages contain the position, velocity, and electrical current measured by each joint. Robot command messages contain the requested job, job step (i.e., each discrete instruction executed by the robots during the execution of a job), a string name of the instruction, and the start/end timestamps. Robot controller messages identify important manufacturing process events (e.g., batch start and end, timestamps when new jobs are received, timestamps when a robot is within proximity of a machining station).

5.4 Machining Stations

The machining stations simulate the actions and industrial communications of a typical machining center [9]. The stations provide access to their internal registers via a Modbus TCP server, giving the PLC and HMI insight into the machining process in real-time (e.g., job progress, machine state, doors open/closed). This data is used primarily by the PLC for supervising the process and using that information to decide which jobs to disseminate to each robot. This data is also accessed and processed by the HMI to display to the operator.

Each station has a native event logger, which can be enabled and disabled with a configuration option. When enabled, the station will log each transition of the machine state. Each entry that is appended to the log file includes: a timestamp, string name of the new machine state, and the value of the station's part counter.

5.5 Servers

The physical servers used to execute the robot controllers and robot drivers have a configuration option to generate log files of server performance metrics. Capturing of this data is performed by a Python script, and is executed independently of the robot processes. A summary of these metrics is shown in Table 6.

The performance logger inherently imparts its load on the server, resulting in the logger recording its own system performance impact. Due to the comparative analysis methodology (see Section 7), the logger's performance impact to the server will not affect performance analysis.

Table 6. Summary of captured server performance data.

Name	Description
Processor timers	Total amount of time across all cores consumed executing user tasks, system tasks, idling, etc.
Disk counters and timers	Counters and timers tracking reads and writes to the disks.
Network counters	Counters tracking transmitted and received packets, bytes, errors, and drops.
Memory counters	Amount of memory available at the time of measurement.

5.6 Networks

Capture files of all raw network traffic are recorded using the TCPDUMP software [10] on a measurement server dedicated to performing network captures and calculations. Each networking device shown in Figure 4 has a dedicated mirror port that forwards all received traffic to the measurement server. Each stream of mirrored packets enters a traffic aggregator, which combines all of the traffic into a single stream before it is forwarded to the measurement server.

Physical network probes are utilized to forward all traffic transmitted to/from the PLC, HMI, and Station 1. Each dedicated probe forwards the captured network traffic to a dedicated Ethernet port on the measurement server. The network traffic is subsequently stored in capture files dedicated to each device.

After an experiment has completed, tools like EDITCAP, CAPINFOS, and TSHARK (components of the WIRESHARK software package) [11] are used to dissect specific metrics from the raw traffic for calculating networking KPI.

5.7 Time Synchronization

A Network Time Protocol (NTP) server is hosted on the workcell router for hosts within the network. The router synchronizes its time with a grandmaster clock located within the testbed network. Each host that performs timestamped logging within the workcell is time synchronized with the NTP server, and is typically accurate to within 1 millisecond of the grandmaster clock.

It is important to measure and record the clock drift between workcell hosts when the logs are used for event correlation, as the clocks in each host tend to deviate from the reference clock (the grandmaster) over time. The NTP performance file from each host, which contains the measured clock drift and offset, is captured to validate that the expected clock accuracy was maintained during the experiment.

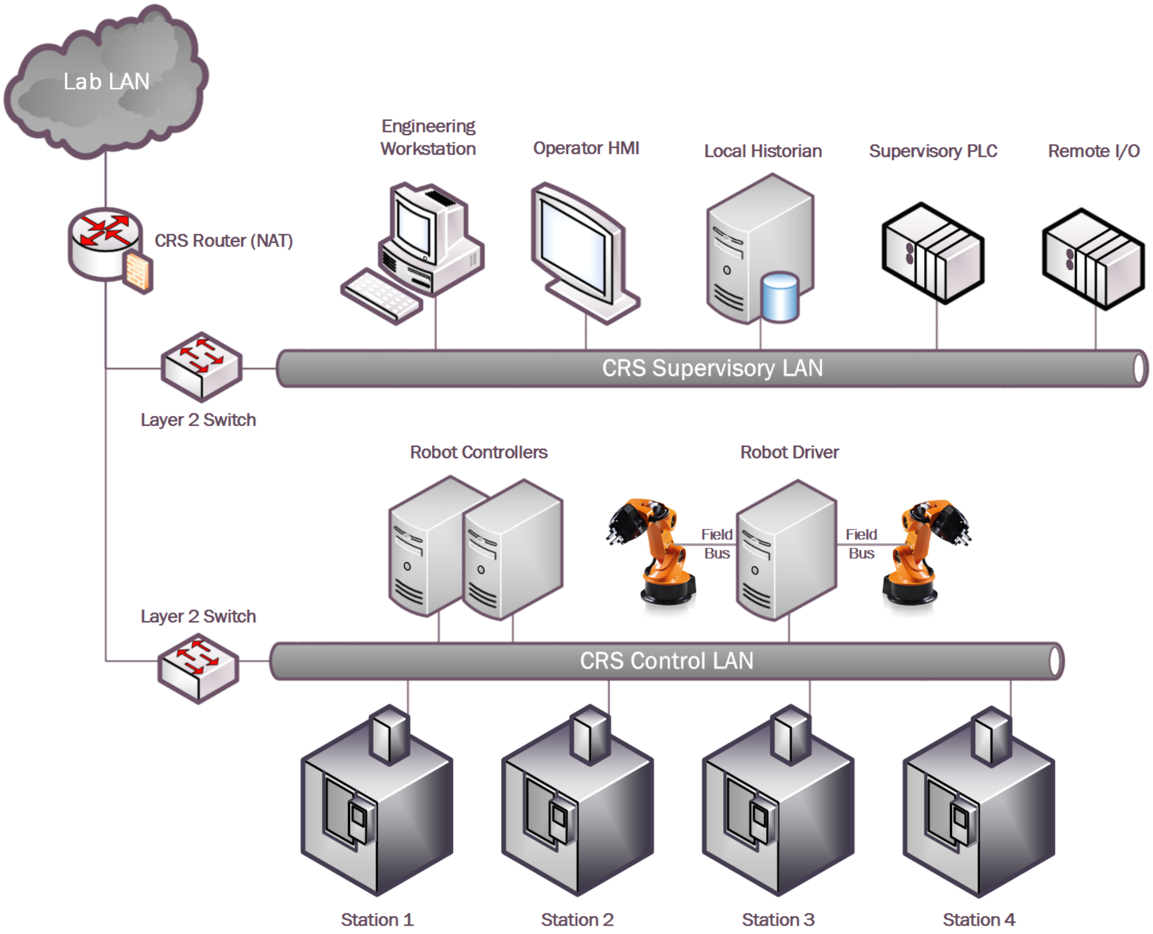

Fig. 4. Network topology diagram of the CRS.

6. Measurement Data Retrieval and Storage

The measurement data files, described in Table 4, are downloaded from each workcell host to a large redundant storage array within the testbed network. All of the data files are organized by tool and configuration, using a descriptive folder naming structure described in Table 7, and subfolder structure for organizing the different types of measurement data. Subfolder locations of each measurement file are also detailed in Table 4.

The measurement data files from successful experiments will be made publicly available via a website on the nist.gov domain after the data has been validated and checked for completeness. The resulting structure of the data will be described in a future publication.

Table 7. Description of the folder naming structure for measurement data.

Name	Values	Description
Testbed system	C, P	Testbed system the measurements were taken: the Collaborative Robotics System (C), or the Process Control System (P).
Security level	L,M,H	Low (L), moderate (M), or high (H) Manufacturing Profile security level.
Technology identifier	000 to 999	Sequential, incrementing identifier indicating the order in which the tool was installed.
Configuration identifier	.0 to .9	Identifier of the technology or tool configuration.
Separator	-	Dash separator.
Technology name	`string`	Descriptive name of the installed technology or tool.
Separator	-	Dash separator.
Configuration name	`string`	Descriptive name of the measured configuration.

Consider the example folder name: `CL004.2-AntiVirus-FullScan`. This folder name indicates it contains data from an experiment on the CRS (C), and it was taken during the implementation of the low security level (L). It was scheduled as the fourth technology to be installed (004) and the second configuration (.2). The technology installed was antivirus software, and the action tested was a "full-scan". More detailed information about the data is contained in the `README` file in the root directory of the folder.

7. Characterization and Analysis of Measurement Data

After each tool has been installed or a configuration measurement is performed, the data from the experiment is analyzed to determine if the manufacturing process performance was negatively impacted. This analysis is performed with an approach called exploratory data analysis (EDA), and is described in great detail in the NIST Engineering Statistics Handbook [12]. This analysis process requires the workcell to be characterized for proper comparative analysis to be performed.

7.1 Workcell Characterization

A characterization procedure (further described in Appendix A) must first be performed to understand the characteristics of the manufacturing process and validate its stability. If the manufacturing process is stable, the response variables (KPI) will have constant means and constant variances over time, and will also have a constant distribution [12]. Two response variables (KPI) were chosen as candidates for detecting performance impacts: part processing time, and workcell cycle time.

The characterization process concluded that the manufacturing process is stable, and it is valid to use EDA for detecting performance impacts to the manufacturing process.

Data from the characterization process was used to calculate the minimum sample size to confidently indicate the process was impacted: 35 parts.

The characterization data was also used to test multiple hypotheses related to the performance and stability of the workcell. For example, do parts produced after the workcell has warmed-up have a shorter production time than if the workcell is cold-started? Of the four hypotheses (see Table 8 in Appendix A), none of the null hypotheses were rejected, further exhibiting the stability of the manufacturing process. Box plots showing the distributions used for testing the hypotheses are shown in Figure 11.

7.2 Comparative Analysis of Measurement Data

Graphical comparisons will be regularly performed using the measurement data. Some example comparisons include: baseline measurements vs. the completed security level implementation, between each sequential technology installations, and between each operational configuration of a single technology. The pre-implementation baseline may also be used to compare the impact between each security level. Two graphical techniques will be utilized consistently for these analyses: bihistograms, and box plots. Both of these graphical methods eliminate the reliance on quantitative t-tests, as more information can be determined visually than can be provided by a single t-test.

A bihistogram is typically used to compare the location, scale, skewness, and outliers between two processes [12]. The bihistogram displays two histograms on a single figure, with one histogram mirrored across the y-axis enabling easier comparative analysis versus two overlaid histograms.

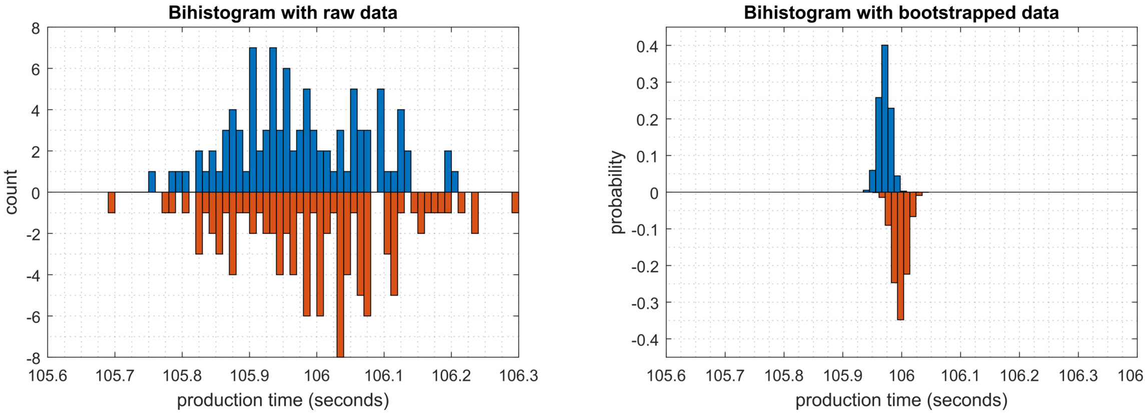

Fig. 5. Example bihistograms of production data. The first bihistogram uses raw production data, while the second uses bootstrapped production data.

The bootstrap method can be combined with the bihistogram to better understand the confidence interval for the mean, as a histogram with a small sample size may not provide a clear representation of the distribution. The method involves randomly selecting samples (with replacement) from the original data set, and calculating a statistic of the resulting subsample [12]. This process is repeated many times, with the statistics computed from each subsample used to create the bootstrap histogram.

The first plot in Figure 5 (left) shows a bihistogram of production times from two 100-part experiments, resulting in sparse histograms, making it difficult to visualize the sample mean. The second plot (right) shows the resulting bootstrap bihistogram from the same production data. The resulting bihistogram shows that the two experiments have similar distributions, and a deviation of means (around 30 milliseconds).

A box plot will be used to compare more than two sets of measurement data. This may include sequential comparisons (as shown in Figure 6), comparisons between unique tool configurations, or comparisons between different combinations of tools. Results from these graphical comparisons can be used to indicate if further testing is required to determine the source of the performance impact.

Other metrics and KPI will be analyzed as-required when implementing the security levels. For example, if the cause of a performance impact is suspected to be the increased network throughput (KPI 6.5 in [3]) on a specific ICS device, the measurement data from the baseline will be characterized and analyzed to determine its feasibility as a response variable. It is not necessary to analyze every KPI for all experiments because the divergence of a KPI from the expected value does not necessarily mean that the deviation has resulted in a performance impact to the manufacturing system.

Some of the tool operations generate short-term actions that complete within the normal processing time of a batch (and sometimes within the processing time of a single part). Tool actions that have this type of behavior will be manually scheduled and initiated at a specific time during production. Performance impacts from these types of tool actions will be evident during post-analysis of the data. For short-term operations, the KPI may not indicate a statistically-significant impact, but its short-term impact may be evident by analyzing high resolution time-series data on a trend plot (or similar time-series plot).

A detailed example of the comparative analysis process is shown in Appendix B. The baseline measurements for the example analysis were obtained from the low security level baseline, and the experimental measurements were created by imparting a network delay on the CRS network. A delay of 10 ms was enforced on all ingress and egress PLC traffic to simulate an impactful tool being installed on the PLC communications path (e.g., a queuing network-based firewall).

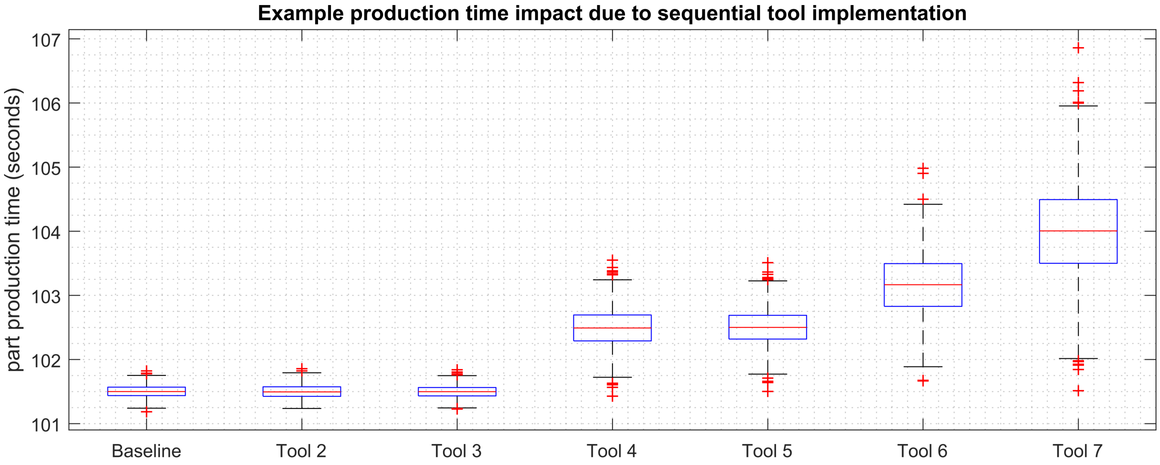

Fig. 6. An example box plot illustrating how increasing performance degradation of a manufacturing system may appear as tools are sequentially added to the system.

8. Conclusion

The implementation and validation of the Manufacturing Profile on the Cybersecurity for Smart Manufacturing Systems project Collaborative Robotics System, and its resulting performance impact measurements, will provide an excellent source of data for developing new cybersecurity guidance for the manufacturing sector, and numerous other industries that utilize ICS. The final results of this research will give industry the confidence it needs to successfully implement cybersecurity technologies within ICS without the worry of negatively impacting their manufacturing processes.

References

[1] Stouffer K, et al. (2017) Cybersecurity Framework Manufacturing Profile. National Institute of Standards and Technology, NISTIR 8183. https://doi.org/10.6028/NIST.IR.8183

[2] Candell R, Zimmerman T, Stouffer K (2015) An industrial control system cybersecurity performance testbed. National Institute of Standards and Technology, NISTIR 8089. https://doi.org/10.6028/NIST.IR.8089

[3] Zimmerman T (2017) Metrics and key performance indicators for robotic cybersecurity performance analysis. National Institute of Standards and Technology, NISTIR 8177. https://doi.org/10.6028/NIST.IR.8177

[4] Tang C (2017) Key performance indicators for process control system cybersecurity performance analysis. National Institute of Standards and Technology, NISTIR 8188. https://doi.org/10.6028/NIST.IR.8188

[5] United States Department of Homeland Security Cybersecurity and Infrastructure Security Agency. Critical Manufacturing Sector. [ONLINE], Available: https://www.dhs.gov/cisa/critical-manufacturing-sector.

[6] Zimmerman T (2018) Quantifying network performance impacts on industrial control systems. National Institute of Standards and Technology, NISTIR 8226. https://doi.org/10.6028/NIST.IR.8226

[7] ROS.org, https://www.ros.org/.

[8] Python.org, https://www.python.org/.

[9] Zimmerman T (2017) Machine Simulator - GitHub, https://github.com/TimZim-NIST/machine_sim.

[10] TCPDUMP/LIBPCAP public repository, http://www.tcpdump.org/.

[11] Wireshark, https://www.wireshark.org/.

[12] NIST/SEMATECH e-Handbook of Statistical Methods (2013), https://www.itl.nist.gov/div898/handbook/.

Appendix A: Characterization of the Workcell

Measuring the performance impact of cybersecurity technologies on the manufacturing process requires detailed knowledge of the process behavior. This knowledge can be obtained through an analysis called characterization, in which data from the manufacturing process is used to find the typical distribution and variation of the chosen KPI. Once the manufacturing process is characterized, performance impacts can be detected and quantified by comparing experimental data to the characterized data or baselines.

For this research, two potential response variables (KPI) were chosen as candidates: the part production time (2.1 in [3]) and the part cycle time (2.2 in [3]). Since parts must progress sequentially through the workcell, it is assumed that performance impacts to any process operation will likely be detected by one of these two KPI.

The first question that must be answered is, how many parts must be produced to accurately determine the mean of the KPI? The larger the sample size obtained during the characterization, the more accurately the process mean can be estimated. However, sample size is also restricted by the amount of time required to produce the sample. For each sample that is produced on the workcell, it can be said that the mean of the sample is within a certain amount of error to the population mean. But how many samples are required? This problem can be formally defined by the equation:

$$\overline{Y} - \delta \leq \mu \leq \overline{Y} + \delta \quad , \tag{1}$$

where μ is the estimated average of the KPI, $\overline{Y}$ the sample mean, and δ is the error of estimation. This is simply formalizing that the sample mean $\overline{Y}$ is within some error $\pm\delta$ of the population mean μ. Because the standard deviation of the manufacturing process can be estimated by using data previously obtained from the workcell, the following equation can be used to find the error of estimation:

$$\delta = \frac{\sigma}{\sqrt{N}} z_{1-\alpha/2} \quad , \tag{2}$$

where σ is the standard deviation of the manufacturing process, N is the sample size, and $z_{1-\alpha/2}$ is the critical value from the normal distribution for $1-\alpha/2$ (i.e., the confidence interval). This equation can be rewritten as an inequality to estimate the minimum quantity of samples required:

$$N \geq \left(\frac{z_{1-\alpha/2}}{\delta}\right)^2 \sigma^2 \quad . \tag{3}$$

To calculate the number of samples required, a confidence interval of 95% was chosen ($\alpha = 0.05$). At this point, the only remaining unknown is the error of estimation. Since the data used to calculate the response variables is captured by the supervisory PLC and its timestamp resolution is limited to 10 milliseconds, it was decided to use an error of estimation of 0.01 seconds. This results in the solution:

$$N \geq \left(\frac{1.96}{0.01}\right)^2 0.11^2 \geq 464 parts \quad .$$

Therefore, to estimate the mean part production time with an error estimation of less than 0.01 seconds with a probability of 95% on a process with an estimated standard deviation of 0.11, a minimum of 464 parts must be produced during the characterization phase. Since the sample size required is so large, it was decided to increase the sample size to 1000 parts.

The sample size was increased because: it allows for the total sample size to be evenly divided into smaller "batches", it enables sub-samples of the characterization data to be selectively combined to test hypotheses of the underlying manufacturing process, and it further increases the accuracy of the estimated mean.

The characterization sample was divided into ten equal batches of 100 parts. The production of each batch was then distributed across five calendar days, with two batches produced per day: one batch in the morning, and one batch in the afternoon. The first four batches were produced on a Thursday and Friday, and the last six batches were produced the following Monday, Tuesday, and Wednesday.

The batch distribution across multiple calendar days allowed for the following hypotheses to be evaluated from the data:

- Part production time is consistent across all batches.
- Part production time does not differ from day-to-day.
- Part production time of batches produced in the morning does not differ from batches produced in the afternoon.
- Part production time of batches produced during a calendar week does not differ from the batches produced during the following calendar week (i.e., after the workcell is shut down for an extended period of time and subsequently restarted).
- Part production time does not differ from the first twenty five parts of a batch to the last twenty five parts of a batch (i.e., the robots do not transfer parts at faster or slower rates depending on when the part was transferred).

At this point, the batches were produced and the production data captured. The production time for each individual part is shown in Appendix C, and is organized by batch number. The production time of the first part from each batch is removed from the data set because its production time is not representative of the later parts. This is a result of the workcell being purged when the manufacturing process is started. This results in a final characterization sample size of $n = 99$ parts per batch, and $N = 990$ parts total for all ten batches combined.

Exploratory Data Analysis

Now that the production data has been obtained, exploratory data analysis (EDA) is employed to gain insight into the manufacturing process. The EDA process is described in great detail in the NIST Engineering Statistics Handbook [12]. Most importantly, it must be determined if the manufacturing process is stable and consistent. If not, it will be impossible to effectively model the manufacturing process and measure any performance impacts.

The expectation for the EDA of the workcell KPI (specifically, the part production time KPI) is to fit a univariate model. This type of model has only two components: a constant term, and an error term. The constant term (which, as of this point, is unknown) represents the fixed location of the distribution, and the error term represents the expected amount of randomness in the process (the variance of which must also be fixed). The model also has other requirements, which will be addressed throughout this section. The univariate model is defined as:

$$Y_i = C + \varepsilon \quad ,$$

where C is the estimated mean and ε is the error term.

Scatter Plot

The resulting scatter plot of part production time vs. part number is shown in Figure 7. From this plot, it is evident there are no large shifts in location, variation, or any significant outliers. The manufacturing process appears stable and consistent (at least, for this response variable).

A horizontal line (red) representing the mean part production time was added to the plot to aid the analysis or variance around the mean. The plot appears to show a larger quantity of parts with values above the mean line than below, and a larger variance for parts with values above the mean. Further analysis is required to verify this hypothesis.

4-Plot

The next step is to verify the process is in-control (i.e., stable), and to determine if the resulting data is normally distributed. This is performed with a 4-plot, which combines four plots: a run-sequence plot, lag plot, histogram, and normal probability plot. The resulting 4-plot from the characterization dataset is shown in Figure 8.

As discussed in the previous section, the run-sequence plot shows there are no large shifts in location, variation, or any significant outliers, suggesting that the process is stable and in-control. The fitted trend line (shown in red) does appear to show a slight upward trend, but it is not statistically significant.

The lag plot is used to determine if the data are random. It is produced by comparing the production time of part Y_i to the production time of part Y_{i+1}. If there is a relationship between the production times of successive parts, an identifiable structure will result in the lag plot. If there is no relationship, the data should appear random and reasonably

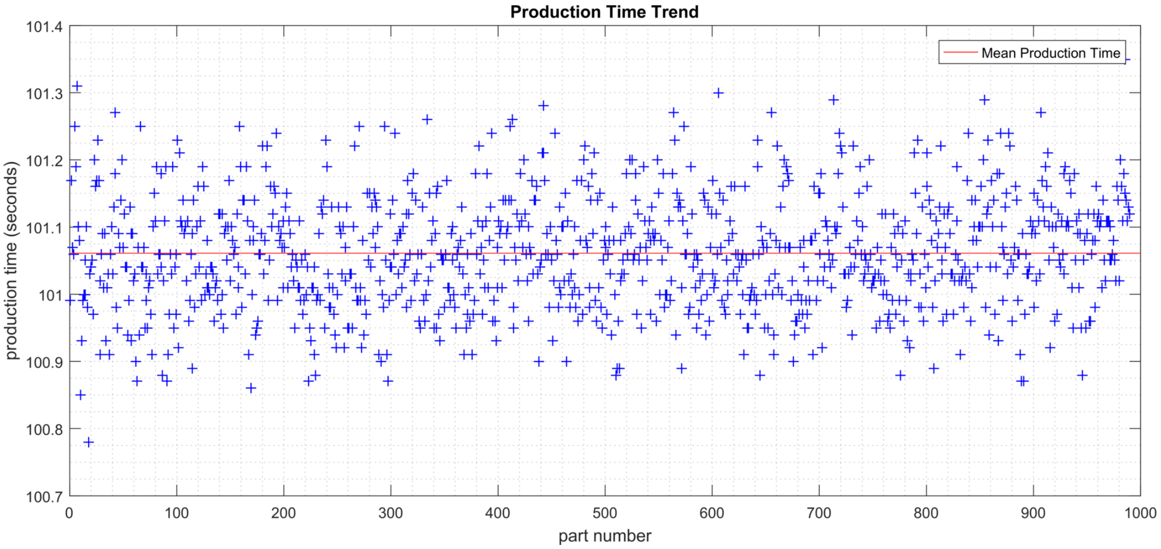

Fig. 7. Scatter plot of the part production time KPI from the characterization process data. The horizontal red line denotes the mean part production time.

symmetric. The lag plot created from the characterization data in Figure 8 does not show any obvious structure and does not show any serial correlation, therefore implying that the data is random.

The histogram plot is used to visualize the distribution of a data set, and to determine which model is best for statistically representing the data. Many features of the distribution can be determined from the plot including: the center and spread, if the data is skewed, and if the data is multimodal. The histogram created from the characterization data implies that the data is normally distributed, is slightly skewed towards the right tail, and is unimodal.

These observations are also confirmed by the normal probability plot, which plots the data set against a theoretical normal distribution. Evidence of the slight skew towards the right tail is visible in the top-right of the plot, which shows the points trending below the red line representing the theoretical normal distribution.

Analysis of the right skew is beyond the scope of this paper, but it is believed that its presence is likely because of the varying performance of the manufacturing control systems and networking equipment (known in this context as operational technology, or OT). There is a theoretical minimum amount of time a part can travel through the process due to the OT limitations, but there is no theoretical limit to the amount of delay that can be imparted by the OT. However, for the purposes of characterizing the manufacturing process, the data is accurately represented by the normal distribution.

Further analysis was performed to verify the randomness of the data. The lag plot included in the four plot (Figure 8) is limited to a “lag of 1” (i.e., the data point Y_i is correlated with the point Y_{i+1}). This does not, however, show if there are correlations between data points Y_i and Y_{i+2}, Y_{i+3}, etc. These relationships can be visualized using an autocorrelation plot, which is shown in Figure 9.

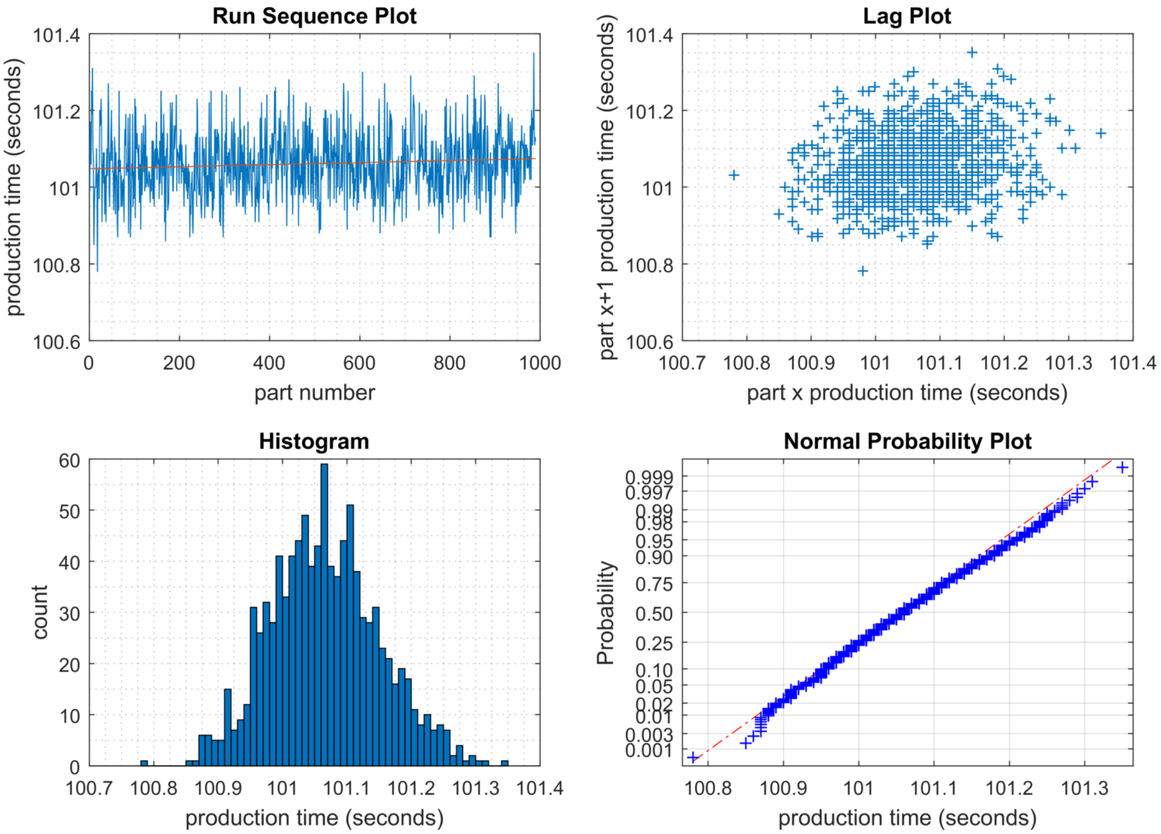

Fig. 8. Four plot of the ‘part production time’ KPI from parts produced during workcell characterization.

The plot shows the autocorrelation for data points from lag 1 to lag 100, along with the significance bands for 95% and 99% significance. There appears to be a statistical significance for lag 1, and some significance for lag 2, 3, and 4 above the 95% confidence band. The data for lag 1 is the same that is visualized in the lag plot from the 4-plot shown in Fig. 8, and was already determined to not have any obviously visible structure.

To further test the statistical significance shown at lag 1, it must be realized that the autocorrelation plot assumes the data is in sequential run order. This becomes an issue when correlating data points from different batches, which happens to be the case with the part production time data set, since it is composed of 10 sequential batches combined into a single dataset. Operationally, the manufacturing process is halted and the workcell is shut down between the batches, meaning there is no inherent relationship between the data points being compared.

To resolve this issue, autocorrelation plots were created for each batch up to lag 25, as shown in Fig. 10. While some of the plots do show levels of correlation greater than the confidence intervals, the correlations are neither consistent between batches nor of great enough significance to imply that the data is not random.

It is hypothesized that the significance at lag 1 through lag 4 in the batches shown in

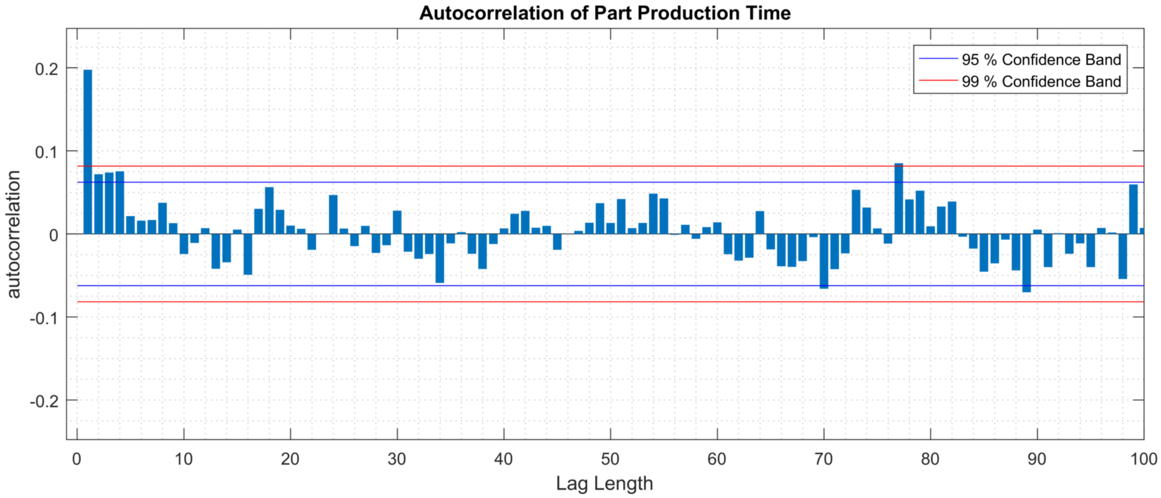

Fig. 9. Autocorrelation plot showing lag 1 to lag 100 of the production time KPI.

Fig. 10 are due to the parallel operations of the workcell and its underlying OT systems. For example, assume there are two parts in the workcell, one at Station 1 and one at Station 2. Since both of those stations are serviced by a single robot, if the part being transferred out of Station 2 is delayed, the part waiting at Station 1 will incur the same amount of delay before it is transferred. Therefore, the effect of communications delays between the OT systems is the likely cause.

However, during normal operating conditions the probability of a part accumulating delay is unlikely (as evident in the data). Future experiments may increase the probability of OT-related delay, resulting in a stronger correlation at lag 1 through 4. Therefore, lag analysis should continue to be performed during the implementation.

Based on the data set from the characterization batches and the subsequent EDA, the process is said to be in-control, stable, and random. This validates the assumption that the process can be modeled using the normal distribution, and the univariate model can be used to model the production time of parts produced in the workcell.

The production time KPI will continue to be used as the primary response variable for indicating a performance impact to the manufacturing system. Any other potential response variables used during the implementation must also undergo the same characterization process.

Batch Evaluations

Distributing the batches across multiple days, and the time of day, enabled multiple hypotheses to be tested during the characterization process. The hypotheses and results are listed in Table 8, and the box plots of the data are shown in Figure 11. None of the evaluations produced statistically significant results, failing to reject any of the null hypotheses, further displaying the stability of the manufacturing process.

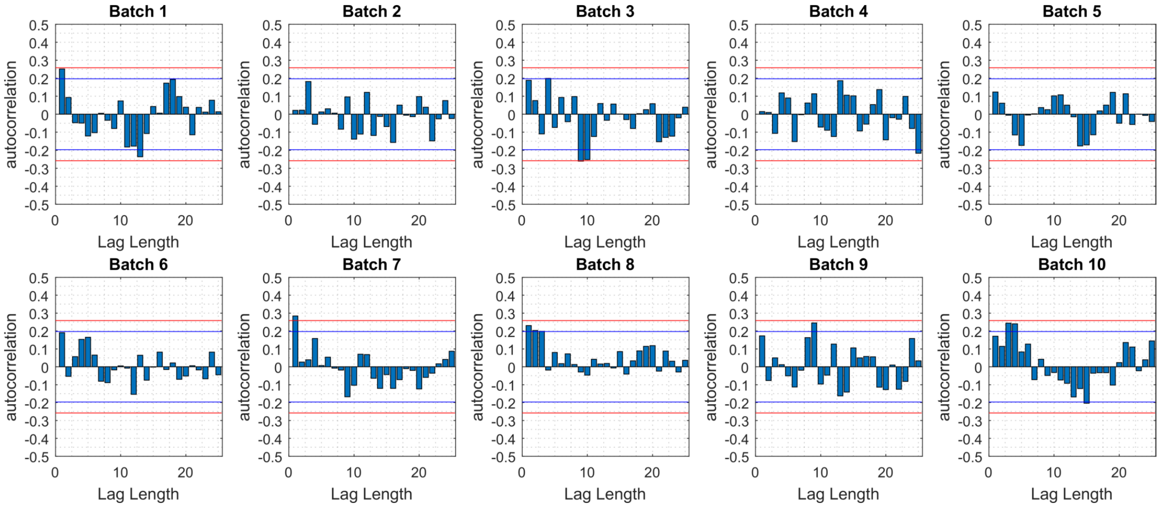

Fig. 10. Autocorrelation plots for each batch produced during the workcell characterization. Each plot includes lag 1 to lag 25.

Table 8. Hypotheses tested on the characterization data.

Plot	Null Hypothesis (H_0)	Reject H_0?
Batch v. Batch	KPI does not change between batches.	NO
Morning v. Afternoon	KPI does not change from morning to afternoon.	NO
Day v. Day	KPI does not change between calendar days.	NO
Day v. Day	KPI does not change after an extended shut down.	NO
First 25 v. Last 25	KPI does not change between production stages.	NO

Sample Size for Cybersecurity Technology Implementations

Just as Equation 3 was used to calculate the sample size required for characterization, the equation can also be used to estimate the bound on the error for a specified sample size. This information can then be used to find the optimal sample size and bound of the error.

A table was created (see Table 9) to determine the bound of the error using Equation 1. A confidence interval α=95 % was used, resulting in the equation:

$$\delta = \frac{\sigma}{\sqrt{N}} z_{1-\alpha/2} = \frac{0.846}{\sqrt{N}} z_{0.025} \quad .$$

After reviewing the results, a sample size of 35 parts was selected for future experiments due to the relatively quick experiment time of 1 hour, and the minimal bound on the error of ±0.028 seconds. An important consideration for the error bound is the 0.01 seconds PLC task cycle time. With the chosen sample size of 35 parts and a task cycle time of 0.01 seconds (e.g., a timer resolution of 0.01 seconds), the bound on the error of estimation can be rewritten in terms of timer periods, or ±3 PLC task cycles. In statistical terms, there is a

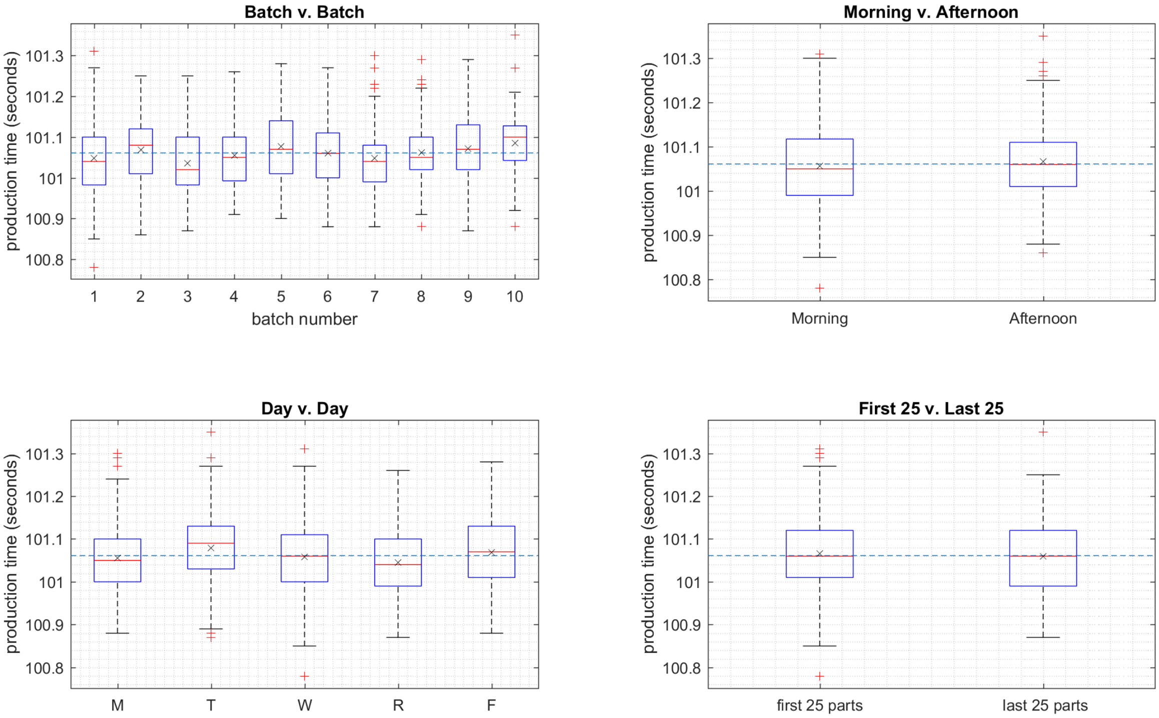

Fig. 11. Box plots of the batch evaluations using production time data from the workcell characterization. The first plot (‘Batch v. Batch’) compares the distribution of part production times between each batch. The second plot (‘Morning v. Afternoon’) compares the distribution of parts produced in the morning vs. the afternoon. The third plot compares the distribution of parts produced on each day. The fourth and final plot shows the distribution of parts produced at the beginning of the batch vs. the end of the batch.

95 % probability (confidence interval) that a sample size of 35 parts covers or contains the true mean.

For proper analysis of the results, it is important to understand how the error estimation increases as the standard deviation changes. This relationship is shown in Figure 12. The plot identifies two examples demonstrating the increasing error of estimation as the standard deviation of the manufacturing process also increases. The plot shows how the error of estimation between the two different values of σ causes the error of estimation (δ) to increase from ±0.028 to ±0.053. Therefore, it is important that the standard deviation be calculated for each KPI to understand the effect it has on the distribution, and if the resulting calculations are statistically significant.

Appendix B performs a comparative analysis of the standard deviation from example production data, and calculates the critical value that indicates a statistically significant change to the standard deviation (with a confidence interval of 99 %).

Table 9. Bound on the error of estimation and the amount of workcell time required to obtain a sample of size N at 95% confidence.

N	$\delta_{production}$	production time (minutes)	N	$\delta_{production}$	production time (minutes)
5	±0.074	35	60	±0.021	79
10	±0.052	39	65	±0.021	83
15	±0.043	43	70	±0.020	87
20	±0.037	47	75	±0.019	91
25	±0.033	51	80	±0.019	95
30	±0.030	55	85	±0.018	99
35	±0.028	59	90	±0.017	103
40	±0.026	63	95	±0.017	106
45	±0.025	67	100	±0.017	110
50	±0.023	71	990	±0.005	820
55	±0.022	75			

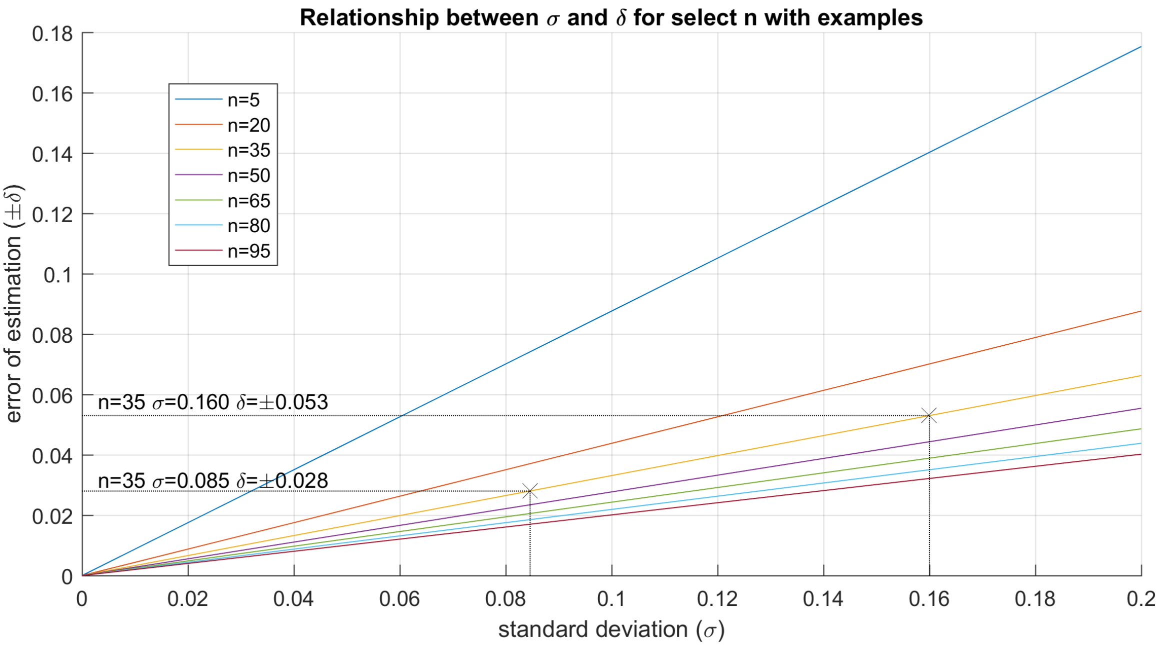

Fig. 12. Plot displaying the relationship between the standard deviation σ, error of estimation δ, and sample size n. The two examples shown demonstrate the increasing error of estimation as the standard deviation of the manufacturing process becomes larger. In this case, an increase of σ to 0.160 causes δ to increase to ±0.053, much larger than the characterized δ of ±0.028.

Appendix B: Example comparative analysis

To validate the methodology described in this report, a performance impact was deliberately imparted on the CRS. More specifically, a constant network delay was forced on all packets received by and transmitted from the PLC. This network link between the workcell and its PLC is critical to maintain because of the supervisory functions it performs: communicating with machining stations to gather operational data, and transmitting jobs to the robot controllers.

To generate the network delay, an IXIA ANUE network emulator was physically connected into the network path between the PLC and the CRS network. The emulator was configured to impart a 10 millisecond delay on all incoming and outgoing traffic. This network performance impact is comparable to a network-based firewall that queues network traffic to be analyzed against a set of rules. This results in a round-trip time (RTT) of 20 milliseconds above the baseline RTT (typically $<$ 1 millisecond).

The workcell was exercised through the normal production process with all measurement systems activated, producing a batch of 35 parts. After the workcell completed the batch, all measurement data was captured and analyzed.

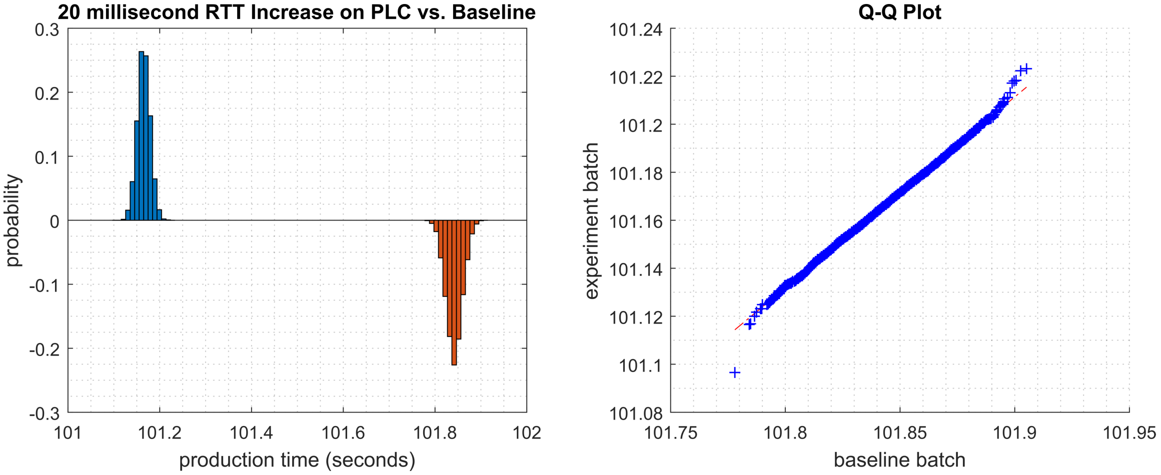

Fig. 13. Bihistogram showing the production time probability of the baseline measurements (blue-colored histogram) and the post-implementation measurements (orange-colored histogram).

The first step of the comparative analysis procedure was to determine if the technology impacted the performance of the manufacturing process. A bihistogram plot was produced using production time data (KPI 2.1 [3]) from the experiment and production time data from a previous baseline. The resulting bihistogram is shown in Figure 13.

The bihistogram shows an obvious performance impact to the manufacturing process: the average time required to produce a part increased from 101.17 seconds during the baseline to 101.84 seconds during the experiment. The bihistogram plot indicates that the null hypothesis (the tool does not affect the performance of the manufacturing process) should be rejected, and the alternate hypothesis (the tool affects the performance of the manufacturing process) should be accepted.

A quantile-quantile (Q-Q) plot is used to determine if the distribution of the two data sets are different, and can be used to determine if there are shifts in location, shifts in scale, changes in symmetry, or the presence of outliers [12]. The similar "normal probability plot" was used during the characterization process, which compares the data set to the normal probability distribution. As with the bihistogram, the Q-Q plot (shown in Figure 13) shows an obvious performance impact (i.e., shift in location). More importantly, the linearity of the plot indicates that the distribution has not changed (i.e., the typical variation of part production times has not been affected by the cybersecurity tool).

It is not obvious from the bihistogram or Q-Q plot if the two data sets have the same standard deviation. This can be determined numerically, as described in §7.3.2 in [12]. First, the critical value from the F table for a significance level of 1% with degrees of freedom $v_1 = 33$ and $v_2 = 33$ is found.

$$F_{1-\alpha,v_1,v_2} = F_{0.99,33,33} = 2.29 \quad .$$

The F statistic is computed using the variance of each data set (in this case $\sigma_b = 0.084$, and $\sigma_e = 0.107$).

$$F = \frac{s_b^2}{s_e^2} = \left(\frac{0.084}{0.107}\right)^2 = 0.616 \quad , \tag{4}$$

where s_b^2 is the variance of the baseline part process times, and s_e^2 is the variance of the experimental part process times. Since F is less than the critical value $F_{1-\alpha,v_1,v_2}$, the alternate hypothesis (the standard deviations of each data set are different) can be rejected. The difference in standard deviations between the baseline measurements and the experimental measurements are, therefore, not statistically significant.

The formula can be rewritten to determine *what* value of standard deviation would result in the null hypothesis being rejected (when the experimental data is compared to the true baseline). This is possible since the experiment methodology has a consistent sample size of 35 parts, and the manufacturing process is known to be stable.

Rewriting Eq. 4 for s_e^2,

$$s_e^2 = F_{0.99,33,33} \cdot s_b^2 \quad ,$$

which is equivalent to,

$$\sigma_e = \sigma_b \cdot \sqrt{F_{0.99,33,33}} \quad ,$$

and finally,

$$\sigma_e = 0.085 \cdot \sqrt{2.29} = 0.129 \quad .$$

Therefore, if the standard deviation of the experimental batch (σ_e) with a sample size of 35, and σ_b of the true baseline ever exceeds a value of 0.129, the null hypothesis is rejected, meaning that the installed technology or tool impacted the performance of the manufacturing process.

It is important to note that the underlying mechanisms allowing the performance impact to propagate to the manufacturing process cannot be determined from the EDA, nor is it meant to be. The EDA process is meant to indicate if the manufacturing process has been impacted, regardless of the underlying mechanism. However, the cause of the impact *is* known: the cybersecurity technology, or its configuration. This is the most important indicator, as most manufacturers consider maintaining production goals and quality goals to be a primary business objective.

In terms of the Manufacturing Profile implementation at NIST, it is important to determine *how* the technology caused the impact. This information will be used as the foundation for future guidance language and to enhance the Manufacturing Profile. A detailed investigation of how the workcell ICS operation was impacted by this experiment is described in NIST IR 8226 [6].

Appendix C: Production Data

Table 10. Part production time data measured during the characterization process. The unit of measure is seconds, with resolution down to hundredths of seconds.

Part No.	Batch 1	Batch 2	Batch 3	Batch 4	Batch 5	Batch 6	Batch 7	Batch 8	Batch 9	Batch 10
1	100.99	101.06	101.01	100.99	100.96	100.98	101.00	100.91	100.98	100.97
2	101.17	101.23	101.11	101.04	100.93	100.95	101.13	101.16	101.03	101.06
3	101.07	100.92	101.07	101.12	101.01	101.05	100.96	101.09	100.98	101.03
4	101.06	101.21	101.13	101.04	101.07	101.08	101.07	100.99	101.13	101.03
5	101.25	101.11	101.15	101.14	101.11	101.07	101.12	101.02	100.96	101.04
6	101.19	101.10	101.01	101.06	101.09	101.10	101.17	101.15	101.06	101.11
7	101.31	101.14	101.11	101.24	101.10	100.97	101.03	101.15	101.01	101.09
8	101.10	101.09	101.09	101.00	101.01	101.06	101.01	101.06	101.16	101.10
9	101.08	101.10	101.04	101.08	101.14	101.08	101.02	101.18	101.21	100.98
10	100.85	101.02	101.06	101.03	101.05	101.03	101.12	100.92	101.04	101.15
11	100.93	101.11	101.02	101.09	101.14	100.94	101.06	101.02	101.14	101.09
12	100.99	100.96	100.95	101.09	101.09	101.13	101.30	101.04	101.12	101.02
13	101.00	100.99	101.05	101.01	101.10	101.08	101.15	101.11	101.08	101.14
14	101.00	101.06	101.04	101.10	101.14	101.00	101.07	101.10	100.98	101.11
15	101.05	100.89	101.01	101.04	101.25	100.88	100.99	101.08	100.89	101.03
16	101.10	101.09	101.11	101.05	101.14	100.89	101.14	101.06	101.02	101.27
17	100.98	101.14	101.01	100.98	101.26	100.96	100.98	101.04	101.11	101.17
18	100.78	100.98	100.99	101.12	100.98	100.89	100.97	101.01	101.07	101.13
19	101.03	101.03	101.05	101.17	100.95	101.08	101.05	101.05	101.09	101.11
20	101.04	101.10	100.97	101.12	101.07	101.02	101.00	101.29	101.13	101.19
21	101.05	101.16	100.99	101.06	101.05	101.03	101.02	101.10	101.22	101.16
22	100.97	101.04	101.06	100.99	101.12	101.05	101.08	101.10	101.05	101.21
23	101.20	101.11	101.00	101.15	101.02	101.03	101.11	101.09	101.03	101.05
24	101.16	100.99	101.04	101.18	101.10	101.09	101.16	101.10	100.96	101.10
25	101.17	101.19	100.87	101.02	101.18	101.18	101.09	101.18	101.12	100.92
26	101.23	101.16	100.95	101.14	101.11	101.06	101.03	101.24	101.07	101.13
27	101.17	101.00	100.93	101.03	101.07	101.10	101.16	101.22	100.97	101.10
28	100.91	101.03	101.01	101.06	101.20	101.20	101.06	101.21	101.12	101.07
29	101.02	101.08	101.01	100.96	101.12	101.09	101.02	101.13	101.13	100.96
30	101.06	101.01	101.04	100.95	101.06	101.20	101.06	101.03	101.09	101.10
31	101.09	101.01	100.91	101.06	101.02	101.05	101.04	101.03	101.03	100.97
32	101.01	101.00	100.88	101.06	100.96	101.06	100.97	100.98	101.03	101.12
33	101.09	101.04	100.99	101.13	101.10	101.14	101.05	100.99	100.97	101.05
34	100.93	101.04	101.00	100.99	101.07	101.18	101.06	101.03	101.02	101.04
35	101.01	101.03	101.09	100.96	101.06	100.97	101.00	101.07	101.06	101.11
36	101.03	100.96	100.99	101.00	101.17	101.12	100.91	101.07	101.19	101.06
37	100.91	101.01	101.13	101.26	101.03	101.05	101.16	101.04	101.11	101.03
38	101.06	101.02	101.12	101.02	101.14	101.07	101.05	100.94	101.05	101.18

39	101.10	100.99	101.14	101.15	101.14	100.95	100.95	101.22	101.02	101.19
40	101.03	101.12	100.95	101.04	101.16	101.02	100.95	101.04	100.91	101.10
41	101.13	101.12	101.23	101.08	101.03	101.19	100.99	101.06	101.02	101.11
42	101.27	100.97	101.19	101.11	100.90	101.11	101.00	101.11	101.03	101.10
43	101.18	101.12	101.03	100.95	101.04	101.07	101.01	101.08	100.98	101.17
44	100.98	101.00	101.01	101.09	101.05	101.05	101.06	101.15	101.04	101.15
45	100.95	100.95	100.98	100.95	101.21	101.09	101.00	101.19	101.02	101.01
46	101.09	101.09	100.98	101.15	101.21	101.11	100.97	101.06	101.17	100.95
47	101.07	101.10	101.13	101.16	101.28	101.08	100.99	101.10	101.24	101.18
48	101.14	101.05	101.01	100.97	101.17	100.97	101.19	101.02	101.03	101.05
49	101.20	101.10	100.97	100.93	101.06	101.13	101.23	101.06	100.96	101.11
50	101.07	101.17	100.94	101.06	101.09	101.09	101.06	101.03	101.14	101.09
51	101.12	101.02	100.92	101.05	101.06	101.08	100.88	101.05	101.01	101.01
52	101.04	101.10	101.00	101.05	100.98	101.09	101.10	101.16	101.08	101.12
53	101.04	101.08	101.13	100.98	101.10	100.96	101.08	101.04	101.10	101.09
54	100.94	101.06	101.10	101.06	101.05	101.20	101.07	101.20	101.14	100.95
55	100.99	101.07	101.09	101.17	100.93	101.15	100.95	101.06	101.08	100.88
56	101.13	100.97	101.03	101.01	101.02	101.00	101.00	101.01	101.04	101.05
57	101.09	101.12	101.10	101.01	101.24	101.08	101.07	101.02	101.10	101.08
58	100.93	101.01	100.92	100.99	101.08	101.05	100.95	101.02	101.10	100.95
59	101.09	101.04	101.02	101.09	101.06	101.18	101.02	101.03	101.16	101.09
60	101.02	101.25	101.07	101.05	100.98	100.98	101.09	100.95	101.10	101.12
61	101.03	101.19	101.13	101.10	101.13	101.00	101.27	101.07	101.20	100.96
62	100.90	100.98	100.95	101.00	101.01	101.00	101.00	101.03	101.29	100.96
63	100.87	101.14	100.95	100.98	101.00	101.00	100.91	100.96	100.98	101.06
64	101.07	101.14	101.03	100.95	101.10	101.07	100.94	101.05	101.12	101.08
65	100.94	101.03	101.03	101.06	101.10	101.11	101.03	101.02	101.23	101.12
66	101.25	101.19	100.95	101.07	101.02	100.99	101.08	101.12	101.13	100.98
67	101.04	101.11	101.10	101.04	101.15	101.02	101.22	101.00	100.94	101.08
68	101.04	100.91	101.22	101.08	100.90	101.18	101.01	101.02	101.04	101.16
69	101.07	101.08	101.01	101.13	101.11	101.27	101.07	101.12	101.17	101.11
70	100.95	100.86	100.99	101.03	101.02	101.23	101.13	101.00	101.10	101.12
71	101.03	101.00	100.99	100.91	100.98	100.94	101.00	100.96	101.14	101.10
72	100.99	101.10	101.25	101.00	101.03	101.15	101.16	100.97	100.97	101.11
73	100.95	101.08	100.99	100.95	101.00	101.13	101.20	100.95	101.04	101.01
74	101.01	101.14	101.01	101.06	101.06	101.14	101.07	101.07	101.18	101.14
75	101.02	100.94	100.92	100.97	101.01	101.13	101.19	100.99	100.93	101.05
76	100.97	100.95	101.08	101.01	101.10	101.01	101.18	101.12	101.18	101.02
77	100.91	100.96	100.94	100.96	100.98	100.89	101.17	101.09	101.24	101.11
78	101.10	101.06	100.99	101.09	101.17	101.11	101.07	100.97	101.10	101.11
79	101.15	101.06	100.97	100.91	100.96	101.25	101.07	101.06	100.99	101.11
80	101.11	101.09	101.15	100.97	101.00	101.06	100.99	101.04	101.07	101.05
81	101.19	101.22	101.05	100.99	101.08	101.06	101.02	101.10	101.06	101.10
82	101.03	101.03	101.15	101.18	101.17	101.09	100.90	100.99	101.13	101.05

83	101.06	101.08	101.04	101.06	100.97	100.97	100.94	100.88	101.15	101.06
84	101.06	101.19	101.13	101.11	101.20	100.98	100.96	101.05	101.03	101.08
85	101.04	101.22	101.09	101.22	101.22	101.19	100.96	101.09	101.24	101.05
86	101.18	100.99	101.15	101.03	101.08	101.12	101.00	100.96	101.22	101.02
87	100.88	100.99	101.07	101.07	101.07	101.04	100.97	101.11	101.05	101.16
88	101.02	101.05	101.14	101.10	101.16	101.06	101.07	100.99	100.96	101.14
89	101.11	101.16	101.13	100.97	101.01	101.01	100.99	100.93	101.04	101.02
90	101.19	101.10	101.10	101.05	100.98	101.06	100.97	101.03	101.07	101.20
91	100.87	101.12	100.91	101.01	101.14	100.95	101.04	100.92	101.16	101.16
92	100.91	101.16	101.02	101.11	101.18	101.05	101.08	101.02	101.14	101.11
93	100.97	101.15	100.90	101.18	101.09	100.97	100.95	101.03	100.99	101.18
94	101.06	101.24	101.00	101.01	101.21	101.02	100.99	101.06	101.06	101.15
95	101.03	100.99	100.98	101.08	100.99	100.99	100.97	101.05	101.06	101.35
96	101.19	101.10	101.25	101.13	100.96	101.03	101.02	101.21	100.91	101.14
97	100.95	101.08	101.05	101.24	101.14	101.15	100.99	101.15	100.87	101.11
98	101.04	101.09	100.91	100.96	101.01	101.14	101.08	101.09	101.09	101.13
99	100.97	101.07	100.87	100.99	101.15	100.94	101.08	101.23	100.87	101.12