



Defending against product-oriented cyber-physical attacks on machining systems

Mohammed S. Shafae¹ · Lee J. Wells² · Gregory T. Purdy³

Received: 30 July 2018 / Accepted: 17 April 2019 / Published online: 22 November 2019
© Springer-Verlag London Ltd., part of Springer Nature 2019, corrected publication 2019

Abstract

Industry 4.0 and its related technologies (e.g., embedded sensing, internet-of-things, and cyber-physical systems) are promising a paradigm shift in manufacturing automation. However, with a continual increase in device interconnectivity, securing these systems becomes crucial. As these systems evolve, opportunities for cyberattacks extend to include attacks that can physically alter parts (Product-Oriented C2P attacks). Fortunately, since these cyber-physical attacks affect the physical world, there exists potential to detect an attack through its physical manifestation. Typically, in manufacturing, quality control (QC) systems are used to detect quality losses or deviations from nominal. This paper proposes that QC tools can be adapted to act as physical detection layers as part of a defense-in-depth strategy (common IT security strategy) that increases the difficulty/cost required for a successful attack. However, effectively designing physical detection layers requires understanding the extent to which attacks can (and cannot) be designed to avoid detection. In response, this paper proposes a machining specific attack design scheme and an attack design designation system (ADDS) that provides the structure to populate a wide variety of potential attacks. To illustrate the importance of applying a defense-in-depth strategy for machining, a case study is conducted with several realistic attacks against an example machining process that collects in-situ process data. Within this case study, the proposed ADDS is employed to systematically describe how these attacks could be designed to avoid detection. Finally, through this exploration, this paper shows how employing process-domain knowledge to understand the effects of Product-Oriented attacks on process physics can further aid in detection layer designs.

Keywords Cyber-physical manufacturing systems · Cyberattacks · Machining · Process monitoring · Quality control

1 Introduction

In modern manufacturing, the growing adoption of Industry 4.0 and its related technologies (e.g., internet-of-things, cyber-physical systems, cloud computing) is introducing an unprecedented level of connectivity for manufacturing cyber-physical systems [1, 2]. Unfortunately, this has also created opportunities for adversaries to commit non-traditional

cyberattacks that aim to inflict harmful physical changes in a product, a process, and/or an entire production system [3]. An example of this type of cyberattacks was conducted on a German steel mill in 2014 [4]. Adversaries gained access to the plant's network, causing multiple physical system component failures and massive physical damage. Another example was demonstrated on an additive manufacturing process by Sturm et al. (2014) [5]. In their attack, an internal void was added into a tensile test specimen by altering its .STL file resulting in a printed part that exhibited a significant loss in strength.

Generally, cyber-physical systems, such as advanced manufacturing systems, can be defined as systems that consist of highly integrated and coordinated cyber and physical entities. Cyberattacks against these systems (referred to as cyber-physical attacks) can target either a cyber or a physical entity to affect either a cyber or a physical entity. As discussed by Yampolskiy et al. (2012; 2013) [6, 7], cyber-physical attacks can be described by the domain that the attack is targeting/

✉ Mohammed S. Shafae
shafae1@email.arizona.edu

¹ Department of Systems and Industrial Engineering, University of Arizona, Tucson, AZ 85721, USA

² Department of Industrial and Entrepreneurial Engineering and Engineering Management, Western Michigan University, Kalamazoo, MI 49008, USA

³ Department Industrial and Systems Engineering, Auburn University, Auburn, AL 36849, USA

influencing (cyber or physical domain) and the domain that is being effected/victimized (cyber or physical domain), which results in the following attack categories: Physical-to-Cyber (P2C), Cyber-to-Cyber (C2C), Cyber-to-Physical (C2P), or Physical-to-Physical (P2P). For instance, the aforementioned attack on the tensile test specimen is an example of a C2P attack in which the influenced entity was cyber (part's .STL file) while the victimized entity was physical (the tensile test specimen).

While understanding and developing defense strategies for each of the aforementioned cyber-physical attack categories is equally as important, this paper focuses on a specific type of C2P attacks, namely, Product-Oriented C2P attacks. In manufacturing, C2P attacks can be generalized as attacks that are either Process-Oriented or Product-Oriented. Process-Oriented attacks aim to disrupt manufacturing processes, which can be accomplished through numerous C2P attacks, such as destroying/incapacitating equipment, canceling incoming raw material shipments, or activating fire suppression systems. Product-Oriented attacks aim to maliciously alter a part's design intent with the hope that the altered part will reach its customer. Product-Oriented C2P attacks can degrade a part's quality and/or decrease its functional performance and/or reliability. Such effects could be catastrophic in terms of consumer safety and financial losses (e.g., increased warranty costs, injury settlements, and/or loss of customer trust) [8, 9]. This work focuses on Product-Oriented C2P attacks against production systems and assumes that all prototypes used during system design have not been subjected to an attack.

Efficiently defending against cyber-physical attacks, such as Product-Oriented C2P attacks, requires a clear understanding of how these attacks could be designed and implemented [10]. For the cyber domain, multiple repositories exist to catalog common attack patterns (designs) targeting software/Information Technology (IT) systems to assist the community in developing more secure systems [11]. Example repositories include the Common Weakness Enumeration (CWE) list [12], the Common Vulnerabilities and Exposures (CVE) list [13], the NIST National Vulnerability Database (NVD) [14], and the Common Attack Pattern Enumeration and Classification (CAPEC) repository [15]. For the manufacturing physical domain, unfortunately, there is little to no focus on understanding how the physical component(s) of Product-Oriented C2P cyberattacks can be designed.

Another aspect of defending against cyber-physical attacks is to adopt specific defensive countermeasures (both detective and protective) for both the cyber and the physical domains of the system. Multiple cyber detective and protective defense layers implemented throughout the IT system have been commonly accepted in practice. The use of multiple defense layers helps to create a more robust security system that results in unattractive (costly) targets for potential adversaries. This

defense strategy is usually referred to as defense-in-depth in traditional cyber-security literature [10, 16]. In the emerging research area on manufacturing security, apart from recent works on attacks detection at the process level using in-situ process monitoring systems, a holistic discussion on potential physical detection layers within the defense-in-depth framework is lacking.

As just mentioned, one of the evolving areas within the manufacturing security field focuses on detecting Product-Oriented C2P attacks at the process level. Monitoring techniques can be developed to detect physical process changes (induced by an attack) by monitoring in-process variables (e.g., vibration, temperature, power consumption). In this research, these variables are treated as "side-channels", which have proven to be a valuable approach to detecting Trojans in integrated circuits [17]. While producing significant results, this research area suffers from three key shortcomings: (1) The majority of current efforts have been limited to case studies implementing arbitrary attacks to demonstrate the use of traditional data-driven methods to detect attack induced process anomalies, (2) There is a lack of systematic understanding of how the physical components of Product-Oriented C2P attacks can be designed, and (3) There has been no focus leveraging physical process knowledge to understand the relationship between attacks and their effects on in-process variables. It should be noted that further details regarding these shortcomings are discussed in Section 2.

In addition to the abovementioned limitations, there is a disproportionately large amount of research in this area being devoted to securing additive manufacturing (AM) technologies. Despite the significant advances in AM over the past decade, machining (subtractive manufacturing) still remains, and will remain for the foreseeable future, the backbone of manufacturing. This is true especially in the case of mass production of low complexity and high precision parts [18]. This is evident in the steady demand for machining technologies, which has a \$7.5 billion expected expenditure on new machine tools in 2018, an increase of 5% compared to estimates for 2017 [19]. Additionally, development of hybrid manufacturing technologies is a growing area of interest as the solution to attain the flexibility of additive manufacturing with the precision of machining technologies [18]. To ensure that all key manufacturing processes are secure, now and in the future, machining infrastructure should be a priority for the industry and the manufacturing security research community.

To address the shortcomings described in this section, four contributions in this paper help advance our understanding of Product-Oriented C2P attacks. First, this paper proposes the adoption of the defense-in-depth strategy/framework to utilize current quality control (QC) resources as physical detection layers to defend against Product-Oriented C2P attacks. Second, an attack design scheme is proposed to systematically describe key design elements of Product-Oriented C2P attacks

against machining. Based on this scheme, a machining specific Attack Design Designation System (ADDS) is proposed to provide a structure to populate a wide variety of potential attacks. Third, to provide practitioners with an example of how use the ADDS to develop relationships between Product-Oriented C2P attacks and physical defense-in-depth strategies, an experimental investigation is conducted. In this case study, several realistic Product-Oriented C2P attacks, designed according to the proposed attack design scheme, are implemented against an example machining process. Finally, the use of a side-channel driven monitoring system is explored to show the effectiveness of employing physical process knowledge to understand the relationship between the attacks and their effects on process physics.

The remainder of this paper is organized as follows. In Section 2, previous relevant research is reviewed. Section 3 introduces the concept/framework of physical defense-in-depth for manufacturing systems. The proposed attack design scheme is discussed in Section 4. The experimental investigation is described in Section 5. The side-channels monitoring system and the experimental investigation results are discussed and reported in Sections 5 and 6, respectively. In Section 7, several remarks draw the paper to its conclusion and potential future work is identified.

2 Literature review

Manufacturing cyber-security is currently receiving significant attention from the research community. With respect to Product-Oriented C2P attacks, the literature can be separated into three basic categories that focus on (1) constructing frameworks or taxonomies, (2) developing attack detection approaches, and (3) demonstrating possible attacks [8]. Due to the relative infancy of this research area, these topics are often discussed separately. However, in order to develop deployable solutions, these areas need to be addressed in tandem. As a first step toward deployable solutions, this paper will bridge this gap for machining processes. Therefore, the following three sub-sections provide a literature review of the three aforementioned research areas, respectively. It should be noted that this literature review only includes research that is applicable to Product-Oriented C2P attacks.

2.1 Framework and taxonomies literature

The focus of this research area is on providing systematic methods to quantify different issues related to cyber-physical security in manufacturing. The results of these efforts have provided researchers with a deeper understanding of manufacturing cyber-physical from a security perspective.

Frameworks. Several frameworks have been proposed regarding the flow of physical or cyber entities through a production system to identify vulnerabilities (or risk). Hutchins et al. (2015) [20] outlined a framework for identifying cybersecurity risks in manufacturing. The focus of the framework was on data transfer within the manufacturing and supply chain environments. Through this framework, several mechanisms for identifying generic and manufacturing-specific vulnerabilities were identified. Chhetri et al. (2017) [21] presented a cross-domain security analysis framework for cyber-physical production systems. Their framework proposed that the relation between the cyber domain model and physical domain information flows can be abstracted using data-driven models. DeSmit et al. (2017) [22] proposed a framework to both identify and quantify vulnerabilities in a manufacturing system. Their work considered not only data transfer through a cyber-physical system, but also included human and physical entities and their interactions.

Taxonomies. Several research efforts have focused on developing cyber-security taxonomies or classification schemes that are specific to the unique nature of manufacturing. Yampolskiy et al. (2016) [23] provided three dimensions of classification with respect to attacks on 3D printers. These dimensions focused on (1) elements in the additive workflow; (2) manipulations of these elements; and (3) the effects of these elements. Pan et al. (2017) [24] discussed two taxonomies for IoT-based production systems to better understand the potential dangers. The first taxonomy focused on cyber-physical attacks against manufacturing processes, whereas the second concentrated on quality inspection measures to counteract the attacks; however, no overlap between these two taxonomies was considered. In response, Elhabashy et al. (2018) [8] presented a taxonomy that focused on the overlap between attacks aimed at both manufacturing and quality control. DeSmit 2017 [25] proposed a taxonomy that classified cyberattacks against manufacturing into four classes: (1) information gathering, (2) enabling, (3) masking, and (4) amplifying. Using this taxonomy, they demonstrated the ability to analyze the impact of an attack based upon eight unique severity metrics.

The research into frameworks and taxonomies has paved the way for understanding the behavior of Product-Oriented C2P attacks. Unfortunately, this collection of work has focused solely on manufacturing from the system-level. In general, framework and taxonomy research needs to be extended to consider the behavior of attacks for specific manufacturing processes. While there exists some degree of overlap across all manufacturing processes, individual process nuances will result in attack-process interactions that cannot be captured in high-level generic frameworks or taxonomies. As part of this research, an attack design scheme for categorizing Product-Oriented C2P attacks on machining is presented in Section 4.

2.2 Attack detection literature

Several approaches to detect the physical manifestation of a Product-Oriented C2P attack have been proposed. In general, these approaches can be separated into two categories: (1) attack detection for additive processes and (2) process independent attack detection.

Attack detection for additive processes. Sturm et al. (2016) [26] investigated the use of impedance signatures (captured by piezo-electric transducers) for in-situ defect monitoring in additive processes. While their paper did not focus specifically on Product-Oriented C2P attacks, the authors did mention the applicability of this approach in cyber-security for additive processes. Chhetri et al. (2016) [27] proposed using process variables, such as acoustic emissions, to detect attacks against Fused Deposition Modeling (FDM) additive processes. Through an experimental setup, the authors showed that acoustic emission can detect a simple attack which changes the velocity, displacement, and movement of the printer's axes. Similarly, Belikovetsky et al. (2017) [28] used acoustic emissions as side-channels to detect attacks that added or removed G-code commands, modified length parameters and extruder's speeds, or reordered G-code commands.

Process independent attack detection. The use of impedance signatures as a side-channel for detection is not limited to additive processes. Vincent et al. (2015) [17] proposed using side-channels as a generic process-independent approach to detect Product-Oriented C2P attacks in production systems. However, their paper only proposed the concept and was never extended to include experimental results or applicable case studies. Wu et al. (2017) [29] demonstrated the use of both vision and acoustic emission sensors to detect a range of attacks against both additive and machining processes.

Excluding the work of Sturm et al. 2016 and Vincent et al. 2015, who proposed the use of novel sensors to detect attacks, the remainder of research in this area has basically demonstrated that existing process monitoring techniques can be used to detect attacks. While these demonstrations are vital to progressing the research field, they have several drawbacks. First, previous attack detection research has focused on implementing arbitrary Product-Oriented C2P attacks that did not have a specific goal with respect to how the manufactured part is altered. Second, no research has explored how an adversary may be able to thwart process monitoring via side-channels through intelligently designed attacks. Third, current efforts do not employ physical process knowledge to understand the relationship between Product-Oriented C2P attacks and their effects on manufacturing process dynamics. This presents a clear loss in the ability to leverage the unique characteristic that Product-Oriented C2P attacks produce physical manifestations. Finally, these research

efforts have, for the most part, relied upon traditionally used data-driven methods for attack detection. These traditional methods were not developed for the purpose of security, and therefore may be inappropriate for attack detection. In fact, it may be possible within an attack, to exploit assumptions made in the development of these methods to evade detection. Altogether, these limitations do not allow researchers to understand the extent to which side-channels can detect a wide range of possible attacks.

In response to these shortcomings, Sections 5 and 6 introduce an experimental investigation study in which six different attacks on a simple geometry are systematically designed following the attack design scheme proposed in Section 4. Second, several attacks demonstrate how an adversary can intelligently leverage physical process knowledge and the assumptions of traditional signal monitoring techniques to thwart detection via side-channels. Third, utilizing the physical process knowledge, each attack impact on the machining process physics is discussed along with its impact on side-channels process data measurements. Last, the use of side-channel monitoring systems is explored with the threat of attacks on product quality in mind. Several detection metrics were selected for this purpose.

In addition to the aforementioned gaps, the attack detection literature has been limited only to attack detection at the process level. However, manufacturing has relied heavily upon QC to ensure products meet design specifications. Numerous QC tools have been designed and deployed across every aspect of manufacturing, from simple human visual inspection plans to mathematically rigorous in-situ process monitoring techniques. To take advantage of these already-available resources, Section 3 proposes the adoption of the defense-in-depth strategy to transition current quality control resources to become physical detection layers to defend against Product-Oriented C2P attacks.

2.3 Attack demonstrations

The majority of the earliest works in cyber-physical security for manufacturing systems focused on generating both academic and industrial awareness. This awareness was created through the use of attack demonstrations that can be categorized as focusing on either machining or additive processes.

Machining attack demonstrations. Wells et al. (2014) [3] discussed some of the cyber-security-related weaknesses existing in production systems through the use of a case study for a machining process. In their case study, the manufacturing of a tensile test specimen on a Computer Numerical Control (CNC) milling machine was attacked through altering the tool path files as part of an undergraduate student project. The purpose of the case study was not only to demonstrate the attack feasibility, but also to assess the diagnostic abilities of

future engineers. Turner et al. (2015) [30] expanded upon the work of [3] and analyzed potential attack surfaces within manufacturing, such as the design tool chain, control, and direct equipment attack surfaces.

Additive manufacturing. Sturm et al. (2014) [5] pointed out existing cyber-related weaknesses in additive manufacturing processes and discussed a variety of potential cyber-physical attacks. This discussion concluded with demonstrating an attack that added an internal void into a tensile test specimen by altering its .STL file. The result was a printed part that exhibited a significant loss in strength. In a manner similar to [3], student groups unknowingly participated in the experiment and were not able to detect the occurrence of the attack, even after destructively testing the part. Expanding upon the work of [5], Zeltmann et al. (2016) [31] investigated two types of additive manufacturing attacks, namely, embedding internal defects and altering print orientations. Ultrasonic inspection and Finite Element Analysis showed that these attacks were difficult to detect and that they had a negative impact on part performance, respectively. Belikovetsky et al. (2016) [32] demonstrated an attack against a 3D printed propeller for a quadcopter. In this demonstration, the propeller's design file was remotely compromised, causing the quadcopter propeller to collapse during flight. In addition to the demonstration, they also identified attack opportunities, analyzed the attack's full chain, and developed a methodology to assess attack difficulty in additive manufacturing.

Moore et al. (2017) [33] demonstrated an additive manufacturing attack by tampering with a FDM printer's firmware. Two different attacks were implemented, one attack affected the printer's control flow and the other attack increased the printer's extrusion rate. Slaughter et al. (2017) [34] presented an attack that compromised an additive manufacturing processes' own quality control system to maliciously attack parts. More specifically, they demonstrated that an attack against a powder bed fusion system could be done by attacking the infrared imaging system used for closed-loop quality control.

While the aforementioned research has successfully brought awareness to the threat of Product-Oriented C2P attacks, attack demonstrations have limited themselves to only consider one or two attack scenarios. However, given the complex nature of manufacturing systems, the number of possible attack scenarios for a given process is almost limitless. The immense amount of possible attack scenarios can be attributed to several unique aspects of manufacturing. First, theoretically speaking, there exists an infinite number of ways to alter a part (e.g., infinite number of possible tool-paths). Second, multiple inputs are necessary for a manufacturing process including product design data, process parameters, and incoming raw materials. Third, for each input into a manufacturing process, multiple exploitation opportunities

exist to maliciously alter the product being produced. As part of this research, Section 5 introduces an experimental investigation case study in which the systematic attack design scheme, proposed in Section 4, is used to design a variety of Product-Oriented C2P attacks, on an example machining process, with varying levels of design sophistication.

3 Defense-in-depth for cyber-physical manufacturing systems

For almost two decades, the strategy of defense-in-depth has been a widely accepted best practice in cyber-security. This strategy relies upon the development of multiple defense layers to create a robust security solution that results in unattractive (costly) targets for potential adversaries [10]. According to the Department of Homeland Security (DHS), the defense-in-depth strategy relies upon both detective and protective layers to impede cyber intruders. Similarly, detect and protect are two of the five core functions of the NIST cyber-security framework [16]. Cyber-physical systems and cyber-physical attacks consist of, and can be uniquely described by, cyber and physical components. Likewise, detective and protective layers can (and should) occur in both cyber and physical forms. Given that the scope of this paper revolves around physical defense in manufacturing, the discussion focuses only on physical defense layers. Information on cyber defense layers is beyond the scope of this work and can be found in resources such as the DHS recommended practices document [10] and the NIST cyber-security framework for critical infrastructures [16].

Over the course of the past several decades, numerous forms of physical protection have been identified and recommended as best practices for enhancing security. Examples of common physical protection include, but are not limited to fences, gates, access cards, locked equipment cabinets, video cameras, lighting, and physical port (e.g., USB) blocks. However, outside of recent work with side-channel detection approaches (Section 2), significantly less attention has been placed on physical detection layers. As part of the defense-in-depth framework, an organization should utilize its available resources to provide effective layers of detection and protection [10]. With respect to Product-Oriented C2P attacks in manufacturing, this begs the question, "What physical resources are available to detect attacks and how can they be leveraged to develop multiple defense layers?" Answering this question requires identifying what these layers would actually be defending against.

As discussed by Wu and Banzhaf (2010) [35], the ultimate goal for detection in cyber-security is to quickly detect threats before they inflict widespread damage. With regard to Product-Oriented C2P attacks, this paper defines widespread damage as producing and ultimately delivering physically

altered products to a customer. From this definition, the answers to the aforesaid question reside in quality control (QC) resources. Over the past century, manufacturing has relied heavily upon QC to ensure goods are produced and delivered as designed [36]. Numerous QC tools have been designed and deployed across every aspect of manufacturing, from simple human visual inspection plans to mathematically rigorous in-situ process monitoring techniques. Unfortunately, QC tools are not designed considering the possibility for a Product-Oriented C2P attack, which may limit their direct application as physical detection layers. More specifically, QC tools may be incapable or highly insensitive to maliciously designed product alterations, especially for alterations that are intelligently designed to circumvent QC. In this paper, QC regimes that have the potential to be developed into physical detection layers are separated into three categories: personnel, inspection, and process. The current status of these QC regimes and their ability to be developed into defense-in-depth detection layers is briefly discussed below.

Personnel. Possibly the most straightforward approach to detecting a Product-Oriented C2P attack is for experienced manufacturing personnel to detect physical alterations to manufactured parts or changes to a manufacturing process. It should be noted that this category is not referring to human inspectors, i.e., humans whose role in a manufacturing system is to inspect and ensure part quality. The category is referring specifically to manufacturing personnel, i.e., any human whose role in the manufacturing system does not necessarily include inspection but does include direct interaction with manufactured parts or processes. Potential personnel within this category include machinists, material transporters, mechanics, packagers, or quality engineers.

While this may appear to be a simple QC regime to develop into a detection layer, it may be quite challenging for parts that are produced in large volumes, exhibit very complex features or geometries, or are very small in size [8]. In addition, the role of operators is shifting toward performing tasks across different systems rather than operating a single manufacturing process. Operators in advanced manufacturing systems have less hands-on interactions with manufacturing processes. Operators rely on software tools such as HMI (human machine interface), SCADA (supervisory control and data acquisition), MES (manufacturing execution systems), and MI (manufacturing intelligence) to produce alarms and warnings regarding process performance.

Inspection. Most manufacturing QC systems rely upon post-production part inspection. These inspections include but are not limited to (1) visual inspection (e.g., human, machine, x-ray); (2) feature measurements (e.g., micrometers, optical comparators, Coordinate Measurement Machines (CMMs), etc.); or (3) gauging (e.g., go/no-go gauges). Given the

specific manufacturing scenario, these inspections can be performed on or off-line using a variety of sampling strategies (e.g., 100% inspection, sub-grouping, first and last). Furthermore, the results of these inspection procedures can be used to accept/reject individual parts of lots and are often used as the basis for statistical process control (SPC) applications.

Developing this QC regime into a detection layer requires overcoming a significant hurdle. Specifically, inspection tools/approaches are not designed to detect (or may be incapable of detecting) the effects of an attack. For instance, inspection systems usually focus on pre-determined (physical or statistical) features, such as Key Quality Characteristics (KQCs) or Principal Components (PCs), respectively. Any Product-Oriented C2P attack that alters a feature that is not inspected may go undetected. Similarly, these approaches only focus on measuring features that should exist. Adding additional features, for example an extra hole, may go undetected. Additionally, in some manufacturing scenarios, not every single manufactured part is inspected, allowing for attacks that only affect non-inspected parts to completely circumvent this QC regime.

Process. Advanced manufacturing systems place significant emphasis on process monitoring in real-time. For instance, for machining processes, vast amounts of research efforts have focused on creating real-time automated process monitoring, diagnosis, and control systems. These methods aim to automatically detect, diagnose, and compensate for process/product anomalies utilizing in-situ sensor measurements of process variables (e.g., cutting forces and vibration) [37]. Similarly, research efforts have also focused on the real-time monitoring of Key Performance Indicators (KPIs) (e.g., process startup times, throughput rates, equipment availability) within a control system to detect anomalies including cyberattacks [38]. It should be noted that this should not be confused with the research revolving around physically detecting attacks against Industrial Control Systems (ICS). In this research area, physical system variables are used to detect anomalies in controller logic, sensor readings, systems states, or control commands [39].

While all three QC regimes are important to develop into detection layers, we believe that developing the Process QC regime has the most potential to significantly increase manufacturing security. However, there are still issues that need to be overcome to transform process monitoring into a detection layer. Specifically, process monitoring focuses on feature extraction and monitoring strategies that may not be sensitive to Product-Oriented C2P attacks. For instance, in machining, features are extracted based upon specific process monitoring objectives; such as, (1) tool conditions, (2) chip conditions, (3) process conditions, (4) surface integrity, (5) machine tool state, or (6) chatter detection [37]. No research

exists that focuses on which features should be extracted when the process monitoring objective is security.

Together these three QC regimes suffer from the same two highly related hurdles that prevent their direct implementation as detection layers. First, all QC tools are developed with specific assumption regarding the nature of the manufacturing process, such as the statistical behavior of the data or possible process shifts. It is possible that these underlying assumptions are either no longer valid in the presence of an attack or the attack itself invalidates the assumptions. Second, the implementation of QC tools provides manufacturers with a false sense of security against Product-Oriented C2P attacks, because they are unaware or oblivious to the fact that attacks can be intelligently designed to circumvent QC systems. Overcoming these hurdles requires re-visioning how QC systems are designed and implemented across all three of these QC regimes.

Lastly, designing an effective and efficient physical detection layer strategy requires understanding how attacks can be designed and to what extent those designs can (and cannot) avoid detection. In response, Section 4 presents an attack design scheme and a designation system (machining specific) that provides the structure to populate a wide variety of potential attacks.

4 Attack design scheme

This section introduces an attack design scheme to systematically describe the key elements of Product-Oriented C2P attacks against machined parts. This scheme provides the first step toward developing a needed body of knowledge that will offer a clear and common understanding of how Product-Oriented C2P attacks on manufacturing systems can be designed and implemented. The scheme is similar to the Common Attack Pattern Enumeration and Classification (CAPEC™) repository for software security applications [15, 40]. The systematic understanding of Product-Oriented C2P attacks against machined parts will assist in the development of secure manufacturing systems utilizing physical defense-in-depth strategies. The elements of the attack design scheme are described in Section 4.1. These elements are captured in the proposed attack design designation system (ADDS) in Section 4.2. ADDS provides the structure to populate the wide variety of potential attacks targeting machining environments to aid in designing better defenses for these systems.

4.1 Attack design scheme elements

The attack design scheme for a Product-Oriented C2P attack in machining includes three elements: (1) the quality integrity category; (2) the design considerations; and (3) the

implementation location. The quality integrity category describes which physical aspect of a machined part is altered by an attack. Design considerations consist of possible system information, gathered by the adversary, which can be used to design an attack with a lower likelihood of detection. The implementation location corresponds to where in the value chain, the attack could be introduced to accomplish a specific objective, but not necessarily where the physical manifestation of the attack occurs in the production system. Here, the implementation location is synonymous with what Yampolskiy et al. [6, 7] referred to as the influenced element and the location of the physical manifestation would correspond to the victim element. Additional details about the attack design scheme elements are provided the following subsections.

4.1.1 Quality integrity category

The attack design scheme captures which physical aspect of a part is altered by an attack through the quality integrity category. For machining processes, three criteria [41] determine if a part adheres to the original design parameters: (1) dimensional accuracy; (2) acceptable surface roughness; and (3) material properties consistent with design intent. These three criteria provide the basis for the three quality integrity categories: geometric, surface, and material quality integrity categories, respectively:

- a. The geometric quality integrity category describes the nominal geometry that can be altered by an attack (e.g. feature size/shape and/or added/removed feature);
- b. The surface quality integrity category describes the nominal surface characteristics—roughness, waviness, and/or lay—that can be altered by an attack; and
- c. The material quality integrity category describes the alteration of material characteristics (hardness or microstructure) as a result of varying the process physical parameters and/or the incoming material characters.

Any attack successfully altering physical part aspect(s) within one or more of these quality integrity categories will result in a part which does not match design intent. With respect to the scope of this work, the focus is on attacks that degrade the part geometric quality integrity for machining processes.

The American Society of Mechanical Engineers (ASME) [42] defines the Geometric Dimensioning & Tolerancing (GD&T) system as “...an essential tool for communicating design intent - that parts from technical drawings have the desired form, fit, function and interchangeability.” The present work adopts the GD&T concepts of type and characteristic to construct the geometric integrity classes and sub-classes needed to populate the individual geometrical physical aspects that

can be potentially altered by an attack. Table 1 shows the geometric integrity classes and sub-classes used to communicate the geometric feature affected in an attack on a machined part.

Together the quality integrity category, class, and sub-class provides the framework to specifically define the geometrical aspect(s) affected through a Product-Oriented C2P attack. The class letter abbreviations and sub-class numerical values, provided in Table 1, are used in the Attack Design Designation System (ADDS). The ADDS is described in Section 4.2 and used throughout Section 5 to demonstrate a number of attacks that can be identified through this system. To illustrate geometric integrity classes and sub-classes, a visual representation of a nominal and two attacked cylinders are shown in Fig. 1.

In this example, two different geometric quality integrity category attacks on a straight cylinder with a keyway can be cataloged by the ADDS. The first attack alters the form class and cylindricity sub-class (G.F.4) by changing the part from a straight cylinder to a barrel shape. The second attack alters the size class and linear dimension circular feature sub-class (G.S.1) by increasing the diameter of the cylinder. It should be noted that this example does not include the attack design consideration in the ADDS which is introduced in the next sub-section.

4.1.2 Attack design considerations

When a manufacturing system is viewed from a defense-in-depth perspective, for an attack to be successful, the design of

Table 1 Geometric Integrity Classes and Sub-classes

Class	Sub-class		
Form	F	Straightness	1
		Flatness	2
		Circularity	3
		Cylindricity	4
Location	L	Position	1
		Concentricity	2
		Symmetry	3
Orientation	O	Angularity	1
		Perpendicularity	2
		Parallelism	3
Profile	P	Profile of a line	1
		Profile of a surface	2
Runout	R	Circular	1
		Total	2
Size	S	Linear dimension circular feature	1
		Linear dimension straight feature	2
		Angular dimension	3



Fig. 1 Nominal (left), attacked barrel (center), and attacked diameter (right) machined cylinder

the attack would need to consider a manufacturing system's defense layers. Design considerations, based on system information and gathered by the adversary, must be incorporated in designing an attack with a lower likelihood of detection. As discussed earlier, the defense-in-depth strategy makes it more challenging or costly for the adversary to pursue an attack. This can be demonstrated in a manufacturing system by considering the different layers that have been developed for detection.

Attack designs need to consider both cyber and physical defense layers. As previously noted, the focus of this paper is on physical design considerations. From the discussion in Section 3, physical detection layers include personnel, inspection, and process. Acquiring information regarding the implementation of these defense layers may allow for considerations in an attack's design that will decrease its probability of being detected. Example considerations are listed in Table 2 and are meant to be an initial population with new entries and/or subsets to be potentially added over time. The numbering system of the design considerations in Table 2 is used in the attack design designation system described in the following sub-section.

The attack design considerations listed in Table 2 are generalized considerations and additional more manufacturing system specific considerations may exist. The following paragraphs provide examples of what an adversary needs to consider to overcome the hurdles created by personnel, inspection, and process detection layers.

Personnel layer. For product considerations, a seasoned/experienced operator(s) who is knowledgeable about a product's nominal geometry may notice geometric alterations. To overcome this hurdle, for example, the magnitude of a feature size alteration can be made relatively small to decrease the likelihood of detection by an operator either at the machine or a downstream process. Similarly, it may be easier to notice attacks on one feature type versus another. For instance, it may be preferable not to perform an attack that results in a nominally symmetrical product becoming asymmetric.

For process considerations, a seasoned/experienced operator(s) who is knowledgeable about the process may cognitively notice (intuition) changes to the process behavior. For instance, if an attack resulted in an increased depth of cut, a seasoned machinist may notice additional vibrations or noise in the process. To overcome this hurdle, another process

Table 2 Design considerations of Product-Oriented C2P attacks in machining

Layers	Consideration	Information required for consideration
Personnel	1 Product detection	The cognitive proficiency of an operator to use their experience/knowledge of a product to detect alterations
	2 Process detection	The cognitive proficiency of an operator to use their experience/knowledge of a process to detect alterations
	3 Human machine interface	Digital information made available to an operator regarding a product or process that could indicate an attack
Inspection	4 Inspection plan	The design parameters of the inspection plan
	5 Key quality characteristics	The geometrical features to be inspected and their specification limits
	6 Inspection procedure	The type(s) of inspection equipment used and how the inspection is performed
Process	7 Industrial control system	How a control system may be used for monitoring the process
	8 in-situ process monitoring	How in-situ process monitoring may be used for monitoring the process
	9 Variation	Common cause process variation

parameter (e.g., cutting speed or feed rate) could be changed to compensate for the increased depth of cut.

For human machine interaction considerations, in modern manufacturing systems, operators rely on software tools such as HMI (human machine interface) and SCADA (supervisory control and data acquisition) to produce alarms and warnings regarding process performance. For instance, an operator responsible for a manufacturing cell encompassing multiple CNC machines may be supported by a process data digital readout showing real-time process status feedback. This panel may include information such as overall average power consumption, average vibration level, and coolant pressure. The knowledge of what and how process information may be supplemented to the operator may aid in designing attacks that overcome this detection layer. For example, an attack resulting in a decreased depth of cut will decrease the overall vibration level reflected on the data panel presented to the operator. The attack may consider compensating for this decrease by increasing the feed rate so that the overall vibration level remains the same. Another way to achieve this is to simultaneously attack the data panel software itself to tamper with the overall vibration level readout on the screen to make it appear normal. The decision of how to implement this consideration relies on the knowledge of the system vulnerabilities and the relative complexity of exploiting them.

Inspection layer. In manufacturing QC systems, the use of post-production inspection offers another layer for detecting product changes from nominal geometry. Considering the specifics of the inspection process such as the inspection plan, the product's key quality characteristics (KQCs), and the inspection procedure helps overcome the inspection detection layer. For the inspection plan, for instance, knowledge of the sampling plan details such as the sampling interval and strategy could aid in performing attacks targeting only parts that are not considered for inspection. For KQCs, they are the pre-determined critical geometrical features and their specification limits selected for quality checks. An attack, for example,

could consider this information and alter a geometric feature not inspected by a large magnitude and yet may go undetected. For the inspection procedure consideration, details such as the measurement device type and the measurement setup help in designing attacks in which altered features could pass inspection. For example, changing the CAD file used to create the inspection program of a CMM to match the attacked geometry will result in the CMM acceptance of this altered geometry.

Process layer. As discussed in Section 3, process monitoring can aid in detecting changes of a process characteristic driven by an attack aimed at altering a product's design intent. Two types of real-time monitoring systems can be utilized for such defense purpose: industrial (machine) control system-based monitoring and/or in-situ process physical variables monitoring system. The knowledge of how any or both of these systems may be used for monitoring the process can offer a venue for overcoming this detection layer hurdle. For an industrial (machine) control-based monitoring system, for example, embedded sensor signals (e.g., tool/spindle position) can be tampered to avoid detecting changes in the cutting path. As for an in-situ process monitoring system, knowledge of what physical signal descriptors (e.g., statistical descriptors, wavelets coefficients, FFT coefficients) that could be extracted and monitored can aid in designing attacks that do not necessarily impact these descriptors. For instance, it is customary to only focus on monitoring signal amplitude descriptors to ensure that a physical quantity such as cutting vibrations remain within allowable limits. Such knowledge can help in designing attacks that do not impact the magnitude of the signal but the local time signatures of the signal (see attack 2 description in Section 5 for further clarification).

For attacks that may change descriptors being monitored, the knowledge of the typical process signal variations could aid in designing attacks that can be hidden within the system's natural variability. For example, to obtain a specific diameter reduction in turning could be accomplished by changing the

depth of cut in either the roughing pass or finishing pass. It is widely known that there is more variation in the process physics of the roughing operation over finishing. Therefore, a larger magnitude diameter reduction could be achieved in the roughing pass with a lower probability of detection over the same diameter reduction in the finishing pass.

Based on the discussion of this section, design considerations must be incorporated in designing an attack with a lower likelihood of detection. The implementation location is the final element in the attack design scheme and is introduced in the next sub-section.

4.1.3 Implementation location

The implementation location is the cyber entity within the value-chain where the attack is introduced to accomplish a desired physical change but does not necessarily need to be where the physical manifestation of the attack occurs. Figure 2 depicts the traditional manufacturing system value-chain for machined products and illustrates the possible implementation locations. It should be noted that part design, used in the following figure, describes the design intent of the part to be created through the value-chain. This term is used since the part design includes various aspects, such as the tolerances of the part which dictate the process parameters used. The use of part design, in this case, is not just the nominal part geometry but instead describes the design intent used as an input at other locations in the value-chain.

An important aspect of the attack implementation location element is that an attack with the same physical manifestation can be introduced at multiple cyber entities (implementation locations) in the manufacturing system value-chain. For example, consider an attack to change a dimension of a part from nominal. This attack could be introduced at three different locations by (1) altering the CAD geometry; (2) altering the roughing G-code; or (3) adjusting the machine offsets in the

machine controller of the process. An attack at any of the three locations could result in the same dimensional change of the part. The numerical numbering system for implementation locations, provided in Fig. 2, will be used in the attack design designation system introduced in the following sub-section.

4.2 Attack design designation system

The Attack Design Designation System (ADDS) is a systematic way to capture unique attack design schemes for Product-Oriented C2P attacks in machining. The quality integrity category, category class, category sub-class, design consideration(s), and implementation location(s) of the attack design scheme are captured through ADDS. This designation system is modeled after the NIST National Vulnerability Database (NVD). The proposed ADDS provides the first step in developing a similar body of knowledge to describe unique attack design schemes in machining. Figure 3 shows the ADDS with the potential values for each element.

The quality integrity category can take values of geometric (G), surface (S), or material (M). The category class uses the letter abbreviations found in Table 1 corresponding to form (F), location (L), orientation (O), profile (P), runout (R), and size (S). The category sub-class for the geometrical integrity category uses the numbering system found in Table 1. The design consideration can take the values found in Table 2. The implementation location corresponds to the manufacturing system value-chain numbering system detailed in Fig. 2. If multiple items need to be recorded in an element of ADDS, they are recorded in ascending numerical order.

5 Experimental investigation

To demonstrate the relationship between Product-Oriented C2P attacks and physical defense-in-depth strategies, an

Fig. 2 Traditional manufacturing system value-chain

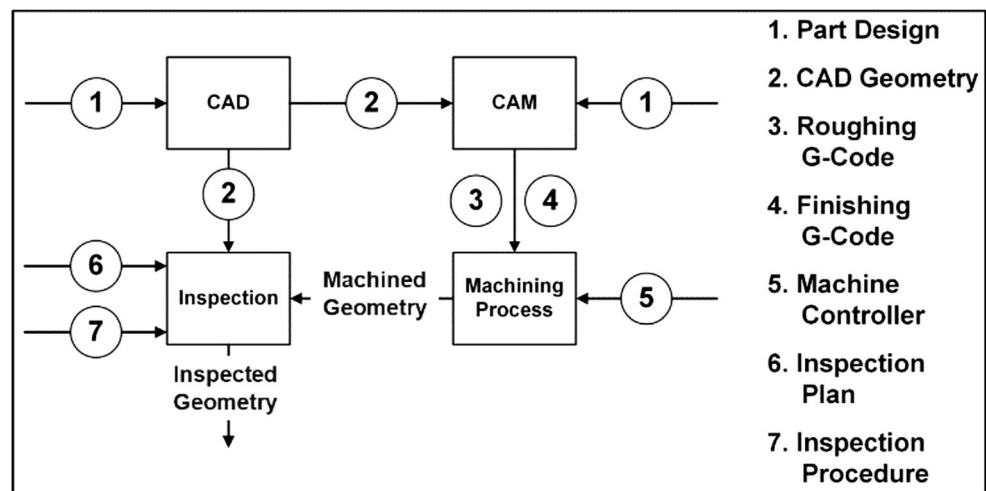
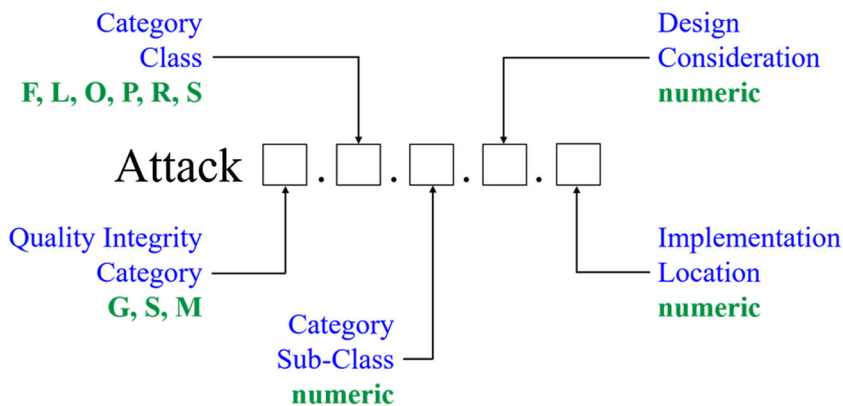


Fig. 3 Attack design designation system for machining



experimental investigation study is conducted. This section introduces an experimental case study conducted with several realistic attacks against a turning process. It should be noted that the machining process chosen was turning; however, another machining process could have been chosen. Additionally, this section discusses the effectiveness of employing physical process knowledge in understanding the physical manifestations of potential attacks through in-situ side-channels measurements of in-process variables.


The attacks presented in this case study are designated following the proposed ADDS. As mentioned in Section 3, designing effective and efficient physical detection layers require understanding how attacks can be designed and to what extent those designs can (and cannot) avoid detection. The ADDS allows for an approach to describe the considerations that are made during an attack design, but does not focus on

the specific details of executing an attack, which are process-dependent. In essence, this section provides an example to practitioners on how the ADDS can populate potential attacks against a system with a given defense-in-depth strategy. Once identified, process-specific details for potential attacks can be determined and their affects against the system can be analyzed to develop a more robust defense-in-depth strategy.

5.1 Test part selection, design, and production

A spool (valve slide) is selected to be the test part for the experimental investigation in this study. It is a critical element of a wide variety of hydraulic valves. This element controls the fluid flow in different paths within a valve to control the opening and closing of the valve whenever needed. Hydraulic valves are a critical component in most of the mechanical and

Fig. 4 Simplified spool design: 3D isometric view (top) and 2D drawing with dimensions (bottom)



mechatronics systems produced today. Several hydraulic valves, for example, are used in landing gear assemblies of modern aircrafts. These valves are included within different components of the landing gear system assemblies such as brakes, retraction, and expansion mechanisms [43]. Spools within hydraulic valves are typically produced using machining processes. Hence, similar to all machined parts, the functional quality of these spools relies on their geometrical, surface, and material integrities, as discussed in Section 4. If the manufacturing system used to produce these critical elements fall victim to a Product-Oriented C2P attack, the consequences could be catastrophic. Depending on the altered geometrical aspect of the part, the impact may vary from failing during assembly with other valve elements, to causing the hydraulic system not to operate as intended, or to causing a sudden failure/damage of the valve and hence the mechanical system (e.g., landing gear of an aircraft). This motivated the choice of a spool as a test part for the experimental investigation to show case the importance of adopting defensive strategies for securing production systems producing such critical products.

For the purpose of this study, a simplified spool has been designed as shown in Fig. 4. Despite the simplicity of this part's geometry, a wide range of geometrical quality integrity attacks have been developed and implemented for this study. Within this study, eleven spools were produced, five nominal (good) parts and six attacked parts that had their geometry altered from nominal under different types of attacks (described in the following sub-section).

A CNC lathe (ECOCA PC-4615E) test-bed was developed to produce these parts and collect the relevant process data (e.g., cutting forces and power). To measure the cutting power as a side-channel of the process physics, this machine was

equipped with a hall-effect power sensor (Load Controls Inc. UPC-FR). The power signals were pre-processed and digitized using a data acquisition system, NI cDAQ-9174, equipped with NI 9229 analog input module. It should be noted that the hall-effect power sensor allows for measuring the spindle electrical power which is directly proportional to the cutting power. Hence, it is not intrusive to the process and does not require any special set-up making it more applicable for use in industrial settings.

The machining operations for this experimental investigation are dry single-point oblique cutting of annealed alloy steel (AISI 4140) stock rods, 58.0 mm diameter and 300.0 mm in length. In each experiment, the workpiece was exposed a distance of 268.0 mm from the chuck and was centered and supported by a tail-stock. All workpieces were pre-machined to a diameter of 53.0 mm over a length of 175.0 mm. Afterwards, roughing and finishing operations were performed following the cutting paths in Fig. 5.

A fresh roughing tool per part (DNMG 15 04 08-PR) was installed to perform the roughing operations, following the tool path illustrated in Fig. 5 (top), in nine cutting cycles as indicated by numbers in the figure below. In each of these cycles, constant cutting velocity, feed rate, and depth-of-cut were used and set to 200 m/min, 0.33 mm/rev and 1.25 mm, respectively. The straight lines represent the actual cutting cycles where the cutting tool moves from the right to the left. Each straight line (cutting cycle) is followed by a dashed dotted-line indicating the return path of the tool after going from left to right to start the following cutting cycle. Following the roughing operations, a single finishing cycle is performed using a fresh finishing tool per part (DNMG 15 04 08-PF), following the tool path shown in Fig. 5 (bottom). In the finishing cycle, constant cutting velocity, feed rate, and

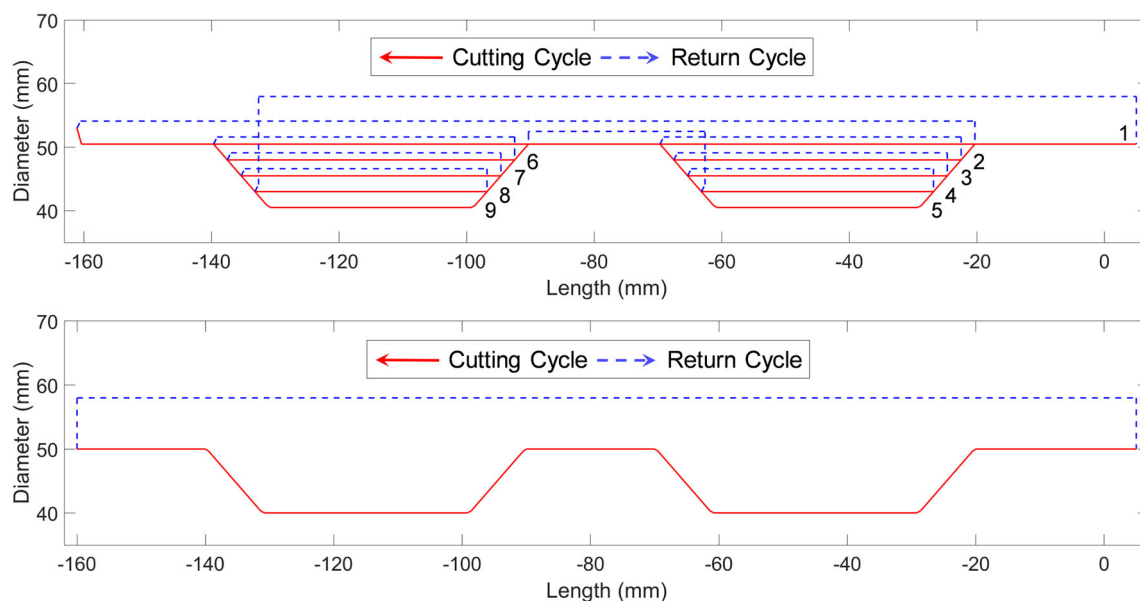


Fig. 5 Nominal roughing operations tool path (top) nominal finishing operation tool path (bottom)

depth-of-cut were used and set to 200 m/min, 0.2 mm/rev and 0.25 mm, respectively. It should be noted that fresh cutting tools are used to avoid any assignable variability (e.g., due to tool wear) in the process other than the attacks induced variability. In the future, further research is needed to relieve this constraint.

5.2 Implemented attack scenarios

This section discusses the details of six different attack scenarios in which the attacks were designed following the attack design scheme proposed in Section 4. These attacks focus on maliciously degrading the spool geometrical integrity. The following sub-sections describe the key design elements of each attack, including considering the underlying manufacturer's defense scenario, how its physical components are implemented, and how it is designated following ADDS. Additionally, utilizing the physical process knowledge, each attack impact on the machining process physics is discussed. Due to the primary focus of this work on the physical components of the attacks, as previously mentioned, it would be always assumed that the adversary has gained access to influence the cyber domain elements required for a specific attack. In other words, the attacks demonstrated within this section are performed along the lines of a white-box ethical hacking effort [44].

5.2.1 Attack scenario 1 (G.S.1.148.56)

Design The objective of this attack is to decrease the spool outside diameters, shown in Fig. 4, by 0.3 mm so that the final diameters of the 50.0 mm sections and the 40.0 mm sections become 49.7 mm and 39.7 mm, respectively. This attack is to be implemented against a manufacturer who has a post-production inspection strategy to inspect 50% of the produced spools (i.e., inspect every other part). Additionally, the manufacturer has deployed an in-situ process monitoring system to monitor the finishing operation's spindle power signatures.

An essential element in the attack design is the system design considerations layered to minimize the detectability of the attack's physical impact on the product and the process. One possible combination, among many others, may include (1) physical change magnitude is chosen to be very small compared to the part original size (0.6% and 0.75% diameter reductions for the large and small diameters sections, respectively) making it hard to be detected by an operator (design consideration #1) who is knowledgeable about a product's nominal geometry; (2) parts to be attacked are those that will not be inspected according to the sampling plan (design consideration #4) so the attack needs to be coordinated not to coincide with an inspection; and (3) the attack to be implemented in a way that only impact the process physics in the roughing operation but not in the finishing operation due the

fact that the manufacturer has adopted an in-situ monitoring systems using finishing process data measurements (design consideration #8).

Implementation. This attack can be implemented by attacking the machine controller (implementation location #5) to positively shift the workpiece coordinate system in the x -direction (i.e., radial direction) by 0.3 mm. This results in a shift of 0.15 mm magnitude between the actual physical centerline of the part and the centerline defined to the controller (typically defined by the operator). Hence, the actual depth-of-cut of the first roughing path will increase by 0.15 mm than the nominal/programmed depth of cut (1.25 mm) while the depth-of-cut of the remaining roughing paths remain unaffected. Moreover, the same shift in the workpiece coordinate system has to be maintained during the finishing operation and hence, the actual finishing depth-of-cut will be similar to the nominal. Therefore, the final diameter will be less by 0.3 mm than the nominal diameter (i.e., radius is less by 0.15 mm). Additionally, the adversary needs to have access to the real-time inspection plan to coordinate which parts to attack and ensure they do not coincide with an inspection (implementation location #6).

Designation number The design and implementation of this attack can be designated as an attack on Geometrical Integrity, Size Class, Linear Dimension of Circular Feature Sub-Class (G.S.1), while accounting for design considerations 1, 4, and 8 as well as implementation locations 5 and 6. Hence, the designation number of this attack is G.S.1.148.56.

Impact on process physics The change in the process physical parameters resulting from this attack is the increase of the first roughing path depth-of-cut by 0.15 mm (actual depth-of-cut is 1.4 mm instead of 1.25 mm). This results in an increase of the estimated Material Removal Rate (MRR) of the first roughing path from $82.5 \text{ cm}^3/\text{min}$ to $92.4 \text{ cm}^3/\text{min}$ (i.e., 12.0% increase) according to the following formula:

$$\text{MRR} = d (\text{depth of cut}) \times f (\text{feed rate}) \times v (\text{cutting speed})$$

It should be noted that such increase in the depth-of-cut and hence the MRR is inevitable when designing such attack. This is true as long as the adversary desires to maintain the needed number of cutting paths similar to the nominal. Otherwise, extra number of cutting paths will be needed and that could be easily detected by the human operator.

From the machining literature, it is well known that physical process variables such as cutting power, vibrations, and forces are directly proportional to the MRR. Hence, there is an opportunity to detect the physical impact of this attack on the machining process. In fact, the increase in the process vibrations (hence, noise) can be noticed by a seasoned machinist

which offers a venue for the attack detection not accounted for in this attack design (design consideration #2). Additionally, the change in the MRR of the first roughing path can be manifested by measuring in-process machining variables such as cutting power and forces. This requires a real-time signal monitoring system that is designed to account for the possibility of a Product-Oriented C2P attack. It should be noted that this process related physical manifestation of the attack will be only reflected in the roughing operations measurements of relevant process variables. From a defense-in-depth perspective, this shows the importance of deploying as many layers as possible to increase the difficulty in designing an undetectable Product-Oriented C2P attack.

5.2.2 Attack scenario 2 (G.S.1.1568.235)

Design. The objective of this attack achieves the same product alteration as in attack scenario 1 while attacking a manufacturer who has adopted a more in-depth quality control strategy in which in-situ process monitoring system was deployed for both roughing and finishing operations. Additionally, this manufacturer has a 100% post-production inspection requirement for the spool outside diameters. This inspection process is to be performed manually by an operator using Vernier calipers.

First, aiming at avoiding the in-situ detection possibility of the attacks' physical impact on the roughing process (i.e., increase of the MRR of the first roughing path), an adversary needs to take this specific information about the monitoring system into consideration (design consideration #8). In particular, the design of the attack needs to consider compensating for the increase in the MRR of the first roughing path due to the increase in the depth-of-cut. This can be done through equating the MRR of the attacked process to the nominal process' MRR by decreasing the feed rate of the first roughing path proportionally to the increase in the depth-of-cut (0.15 mm). This leads to a reduced feed rate value of 0.2946 mm/rev for attack scenario 2 and a similar MRR to the nominal process (i.e., similar cutting power and force requirements). It should be noted that both attacks 1 and 2 consider the in-situ process monitoring detection layer of the manufacturer's defense system. However, the specific consideration within each attack regarding this layer is a function of how the layer is designed.

Second, to avoid detection through the 100% manual post-production inspection, the specification limits on the parts outside diameters as a KQC will be altered to allow for attacked parts to be accepted by the inspection operators (design considerations #5 and 6). Finally, similar to attack scenario 1, the physical change magnitude in this attack is relatively small making it hard to be cognitively detected by the operator (design consideration #1).

Implementation. Similar to attack scenario 1, the machine controller need to be attacked to shift the workpiece coordinate system (implementation location #5). Additionally, the roughing G-code (implementation location #3) will be altered to change the feed rate of the first roughing path only. To implement the alteration of the specifications limits on the part outside diameter, the KQCs specifications provided to the inspection station can be altered (implementation location #2).


Designation number. The design and implementation of this attack can be designated as an attack on Geometrical Integrity, Size Class, Linear Dimension of Circular Feature Sub-Class (G.S.1), while accounting for design considerations 1, 5, 6, and 8 as well as implementation locations 2, 3, and 6. Hence, the designation number of this attack is G.S.1.1568.235.

Impact on process physics. In this attack, decreasing the feed rate of the first roughing path may succeed in mitigating the impact that attack scenario 1 had on the magnitude of process physical variables such as cutting power. However, this decrease in feed rate will result in increasing the cutting time needed for the first roughing path and consequently the total completion time of the roughing process. Such a change in cutting time can be detected by monitoring the time needed for performing the first roughing path utilizing the measurements of the in-situ process monitoring system (another specific consideration within the generalized design consideration #7) or through monitoring the machine controller execution time of the whole roughing operation (design consideration #6). Similar to attack scenario 1, this change will only impact the process variable measurements for the roughing operation. It is worth noting how it is inevitable in this case to compensate for the change in the amplitude of in-process variable measurements without affecting the time needed to perform the process and vice-versa.

5.2.3 Attack scenario 3 (G.O.1.13569.2)

Design. The objective of this attack is to increase the angles of the four tapered sections of the spool, shown in Fig. 4, by 5° so that the final angles become 35° instead of 30°, see Fig. 6. The manufacturer to be attacked has a quality control strategy that include 25% part inspection plan using a CMM to check the angles of the four tapered sections. Additionally, the machine power consumption is acquired through the machine controller utilizing the embedded power sensor in the CNC machine. The average power consumption is then compared to the nominal average power consumption previously collected from controlled experiments. This information will be used to alert the operator if the power consumption of the machine goes beyond the pre-determined threshold for machine power consumption per part.

Fig. 6 Recess dimensions: nominal part (left) and attacked part (right)



To minimize the detectability of the physical impact of the attack on the product, the adversary can consider a combination of attack design considerations. First, physical change magnitude of the angles is chosen to be relatively small so it may not be detected by the machine operator (design consideration #1). Second, the geometrical quality specifications used to program the CMM inspection procedure will be altered (design considerations #5 and 6). Finally, increasing the angles by only 5° will lead to a slight increase in the MRR (impact the cutting power) while performing the taper turning operations which are performed in a very small time compared to machining the other portions of the part. The impact of this change on the overall average power consumption of the machine is very limited and may be within the allowable variation threshold (design considerations #9). This will also subvert the possibility of the attack being detected by the operator through HMI power consumption information (design considerations #3).

Implementation This attack can be implemented by attacking the CAD geometry (implementation location #2) to alter the angles to the desired attack value and consequently lead to a different G-code to be produced by the CAD/CAM engineer. Additionally, changing the CAD geometry results in a CMM inspection program based on the altered geometry and hence, the altered parts will be accepted by the CMM.

Designation number The design and implementation of this attack can be designated as an attack on Geometrical Integrity, Orientation Class, Angularity Sub-Class (G.O.1), while accounting for design considerations 1, 3, 5, 6, and 9 and implementation location #2. Hence, the designation number of this attack is G.O.1.13569.2.

Impact on process physics. This attack results in changing the relative location of the tool and the workpiece as a function of time compared to the nominal cutting path. Changing the


recess angle from 30° to 35° results in increasing the total surface length in the feed direction of each recess from 52.6303 to 53.0741 mm, including a change from 31.6 to 34.5 mm in the straight portion of the recess surface, as shown in Fig. 6. This change in the cutting length will result in increasing the cutting time for each of the four roughing paths of each recess. Consequently, the total cutting time of the attacked part will be longer than the nominal part. Additionally, due to the change of the recess angles, the gradient of gradual increase in the depth-of-cut in this taper-turning operation will increase leading to an increase in the MRR and hence increase in machining process variables such as cutting power. These process-related physical manifestations of the attack should be seen in relevant in-situ process data measurements (e.g., cutting power) for both the roughing and finishing operations.

5.2.4 Attack scenario 4 (G.L.1.147.346)

Design The objective of this attack is to shift the location of the 50 mm diameter middle section by 2 mm to the right, as shown in Fig. 7. This attack is to be implemented against a manufacturer who has a post-production inspection strategy to inspect 25% of the produced spools (i.e., inspect every fourth part) for the correctness of the location of the middle section. Additionally, the manufacturer has deployed an industrial control monitoring system to utilize machine controller data to monitor various machine and process status metrics including average power consumption of the machine and the execution time of G-code programs.

Aiming to avoid detection of this attack impact by an operator, who is knowledgeable about the product and/or in post-production inspection, similar to attack scenario 1, the adversary chose a relatively small physical change magnitude (design consideration #1). The attack was also coordinated to attack only parts that will not be inspected (design

Fig. 7 2D drawing with dimensions under attack 4



consideration #4). In essence, although this attack decreases the needed time to cut the first recess while decreasing it for the second, the total completion time of the process remains the same. Hence, the execution time of the G-code program does not change. Moreover, the geometrical change in this attack does not involve changing any of the process independent cutting parameters. Hence, the average amplitude of a dependent variable such as cutting power remains constant.

Implementation. This attack can be implemented by attacking the G-code for both the roughing and finishing operations (implementation locations #3 and 4). Furthermore, similar to attack scenario 1, the adversary needs to have access to the real-time inspection plan (implementation location #6) to realize design consideration #4.

Designation number. The design and implementation of this attack can be designated as an attack on Geometrical Integrity, Location Class, Position Sub-Class (G.L.1), while accounting for design considerations 1, 4, and 7, and implementation locations #3, 4, and 6. Hence, the designation number of this attack is G.L.1.147.346.

Impact on process physics. This attack results in changing the relative location of the tool and the workpiece as a function of time compared to the nominal cutting path. In particular, it decreases the length of the first recess in the feed direction from 50 to 48 mm while it increases the length of the second-recess from 50 to 52 mm. For the roughing operations, this results in decreasing the total estimated cutting time of cutting the first recess from 23.057 to 22.046 s (estimated machining times were obtained from the CAM software). Additionally, the total estimated cutting time of cutting the second recess will increase from 23.057 to 24.068 s. As for the finishing operations, this attack will lead to decreasing the total estimated cutting time of cutting the first recess from 11.739 to 11.362 s. Additionally, the total estimated cutting time of cutting the second recess will increase from 11.739 to 12.116 s. As discussed in the attack design, the total

estimated time required for machining the attacked part will not change from the nominal part. It should be noted that the local variations in the cutting time for different geometrical features such as the recesses can be manifested by measurements of in-process machining variables such as cutting power. However, this requires a real-time signal monitoring system that is designed with the possibility that the signal may vary not only in amplitude but also in time.

5.2.5 Attack scenario 5 (G.L.1.1478.346)

Design. The objective of this attack is to achieve the same product alteration as in attack scenario 4 while attacking a manufacturer who has adopted a more in-depth quality control strategy in which in-situ process monitoring system was deployed for both roughing and finishing operations. This is in addition to all the other quality control measures used in attack scenario 4.

Aiming at avoiding the detection of this attack impact by an operator who is knowledgeable about the product, in post-production inspection, and/or the machine controller-based system, similar to attack scenario 4, the adversary needs to adopt the same design considerations numbers 1, 4, and 7. As for the in-situ detection possibility of the attacks' physical impact on the process (i.e., the local variation in the time required to cut the recesses as described in attack scenario 4), an adversary needs to take the in-situ monitoring system added in this scenario into consideration (design consideration #8). In this particular case, the design of the attack needs to consider compensating for the change in the cutting time of machining the recesses due to the change in the cutting length per each cutting path. This can be done through equating the cutting time of the nominal process by varying the feed rate from 0.33 mm/rev for roughing and from 0.2 mm/rev for finishing according to each diameter level. This can be done according to the following formula,

$$t = \frac{v \times l}{\pi \times D \times f}$$

where v is the cutting speed, l is the cutting length, D is the cutting diameter, and f is the feed rate. Table 3 shows the altered roughing and finishing feed rates under attack 5 to equate the required times for machining attacked part to the nominal parts.

Implementation. The implementation of this attack is fairly similar to attack 4. The only difference is to include the feed rate changes reported in the above table in the attack on the roughing and finishing G-codes.

Designation number. By adding the design consideration #8 in addition to the designation of attack 4, the designation number of attack 5 becomes G.S.1.1478.346.

Impact on process physics. In this attack, varying the feed rate following the pattern shown in Table 3 may succeed in mitigating the variation in local time signatures of process physical variables measurements. However, this variation of the feed rate will result in varying the local amplitude signatures of physical variables measurements. It is also worth mentioning that the overall average amplitude of process data signals is supposed to remain similar to the nominal parts. However, the local time variation can be detected by monitoring the local amplitude signatures of physical side-channel measurements. Additionally, varying the feed rate in this manner may not be completely successful in matching the time signatures to the corresponding nominal signatures due to the inability to accurately mimic/predict controller behavior.

5.2.6 Attack scenario 6 (G.F.1.15.34)

Design. The objective of this attack is to change the shape form of the two 40.0 mm diameter recess sections from a straight cylinder to a barrel. The resulting barrels will have a 40.3 mm diameter at its center and a 40 mm diameter at both ends. The manufacturer to be attacked has a quality control strategy that implements a 100% post-production inspection plan for only the spool’s larger outside diameters (Ø50 mm sections). This inspection process is to be performed manually by an operator using Vernier calipers.

To minimize the detectability of the physical impact of the attack both on the product and the process, the adversary can

consider a variety of combinations of attack design considerations. One combination includes (1) physical change magnitude to be relatively small to reduce the probability of being detected by the machine operator (design consideration #1) and (2) attacking features that are not typically inspected (i.e., not included in the inspection procedure) as these features may not be considered critical by the manufacturer (design consideration #5). In this specific example, tight tolerances are not required for the attacked features and hence a wider dimensional variation would be acceptable. The adversary can make use of this information and alter the geometry of these features in a way that affects the function while these features will not be inspected.

Implementation. This attack can be implemented by attacking the G-code for both the roughing and finishing operations (implementation locations #3 and 4).

Designation number. The design and implementation of this attack can be designated as an attack on Geometrical Integrity, Form Class, and Straightness Sub-Class (G.F.1), while accounting for design considerations 1 and 5, and implementation locations #3 and 4. Hence, the designation number of this attack is G.F.1.15.34.

Impact on process physics. A typical CAM software package would generate a tool-path in which the cutting parameters remain similar in as many cutting paths as possible. As described earlier in 5.1, there are four nominal cutting paths to each recess roughing operation. The cutting parameters for these four paths are constant across all different paths (depth-of-cut is 1.25 mm, cutting velocity is 200 m/min, and feed rate is 0.33 mm/rev). However, to produce the spool with the variation intended under attack scenario 6, the CAM software will change only the last roughing path of each recess (roughing paths 5 and 9) by varying the depth-of-cut following the rate of diameter change for the barreled sections. Accordingly, the largest depth-of-cut will be 1.25 mm at both ends of the recess while it decreases gradually going toward the midpoint of the concave line until it reaches the smallest depth-of-cut value being 1.10 mm. Figure 8 shows the cross section of the total removed material. This gradual variation in the depth-of-cut for roughing paths 5 and 9 will lead to a similar variation in the MRR.

Table 3 Altered roughing and finishing feed rates under attack 5

Path#	First recess				Finishing	Second recess				
	Roughing					Roughing				Finishing
	1	2	3	4		1	2	3	4	
Feed rate (mm/rev)	0.315	0.314	0.312	0.309	0.187	0.345	0.346	0.348	0.351	0.213






Fig. 8 A schematic illustration of depth-of-cut variation for roughing paths 5 and 9 under attack 6

As discussed previously, variations in the MRR can be reflected in the measurements of in-process physical variables such as cutting power and forces. Hence, this attack can be detected with proper side-channel monitoring system.

6 Side-channel measurements analysis and discussion

In this section, the side-channel measurement signatures of the spindle power for all experimental parts are presented and analyzed. For brevity, this section primarily focuses on roughing operations power signatures. As described in Section 5, among the eleven test parts, five parts were produced nominally, shown in Fig. 4. The corresponding measured power signatures for these five parts are shown in Fig. 9. The roughing operations for these parts were performed in nine cutting cycles, following the tool path illustrated in Fig. 5 (top). These nine cycles are clearly reflected in the spindle power signatures, shown in Fig. 9, where vertical lines are used to mark the start and the end of each of these cycles, respectively.

Figure 9 shows clearly that the measured power signatures successfully depict the spool geometrical features. For instance, the power measurements, shown in Fig. 9, show a clear spike at the beginning of each of the cutting cycles two through nine. This can be attributed to the gradual increase in the material removal rate to perform the α -taper turning operation in the two recesses within the spool. Several functional features (e.g., mean, maximum, duration, and slope) of this spike can be correlated to the part geometry, tool geometry, and the cutting parameters used to produce a feature. Additionally, the horizontal portions of the power signal around the 5 KW level correspond to the straight turning

operations performed in each of the nine roughing cycles. As discussed earlier, the power level and completion time of each cycle can be correlated to the process independent parameters such as feeds, speeds, and tool geometry. Therefore, changes in a process power signature compared to the nominal/baseline signatures, while producing a specific part, can be used as an indicator of a geometrical change induced by a Product-Oriented C2P attack.

To investigate the assignable variations in the power signature associated with different attack scenarios, five detection metrics were extracted for each cutting cycle from the measured power signals of all test parts. These metrics are:

1. Average power level of the steady-state time-window (shown in Fig. 9 as lightly colored rectangles on the signal) of the straight turning portions of each cycle;
2. Overall average power level for all nine cutting cycles including both straight and taper turning portions;
3. Maximum power level within each cutting cycle;
4. Cutting time within each cutting cycle; and
5. Overall part's production time including return cycles.

Values of detection metrics extracted from the attacked parts' signatures are compared against the 95% prediction-intervals estimated for the same metrics extracted from the five nominal parts, using the t-distribution based on the nominal parts sample size and a 5% significance level. These comparisons against the prediction-interval for each detection metric are to test whether or not a detection metric for a new part belongs to the in-control population. Table 4 shows the detection test results for each attack based on the three per-cycle metrics (metrics 1, 3, and 4) and the two overall metrics (metrics 2 and 5) resulting in 29 tests for each attack (new sample).

Fig. 9 Spindle power signatures of the five nominal parts

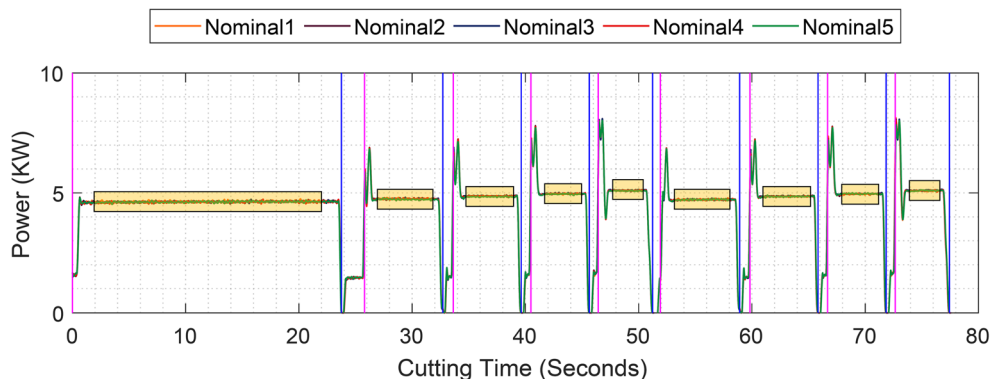


Table 4 Detection Test Results (symbols X, O, and # indicate true positive, true negative, false positive, and false negative decisions, respectively)

Detection Metric	Cutting Cycle #	Detection Test Results for Attack #					
		1	2	3	4	5	6
Steady-State Average Spindle Power (KW)	1	X	O
	2	X	.
	3	X	.
	4	X	.
	5	X	X
	6	.	.	O	O	X	O
	7	.	.	O	.	X	.
	8	X	.
	9	X	X
Overall	X	.	.	.	O	X	
Maximum Spindle Power (KW)	1	X	.	.	.	O	O
	2	.	.	X	.	X	.
	3	.	.	X	.	#	.
	4	O	O	X	O	X	.
	5	.	.	X	.	#	O
	6	O	.	X	O	X	.
	7	O	.	X	.	X	.
	8	.	.	X	.	X	.
	9	.	.	X	.	X	.
Cutting Time (Seconds)	1	.	X
	2	.	.	.	X	O	.
	3	.	.	X	X	O	.
	4	.	.	X	X	O	.
	5	.	.	X	X	O	.
	6	.	O	X	X	O	.
	7	.	.	X	X	.	.
	8	.	.	X	X	.	.
	9	.	.	X	X	.	.
Overall	.	X	X	.	.	X	

Section 5 provided a detailed discussion on the physics behind the attacks and how each attack’s physical impact on the process can be characterized through a side-channel in-situ sensor measurements of process data such as cutting power. Accordingly, the detection metrics that should have signaled that an attack had occurred have been identified; these are indicated by the colored cells in Table 4. The decisions in this table confirm our process physical knowledge driven characterization with a collective true positive rate of 95.9% (assuming test independence). The following paragraphs summarize, for each attack the power signal results confirming the conclusions drawn in Section 5.

Attack 1 The increase of the depth-of-cut of the first roughing path resulted in a significant increase in the steady-state average and maximum power levels of the first segment of the signal. Additionally, this attack increased the overall average power level of the whole roughing operation.

Attack 2 The decrease of the feed rate of the first roughing path, to mitigate the power level increase due to the increased depth-of-cut, resulted in a significant increase in the time of the first segment of the signal as well as the overall completion time of the roughing operations. It should be noted that the estimated reduction in the feed rate did not completely mitigate the increase in the average power level of the first roughing cycle. One potential reason is the inability to accurately mimic/predict controller behavior when estimating the shifted feed rates.

Attack 3 The increase of the recesses taper angle resulted in increasing both the power and time requirements for cutting these four tapers. This change was clearly characterized by the increase in the maximum power level for cycles two through nine of the roughing operations as well as the completion time for these cycles.

Attack 4 The dislocation of the middle section of the spool shifted the completion time of cutting cycles two through nine. This is detected by the decrease in the cutting time of the segments two through five and the increase of the time in segments six through nine.

Attack 5 The variation of the feed rate for cutting cycles two through nine outlined in Table 3, to mitigate the local time variation due to the dislocation of the middle spool section, resulted in reducing the steady-state average power level for segments two through five while increasing it for segments six through nine. Additionally, the maximum power was shifted similarly for seven out of the nine cycles. It should be noted that the estimated variation in the feed rate did not completely mitigate the local variation in the completion time of different cutting cycles. One potential reason, as discussed in Section 5, is the inability to accurately mimic/predict controller behavior when estimating the shifted feed rates.

Attack 6 The change in the depth-of-cut profile, shown in Fig. 8, was reflected in the power signature leading to a clear drop in the average power level of the corresponding signal segments number five and nine. This drop in the power of these two cutting cycles reduced the overall average power level. Additionally, the total completion time of the roughing process was increased due to the increase in the cutting length through changing the cutting path from a straight line to a concave line. However, the local completion time of cycles five and nine did not change significantly compared to the nominal values. This showcases a possible attack induced signal shift that can be hidden within the system variation as discussed in attack design consideration number 9 in Section 4.

7 Conclusions and future work

At its heart, this paper proposed the adoption of the defense-in-depth strategy to transition current quality control (QC) resources to become physical detection layers to defend against Product-Oriented C2P attacks. The discussion regarding the adoption of the defense-in-depth strategy has revealed that to truly defend advanced manufacturing systems from these attacks requires a collective effort across the manufacturing environment from shop floor operators to product/process designers and engineers.

To demonstrate the benefits that can be obtained through the incorporation of the defense-in-depth strategy for machining processes, this paper proposed an attack design scheme and a designation system (ADDS) to not only describe but also understand how Product-Oriented C2P attacks can be designed. This provides the first step toward developing a needed body of knowledge that will provide a clear common understanding of how Product-Oriented C2P attacks on manufacturing systems for machined products can be designed and implemented, similar to the Common Attack Pattern Enumeration and Classification (CAPECTM) repository for software security applications. Through the use of the ADDS, machining process/product designers now have the opportunity to understand the wide variety of potential attacks that could affect their systems. Utilizing this knowledge, designers can leverage available QC resources to increase the difficulty/cost an adversary has to overcome to successfully implement a Product-Oriented C2P attack. One of the future efforts is to apply ADDS for real-world systems.

To highlight the importance of implementing a defense-in-depth strategy, this paper demonstrated several realistic Product-Oriented C2P attacks against a machined spool valve. These demonstrations show that manufacturing systems that implemented numerous quality control tools, that have not been designed to account for attacks, can be compromised. Furthermore, these demonstrations illustrated the importance of incorporating numerous detection layers to impede/prevent attacks. For example, it was shown that simultaneously monitoring the local amplitude and time signatures of machine power significantly reduces the potential for intelligently designed attacks to circumvent in-situ process monitoring. These demonstrations set the foundation for developing process physics-driven machine learning tools for monitoring the ever-increasing real-time big data of manufacturing processes for security purposes.

This work focused on the discussion of using side-channels as an approach to increase the Product-Oriented C2P detection capabilities within the quality control process regime. Further research needs to focus on (1) developing process monitoring techniques that are capable of monitoring for both traditional process anomalies and Product-Oriented C2P attacks driven anomalies (necessary for developing deployable solutions in

real manufacturing facilities); (2) identifying features that are sensitive to a wide range of attacks; and (3) identifying features that are the most sensitive to specific attacks. For the second research area, it is needed to develop robust/generic strategies for using side-channels as detection layers. The third research area is needed to develop machining-scenario specific uses for side-channels. More specifically, if it is known that a system is highly susceptible to a specific attack or that an attack against a specific part feature could cause significant damage, it may be desirable to incorporate specific side-channel features. Furthermore, this paper did not suggest any specific features to extract (e.g., statistical features, wavelet coefficients, principal components) for side-channel data streams. It should be re-emphasized that previously identified side-channel features for process monitoring may not be optimal for detecting Product-Oriented C2P attacks. Finally, this paper did not focus on nor discuss the important trade-offs between the occurrence of false alarms and true-positives. In theory, an infinite number of side-channel features could be monitored; unfortunately, the false alarms costs of such a system would be impractical. Further research needs to address the cost of implementing a depth-in-defense strategy, from both false alarm and resource requirement viewpoints.

Finally, we fully acknowledge that machining is but a small aspect within manufacturing. As such, we envision the proposed ADDS to be an initial step toward a larger approach to cataloging and understanding attacks against manufacturing systems. Further research in this area will focus on the development of a Manufacturing Attack Designation System (MADS), which will incorporate how attacks can be designed considering all manufacturing aspects, including attacking assembly systems, MRP/ERP systems, and supply chains. Together, MADS and CAPEC will fully define the aspects of any Product-Oriented C2P attacks. This knowledge will provide the basis for designing future secure manufacturing systems, which is a fundamental requirement for the successful transition to Industry 4.0.

Acknowledgments The authors would like to thank the Center for Innovation Based Manufacturing (CIBM) at Virginia Tech and its director, Dr. Jaime Camelio, for utilizing the center's manufacturing equipment, sensors, and DAQ system used in running the experimental investigation. Also, the authors thank SANDVIK COROMANT for providing the cutting tools used in the experimental investigation.

Funding information This research work was partially supported by the National Science Foundation (NSF) and Department of Homeland Security (DHS) under grant no. CNS 1446804.

References

1. Jazdi N (2014) Cyber physical systems in the context of Industry 4.0. In: IEEE International Conference on Automation, Quality and Testing, Robotics. IEEE, pp 1–4

2. Ren L, Zhang L, Tao F, Zhao C, Chai X, Zhao X (2015) Cloud manufacturing: from concept to practice. *Enterp Inf Syst* 9(2):186–209
3. Wells LJ, Camelio JA, Williams CB, White J (2014) Cyber-physical security challenges in manufacturing systems. *Manufact Lett* 2(2):74–77. <https://doi.org/10.1016/j.mfglet.2014.01.005>
4. Lee RM, Assante MJ, Conway T (2014) German steel mill cyber attack. *Industrial Control Systems*. SANS Institute
5. Sturm LD, Williams CB, Camelio JA, White J, Parker R (2014) Cyber-physical vulnerabilities in additive manufacturing systems. In: 25th Annual Solid Freeform Fabrication Symposium, Austin, TX
6. Yampolskiy M, Horvath P, Koutsoukos XD, Xue Y (2012) Sztipanovits J Systematic analysis of cyber-attacks on CPS-evaluating applicability of DFD-based approach. In: the 5th International Symposium on Resilient Control Systems (ISRCS). IEEE, pp 55–62
7. Yampolskiy M, Horvath P, Koutsoukos XD, Xue Y (2013) Sztipanovits J Taxonomy for description of cross-domain attacks on CPS. In: Proceedings of the 2nd ACM international conference on High confidence networked systems. ACM, pp 135–142
8. Elhabashy AE, Wells LJ, Woodall WH, Camelio JA (2018) A cyber-physical attack taxonomy for production systems: a quality control perspective. *J Intell Manuf*:1–16
9. Sturm LD, Williams CB, Camelio JA, White J, Parker R (2017) Cyber-physical vulnerabilities in additive manufacturing systems: a case study attack on the STL file with human subjects. *J Manuf Syst* 44:154–164
10. Fabro M, Gorski E, Spiers N (2016) Recommended practice: improving industrial control system cybersecurity with defense-in-depth strategies. DHS Industrial Control Systems Cyber Emergency Response Team
11. Blackwell C, Zhu H (2014) *Cyberpatterns: unifying design patterns with security and attack patterns*. Springer
12. Mitre-Corporation (2018) Common weakness enumeration (CWE). Mitre Corporation. <http://cwe.mitre.org/index.html>. Accessed 05/04/2018
13. Mitre-Corporation (2018) Common vulnerabilities and exposures (CVE®). Mitre Corporation. <http://cve.mitre.org/index.html>. Accessed 05/04/2018 2018
14. Mell P (2005) The national vulnerability database. NIST Presentation
15. Mitre-Corporation (2017) Common attack pattern enumeration and classification (CAPEC™). Mitre Corporation. <https://capec.mitre.org/>. Accessed 05/04/2018 2018
16. NIST (2018) Framework for improving critical infrastructure cybersecurity, Version 1.1. 1.1 edn. National Institute of Standards and Technology
17. Vincent H, Wells L, Tarazaga P, Camelio J (2015) Trojan detection and side-channel analyses for cyber-security in cyber-physical manufacturing systems. *Procedia Manufactur* 1:77–85. <https://doi.org/10.1016/j.promfg.2015.09.065>
18. Huang Y, Leu MC, Mazumder J, Donmez A (2015) Additive manufacturing: current state, future potential, gaps and needs, and recommendations. *J Manuf Sci Eng* 137(1):014001
19. Kline S, Guckes ACM, Schafer J (2017) Machine tools. 2018 Capital Spending Survey Results
20. Hutchins MJ, Bhinge R, Micali MK, Robinson SL, Sutherland JW, Dornfeld D (2015) Framework for identifying cybersecurity risks in manufacturing. *Procedia Manufactur* 1:47–63
21. Chhetri SR, Wan J, Al Faruque MA (2017) Cross-domain security of cyber-physical systems. Design Automation Conference (ASP-DAC), 2017 22nd Asia and South Pacific, IEEE:200–205
22. DeSmit Z, Elhabashy AE, Wells LJ, Camelio JA (2017) An approach to cyber-physical vulnerability assessment for intelligent manufacturing systems. *J Manuf Syst* 43:339–351
23. Yampolskiy M, Skjellum A, Kretzschmar M, Overfelt RA, Sloan KR, Yasinsac A (2016) Using 3D printers as weapons. *Int J Crit Infrastruct Prot* 14:58–71
24. Pan Y, White J, Schmidt DC, Elhabashy A, Sturm L, Camelio J, Williams C (2017) Taxonomies for reasoning about cyber-physical attacks in IoT-based manufacturing systems. *Int J Interact Multimed Artific Intel* 4(3)
25. DeSmit Z (2017) Cyber-physical security in advanced manufacturing Doctoral Dissertation, Virginia Tech
26. Sturm LD, Albakri M, Williams CB, Tarazaga P (2016) In-situ detection of build defects in additive manufacturing via impedance-based monitoring. In: Paper presented at the Proceedings of the 27th Annual International Solid Freeform Fabrication Symposium. An Additive Manufacturing Conference, Austin, pp 8–10
27. Chhetri SR, Canedo A, Al Faruque MA (2016) KCAD: kinetic cyber-attack detection method for cyber-physical additive manufacturing systems. Paper presented at the International Conference On Computer Aided Design (ICCAD '16), Austin, 7-10
28. Belikovetsky S, Solewicz Y, Yampolskiy M, Toh J, Elovici Y (2017) Detecting cyber-physical attacks in additive manufacturing using digital audio signing arXiv preprint arXiv:170506454
29. Wu M, Song Z, Moon YB (2017) Detecting cyber-physical attacks in CyberManufacturing systems with machine learning methods. *J Intell Manuf* 30:1–13. <https://doi.org/10.1007/s10845-017-1315-5>
30. Turner H, White J, Camelio JA, Williams C, Amos B, Parker R (2015) Bad parts: are our manufacturing systems at risk of silent cyberattacks? *IEEE Secur Priv* 13(3):40–47
31. Zeltmann SE, Gupta N, Tsoutsos NG, Maniatakos M, Rajendran J, Karri R (2016) Manufacturing and security challenges in 3D printing. *J Miner, Met Materi Soc (JOM)* 68(7):1872–1881. <https://doi.org/10.1007/s11837-016-1937-7>
32. Belikovetsky S, Yampolskiy M, Toh J, Elovici Y (2016) Dr0wned-cyber-physical attack with additive manufacturing arXiv preprint arXiv:160900133
33. Moore SB, Glisson WB, Yampolskiy M Implications of malicious 3D printer firmware. In: Proceedings of the 50th Hawaii International Conference on System Sciences. HICSS, Waikoloa Village, pp 6089–6098
34. Slaughter A, Yampolskiy M, Matthews M, King WE, Guss G, Elovici Y (2017) How to ensure bad quality in metal additive manufacturing: in-situ infrared thermography from the security perspective. In: Proceedings of the 12th International Conference on Availability, Reliability and Security. ACM, p 78
35. Wu SX, Banzhaf W (2010) The use of computational intelligence in intrusion detection systems: a review. *Appl Soft Comput* 10(1):1–35
36. Evans JR, Lindsay WM (2013) *Managing for quality and performance excellence*. Cengage Learn
37. Teti R, Jemielniak K, O'Donnell G, Dornfeld D (2010) Advanced monitoring of machining operations. *CIRP Annals-Manufact Technol* 59(2):717–739
38. Tang C (2017) Key performance indicators for process control system cybersecurity performance analysis. US Department of Commerce, National Institute of Standards and Technology
39. Urbina DI, Urbina DI, Giraldo J, Cardenas AA, Valente J, Faisal M, Tippenhauer NO, Ruths J, Candell R, Sandberg H (2016) Survey and new directions for physics-based attack detection in control systems. US Department of Commerce, National Institute of Standards and Technology
40. Barnum S (2008) Common attack pattern enumeration and classification (capec) schema description. Cigital Inc, https://capec.mitre.org/documents/documentation/CAPEC_Schema_Description_v1.3
41. Kuttolamadom MA, Mears ML, Kurfess TR (2012) On the volumetric assessment of tool wear in machining inserts with complex

- geometries—part 1: need, methodology, and standardization. *J Manuf Sci Eng* 134(5):051002
42. ASME (2009) Dimensioning and tolerancing, p Y145
 43. Federal-Aviation-Administration (2012) Aircraft landing gear systems. In: Aviation maintenance technician handbook - airframe, vol 1. U.S. Department of Transportation, Flight Standards Service
 44. Caldwell T (2011) Ethical hackers: putting on the white hat. *Netw Secur* 2011(7):10–13

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.