Watching and Safeguarding Your 3D Printer: Online Process Monitoring Against Cyber-Physical Attacks

YANG GAO, University at Buffalo, the State University of New York, USA
BORUI LI, Binghamton University, the State University of New York, USA
WEI WANG, WENYAO XU, CHI ZHOU, ZHANPENG JIN, University at Buffalo, the State University of New York, USA

The increasing adoption of 3D printing in many safety and mission critical applications exposes 3D printers to a variety of cyber attacks that may result in catastrophic consequences if the printing process is compromised. For example, the mechanical properties (e.g., physical strength, thermal resistance, dimensional stability) of 3D printed objects could be significantly affected and degraded if a simple printing setting is maliciously changed. To address this challenge, this study proposes a model-free real-time online process monitoring approach that is capable of detecting and defending against the cyber-physical attacks on the firmwares of 3D printers. Specifically, we explore the potential attacks and consequences of four key printing attributes (including infill path, printing speed, layer thickness, and fan speed) and then formulate the attack models. Based on the intrinsic relation between the printing attributes and the physical observations, our defense model is established by systematically analyzing the multi-faceted, real-time measurement collected from the accelerometer, magnetometer and camera. The Kalman filter and Canny filter are used to map and estimate three aforementioned critical toolpath information that might affect the printing quality. Mel-frequency Cepstrum Coefficients are used to extract features for fan speed estimation. Experimental results show that, for a complex 3D printed design, our method can achieve 4% Hausdorff distance compared with the model dimension for infill path estimate, 6.07% Mean Absolute Percentage Error (MAPE) for speed estimate, 9.57% MAPE for layer thickness estimate, and 96.8% accuracy for fan speed identification. Our study demonstrates that, this new approach can effectively defend against the cyber-physical attacks on 3D printers and 3D printing process.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

Additional Key Words and Phrases: Cyber-Physical Security, Online Process Monitoring, Sensor Fusion, 3D Printing

ACM Reference Format:
Yang Gao, Borui Li, and Wei Wang, Wenyao Xu, Chi Zhou, Zhanpeng Jin. 2018. Watching and Safeguarding Your 3D Printer: Online Process Monitoring Against Cyber-Physical Attacks. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 3, Article 108 (September 2018), 27 pages. https://doi.org/10.1145/3264918

Authors' addresses: Yang Gao, University at Buffalo, the State University of New York, School of Computer Science and Engineering, Buffalo, NY, 14260, USA, ygao36@buffalo.edu; Borui Li, Binghamton University, the State University of New York, Binghamton, USA; Wei Wang, Wenyao Xu, Chi Zhou, Zhanpeng Jin, University at Buffalo, the State University of New York, School of Computer Science and Engineering, Buffalo, NY, 14260, USA.

2474-9567/2018/9-ART108 $15.00
https://doi.org/10.1145/3264918

1 INTRODUCTION

Additive manufacturing (AM), also commonly known as 3D printing, is an emerging and evolving cyber-physical system that builds three-dimensional objects by accumulating layer-upon-layer of materials, as opposed to traditional manufacturing process of material removal such as casting and milling. With the advances in materials and mechanical technologies, AM has gained significant growth and development in the last few decades. According to the Wohlers Report 2016 [58], the AM industry has a 25.9% CAGR (Corporate Annual Growth Rate) and is expected to exceed 20 billion by 2020.

With the significant advantages in short time-to-market, freedom to design, reduced tooling costs and material diversity, 3D printing has been applied in a variety of industries, ranging from many safety- and mission-critical applications to consumer-grade, user-friendly products, including biomedical fabrication, architecture and construction, and aerospace and automotive manufacturing. SpaceX launched its Falcon 9 rocket with a 3D printed part in its engine [44]. General Electric intended to apply its first 3D-printed parts in their aircraft engine platform [39]. The U.S. Navy has very recently unveiled the first proof-of-concept, 3D-printed submersible hull, which was 90 percent cheaper and produced within a few days instead of the conventional 3-5 months [30]. Disney has extensively used 3D printing and digital fabrication in designing more responsive and interactive media and entertainment products [23]. More broadly, it has seen a growing trend that 3D printing has been integrated and incorporated with the advanced wearable technologies and smart devices to create a pervasive and ubiquitous computing environment [4, 25]. Numerous efforts have been made to promote the vision of personal fabrication and supporting a growing DIY community building interactive objects and machines [40].

As 3D printing becomes more popular and widely accessible, some highly sensitive products are getting involved, which makes 3D printer a potentially significant target of attacks and the cyber-physical vulnerabilities of 3D printing process become non-negligible [8, 10, 16, 54, 60]. A cyber-physical attack could be designed to lower production efficiency or damage the printer itself. More insidiously, it can compromise a printing process to fabricate defective or faulty parts that may pass inspection but eventually fail during the real operation, and thereby, result in serious injuries and catastrophic consequences [50, 51]. For parts produced by 3D printing technology, previous research has shown that an attack can alter design files or process parameters (e.g., tool paths) to bring a part out of specification and such attacks can disrupt the product design process and adversely affect a product's design intent, performance, or quality [56]. To this aim, some preliminary studies have explored the effects of altered printing orientation and undetectable defects in the interior of 3D-printed parts, using sophisticated ultrasonic testing [64]. Similar study about the influences of printing orientation and layer thickness on material properties has also been conducted at NASA [57]. As shown in Figure 1, traditional defenses against cyber-physical attacks in 3D printing can be roughly divided into two groups: pre-process software validation, and post-process quality control. For existing software techniques [11, 41], they still can't guarantee a high security level due to the stealth and persistence of firmware attacks. For instance, the buggy codes and third-party libraries included in firmware image files remain vulnerable and could be exploited by attackers. Thus, a firmware-independent integrity inspection is necessary.

On the other side, unfortunately, the traditional quality control in AM is largely limited to offline processing. Existing mitigation strategies, such as weighting test and redundancy, were intended to detect accidental manufacturing defects based on statistical analysis for large volume production, not suitable for those deliberately hidden flaws without any geometry/weight change and small-batch AM fabrications. For contact inspections [29], it would lead to high scrap rates [38]. Other existing visual and measured inspection methods (e.g., 3D laser scanning) usually demand high equipment and computational costs [1] and are difficult to deploy in measurement of AM parts with free-form geometries and support structures attached. Ultrasonic-based measurement [17] is designed to detect surface displacement and shell thickness, which is inaccurate to detect complicated

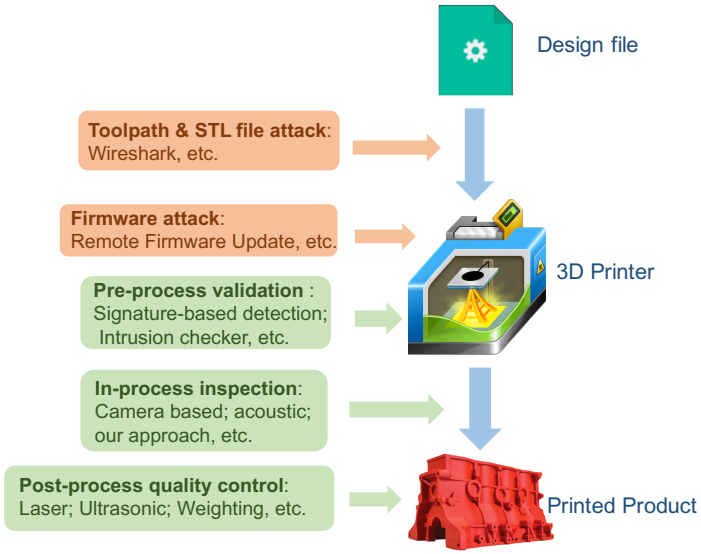

Fig. 1. Taxonomy of attacks and defenses in additive manufacturing process chain.

and imperceptible inner structure flaws [63]. In addition, post-process inspection has a relative high time cost especially for refined 3D printed products.

Unlike the large manufacturing enterprises that often use expensive industrial powder-based 3D printers (e.g., SLS), more affordable consumer-grade FDM (Fused Deposition Modeling) 3D printers are still the most widespread desktop technology and thus are favored by small businesses and end users. Moreover, 3D printers, as a potential ubiquitous computing technology available in people's daily lives [5, 34, 42], have also brought the concerns in privacy and security. In other words, those new ubiquitous computing components are more vulnerable to cyber-physical attacks that can lead to stealing or corrupting designs, and also can't deploy the high-resolution commercial-grade monitoring equipment (e.g., laser scanners or ultrasonic displacement sensors) because of their prohibitive costs. By leveraging heterogeneous sensing and computing devices, this paper aims to provide the end users with a ubiquitous solution to address the unique security demands that could be deployed in general desktop FDM printers. The proposed approach can effectively detect not only the surface abnormalities in a 3D-printed object, but also the malicious modifications in each infill layer. Specifically, we seek to address the potential cyber-physical threats on the 3D printing process chain, and analyze the specific printing attributes, including the infill path, printing speed, layer thickness, and fan speed, all of which have significant impact and influence on the quality of 3D printed products. We also develop a multiple sensor fusion model to extract and estimate the printing attributes. Experimental results show that, for a complex 3D printed design, by analyzing the multi-faceted, real-time measurement and status information collected from the Inertial Measurement Unit (IMU) sensor and video camera, our method can successfully reconstruct the printing attributes with a 4% Hausdorff distance compared with the ground truth for the infill path, 3.12% MAPE (Mean Absolute Percentage Error) for the printing speed, 9.57% MAPE for the layer thickness, and 96.8% accuracy for fan speed identification.

To the best of our knowledge, this study is the first of its kind to explore an effective and ubiquitous defense mechanism against cyber-physical attacks on 3D printers from both kinetic and thermodynamic prospects. Our contributions are summarized as follows:

- We explore the cyber-physical attack threats in the 3D printing process chain, particularly on the firmwares of 3D printers, and formulate the attack models.

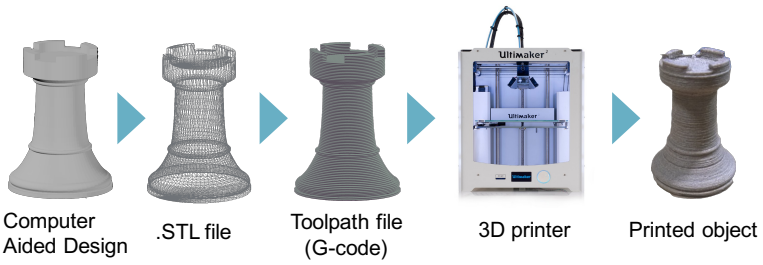

Fig. 2. 3D printing process chain.

- We propose and develop a real-time low-cost pervasive monitoring approach, capable of detecting and defending against the cyber-physical attacks on consumer-grade 3D printers. The in-process defense mechanism is based on the reconstruction of printing attributes and analysis of multi-faceted measurement and status information from the IMU sensor and camera.
- We validate the effectiveness of the proposed, multiple sensor fusion based in-process monitoring system on different types of printers in a real case study.

The rest of this paper is organized as follows: We introduce the background of 3D printing and explore the potential cyber-physical attacks on 3D printing in Section 2. Then we investigate and formulate the specific cyber-physical attack models on four crucial printing attributes in Section 3. Based on the proposed attack models, we propose a multiple sensor fusion based online process monitoring system in Section 4. Section 5 presents the evaluation results for the reconstruction of printing attributes including the infill path, printing speed, layer thickness, and a real case study of cyber-physical attacks on 3D printers. The limitations and future work are discussed in Section 6.

2 BACKGROUND AND RELATED WORK

2.1 Overview of 3D Printing Process Chain

Currently, there are many existing technologies applied in 3D printing, such as Fused Deposition Modeling (FDM), Selective Laser Sintering (SLS), Electron-Beam Melting (EBM), and Stereolithography (SLA). It was reported that [58] over 278,000 desktop 3D printers were sold worldwide in 2015 and most of them were based on FDM technology, more affordable and accessible compared with other methods. Meanwhile, thanks to the unrelenting efforts of the community on low-cost, open-source 3D printers [3, 37], such as the RepRap project [28] and the LulzBot 3D printers, it has seen a boom in entry-level 3D printing machines and the cost of 3D printers has decreased dramatically. Since 2010, the price of 3D printers that used to cost $20,000 has dropped to a level of $1,000 or less. According to a recent forecast report [22], the low-cost, sub-$1000 desktop 3D printers will continue to be a major driving force for growth and is expected to grow at a rate of 12% into 2020. Therefore, in this paper, we will focus on the FDM 3D printers.

Figure 2 shows a complete 3D printing process consisting of five steps. First, Users create a digital 3D model using Computer Aided Design (CAD) software and convert it into a stand StereoLithography (STL) file which is widely used in rapid prototyping. Then, during the Computer Aided Manufacturing (CAM) process, a layer description toolpath file (G-code) is generated by performing operations that includes slicing, path-planning, and support generation, by using 3D printing slicing tools such as Cura and Slic3r. Those toolpath instructions will be sent to the 3D printer to instruct the firmware to control the motor movement, the fan speed, and the extruder's printing speed in order to fabricate the desired product.

G-code is the *de facto* machine code in 3D printers to describe the object's geometry information and the printing process information. Specifically, as shown in Figure 4, the geometry information is specified by the

X, Y coordinates in the instructions, while the printing settings and run-time process parameters (e.g., nozzle printing speed with/without extrusion) are described as instruction options (e.g., E—, G—, Z—, etc.). The process information is independent of the original design and only affects the printing process.

2.2 Cyber-Physical Attacks on 3D Printing

Cyber-physical security is not new for traditional manufacturing. Many research have been performed focusing on Supervisory Control and Data Acquisition (SCADA) [2, 56]. While additive manufacturing shares the similarity of cyber-physical security issues, it possesses some unique challenges. The 3D printers are highly integrated with the production cycle, which can be attacked anywhere from the initial design, data transmission, fabrication, to the final quality inspection and product verification. In addition, as cyber-manufacturing or IoT-based manufacturing becomes a new trend for 3D printing, it makes 3D printers expose to more and more cyber-attacks [6, 24, 56]. In this section, we will explore and discuss the potential cyber-attacks in the 3D printing process chain.

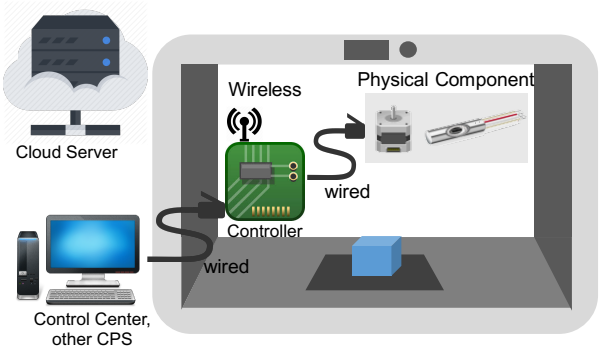

Fig. 3. A Cyber-physical model for 3D printer

For example, if a 3D printer is compromised by an adversary or even if a simple printing setting is maliciously changed, the mechanical properties (e.g., physical strength, thermal resistance, dimensional stability) of 3D printed objects could be significantly affected and degraded. In case the manufactured parts are used in jet engines or other safety-critical systems, they could endanger human life, disrupt critical infrastructure and produce significant economic and societal impacts [62].

As shown in Fig 3, A simplified cyber-physical model can consist of three parts: (1) the physical components that perform physical thermal and mechanical operations (such as the step motors and heater in the 3D printer), (2) the cyber-physical components that convert the cyber instructions into physical commands (such as the micro-controller), and (3) the cyber components that compile and store the instructions and design files (such as

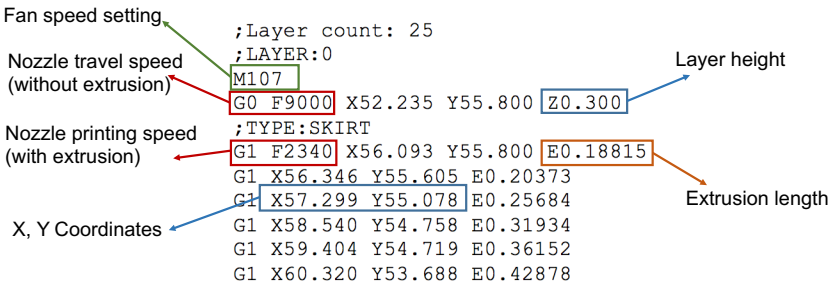

Fig. 4. An example of the main body in G-code

the PC or cloud server connected via USB or Internet). Accordingly, the cyber-physical attacks and defenses on 3D printing may include the following three perspectives:

Wireless communication between the server and the controller. In the cyber domain, eavesdropping and data injection are the two common types of cyber attacks. Traditional cyber-defense mechanisms including encrypted command (G-codes), cypher block chaining encryption, and cryptographic checksum would be useful. In the physical domain, leakage of side-channel information such as acoustic or vibration profile represents another source of possible attacks [43]. Accordingly, frequency hopping and mix of random noises to normal operations could be potential ways to prevent information leakage.

Controller. Compromising controller means accepting adversary commands which are not part of the defined application layer protocol, or interpreting user-commands in an appropriate way. The consequence can be applied to various controlled physical units. External sensors independent from the controller or the firmware as well as data evaluation can significantly defend against such attacks.

USB communication between the PC and the controller. Connected device infection is another type of cyber-physical attack. It is very common to connect various external devices via the USB interface for the flash, erase operations, or updating the firmware [61]. This enables all kinds of attacks via USB connection when the controller runs in the normal mode, while the malicious code is active. However, presence of firewall and anti-virus software on the maintenance computer could effectively address such attacks.

Specifically, aiming at the major components that are vulnerable to attacks in the 3D printing chain, we will explore the potential security threats on the following components: STL files, G-code, and 3D printer.

Security Threats on STL Files. A .STL file is generated from a precise CAD model using tessellation process. After the STL generation, the model is no longer composed of multiple complex mathematical equations, instead it represents an aggregate of many geometry triangular facets. Potential attacks on this stage could be to change the coordinates of the vertices which may modify the product surfaces and inner structure, or add additional vertices inside the model to create new features. However, the design model represented in the STL file still needs to be sliced into layer descriptions (i.e., the toolpath file) before being sent to the 3D printer. Operators or users may still be able to observe the defective layers during the Computer Aided Manufacturing (CAM) process.

Security Threats on Toolpath Files (G-code). Upon receiving the STL file, CAM will generate the toolpath instructions layer by layer. These instructions control the precise movement of the stepper motors on X-Y-Z axes and specify the extrusion settings (e.g., extrusion rate, heating temperature, fan speed, etc.).

As many 3D printers have become to rely on remote printing, which would leave some insecure communication ports. Thus, the toolpath file could be intercepted and modified through the slicing software on the PC and the wireless transmission. For example, Turner *et al.* [52] investigated the toolpath file attacks by using WireShark, and demonstrated that it is possible to intercept the printing information including temperature setting, fan speed, and even the layer description during the transmissions of data and commands from the computer to the 3D printer. However, leveraging some advanced encryption algorithms, the toolpath file transmission can be better secured and protected. In addition, the toolpath file is not the final step in the entire 3D printing process chain, which means that it can still be verified by code profiling before fabrication [14].

Security Threats on 3D Printers. The last step in the 3D printing process chain is the 3D printer itself. This stage of process is mainly vulnerable to the firmware attacks [33]. Even though the STL files and G-codes can be guaranteed as authentic through software inspection, if the controller or firmware is compromised, the 3D printer could still interpret the original instructions into some inaccurate or undesirable operational commands. Thus, it is mandatory to secure the firmware against malicious changes. The firmware not only specifies the geometry information about the design, but also controls the specific manufacturing process settings. For example, the attacker can manipulate the firmware to modify the actual printed coordinates of the infill layers in order to cause the layers to be misaligned, weaken the interior structure, or create the void cavities. Moreover, the attacker

can also increase the nozzle moving speed to an abnormal level, which will result in the under-extrusion of the printed layer. Cui *et al.* [15] exploited the firmware vulnerability and showed that their RFU (Remote Firmware Update) can allow arbitrary injection of malware into the printer's firmware. Chhetri *et al.* [13] implemented the zero-day kinetic cyber-attack onto the 3D printer's firmware, and demonstrated the feasibility of changing some printing parameters of the 3D printed objects.

2.3 Existing Defense Mechanisms on 3D Printing

Some prior work have already employed sensors for in-process monitoring of 3D printing [17, 31, 38]. Rao *et al.* [38] performed an online surface quality monitoring system by multiple sensors. More recently, Straub has conducted a series of research studies and developed a camera-based inspection system [47] to defend against attacks that adversely change printing orientation [48] and infill level [45], as well as to detect the introduced defects [46] and misuse of printing material [49]. However, most of those existing work primarily focused on failure detection, printing optimization, and inspection of geometry and surface. Those online defect inspection and detection systems are not sufficient and reliable to discover model-based attacks. Due to the unique layer-upon-layer process of 3D printing, the product's physical properties such as compressive strength and stiffness can be modified by changing the interior structure without affecting the exterior surface of the product. For example, voids can be placed inside a crucial part during the 3D printing process [50, 51], or the extruded material properties of the infill layers can also be changed without modifying the shell layers. Because a void is completely enclosed inside the model after manufacturing, it is undetectable by common dimensional or visual measurements. Wu *et al.* [59] proposed a camera-based detection method against cyber-physical attacks that maliciously modify the StereoLithography (STL) file. However, their detection model is limited to the void attack in a simple design with the same shell and infill patterns for all the layers. Images from the top view can not precisely reflect the overlapped printing path for complex designs. Chhetri *et al.* [13] used machine learning technique to infer printing traces based on acoustic information emitted by the printer's motors. Their method demands a great amount of sensor data for *a priori* training purpose. Moreover, in real scenarios, the sequential analysis of acoustic information can be easily affected by environmental noises. Bayens *et al.* [6] proposed a three-layers malicious fill patterns detection method by using the microphone, IMU sensors, and the CT scanner. However, the acoustic layer and scanning layer are model-dependent and designed for large-volume production. In addition, beside the kinetic attacks that the paper assessed (e.g., infill patterns), the thermodynamic behaviors (e.g., cooling the filament) will also be potential attack targets.

3 CYBER-PHYSICAL ATTACK MODELS

3.1 Attack Models

According to Figure 4, each command line in the G-code specifies the details of every movement of the nozzle, including the X- and Y-axis coordinates, nozzle traveling and printing speeds, material extrusion length, and height of each layer. As discussed in Section 2.2, the firmware in a 3D printer could be one of the primary targets of cyber-physical attacks which may adversely change the way to interpret and execute the cyber-intellectual property (cyberIP) information received from the G-code.

Specifically, the *kinematic* and *thermodynamic* properties are the two major targets of firmware attacks during the 3D printing process. The *kinematic properties* specify the behaviors of the extruder in sending the correct amount of filament to the hot end and then extruding the melted filament down in thin layers following the designed toolpath, while the *thermodynamic properties* regulate the control for heating, melting, and cooling. From the attacker's perspective, manipulating the kinematic and thermodynamic properties of supporting layers inside the shell would be an effective and almost imperceptible way of attacks that can lead to disastrous consequences. Therefore, in this study, we primarily focus on the attacks on both kinematic and thermodynamic properties and

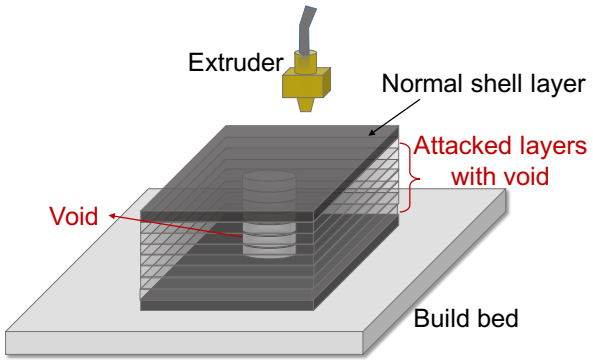

Fig. 5. An example of the void attack.

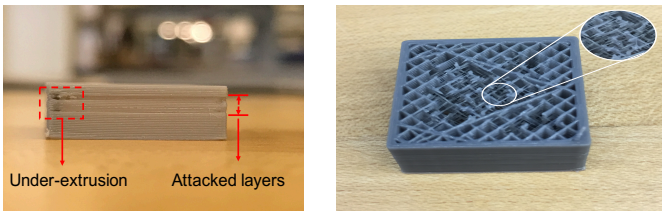

(a) Printing speed attack of shells (b) Printing speed attack of infill layers

Fig. 6. An example of the printing speed attack

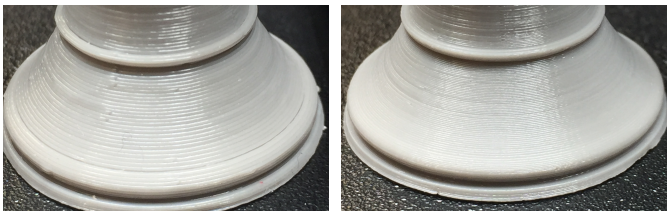

(a) Attacked model with 0.2 mm layer thickness (b) Original model with 0.05 mm layer thickness

Fig. 7. An example of the layer thickness attack

specifically aim to investigate four different types of firmware attacks based on the cyber-IP information of the 3D design model.

Infill Path Attack. Different from the traditional material removal manufacturing processes, the infill structure is not only one of the most significant components for determining the mechanical performance and physical property of 3D printed objects, but also the most challenging part to measure and detect. All these characteristics make the infill structure a perfect target of cyber-physical attacks. Figure 5 illustrates a void attack where a void cavity is created within the printed structure by altering the infill paths for multiple layers. This presence of the cavity will weaken the strength of the manufactured product [9]. For a sensitive component in a complex system, if the voids are located near the load-bearing region, it may cause the direct failure of the printed object and result in catastrophic consequences to the entire system.

Printing Speed Attack. Printing speed is another crucial attribute that influences the printing quality. Attackers can change each layer's printing speed combined with the extrusion rate, to cause under-extrusion or over-extrusion. As shown in Figure 6, the physical and mechanical properties of 3D printed objects are observably altered and degraded by the varying printing speed. Through modifying the printing speed of some randomly or selectively chosen infill layers, the adversary's attack behaviors are more hidden and hard to detect. In addition,

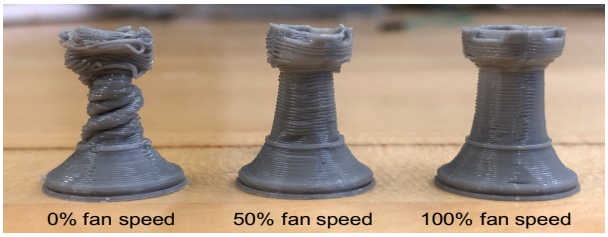

Fig. 8. An example of the fan cooling attack.

the critical speed for under-extrusion or over-extrusion is determined by many factors including the temperature, model structure, material property and the printer itself, and there isn't a generic fixed threshold to limit the printing speed. Thus, a real-time, on-the-fly printing speed monitoring is strongly demanded to guarantee the print quality.

Layer Thickness Attack. The layer thickness is a major attribute that determines the surface roughness and surface texture of 3D printed objects, as shown in Figure 7. It can be thought of as a resolution of the layered manufacturing process [26]. By maliciously changing the layer height values in G-code file, the surface roughness and geometry size of 3D printed objects can be easily re-defined. Although most of time these attacks are visibly detectable and can be resolved through post-processing like polishing, it is still quite challenging to visually identify all of the potential damages if the victim layers are located in the region with complex geometries (e.g., corner swell and ringing) or are sparsely distributed throughout the vertical axis of 3D printed objects.

Fan Cooling Attack. 3D printing is actually the process of filaments remodeling, from the solid state to the liquid state during heating, and back to the solid state again during cooling. The heating and cooling are the two dominant factors that affect the thermal and physical properties of the printed products. Given the critical role of the temperature control, most of existing 3D printers are equipped with built-in nozzle temperature monitoring module and users can also monitor the temperature by setting up an infrared thermometer. However, the cooling process controlled by the fan is less well-defined and often difficult to assess, which may become another unperceivable target of malicious attacks. As an example, Figure 8 shows that different fan speeds will significantly affect the printing quality and all these influences can be hidden when occurring inside a printed object.

3.2 Definitions

Definition 1: Cyber IP. For a 3D printing process, let $\mathbb{C}$ denotes the whole cyber-intellectual property information for the design model. c_i denotes each independent cyber-IP attribute that determines and influences the quality of 3D printed objects, such as the printing speed, the infill path, the layer thickness, and the printing temperature. Specifically, we define C as the complete information set of the design and the 3D printing settings described in G-code.

$$C = \{c_1, c_2, ..., c_k\}, \forall c_i \in \mathbb{C}, C \subseteq \mathbb{C} \tag{1}$$

Definition 2: Cyber-Physical Attacks. Let f be the attacking method. u_i represents the modified IP attribute c_i by adversaries and U denotes the complete set of the modified IP attributes.

$$U = \{u_1, u_2, ..., u_m\}, u_i = f(c_i), \tag{2}$$

Definition 3: Physical Observations. Let P be the variable set of all the physical information that can be measured and collected from the 3D printer. Each p_i is considered as the needed observations combination to reconstruct corresponding cyber-IP information as$\{c_1, c_2, ..., c_n\}$.

$$P \supseteq \{p_1, p_2, ..., p_n\} \tag{3}$$

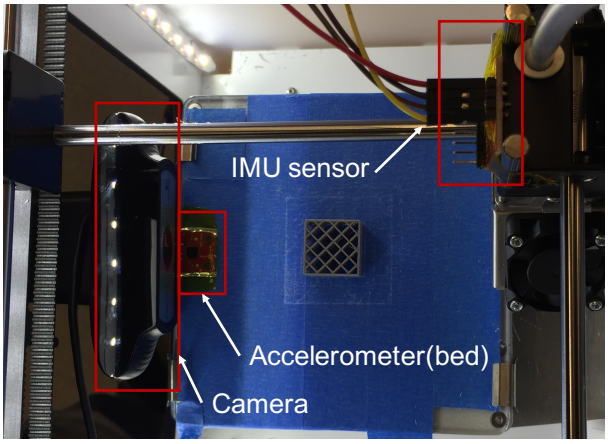

Fig. 9. Setup of multiple sensors for real-time process monitoring (top-down view)

Definition 4: Reconstruction Function. Let G indicates the set of multiple reconstruction functions for physical observations P. And $g(.)$ is each individual analysis function and is determined through the analysis of the mapping relation between each specific cyber-IP attribute and corresponding physical observations.

$$G = \{g_1(.), g_2(.), ..., g_k(.)\} \tag{4}$$

3.3 Problem Formulation

Formulation 1: IP Reconstruction. The purpose of IP reconstruction is to extract the mechanical and printing attributes from the collected physical observations P and thus estimate the cyber-IP information of the model that the 3D printer actually prints. We define S be the estimation set for various type of physical observations through reconstruction analysis:

$$S = \{s_1 \leftarrow g_1(p_1), s_2 \leftarrow g_2(p_2), ..., s_k \leftarrow g_k(p_k)\} \tag{5}$$

Formulation 2: Cyber-Physical Attack Assessment. We evaluate the reconstructed IP information S against the ground truth cyber-IP information C and assess the cyber-physical attacks based on the tolerance $\Delta\varepsilon$.

$$O(S, C) = \begin{cases} > \Delta\ \varepsilon \Rightarrow \text{Attack} \\ \leq \Delta\ \varepsilon \Rightarrow \text{Normal} \end{cases} \tag{6}$$

4 MULTI-SENSOR MONITORING SYSTEM

In this section, we propose a multi-sensor-based, real-time, online process monitoring system to detect cyber-physical attacks. As discussed above (see Section 3.1), this study specifically focuses on the attacks that adversely alter the kinematic and thermodynamic behaviors of the 3D printer, including the infill path, printing speed, layer thickness, and fan cooling. As shown in Figure 9, we instrumented the Ultimaker 2 Go desktop FDM 3D printer with multiple sensors, including an Inertial Measurement Unit (IMU) sensor which is a very common sensor for kinetic estimation [7, 21], an accelerometer, and a camera. (Please refer to Table 1 for sensor configuration details.) The rationale of choosing the appropriate sensors is to provide the needed combination of measurements and observations (i.e., p_i in Equation 4) to reconstruct the corresponding cyber-IP attributes (i.e., s_i in Equation 5), while minimizing the overhead costs and influence on the printing process. Those chosen sensors are light-weight, ultra-portable, and non-intrusive, which ensures that the precision of printing process will not be affected by the proposed defense configuration.

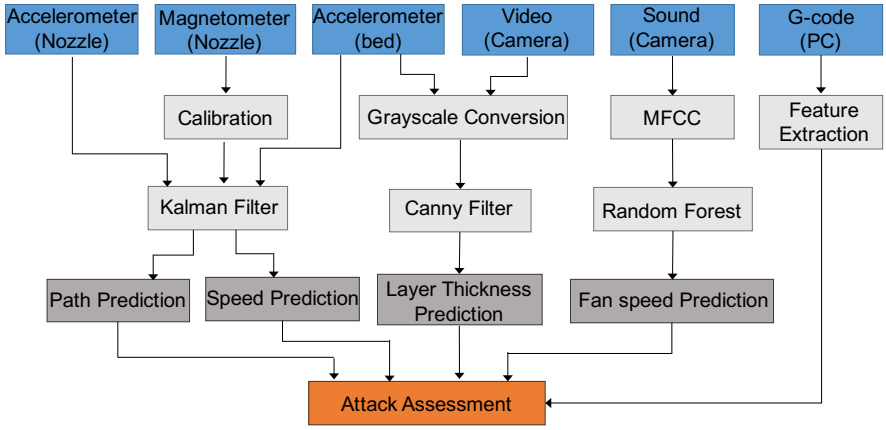

Fig. 10. In-process monitoring methodological flow against cyber-physical attacks on 3D printers

Table 1. Sensor setup and configurations based on different printing characteristics

Sensor Type	Location	Measured Variable	Model	Vendor
Accelerometer	Extruder	Extruder acceleration	LSM9DS1	STMicroelectronics
Accelerometer	Build Bed	Bed vibration	ADXL335	Analog Devices
Magnetometer	Extruder	Extruder Magnetic intensity	LSM9DS1	STMicroelectronics
Video camera	Build Bed	Video images and sound	C960	Logitech

4.1 Multiple Sensor Fusion Model

In the proposed sensor-fusion based online process monitoring system (Figure 10), there are two dimensions of attack analysis: the horizontal and vertical dimensions. Horizontal dimension represents the monitoring of the printing attributes in the X-Y plane, which include the printing infill path and the printing speed. In this dimension, since there is no nozzle rotation, only the acceleration and magnetic intensity of the extruder are collected and fused to compensate and eliminate errors. Specifically, the magnetic intensity data is collected by the magnetometer and fed into a calibration stage to obtain the estimated extruder movement; and the data collected from the two accelerometers are processed by the Kalman filter combined with the calibrated magnetic data. The outputs of the Kalman filter are the estimates of the real-time printing path and speed. In the vertical dimension, a high resolution camera is deployed to monitor the side view of the 3D printed object during the printing process and capture the precise variations of the layer thickness. Besides the sensory information described above for the kinetic attacks, the acoustic channel information recorded by the camera is also acquired to assess the cooling attack.

In summary, we acquire the acceleration, magnetic field, image and acoustic information from these three sensors/devices:

- **Extruder acceleration**: The extruder's movement indicates the actual infill path during the printing process. We use a three-axis accelerometer with a high sampling rate to record the acceleration of the extruder in the real-time manner.
- **Extruder magnetic field intensity**: Due to the high printing speed (10 mm/s to 120 mm/s) and the limited printing space (200 mm × 200 mm), it is insufficient to achieve a high tracking accuracy solely based on the accelerometer. Therefore, we also use the magnetic field intensity of the extruder to help correct the predicted infill path.

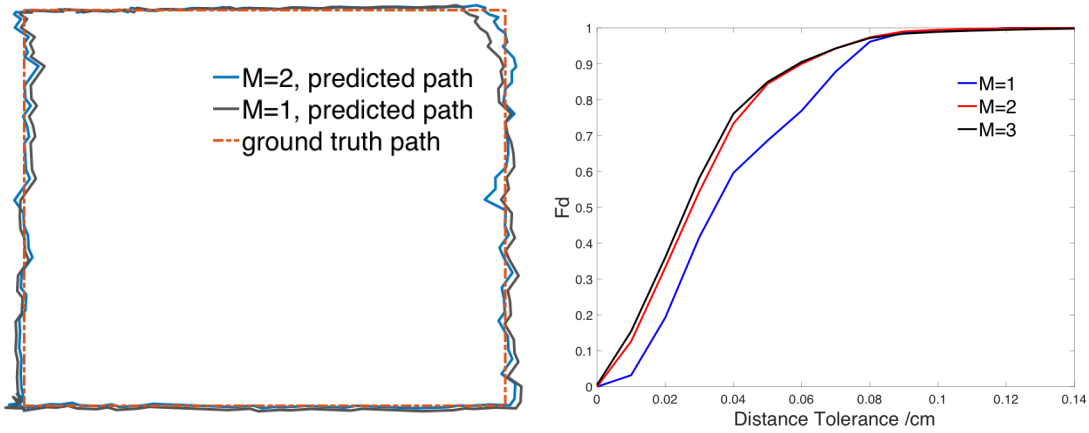

(a) Comparison between predicted square path when M=1 and M=2

(b) Prediction accuracy (DTCDF) for M=1, M=2, and M=3

Fig. 11. Comparison of calibration results for test square shape ($40mm \times 40mm$) under different M. The definition of DTCDF in (b) is discussed in Section 5.2

- **Build bed vibration**: Most desktop 3D printers, including our testing printer Ultimaker 2 Go, apply the lifting of the build bed as the Z-axis movement for the deposition of layers. By measuring the vibration of the build bed, we can monitor the time interval between each layer during the printing.
- **Product side-view images**: The layer thickness for infill and shell is the same for 3D printing. Thus, surface roughness can be measured as the indication of the layer thickness attack. A high resolution video camera is used to monitor the surface roughness of the printed object on the fly.
- **Cooling fan noises**: A spinning cooling fan with higher than 4000 rpm can't be visually inspected and also generate many side channel information. Because the mechanical vibration and magnetic emission are often dominated by environmental factors, we examine the acoustic channel in order to precisely evaluate the fan speed variations.

4.2 Reconstruction of Infill Path and Printing Speed

4.2.1 Calibration Matrix. The earth can be simplified as a rock filled with ferro-magnetic minerals, emitting a strong magnetic field signal that can be captured by a magnetometer. Thus, a magnetometer is widely used for direction navigation [19]. By analyzing the three axes of the magnetic field intensity from the magnetometer, we can get a relatively clear, but not accurate enough, direction for the extruder's movement. However, the alignment of magnetic field varies from place to place, and the readings of the magnetometer could be affected by the environment, such as the stepper motors in the 3D printer [20].

Because the extruder traverses in the horizontal plane, only the X and Y axes are involved in the magnetic field intensity analysis. The relationship between the raw measured magnetometer data and the true path on the X-Y plane can be considered as:

$$\begin{bmatrix} \alpha_x \\ \alpha_y \end{bmatrix} = f(\begin{bmatrix} \beta_x \\ \beta_y \end{bmatrix}, \begin{bmatrix} \gamma_x \\ \gamma_y \end{bmatrix}) \tag{7}$$

where $\begin{bmatrix} \alpha_x \\ \alpha_y \end{bmatrix}$ is the measured vector in the X and Y axes, $\begin{bmatrix} \beta_x \\ \beta_y \end{bmatrix}$ is the real earth magnetic field signal, $\begin{bmatrix} \gamma_x \\ \gamma_y \end{bmatrix}$ is the interference caused by the stepper motors and other external environmental factors, which also vary based on the specific location. We consider these factors would be constant during the entire printing process.

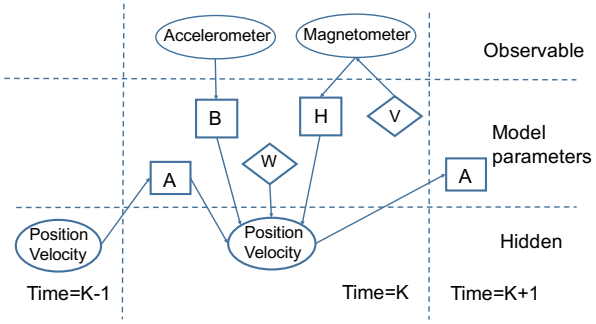

Fig. 12. Kalman filter model for our trace detection.

To obtain the real earth magnetic field signal, a more accurate calibration for the magnetometer reading is needed. However, because the magnetic interference is non-linear, we decide to equally divide the entire printing platform into M^2 regions to obtain the approximation. For each region, we apply an 3×3 affine transformation matrix $M_{transfer}$ [35] for mapping the relationship between $\begin{bmatrix}\alpha_x \\ \alpha_y\end{bmatrix}$ and $\begin{bmatrix}\beta_x \\ \beta_y\end{bmatrix}$, as shown below:

$$\begin{bmatrix}\beta_{x1} & \beta_{y1} \\ \beta_{x2} & \beta_{y2} \\ \beta_{x3} & \beta_{y3}\end{bmatrix} = M_{transfer} \begin{bmatrix}\alpha_{x1} & \alpha_{y1} \\ \alpha_{x2} & \alpha_{y2} \\ \alpha_{x3} & \alpha_{y3}\end{bmatrix} \tag{8}$$

$$M_{transfer} = \begin{bmatrix}1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1\end{bmatrix} \begin{bmatrix}S_x & 0 & 0 \\ 0 & S_y & 0 \\ 0 & 0 & 1\end{bmatrix} \begin{bmatrix}1 & sh_y & 0 \\ sh_x & 1 & 0 \\ 0 & 0 & 1\end{bmatrix} \begin{bmatrix}cos(q) & sin(q) & 0 \\ -sin(q) & cos(q) & 0 \\ 0 & 0 & 1\end{bmatrix} \tag{9}$$

where $\{(\beta_{x1}, \beta_{y1}), (\beta_{x2}, \beta_{y2}), (\beta_{x3}, \beta_{y3})\}$ and $\{(\alpha_{x1}, \alpha_{y1}), (\alpha_{x2}, \alpha_{y2}), (\alpha_{x3}, \alpha_{y3})\}$ are the corresponding three non-linear points on the ground truth path and the measured path to define the transformation matrix. S_x and S_y specify the scale factors along the X and Y axes. sh_x and sh_y denote the shear factors along the X and Y axes. t_x and t_y indicate the displacement along the X and Y axes. q represents the rotation angle.

For a M^2 segmented calibration, the complexity is approximately $O(M^2n^3)$ by using the Gaussian Elimination [53], where $n = 3$. The calibration performance is shown in Figure 11. It is observed that, as M increases, the calibrated path is more close to the ground truth path. When $M = 2$, the calibration can achieve a high accuracy and remain a relative low computation complicity. The evaluation results for the infill path prediction in Section 5 are all based on segmentation for $M = 2$.

4.2.2 Kalman Filter. Because the sensor noises from the accelerometer cannot be completely eliminated, and the high precision requirement for 3D printing, the accelerometer data cannot be directly used to predict the infill path. On the other hand, although the magnetometer data has been calibrated, there still are a rather high level of noises from the sensor itself. Thus, the Kalman Filter [36, 55] is adopted in this study, which produces the estimated location of current state based on the estimates and covariance matrix from the last state and the current observation (the calibrated magnetic field signal). The basic model is shown in Fig. 12.

We develop an extruder tracking algorithm based on the fusion model of the three sensors. First, according to the movement and vibration of the build bed in the vertical dimension, we calculate the time interval of the

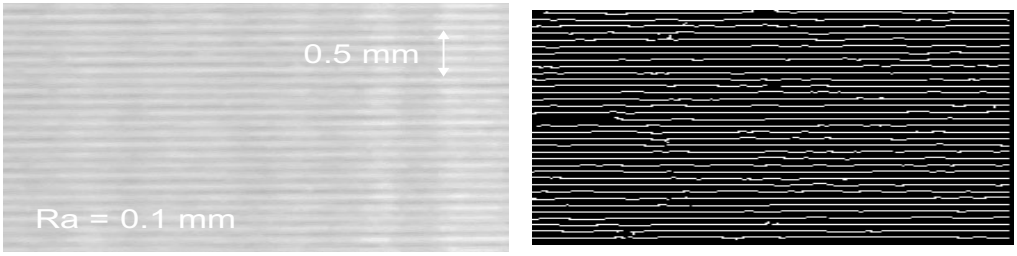

(a) (b)

Fig. 13. (a) The cropped 0.1 mm raw gray-scale image, (b) image processed by the Canny filter

extruder's movement between each layer for segmentation. Second, we use the extruder's magnetic field intensity as the observation, and its acceleration as the control-input vector. (Please refer to Algorithm 1).

Algorithm 1: Extruder Tracking Algorithm

Input: T: sampling period of accelerometer and magnetometer;
z_k: extruder magnetometer data in time series;
u_k: extruder accelerometer data in time series;
b_k: build bed accelerometer data in time series;
w_k: process white noise; v_k: sensor white noise;
Output: v_k: estimated extruder moving speed;
p_k: estimated extruder moving path
$A \leftarrow T$; $B \leftarrow T$; H= [1 0] // Get transition matrix;
$Q_k \leftarrow w_k$; $R_k \leftarrow v_k$ // Get noise covariance;
if $b_k > ThresH_{layer}$ **then**
  Record time k;
end
Segment z_k and u_k for each layer based on k;
foreach *layer* **do**
  $x_0 = 0$; $P_0 = 0$ //initiate ;
  foreach k **do**
    y_k = z_k - H $*$ x_k // compute the innovation vector;
    S_k = H $*$ P_k $*$ H^T + Q_k // compute the covariance of innovation;
    K_k = P_k $*$ H^T $*$ S_k^{-1} // compute the Kalman gain;
    x_k = x_k + K_k $*$ y_k // update the state estimate;
    P_k = (I - K_k $*$ H) P_k// update the error covariance ;
    x_{k+1} = A $*$ x_k + B $*$ u_k+w_k //predict the next state ;
    P_{k+1} = A_k $*$ P_k $*$ A_k^T + Q_k //predict the next error covariance ;
  end
end

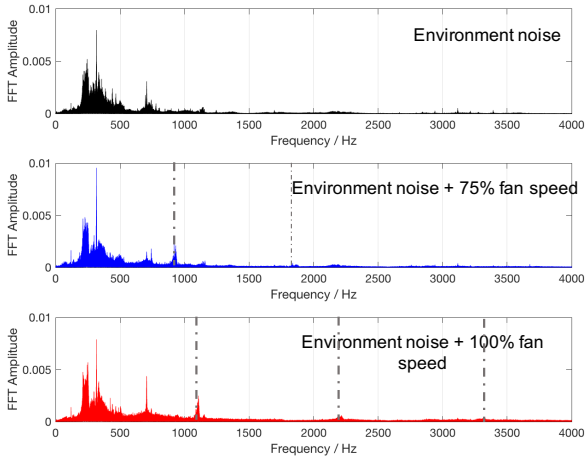

Fig. 14. Noise spectrum in the frequency domain over different fan speeds and the corresponding BPFs.

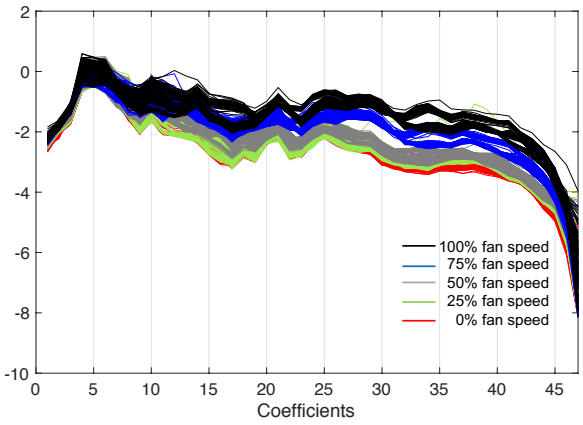

Fig. 15. Extracted coefficients for different fan speeds.

4.3 Layer Thickness Analysis

4.3.1 Gray-scale Pictures. To better evaluate the surface texture and roughness of 3D printed objects and obtain the highest image resolution, the position and focus of the camera have been optimally adjusted. A set of gray-scale images are extracted from the video.

4.3.2 Canny Filter. As shown in Figure 13(a), the edges of layers in the raw gray-scale image becomes more blurry when the layer thickness R_a is 0.1 mm. To precisely extract the vague layer edges, we apply the Canny Filter which is a robust edge detection algorithm [12]. Different from other edge detection methods, the Canny Filter uses two thresholds to determine the strong and the weak edges. The weak edge can be saved if it is connected to some strong edges. Otherwise, it will be filtered out. Due to this unique feature extraction process, the Canny Filter can not only filter out noises but also detect the true weak edges. Figure 13(b) shows the detected layer edge by the Canny Filter. The feature of layer edges and number of pixels between each layer are extracted.

4.4 Fan Speed Estimation

4.4.1 Acoustic Channel Analysis. We compare the spectra for the environmental noises (including noises caused by printing and background), environmental noises plus 75% fan speed, and environmental noises plus 100% fan speed in the frequency domain using FFT (Fast Fourier Transform). In the machine vibration spectra, the interval between each high power peak caused by the fan rotation is called the BPF (Blade Passing Frequency) which can be defined as below:

$$BPF = \frac{N * speed}{60} \tag{10}$$

where N denotes the number of blades of the cooling fan, *speed* means the fan's rotation speed (rpm). It can be seen from Figure 14 that, the intervals of peaks in the noise spectrum fit the simulated BPFs based on the equation for both 75% and 100% fan speed settings. However, even the noise spectrum contains the information of BPFs, which are mixed within the unpredictable environmental noise, we can not directly extract BPFs and calculate the fan speed. Therefore, we adopt a feature extraction method — MFCC (Mel-frequency cepstral coefficients) — to reduce the noise and meanwhile remain the important information of the BPFs.

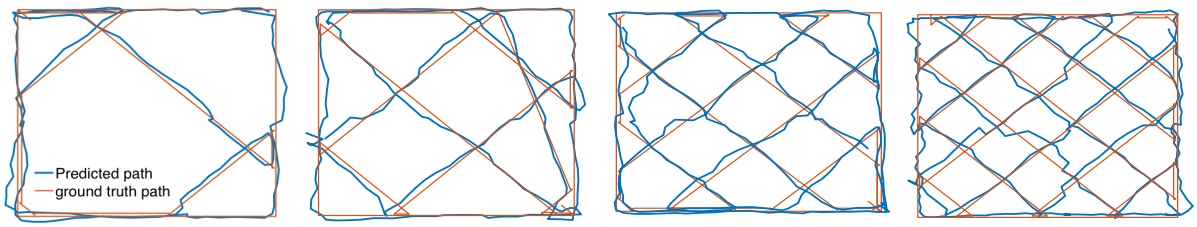

(a) infill rate 5% (b) infill rate 10% (c) infill rate 15% (d) infill rate 20%

Fig. 16. The single-layer predicted paths for a square shape with 5%, 10%, 15% and 20% infill rate respectively.

4.4.2 MFCC. To enhance the high-frequency component, the raw acoustic information will pass through a high-pass filter, and then calculate the frequency spectrum. Different from the original Mel filter bank which is a set of Mel-scaled triangular filters extracting low-frequency information caused by the vocal cords and lips [18, 32], we customize a new filter bank with uniform scale that can extract information evenly distributed over the entire frequency domain. After feature reduction and extraction, the coefficients will serve as the features through a *log* function, as shown in Fig 15, different fan speed settings thus can be discriminated.

4.4.3 Random Forest. It is also observed that the smaller coefficients (e.g., <10) which represent the low-frequency information contributed less to discrimination of the fan speeds. Considering the different feature importance for the fan speed classification, we apply the random forest as the classifier that can filter out the less important features by randomly training multiple weak learners.

5 EVALUATION

In this section, we first analyze the performance of the detection accuracy for four printing attributes respectively and evaluate the system's resistance against the real-world attacks. As a ubiquitous solution for protecting generic consumer-grade FDM 3D printers, the universality of our proposed approach is evaluated and tested on two different types of 3D printers.

5.1 Experimental Setting

As shown in Figure 9, we employ the Ultimaker 2 Go, one of the most widely used desktop FDM 3D printers, as our experimental testbed with the open-source Marlin firmware. Arduino platform is applied to root and modify the firmware. A Quarter-sized versatile LSM9DS1 IMU sensor is equipped on the extruder to record the motion information. It houses a 3-axis accelerometer with a ±16 G scale range and a 3-axis magnetometer with a ± 16 Gauss scale range. We also use the ADXL335 accelerometer attached on the build bed to monitor its movement and vibration along the Z-axis. In addition, the Logitech Pro 960 webcam is employed to monitor the surface texture and roughness of the 3D printed objects with 1080P resolution and 30 Frame Per Second (FPS), and meanwhile record the sound with 44,100 Hz sampling rate.

To collect all the motion information at the same time, We use the Arduino Mega 2560 micro-controller board with a 10-bit analog-to-digital converter. LSM9DS1 is connected to the Arduino Mega 2560 with Inter-Integrated Circuit (I^2C) protocol. ADXL335 is communicated with the Arduino Mega 2560 via three analog input pins. During the monitoring process, the acceleration and magnetic field data are recorded at 50 Hz sampling frequency to ensure a high quality signal with an appropriate computational load.

5.2 Prediction Accuracy

Algorithm 2: DTCDF Evaluation Algorithm

Input: $\{GX_i, Gy_i\}$: ground truth points in X and Y axis;
$\{EX_j, EY_j\}$: Estimated path points in X and Y axis;
DT_{max}: model's dimension
Output: $\{DT_d, F_d\}$: distribution function under different error distances ;
foreach $\{GX_i, Gy_i\}$ **do**
  foreach $\{EX_j, Ey_j\}$ **do**
    $distance_j = \sqrt[2]{(GX_i - EX_j)^2 + (GY_i - EY_j)^2}$ // calculate the euclidean distance;
  end
  $distance_i \leftarrow Min(distance_j)$ // Obtain the minimum distance ;
end
foreach $DT = 0 \rightarrow DT_{max}$ **do**
  count = 0 ;
  foreach $distance_i$ **do**
    if $distance_i < DT$ **then**
      count ++ ;
    end
  end
  F_d = count / i ;
end

5.2.1 Infill Path. To evaluate the prediction accuracy of our method, we select a representative and realistic design model with different infill rates for each layer from 5% to 20%, which covers the normal range for different stiffness requirements. As shown in Figure 16, the predicted paths can match and overlap with the ground truth paths within a small error range.

As the extruder's motion is inconsistently accelerated, the collected data points in the predicted path is not uniformly distributed. Thus, the traditional curve or image similarity metrics cannot appropriately reflect the geometry errors and infill path attacks. We propose the Distance Tolerance Cumulative Distribution Function (DTCDF), based on the Hausdorff Distance which describes the largest distance between sample points in two curves [27]. The definition of DTCDF is formulated as:

$$\begin{aligned} \mathbb{E} &= \{(EX_1, EY_1), ..., (EX_M, EY_M)\} \\ DT(\{GX_i, GY_i\}) &= Hausdorff(\{GX_i, GY_i\}, \mathbb{E}) \\ F_{DT}(d) &= P(DT <= d) \end{aligned} \tag{11}$$

where DT is the distance tolerance, and $d \in [DT_{min}, DT_{max}]$. *Hausdorff* is the Hausdorff Distance function, where N is the number of sample points in each layer for the ground truth path, and M is the number of sample points in each layer for the predicted path. $\{GX_i, GY_i\}$ is the sampling points in the ground truth path and $\mathbb{E}$ is the set of sampling points $\{EX_i, EY_i\}$ in the estimated path. P is the probability function. The detailed DTCDF calculation algorithm can be found in Algorithm 2. The average Hausdorff Distance of 50 layers for difference infill rates are shown in Table 2.

5.2.2 Printing Speed. Printing speed attack is another major type of cyber-physical attacks. By increasing the infill layer's printing speed, it would cause the filament to drag or slip and thus weaken the inner structure without changing the infill paths and outside shells. To evaluate the accuracy for speed prediction, we apply 4 test

Table 2. Prediction Accuracy of Infill Path at Different Infill Rates

Error	5%	10%	15%	20%
Hausdorff Distance (mm)	0.735	0.807	0.856	0.869

Table 3. Prediction Accuracy of Infill Path at Different Printing Speeds

Error	30 mm/s	60 mm/s	90 mm/s	120 mm/s
MAPE (%)	15.16%	4.24%	2.58%	2.05%
MAD (mm/s)	4.548	2.54	3.93	2.46
STD	2.81	1.39	1.54	2.24

Table 4. Prediction Accuracy of Layer Thickness

Error	0.05 mm	0.10 mm	0.15 mm	0.20 mm	0.25 mm
MAPE (%)	15.49%	4.84%	2.94%	1.65%	1.54%
MAD (μm)	7.75	4.84	4.41	3.30	3.85
STD	0.989	0.485	0.466	0.461	0.459

cube models (each has 50 layers) with different infill printing speeds, including 30 mm/s, 60 mm/s, 90 mm/s, and 120 mm/s. For the same nozzle size and filament flow rate, the higher speed means the lower printing quality and less inner support for product's shells. The corresponding prediction accuracy (represented in terms of various error metrics), including the Mean Absolute Percentage Error (MAPE), Mean Absolute Deviation (MAD), and Standard Deviation (STD), is listed in Table 3. In the table, we can observe that the largest error is under low printing speed at 30 mm/s. This is under our expectation, because the step motor would have larger vibration or called resonance during the low rotation frequency, and those introduced vibration would increase the system noise to the IMU sensor.

$$MAPE = \left[\frac{1}{N}\sum_{i=1}^{N}\left|\frac{x_i - \hat{x_i}}{x_i}\right|\right] \times 100 \tag{12}$$

$$MAD = \frac{1}{N}\sum_{i=1}^{N}|x_i - \hat{x_i}| \tag{13}$$

$$STD = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\varepsilon_i - \bar{\varepsilon_i})^2}, \quad \varepsilon_i = |x_i - \hat{x_i}| \tag{14}$$

where x_i is the actual value, $\hat{x_i}$ is the predicted value, ε_i is the error between the actual value and the predicted value, $\bar{\varepsilon_i}$ is the mean of errors, and N is the number of sample points in the predicted path.

5.2.3 Layer Thickness. A 6th-order polynomial curve fitting (see Figure 17) is used to describe the relationship between the number of pixels between two neighboring edges and the layer thickness. The average number of pixels with different layer thickness is calculated based on over 500 layers extracted from 10 individual models. Based upon the measured number of pixels, the estimated layer thickness can be obtained as follows:

$$R_e = P_1{N_p}^6 + P_2{N_p}^5 + P_3{N_p}^4 + P_4{N_p}^3 + P_5{N_p}^2 + P_6{N_p}^1 + P_7 \tag{15}$$

where R_e is the estimated layer thickness, N_p is the number of pixels between two neighboring edges, P_1=1.56×10^{-9}, P_2=-1.84×10^{-7}, P_3=9.17×10^{-6}, P_4=-2.24×10^{-4}, P_5=0.0025, P_6=0, P_7=0. The errors for different layer thickness are

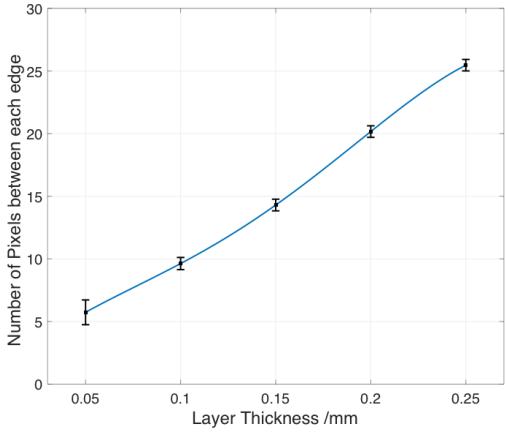

Fig. 17. Curve regression with the STD error bar for the relationship between the average number of pixels between two neighboring edges and the ground truth layer thickness.

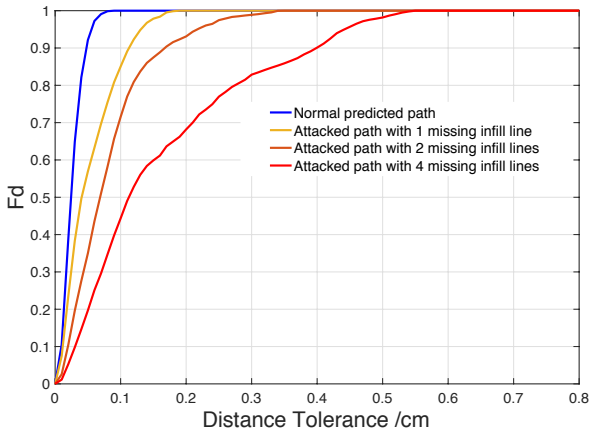

Fig. 18. Average DTCDF curves for 20 normal layers and 20 attacked layers.

listed in Table 4. From the table, we can observe that, due to the resolution limit of the deployed camera, the prediction accuracy drops significantly when the layer thickness increases from 0.05 mm to 0.25 mm.

5.2.4 Fan Speed. To evaluate the performance of estimating fan speed based on the recored acoustic information, especially focusing on model-free and real-time, we record the sound for five different fan speed settings, and five different models (10 layers for each model) for each setting, and repeat 5 times for each model. Thus, after segmentation by each layer, we randomly selected 1000 samples from all five models for training and 250 samples for testing. Each sample will hold 47 features. 50 weak learners were trained for the random forest, the overall identification accuracy is 96.8%.

5.3 Attack Assessment

This section mainly discusses the performance of detecting the firmware attacks on four different printing attributes based on the ground truth information from the G-code file.

5.3.1 Infill Path. To simulate the infill path attacks, we attack the motion control part of firmware in the 3D printer to modify the extruder's target X, Y coordinates received from the G-code file. By skipping some coordinates in the infill layers, the void cavity is formed within the infill patterns. The predicted printing paths for the normal layer and the attacked layer are shown in Figure 19.

We calculate the DTCDF for each individual layer between the predicted paths and the designed paths from G-code to identify the victim layer under attack. Figure 18 describes the DTCDF of the predicted path for normal and attacked layers. It can be observed that, for the normal layers, the Hausdorff Distance is around 0.8 mm, only 4% of the design dimension, which is much smaller than the attacked layers. In addition, as the number of missing infill lines increases, the Hausdorff Distances become larger and the slop of the CDF curves also become more gradual.

5.3.2 Printing Speed. To evaluate the proposed detection mechanism for printing speed attacks, we apply a 20 mm × 20 mm square model with 100 layers and 25% infill rate as our testing target. We set the initial printing speed at 30 mm/s as a relative low values to have a better comparison with the scenario under attack. From the

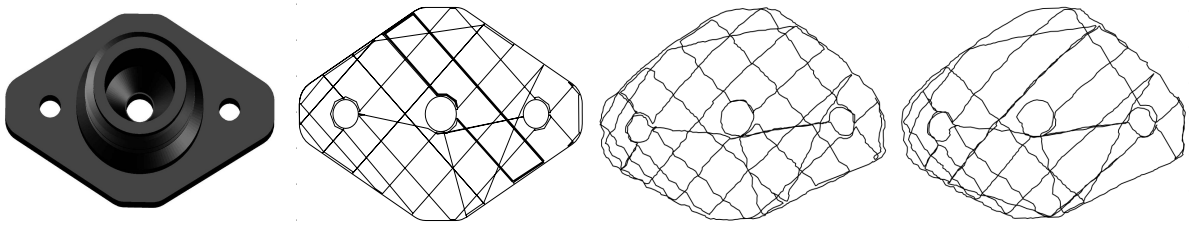

(a) Thread base model (b) Ground truth printing path (c) Estimated printing path for the normal layer (d) Estimated printing path for the attacked layer

Fig. 19. The single-layer predicted paths for a thread-base model. Bold line in (b) indicates the location of missing paths in the attacked layer.

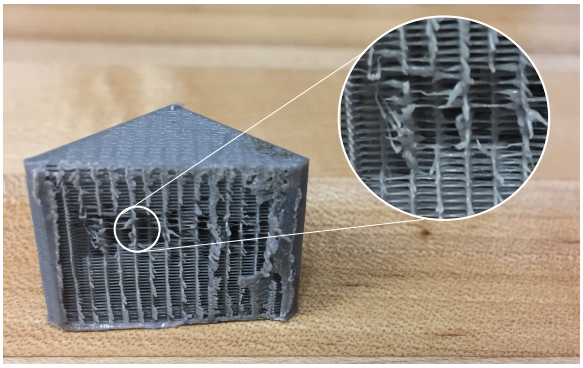

Fig. 20. Cross-sectional view of the attacked model with abnormal infill layer printing speed. The middle layers are clearly under-extrusion and almost form a cavity.

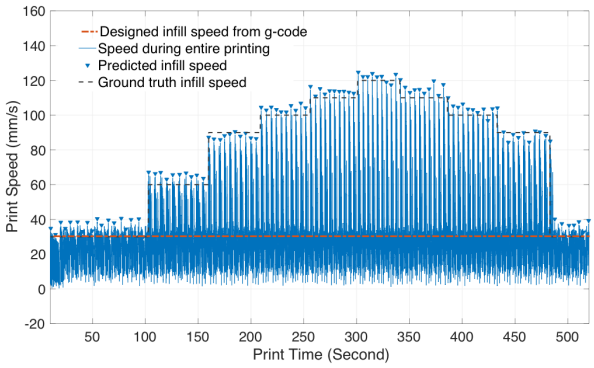

Fig. 21. Prediction of speed curve during the printing process. Y-axis denotes the square root of printing speed.

attacker's perspective, to make the attack in a more imperceptible way, the nozzle moving speed is gradually changed from 30 mm/s to 120 mm/s. According to the inner structure of the 3D printed object as shown in Figure 20, it can be observed that, the first (bottom) 40 layers are finely printed, while the middle 20 layers are under-extruded and disrupted, which will definitely result in a significantly weakened support for the top and side shells.

The predicted speed curve is displayed in Figure 21. The maximum speed points for infill layers generally fit the ground truth speeds. According to the prediction, it can also be seen that, there are abnormal high speed layers around above 110 mm/s from 250 second to 375 second, which match the damaged layers shown in Figure 20. In addition, between each abnormal peak, the speed remains at a low level, which indicates that the shells in the attacked layers are still printed at regular speed. For the overall prediction accuracy of infill layer printing speed, the MAPE is 3.12%, and the normalized mean square error (NMSE) is 3.057.

5.3.3 Layer Thickness. In this subsection, we evaluate the prediction accuracy of surface texture and roughness by randomly locating the attacked layers in a 49-layer 3D printed design. The original thickness is set to 0.1 mm, and we intentionally modify the Z-axis stepper motor's movement to 0.2 mm/step for a few randomly selected layers. The gray-scale images for the normal printed surface and the attacked printed surface captured by the camera are shown in Figure 24, which shows clear, observable differences in the texture patterns. Leveraging the Canny edge detector and the proposed layer thickness fitting model, Figure 22 presents the prediction of the

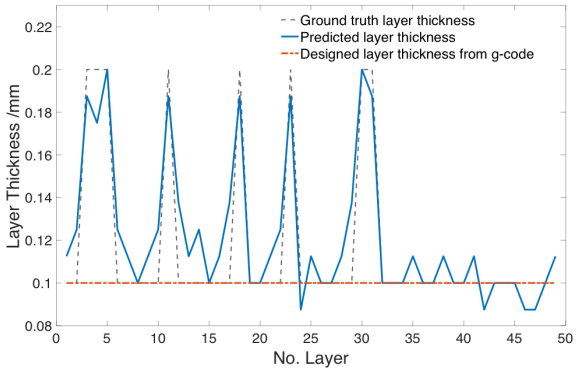

Fig. 22. Demonstration of predicted layer thickness among 49 layers based on attacked mode.

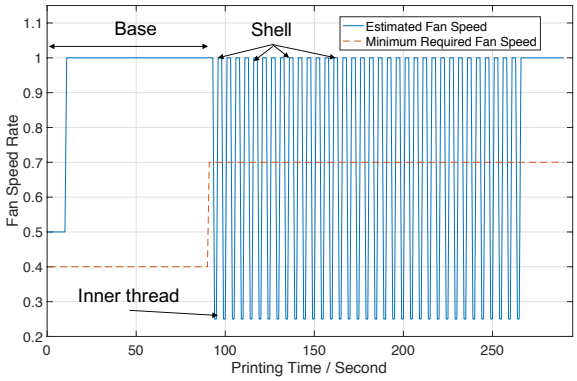

Fig. 23. Demonstration of predicted fan speed rate among attacked inner-thread model over 290 seconds.

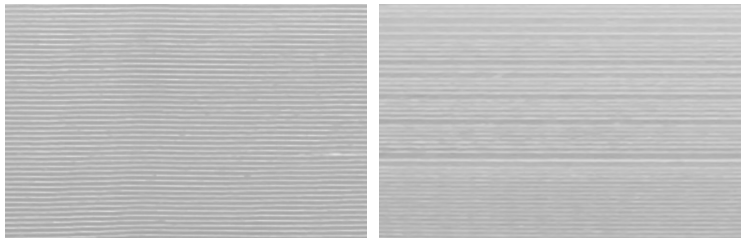

(a) (b)

Fig. 24. (a) Original model with 0.1mm layer thickness surface roughness. (b) attacked model with random distributed 0.2mm layers surface thickness

thickness for each printed layer, compared with the original and the adversely altered layer thickness. In general, the predicted layer thickness matches the ground truth across the entire 49 layers , with an MAPE of 9.57% and an MAD of 10.5 μm.

5.3.4 Fan Speed. To get a better presentation of hidden attack for modifying fan speed, we choose a classic inner-thread model with dimension of 15 mm X 15 mm X 9.75 mm. By only modifying the fan speed setting from the normal 100% to 25 % in the inner thread layers, the filament stacking for the inner thread layers are no longer smooth, however, the attacked printed model is hard to detected from the outside visual inspection, as shown in Figure 25. In our fan speed prediction solution, we record the acoustic file during the entire printing process, and to ensure a good detecting resolution, we segment and predict the fan speed based on the acoustic channel information by every second. As shown in Figure 23, starting from 90 second, the estimate fan speed dropped from 100% to 25% periodically, and we can assess the attacking behavior based on the comparison between the estimated fan speed and the minimum require fan speed which is calculated by the printing time for the corresponding layer.

5.4 Ubiquitous Monitoring

To better demonstrate that the proposed monitoring system can be applied on the generic consumer-grade FDM 3D printers as a ubiquitous approach, in addition to the Ultimaker 2 printer, we also deployed our system on another popular desktop FDM 3D printer — Lulzbot Mini (as shown in Figure 26). The sensor configuration and

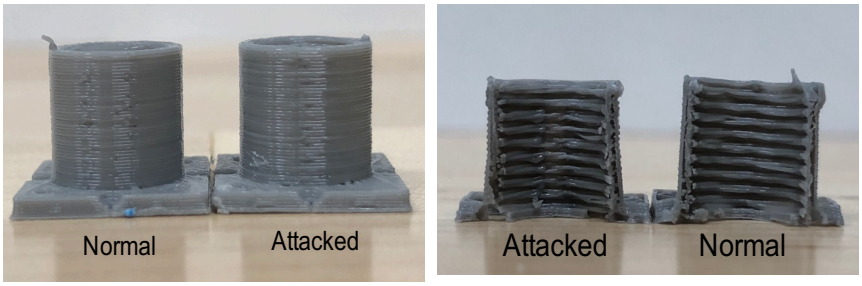

(a) Shell comparison (b) Inner thread comparison

Fig. 25. Hidden fan speed attack for a inner-thread model.

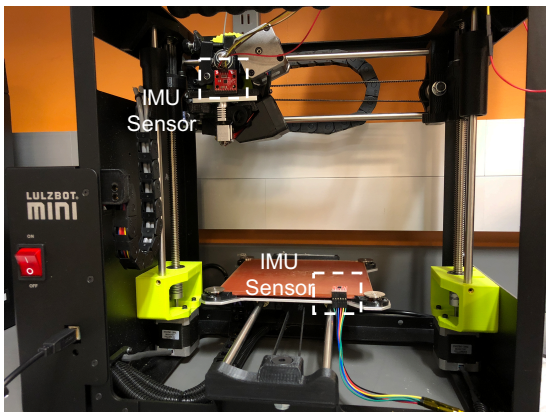

Fig. 26. Sensor deployment on Lulzbot 3D printer

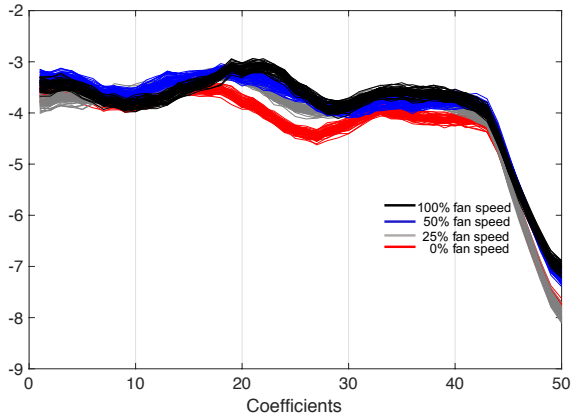

Fig. 27. Extracted coefficients for different fan speeds on Lulzbot 3D printer

deployment are quite similar to the case on the Ultimaker 2, including an IMU sensor attached to the filament extruder, an IMU sensor fixed on the printing bed, and a camera. We then evaluated the two major attributes for hidden attacks (i.e., infill path and fan speed).

5.4.1 Fan Speed. To evaluate the fan speed estimation on the new printer, we recorded 50 sound files for multiple printing tasks under different fan speed settings (40 files for the training and the rest 10 file for testing). Each recording file lasts for one minute. We extracted the MFCC features for each different fan speed as shown in Figure 27, and it can be clearly observed that those coefficients can be effectively differentiated. We can also achieve 100% identification accuracy using the random forest classifier.

5.4.2 Infill Path. Same as the evaluation on Ultimaker 2, we also used the 20mm cube with grid infill patterns as the test-bench model. As shown in Figure 28, we evaluated the errors between the estimated path and the ground truth path under different infill rates and printing speeds (the default infill rate is 20% and printing speed is 60 mm/s). It can be observed that, although the two printers adopt different slicing software and extruder design, the estimation errors are still at the similar level. Those results of two major printing attributes show that our proposed system can also achieve good estimation on the different type of printers, which could be considered as a ubiquitous monitoring system among consumer-level FDM printers.

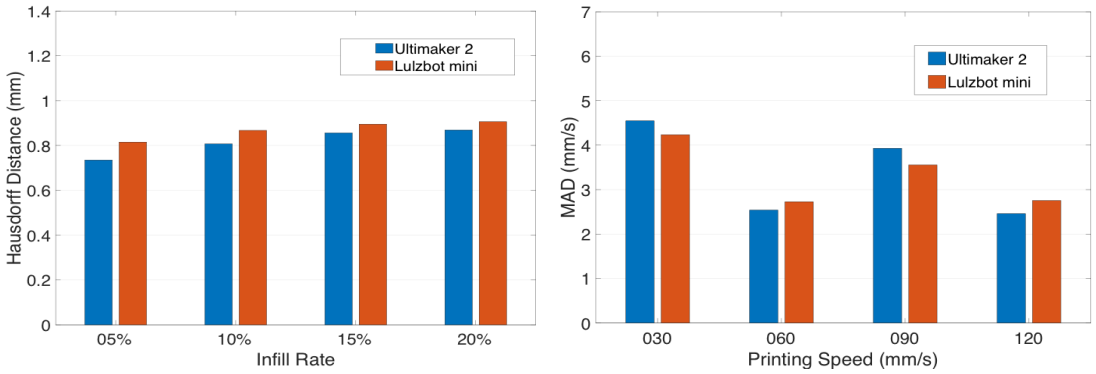

(a) Errors under different infill rates

(b) Errors under different printing speeds

Fig. 28. Comparison of infill path errors between two different types of 3D printers

6 DISCUSSIONS

In this section, we discuss the limitations of the proposed defense mechanism and the future work.

Trivial variance evaluation. For 3D printing and especially for consumer-grade 3D printers, it is unavoidable that the printed objects may come with some invisible or graceful variances, resulted from the mechanical and thermal variances, the process variations, the imperfection of printing filament, or the user's inexperienced operations. In this case, the modifications on 3D printed objects can be categorized into two groups: significant distortions (which may be resulted from malicious attacks or mechanical defects, either of which is not acceptable from the quality perspective) and trivial variances (which may be resulted from normal process variations within the acceptable quality tolerance). As for the significant distortions that may damage the object's functionality and geometry, as shown in the evaluation part in this paper, we can achieve 100% detection rate. In contrast, the trivial variances which might slightly deviate the original parameter settings within the acceptable tolerance level will not threat or sacrifice the quality of the product (i.e., requirements of surface toughness, stress, or color variances). Thus in this paper, to better show the potential threats and damages resulted from those attacks, we only consider those harmful significant distortions (no matter for whatever type of printed objects) with 100% correctly identification rate, and 0% false positive and false negative rates. In the future work, to better assess the impacts of the attacks, especially for those trivial modifications, we need to investigate the relationship between the degree of modifications and the actual changes of the physical and geometry characteristics, which is definitely specific to the 3D design model.

Travel movement. Travel movement means the extruder moves between two points without printing task. This non-printing movement is determined by the slicing software to optimize the printing path. Usually, the travel movement starts from the end point of the previous layer to the start point of the next layer, and the two points are designed to overlap or locate closely. In our study, the travel movements in the test models are short and negligible compared with the prediction results. But for some complicated or special designs, a motion sensor (e.g., infrared camera) is needed to detect the extrusion length.

Magnetic field interference. In our magnetic field calibration stage, we assume that the magnetic interference stays constant during the entire printing process, or the variation can be negligible. Given the fact that the magnetometer is highly sensitive to interfering local magnetic fields and distortions, the reading would be affected if there are electric devices proximate to the 3D printer or metallic objects passing through the generated field. To address this issue, we plan to develop a dynamic magnetic calibration method. By setting up some fixed test points of magnetometers around the printer, we can sense the changes of the magnetic field intensity and adjust the transformation matrix in the real-time manner.

Location of infill modification. As discussed in Section 3, the inner structure of 3D printed objects is very important for printing quality, especially for some sensitive products. Another potential infill layer attack is to slightly change the infill path that is located near the crucial load-bearing structures (e.g., corners, notches, ribs, and gussets), which is more imperceptible and difficult to be detected. A trivial modification of load-bearing components will significantly affect the mechanical and physical properties of the 3D printed objects. In the future work, we will evaluate the infill layer attacks based on, not only the extent and amount of infill paths that are adversely modified, but also the specific location of the modified infill paths according to their level of sensitivity and criticalness to the structural integrity and rigidity of 3D printed design.

Universality of system deployment and calibration. As discussed in the Section 1, the main objective of this study is to provide personal users with an effective and efficient way to protect their designs away from attacks. The sensors in this work are light-weight, non-intrusive, and with small form factor (2 cm by 2 cm). Generally, for most of the desktop FMD 3D printers on the market, the sensor deployment would be similar. The IMU sensors should be attached to the printing extruder's shell, and the camera should be set near the printing bed (approximate 10 cm in our setting). Because the magnetic sensor is sensitive to environmental noises, including the earth and surrounding electrics. The users would have to calibrate the sensor by themselves through the following steps: Firstly, the user needs to print several simple objects with basic shapes (i.e., square and triangular) and record the corresponding magnetic field intensity. Then, our algorithm would generate an update transformation matrix based on the printed path of basic shapes and the recorded magnetic paths. This calibration matrix would store the local magnetic interference that can help to reduce and mitigate the path predication errors in the future printing.

Potential deployment of smartphones. The discussed sensors including IMU, camera, and audio recorder are all available in most existing smartphones. This work lays the foundations and provides preliminary supports for using the most affordable and accessible smart devices such as smartphones to monitor and reconstruct the crucial printing parameters. In the future study, we would like to explore the use of built-in smartphone sensors and develop the app to realize the defense mechanism, which would be much easier and more convenient for personal users to calibrate and to use.

7 CONCLUSION

3D printing has demonstrated superior advantages over the traditional manufacturing, and has been widely adopted in many industries. However, it also brings some unknown risks especially from the perspective of security vulnerabilities. In this paper, we first explore the potential cyber-physical attacks on the 3D printing process chain. A real-time, online process monitoring approach is proposed to defend against the cyber-physical attacks on 3D printers and particularly the adverse modifications of critical printing attributes specified by the firmware. By analyzing the data collected from a set of sensory devices including the accelerometer, magnetometer, and video camera, our method can continuously monitor and verify four most significant printing attributes that could be maliciously modified by an adversary. According to our experimental results, the infill path, printing speed, layer thickness, and cooling fan speed are reconstructed and estimated with a high accuracy, and the abnormal printing behaviors can be properly detected. We expect that our work can lay a foundation for the exploration of effective defense solutions for increasingly popular 3D printing.

REFERENCES

[1] O. Akyol and Z. Duran. 2014. Low-cost laser scanning system design. *Journal of Russian Laser Research* 35, 3 (May 2014), 244–251.
[2] S. Amin, X. Litrico, S. Sastry, and A.M. Bayen. 2013. Cyber security of Water SCADA systems — Part I analysis and experimentation of stealthy deception attacks. *IEEE Transactions on Control Systems Technology* 21, 5 (2013), 1963–1970.
[3] G. C. Anzalone, C. Zhang, B. Wijnen, P. G. Sanders, and J. M. Pearce. 2013. A Low-cost open-source metal 3-D printer. *IEEE Access* 1 (Dec. 2013), 803–810.

[4] M. Baker and J. Manweiler. 2014. From 3D printing to spy cats. *IEEE Pervasive Computing* 13, 4 (2014), 6–9.
[5] Rafael Ballagas, Sarthak Ghosh, and James Landay. 2018. The design space of 3D printable interactivity. *Proceedings of ACM Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2, Article 61 (July 2018), 21 pages. DOI: http://dx.doi.org/10.1145/3214264
[6] C. Bayens, T. Le, L. Garcia, R. Beyah, M. Javanmard, and S. Zonouz. 2017. See no evil, hear no evil, feel no evil, print no evil? Malicious fill patterns detection in additive manufacturing. In *Proceedings of the 26th USENIX Security Symposium*. 1181–1198.
[7] Abdelkareem Bedri, Richard Li, Malcolm Haynes, Raj Prateek Kosaraju, Ishaan Grover, Temiloluwa Prioleau, Min Yan Beh, Mayank Goel, Thad Starner, and Gregory Abowd. 2017. EarBit: Using wearable sensors to detect eating episodes in unconstrained environments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 37.
[8] S. Belikovetsky, M. Yampolskiy, J. Toh, and Y. Elovici. 2016. dr0wned-cyber-physical attack with additive manufacturing. *arXiv preprint: 1609.00133* (2016).
[9] J. T. Belter and A. M. Dollar. 2015. Strengthening of 3D printed fused deposition manufactured parts using the fill compositing technique. *PLOS One* (2015), 1–19. DOI: http://dx.doi.org/10.1371/journal.pone.0122915
[10] S. M. Bridges, K. Keiser, N. Sissom, and S. J. Graves. 2015. Cyber security for additive manufacturing. In *Proceedings of the 10th Annual Cyber and Information Security Research Conference (CISR)*. ACM, 1–3.
[11] C. Byung-Chul, L. Seoung-Hyeon, N. Jung-Chan, and L. Jong-Hyouk. 2016. Secure firmware validation and update for consumer devices in home networking. *IEEE Transactions on Consumer Electronics* 62, 1 (2016), 39–44.
[12] J. Canny. 1986. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6 (1986), 679–698.
[13] S. R. Chhetri, A. Canedo, and M. A. Al Faruque. 2016. KCAD: Kinetic cyber attack detection method for cyber-physical additive manufacturing systems. In *Procedings of the 35th International Conference On Computer-Aided Design (ICCAD)*. ACM, 74.
[14] Jiska Classen, Daniel Wegemer, Paul Patras, Tom Spink, and Matthias Hollick. 2018. Anatomy of a vulnerable fitness tracking system: dissecting the fitbit cloud, App, and firmware. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 5.
[15] A. Cui, M. Costello, and S.J. Stolfo. 2013. When firmware modifications attack: a case study of embedded exploitation. In *Proceedings of the 20th Network and Distributed System Security Symposium (NDSS'13)*. 1–13.
[16] Q. Do, B. Martini, and K.K.R. Choo. 2016. A data exfiltration and remote exploitation attack on consumer 3D printers. *IEEE Transactions on Information Forensics and Security* 11, 10 (2016), 2174–2186.
[17] S. K. Everton, M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare. 2016. Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing. *Materials & Design* 95 (2016), 431–445.
[18] Petko Georgiev, Sourav Bhattacharya, Nicholas D Lane, and Cecilia Mascolo. 2017. Low-resource multi-task audio sensing for mobile and embedded devices via shared deep neural network representations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 50.
[19] F. Goldenberg. 2006. Geomagnetic navigation beyond the magnetic compass. In *Proceedings of Position, Location, And Navigation Symposium (PLANS)*. IEEE, 684–694.
[20] B. Gozick, K. P. Subbu, R. Dantu, and T. Maeshiro. 2011. Magnetic maps for indoor navigation. *IEEE Trans. Instrum. Meas.* 60, 12 (2011), 3883–3891.
[21] Andreas Grammenos, Cecilia Mascolo, and Jon Crowcroft. 2018. You are sensing, but are you biased?: A user unaided sensor calibration approach for mobile sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 11.
[22] T. Greene. 2016. *U.S. 3D Printer Forecast, 2016–2020: New 3D Print/Additive Manufacturing Technologies Fuel Growth.* Technical Report US41333516. IDC Research, Inc., Framingham, MA.
[23] M. Gross. 2013. Creating the magic with information technology. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. 1–2.
[24] A. Hojjati, A. Adhikari, K. Struckmann, E. Chou, T.N. Tho Nguyen, K. Madan, M.S. Winslett, C.A. Gunter, and W.P. King. 2016. Leave your phone at the door: Side channels that reveal factory floor secrets. In *Proceedings of the 23rd ACM Conference on Computer and Communications Security (CCS)*. ACM, 883–894.
[25] J. Hong and M. Baker. 2014. 3D Printing, Smart Cities, Robots, and More. *IEEE Pervasive Computing* 13, 1 (2014), 6–9.
[26] J. U. Hou, D. G. Kim, and H. K. Lee. 2017. Blind 3D mesh watermarking for 3D printed model by analyzing layering artifact. *IEEE Transactions on Information Forensics and Security* 12, 11 (Nov. 2017), 2712–2725.
[27] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge. 1993. Comparing images using the Hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15, 9 (1993), 850–863.
[28] R. Jones, P. Haufe, E. Sells, P. Iravani, V. Olliver, C. Palmer, and A. Bowyer. 2011. RepRap - the replicating rapid prototyper. *Robotica* 29, 1 (January 2011), 177–191. DOI: http://dx.doi.org/10.1017/S026357471000069X
[29] J. P. Kruth, M. C. Leu, and T. Nakagawa. 1998. Progress in additive manufacturing and rapid prototyping. *CIRP Annals-Manufacturing Technology* 47, 2 (1998), 525–540.

[30] A. Liptak. 2017. The US Navy 3D printed a concept submersible in four weeks. https://www.theverge.com/2017/7/29/16062608/us-navy-3d-printing-submersible-manufacturing-military. (July 29 2017). [Online; accessed 20-July-2018].
[31] J. Mireles, C. Terrazas, F. Medina, R. Wicker, and E. Paso. 2013. Automatic feedback control in electron beam melting using infrared thermography. In *Proceedings of the Solid Freeform Fabrication Symposium*.
[32] Mark Mirtchouk, Drew Lustig, Alexandra Smith, Ivan Ching, Min Zheng, and Samantha Kleinberg. 2017. Recognizing eating from body-Worn sensors: combining free-living and laboratory data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 85.
[33] S. B. Moore, W. B. Glisson, and M. Yampolskiy. 2017. Implications of malicious 3D printer firmware. In *Proceedings of the 50th Hawaii International Conference on System Sciences*.
[34] S. Mueller. 2018. Toward direct manipulation for personal fabrication. *IEEE Pervasive Computing* 17, 1 (Jan 2018), 75–81. DOI: http://dx.doi.org/10.1109/MPRV.2018.011591064
[35] K. Nomizu and T. Sasaki. 1994. *Affine differential geometry: geometry of affine immersions*. Cambridge University Press.
[36] Kazuya Ohara, Takuya Maekawa, and Yasuyuki Matsushita. 2017. Detecting state changes of indoor everyday objects using Wi-Fi channel state information. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 88.
[37] J. M. Pearce, C. M. Blair, K. J. Laciak, R. Andrews, A. Nosrat, and I. Zelenika-Zovko. 2010. 3D printing of open source appropriate technologies for self-directed sustainable development. *J. Sustain. Development* 3, 4 (2010), 17–29.
[38] P. K. Rao, J. P. Liu, D. Roberson, Z. J. Kong, and C. Williams. 2015. Online real-time quality monitoring in additive manufacturing processes using heterogeneous sensors. *Journal of Manufacturing Science and Engineering* 137, 6 (2015), 061007.
[39] GE Global Research. 2016. 3D printing creates new parts for aircraft engines. http://www.geglobalresearch.com/innovation/3d-printing-creates-new-parts-aircraft-engines/. (2016). [Online; accessed 1-August-2017].
[40] A. Schmidt, T. Döring, and A. Sylvester. 2011. Changing how we make and deliver smart devices: when can I print out my new phone? *IEEE Pervasive Computing* 10, 4 (2011), 6–9.
[41] D. M. Shila, P. Geng, and T. Lovett. 2016. I can detect you: Using intrusion checkers to resist malicious firmware attacks. In *Proceedings of the IEEE Symposium on Technologies for Homeland Security (HST)*. IEEE, 1–6.
[42] Chen Song, Zhengxiong Li, Wenyao Xu, Chi Zhou, Zhanpeng Jin, and Kui Ren. 2018. My smartphone recognizes genuine QR codes!: Practical unclonable QR code via 3D printing. *Proceedings of ACM Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2, Article 83 (July 2018), 20 pages. DOI: http://dx.doi.org/10.1145/3214286
[43] Chen Song, Feng Lin, Zhongjie Ba, Kui Ren, Chi Zhou, and Wenyao Xu. 2016. My smartphone knows what you print: Exploring smartphone-based side-channel attacks against 3D printers. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 895–907.
[44] SpaceX. 2014. SpaceX lauches 3D-printed part to space, creates printed engine chamber. http://www.spacex.com/news/2014/07/31/spacex-launches-3d-printed-part-space-creates-printed-engine-chamber-crewed/. (2014). [Online; accessed 20-July-2018].
[45] J. Straub. 2017. 3D printing cybersecurity: detecting and preventing attacks that seek to weaken a printed object by changing fill level. In *Proceedings of SPIE, Dimensional Optical Metrology and Inspection for Practical Appl. VI*, Vol. 10220. 1–15.
[46] J. Straub. 2017. An approach to detecting deliberately introduced defects and micro-defects in 3D printed objects. In *Proceedings of SPIE, Pattern Recognition and Tracking XXVIII*, Vol. 10203. 1–14.
[47] J. Straub. 2017. A combined system for 3D printing cybersecurity. In *Proceedings of SPIE, Dimensional Optical Metrology and Inspection for Practical Appl. VI*, Vol. 10220. 1–13.
[48] J. Straub. 2017. Identifying positioning-based attacks against 3D printed objects and the 3D printing process. In *Proceedings of SPIE, Pattern Recognition and Tracking XXVIII*, Vol. 10203. 1–13.
[49] J. Straub. 2017. Physical security and cyber security issues and human error prevention for 3D printed objects: detecting the use of an incorrect printing material. In *Proceedings of SPIE, Dimensional Optical Metrology and Inspection for Practical Appl. VI*, Vol. 10220. 1–16.
[50] L. D. Sturm, C. B. Williams, J. A. Camelio, J. White, and R. Parker. 2014. Cyber-physical vulnerabilities in additive manufacturing systems. *Context* 7, 2014 (2014), 951–963.
[51] L. D. Sturm, C. B. Williams, J. A. Camelio, J. White, and R. Parker. 2017. Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the .STL file with human subjects. *Journal of Manufacturing Systems* 44, 1 (2017), 154–164.
[52] H. Turner, J. White, J. A. Camelio, C. Williams, B. Amos, and R. Parker. 2015. Bad parts: Are our manufacturing systems at risk of silent cyberattacks? *IEEE Security & Privacy* 13, 3 (2015), 40–47.
[53] L. G. Valiant. 1979. The complexity of computing the permanent. *Theoretical Computer Science* 8, 2 (1979), 189–201.
[54] H. Vincent, L. Wells, P. Tarazaga, and J. Camelio. 2015. Trojan detection and side-channel analyses for cyber-security in cyber-physical manufacturing systems. *Procedia Manufacturing* 1 (2015), 77–85.
[55] Chuyu Wang, Jian Liu, Yingying Chen, Lei Xie, Hong Bo Liu, and Sanclu Lu. 2018. RF-Kinect: A wearable RFID-based approach towards 3D body movement tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 41.
[56] L. J. Wells, J. A. Camelio, C. B. Williams, and J. White. 2014. Cyber-physical security challenges in manufacturing systems. *Manufacturing Letters* 2, 2 (2014), 74–77.

[57] R. Whited. 2017. *Failure Analysis of 3D Printed Parts.* Technical Report KSC-E-DAA-TN41114. NASA Kennedy Space Center, Cocoa Beach, FL.
[58] T. T. Wohlers and T. Caffrey. 2016. *Wohlers Report 2016: 3D printing and additive manufacturing state of the industry annual worldwide progress report.* Wohlers Associates.
[59] M. Wu, Z. Song, and Y. B. Moon. 2017. Detecting cyber-physical attacks in CyberManufacturing systems with machine learning methods. *Journal of Intelligent Manufacturing* (2017), 1–13.
[60] M. Yampolskiy, T. R. Andel, J. T. McDonald, W. B. Glisson, and A. Yasinsac. 2014. Intellectual property protection in additive layer manufacturing: Requirements for secure outsourcing. In *Proceedings of the 4th Program Protection and Reverse Engineering Workshop*. ACM, 7.
[61] Mark Yampolskiy, Peter Horvath, Xenofon D Koutsoukos, Yuan Xue, and Janos Sztipanovits. 2012. Systematic analysis of cyber-attacks on CPS-evaluating applicability of DFD-based approach. In *Proceedings of the 5th International Symposium on Resilient Control Systems (ISRCS)*. IEEE, 55–62.
[62] M. Yampolskiy, A Skjellum, M Kretzschmar, R. A. Overfeit, K. R. Sloan, and A. Yasinsac. 2016. Using 3D printers as weapons. *International Journal of Critical Infrastructure Protection* 14 (2016), 58–71.
[63] L. Yang, K. Hsu, B. Baughman, D. Godfrey, F. Medina, M. Menon, and S. Wiener. 2017. Additive manufacturing of metals: the technology, materials, design and production. (2017).
[64] S. E. Zeltmann, N. Gupta, N. G. Tsoutsos, M. Maniatakos, J. Rajendran, and R. Karri. 2016. Manufacturing and security challenges in 3D printing. *JOM* 68, 7 (July 2016), 1872–1881.

Received February 2018; revised July 2018; accepted September 2018