Sabotaging Material Extrusion-Based 3D Printed Parts through Low-Magnitude Kinetic Manipulation Attacks

MUHAMMAD HARIS RAIS, Virginia State University, Petersburg, VA, USA
MUHAMMAD AHSAN and IRFAN AHMED, Virginia Commonwealth University, Richmond, VA, USA

The increasing ubiquity of material-extrusion-based additive manufacturing is motivating cybersecurity researchers to explore its offensive and defensive landscape. Being a physical system, 3D printers have non-zero tolerance specifications for precision and trueness parameters. While a single-bit change in a digital data file is sufficient to fail its integrity and is easily detected through methods such as hashing, the printing process (and subsequently the printed object) remains compliant within the tolerance zone. This study systematically analyzes the material extrusion process and identifies four attack opportunities where low-magnitude kinetic cyberattacks exploit the physical process compliance zone to sabotage the printed part's mechanical properties. The attacks are demonstrated on ASTM-compliant tensile and flexure bars through a man-in-the-middle attack scenario by hijacking the network layer communication between the 3D printer and the printer control machine. The physically stealthy attacks did not produce any evident deformation in the parts' dimensions and mass, while the destructive tests confirm that they are still effective in modifying the tensile and bending strength by up to 25%. The effectiveness of the attacks in bypassing the defenses is assessed by implementing one of the leading detection schemes described in the current literature. The attacks were either not detected at all or detected with a significantly high false negative rate at various attack magnitudes.

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

Additional Key Words and Phrases: Extrusion-bonding attacks, Kinetic attacks, Sabotage, Fused Filament Fabrication

ACM Reference format:
Muhammad Haris Rais, Muhammad Ahsan, and Irfan Ahmed. 2025. Sabotaging Material Extrusion-Based 3D Printed Parts through Low-Magnitude Kinetic Manipulation Attacks. *ACM Trans. Cyber-Phys. Syst.* 9, 1, Article 5 (January 2025), 26 pages.
https://doi.org/10.1145/3704735

1 Introduction

Additive manufacturing methods are commonly used in various industrial sectors, including aviation, automobile, and healthcare. Due to its compelling advantages, such as a faster development

M. H. Rais completed this work while a PhD student at Virginia Commonwealth University. He helped revise the manuscript at Virginia State University.

Authors' Contact Information: Muhammad Haris Rais (corresponding author), Virginia State University, Petersburg, VA, USA; e-mail: mrais@vsu.edu; Muhammad Ahsan, Virginia Commonwealth University, Richmond, VA, USA; e-mail: ahsanm5@vcu.edu; Irfan Ahmed, Virginia Commonwealth University, Richmond, VA, USA; e-mail: iahmed3@vcu.edu.

ACM 2378-9638/2025/1-ART5
https://doi.org/10.1145/3704735

cycle, mass customization, and complex object printing capability, Industry 4.0 considers additive manufacturing as an essential component [9]. Material extrusion is the most common of the seven additive manufacturing methods defined in ISO/ASTM standard 52900 [14]. Material extrusion-based **fused filament fabrication (FFF)** technology is anticipated to attract more attackers after the incorporation of metal-infused filaments [7]. In response to the increased incentive to attackers, cybersecurity researchers have proposed various techniques to secure the FFF-based printing process [10, 40]. A lot of research is available in securing cyber-physical systems using conventional methods [1, 4, 17, 18, 27, 41]. In addition to such solutions, researchers have utilized physical domain knowledge to detect attacks and anomalies [2]. Monitoring the physical process through side channels offers better coverage of the cyber-physical process than conventional cyber-domain monitoring. For example, if an attacker hijacks the network communication and manipulates the G-code file sent to the printer, a physical process monitoring solution shall still detect such an attack.

Monitoring the printing process in the physical domain has its own challenges. The performance of a physical process monitoring solution depends on the quality of the sensing equipment, deployment proficiency, algorithmic errors, and environmental factors (such as changing background sound and lighting conditions). It is an active research area, and the literature review shows that the detection horizon is continuously improving. For instance, Rais et al. [34] claimed to reliably detect a 1 mm deviation in the toolpath with zero false positives and false negatives in a set of objects. As the detection horizon improves, it will overlap with the printer specifications tolerance zone. Unlike a digital artifact, where a single-bit change is also not acceptable, a physical process is considered compliant within the tolerance zone. Even if a monitoring scheme is capable of detecting tiny deviations, reducing the anomaly threshold below the printer's trueness value will likely result in a significant increase in false positives.

If a smart attacker keeps the attack magnitude within the tolerance of the printing process, the attack can likely circumvent the threshold-based detection schemes. One may hypothesize that there should be no reason to worry if the process is progressing within the specified green zone. It is not ascertained if these low-magnitude deviations can consistently and negatively influence the printed parts' properties. As the printers are not designed, nor have their specifications been finalized after considering the impact of machine deviations on different printed objects, it is reasonable to doubt the above hypothesis. FFF characteristics also play an important role in the mechanical properties of the printed object. However, all the characteristics may not offer a good opportunity for a stealthy attack. For instance, changing the build orientation, as presented in Figure 1, significantly reduces the tensile strength but completely changes the toolpath (the printing sequence), making it a simple-to-detect attack for the current schemes defined in the related work.

In this study, the authors examine the FFF printing process to identify minimal kinetic manipulation opportunities (within the printer specifications tolerances) targeting the extrudates bonding to degrade the tensile and/or flexure strength of the printed parts. The paper presents four attacks exploiting extrudates bonding at critical locations. The first two attacks relate to inducing bonding weakness within the infill structure. The third attack attempts to weaken the bonding between the infill and the wall structure. The fourth attack exploits the printing bed kinetics to manipulate the interlayer bond. This study uses a **man-in-the-middle (MiTM)** attack vector to inject the proposed attacks into the printing process by hijacking the G-code file in flight.

To evaluate the effectiveness of the proposed attacks, an experiment is designed with attack magnitudes ranging from 0.015 mm to 0.2 mm (around and below the trueness specifications). Tensile and three-point bending tests are conducted for the attacked and non-attacked samples to measure the attacks' impact on the tensile and flexure strengths. The results confirm that planned malicious deviations within the above-mentioned range are sufficient to compromise the

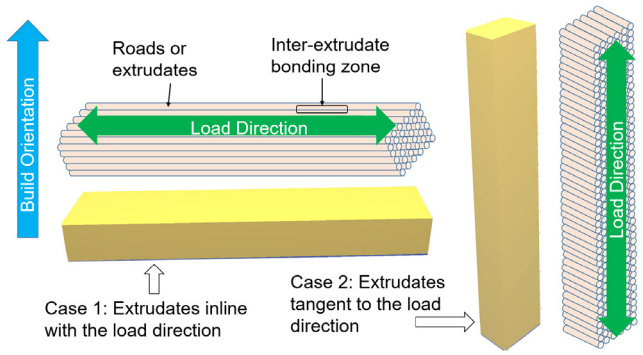

Fig. 1. Impact of FFF characteristics on the mechanical strength explained through an example of changing build orientation.

mechanical strength of the printed parts. Finally, we determine the evasion potential of the attacks by implementing a defensive solution presented in a previous study [34], which incorporates highly refined thresholds for detecting kinetic attacks.

2 Related Work

This section briefly presents the sabotage attacks where the introduced defects are concealed in the final printed object. The researchers have demonstrated sabotage attacks by inducing defects at the designing stage of the printing process [6, 35, 43]. A few other researchers have explored attacks on the G-code instructions (post-designing stage), either through MiTM between the control machine and the printer [33] or by manipulating the slicer memory [20]. The literature also shows limited effort in manipulating the printer's firmware or bootloader to inject defects in the printed parts [28]. Rais et al. [31] proposed two low-magnitude attacks targeting the infill structure to reduce a part's strength. Their paper focused solely on the manipulation of the infill structure. In contrast, our study focuses on the kinetic process to identify all potential attack opportunities, whether they fall within or outside the infill structure. These would qualify as low-magnitude attacks, defined as attacks where the deviation stays within the printer's trueness specifications. Instead of restricting to infill structure, this work systematically examines all sub-structures and inter-substructure bonding instances to identify low-magnitude attack opportunities. Moreover, this study only includes low-magnitude kinetic attacks targeting the bonding between adjacent extrudates, while the cavity attacks and thermodynamic attacks are not in the scope of this work.

As this study aims to find process deviations that can fly below the detection horizon, we discuss the existing attack detection schemes to identify the detection horizon. Chhetri et al. [8] proposed the use of audio emissions to detect kinetic anomalies in the print object. Belikovetsky et al. [5] used the fingerprinting method to authenticate the printed part by generating a master audio profile and using it for the next printed parts. A similar technique was adopted by Gatlin et al. [13] wherein instead of using audio, electric current signals were used to generate a master profile. A deviation beyond the threshold was categorized as anomalous behavior. Gao et al. [12] acquired data through Inertial Measurement Unit sensors and cameras. Using mathematical modeling and image processing, they were able to detect significant geometry distortions due to anomalies in the cooling process. Wu et al. [38, 39] employed static and moving camera techniques to capture and train images on the machine learning algorithm to detect infill deviations in the print geometry. Rais et al. [34] adopted a multi-sensing technique and utilized optical encoders and thermal sensors

to estimate the printing state accurately. Their proposed framework, *Sophos*, transforms G-code instructions through spatiotemporal modeling and compares it with the sensor values. We utilize the above-mentioned studies to estimate attack detection's state of the art.

3 Methodology and the Proposed Attacks

This section first presents the criteria formalized for a successful attack, followed by the existing attack detection horizon identified through literature. Then the precision and trueness values of common FFF printers are reviewed to identify the limits of attack magnitude. Once these constraints are established, compliant attack opportunities in the FFF process are analyzed. Finally, the four proposed attacks are described in the section.

3.1 Defining Success Criteria for the Proposed Attacks

This study hypothesizes that malicious low-magnitude kinetic deviations can noticeably and consistently modify a printed part's mechanical properties while maintaining the geometry and shape. In this context, low-magnitude deviations are considered the ones that are within the printer specifications tolerances. The success of the proposed attacks is based on the validity of the hypothesis, which was examined using the following criteria.

(1) Resultant modifications in the printed parts should remain within the tolerance of a typical FFF printer's specifications
(2) Attacked parts should statistically maintain the shape, dimensions, and weight
(3) Mechanical strength of the attacked parts should be consistently reduced
(4) Attacks should be able to evade the existing state-of-the-art detection schemes discussed in Section 2

3.2 Existing Attack Detection Horizon

The best results in detecting process deviation reported in the literature presented in Section 2 are 1 second per layer for timing profile, 0.05 mm for layer thickness, 1 mm^2 single area mismatch with at least 0.3 mm length per axis, and 5°C variation in the nozzle and printing bed thermal profile. These values constitute the current detection horizon for attacks on the FFF process.

3.3 Precision and Trueness Values of Common FFF Printers

Some vendors and researchers have reported the precision and trueness values of 3D printers, which have been utilized by Rais et al. [31]. For instance, Stratasys, Ltd found a 130 µm tolerance for 95% of parts printed on Fortus 360 mc/400 mc printers [16]. Other studies have observed similar values, such as Kim et al.'s findings of 99 ± 14 µm and 188 ± 14 µm for FFF printers [19], and Msallem et al.'s measurements of 160 ± 9 µm and 50 ± 5 µm for an Ultimaker 3 Ext FFF printer [26]. These studies suggest that low-magnitude variations within these reported values could be exploited by attackers, likely evading detection systems as expected printer behavior.

3.4 Examining FFF Process for Available Attack Opportunities

Figure 2 illustrates a sliced version of a single internal layer of a rectangular prism, highlighting two main components: walls and infill structure. For a solid load-bearing part, a common choice of infill pattern is "lines" or "rectilinear." As depicted in Figure 2, two infill lines are connected by small segments, whose length is proportional to the nozzle diameter. Manipulating the placement and the size of these connecting segments presents opportunities for low-magnitude attacks.

As the molten filament is extruded from a printer's nozzle, it either interacts with the printing bed or with the already extruded filament. Heat energy from the latest extrusion is used in melting

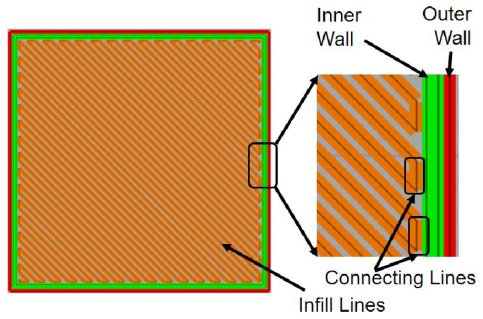

Fig. 2. Components of an internal layer in FFF printing.

(wetting) a small part of the existing filament in its proximity. The process is referred to as interdiffusion. As our aim is to identify kinetic attack opportunities with attack magnitudes within the order of the printer's specifications, our zone of interest is practically restricted to the adjacent extrudates only. The attacks' impact on the parts is measured through destructive mechanical tests, including tension and flexure strength tests.

3.5 Proposed Attacks

Analysis of the printing operation at any instance in time shows that the most recently extruded filament interacts with the existing extrudates belonging to the wall or infill structure of the same layer and the adjacent layers. The extrudates bonding phenomenon can be exploited by subtle kinetic deviations. The following subsections present four feasible attacks by exploiting these extrudates interactions. Attack 1 and Attack 2 are based on our previous work [31] that focuses on infill-related attack issues. These attacks were initially proposed with a focus on infill structure vulnerabilities. The inclusion of these attacks is crucial as they complete the portfolio of kinetic manipulation attacks, providing a broader perspective on how subtle kinetic process modifications can impact the mechanical properties of 3D printed parts.

3.5.1 Attack 1: Infill Lines Spacing Attack. Bonding between two spatially adjacent infill lines influences the overall strength of a solid part. In this attack, two consecutive extrudates from infill lines are separated by increasing the length of the connecting segment by a small fraction. Figure 3 presents one instance of this attack. The attacked connecting segment length d_a is increased by Δd_a, which is a fraction of the original segment length d_o. The length of two adjacent connecting segments d_{c_1} and d_{c_2} is reduced by Δd_{c_1} and Δd_{c_2}, respectively. Equations (1) and (2) present the relationship and constraints of the attack variables. K_s ranging from 0 to 1 is the stealth factor against any visible deformation.

$$0 < \begin{Bmatrix} \Delta d_{c_1} = d_o - d_{c_1} \\ \Delta d_{c_2} = d_o - d_{c_2} \end{Bmatrix} < (1 - K_s) * d_o/2, \tag{1}$$

$$\Delta d_a = \Delta d_{c_1} + \Delta d_{c2} = d_a - d_o, \tag{2}$$

3.5.2 Attack 2: Infill Vertices Spacing Attack. This attack also targets the bonding between consecutive infill extrudates within a layer. Instead of reducing the overlap across the two consecutive infill lines (as in Attack 1), this attack manipulates only one edge of the targeted part as presented in Figure 4. An inverse wedge is produced by reducing the length of two consecutive connecting segments at one edge. Depending upon the attack magnitude, the attack may only reduce the

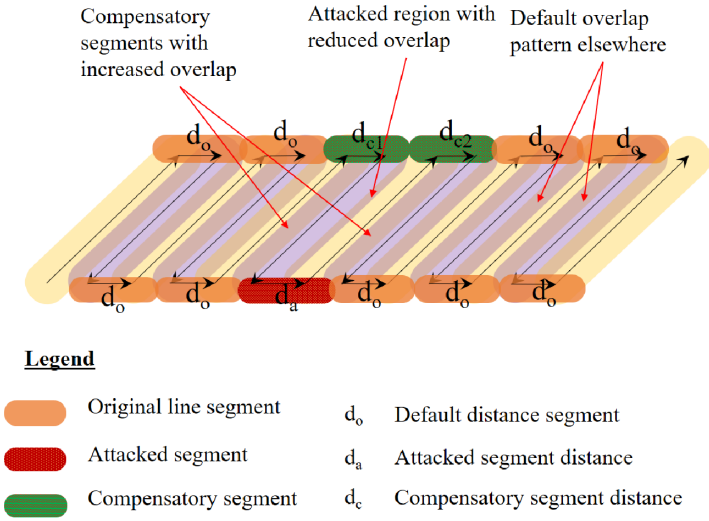

Fig. 3. Infill lines spacing attack representation.

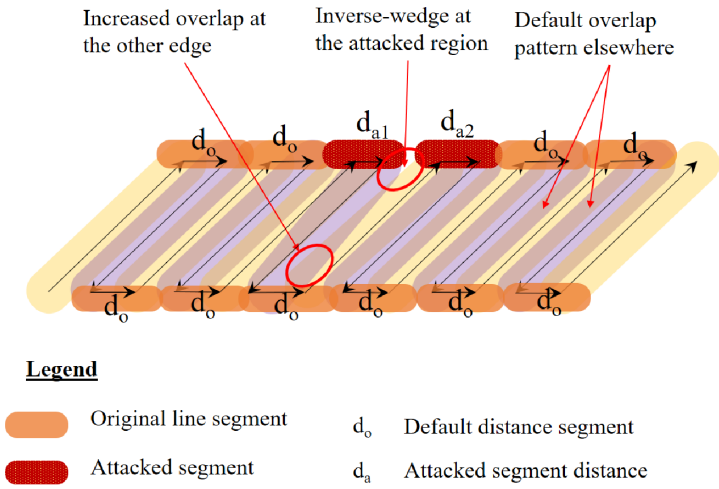

Fig. 4. Infill vertices spacing attack representation.

overlap (for smaller magnitudes) and a visible inverse wedge (for higher magnitudes). As the attack targets internal layers only, it is concealed in the final part for all magnitudes. The attack only manipulates the vertices of the connecting segments at one edge. This attack causes a minimal deviation in the local raster angle and the length of the two consecutive infill lines involved in the attack. Equation (3) represents the change in the length of the infill-lines, and Equation (4) represents the change in the raster angle,

$$d_{IF_a} = \sqrt{d_{IF_o}^2 - 2 * \Delta d_s * sin(\theta_o) * d_{IF_o} + \Delta d_s^2}, \tag{3}$$

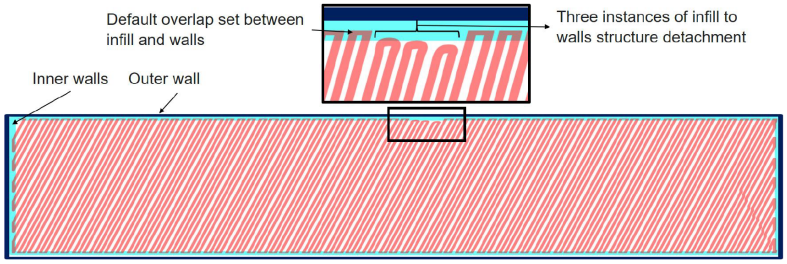

Fig. 5. Infill structure to walls bonding attack.

$$\theta_{IF_a} = tan^{-1} \frac{(d_{IF_o} * sin(\theta_o) + \Delta d_s)}{d_{IF_o} * cos(\theta_o)}, \tag{4}$$

where d_{IF_a} and d_{IF_o} are the modified and the original length of the infill lines, Δd_s is the difference between the original connecting segment length d_o and the modified length d_a, and θ_o *and* θ_a represent the default and the modified raster angles. Considering a 15 mm infill line configured at a raster angle of 45°, a 0.1 mm decrease in the connecting segment length will result in Δd_{IF} (change in infill line length) of around 0.07 mm, and the change in raster angle of 0.2°. These minimal changes are within the printer tolerances and beyond the capability of the existing attack detection schemes. The attack can be accomplished using different values of d_{a_1} and d_{a_2}. Interestingly, the changes in IF_1 and IF_2 have opposite polarity. If one decreases, the other increases, and vice versa. As the attack instances are launched over multiple layers, this polarity reversal helps in canceling out (instead of accumulating) the difference in the original and the attacked printing profile, making it more challenging for the attack detection schemes.

3.5.3 Attack 3: Infill and Wall Structure Bonding Attack. Unlike the previous two attacks, this attack targets fusion between the infill structure and the walls. Slicer software offers a choice to print the infill before or after the internal walls. In either case, these two constituents of internal layers are temporally displaced. If the infill is printed first, the later printed extrudate of the internal wall will interact with the infill structure creating a bond by interdiffusion. This attack manipulates the bonding strength between the infill and wall structure at the point of attack by reducing the overlap between the two regions. Figure 5 presents a typical attack with three instances of varying magnitudes increasing from left to right. Each attack instance is executed by modifying the end vertices of two consecutive toolpath instructions; the first instruction prints the preceding infill line, and the second one prints the targeted connecting segment. The length of the connecting segment is not changed, while the infill line segment length is decreased by the magnitude of the attack (typically a small fraction of a millimeter).

3.5.4 Attack 4: Inter-Layer Bonding Attack. As the filament is extruded out of the nozzle, it also interacts directly with the material from the previous layer. The impact of interlayer bonding on object strength is a well-researched topic [11]. This attack induces interlayer bonding weaknesses in the printed part. After printing a layer, if the bed is lowered more than the designed value without increasing the filament flow rate, the transferred heat and the pressure exerted by the new extrusion on the existing layer are reduced. For instance, if the printing bed is lowered by 0.2 mm for the nth layer against the designed layer thickness of 0.1 mm, the bonding between nth and $(n-1)$th layer will be poor. However, it creates an obvious mark on the sides of the part. To conceal the poor bonding mark, this attack exploits the

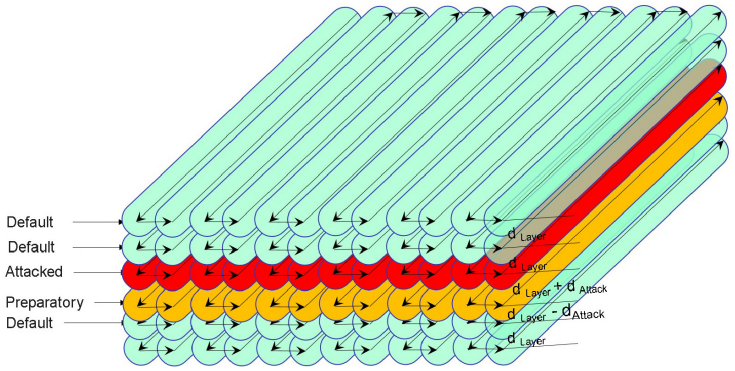

Fig. 6. Manipulation of inter-layer distance for the infill structure.

interlayer bonding for the infill structure only. The wall structure is still printed using the default profile. Figure 6 elaborates an attack scenario showing the infill structure. The layer thickness of the attacked layer (say nth) layer is modified to $d_{layer} + d_{attack}$ where $0 < d_{attack} < d_{layer}$. The increase in layer thickness for the nth layer is pre-compensated at $(n-1)$th layer by reducing the layer height for the infill structure to $d_{layer} - d_{attack}$. One may argue that the attack may incur some deformation in the layers other than the attacked ones. While it is plausible, it is not a concern for the attacker. If the compensatory moves are within the attack detection thresholds and do not cause obvious deformation (being within the infill structure only), they do not conflict with the attack success criteria. The attack is more effective against parts with walls printed prior to the infill. The inner edge of the wall structure will partially hold the elevated infill lines resulting in less pressure and transfer of energy to the lower layer.

4 Experimental Details

4.1 Attack Vector

The 3D printing process generally comprises three stages: designing, slicing, and printing. These stages have multiple components associated with them, with each component susceptible to compromise. For example, researchers have demonstrated many vulnerabilities [3] including in control PCs [6], slicing software [20], communication channels [21, 36], firmware [15, 24, 28, 30], and other elements [23, 37]. Such vulnerabilities and weaknesses open up opportunities for an adversary to sabotage the printing process.

Communication between the designing and slicing stages occurs entirely within the digital domain, where IT best practices, such as fully encrypted communication between independent systems, are typically implemented. In contrast, the printing stage, which involves the 3D printer, often lacks standard encryption and authorization practices [22]. As a result, attackers can target the network link between the control machine and the printer to manipulate printing instructions via an MiTM attack. Whilst most of the devices support network communication, some standalone printers communicate through a serial connection and/or a USB/SD-card file transfer. With the network-connected printers allowing for the remote injection of the malicious file via MiTM, serial and USB manipulation require different attack considerations, wherein the attacker compromises other components, e.g., control-PC or firmware [6, 28, 30].

Algorithm 1: Low-Magnitude Extrudates Bonding Kinetic Attacks

Input: Network traffic b/w printer and controller, Attack$_{param}$, A$_{No}$
Output: G-code$_{Attacked}$
Attack$_{\text{param}}$: { A$_{instances}$, A$_{Mag}$, A$_{Loc}$, A$_{Layers}$ }
Launch ARP Poisoning Attack
Sniff printer - controller communication
if Controller sends G-code to printer **:**
    Extract G-code file → G-code$_{Original}$
    G-code$_{Attacked}$ ← **Attack[A$_{No}$]-function**(Attack$_{param}$, G-code$_{original}$)
Send G-code$_{Attacked}$ to printer via MiTM
Manage communication
Attack-1-function(Attack$_{param}$, G-code$_{original}$) **:**
(Infill lines spacing attack)
while A$_{Loc}$ ∉ Infill-structure: shift A$_{Loc}$
∀ i ∈ A$_{Layers}$:
    a ← Search nearest connecting segment to A$_{Loc}$
    Calculate new x and y coords, such that:
        No change in the slope for any infill or segment
        $|d_{c_1}|$ ← $|d_{c_1}|$ - $|A_{Mag}|$; $|d_a|$ ← $|d_a|$ + $|A_{Mag}|$;
        $|d_{c_2}|$ ← $|d_{c_2}|$ - $|A_{Mag}|$; No change in $|Infill_1|$ & $|Infill_2|$
    ∀ j ∈ Attacked commands :
        Compute new G-code(j)
        Update G-code(j) in G-code$_{Attacked}$
return G-code$_{Attacked}$
Attack-2-function(Attack$_{param}$, G-code$_{original}$) **:**
(Infill vertices spacing attack)
while A$_{Loc}$ ∉ Infill-structure: shift A$_{Loc}$
∀ i ∈ A$_{Layers}$:
    a_1 ← Search nearest connecting segment to A$_{Loc}$
    Calculate new x and y coords, such that:
        No change in the slope of segments (slight change for infill lines)
        $|d_{a_1}|$ ← $|d_{a_1}|$ - $|A_{Mag}|$;
        $|d_{a_2}|$ ← $|d_{a_2}|$ - $|A_{Mag}|$; (Infill-lines magnitude will slightly change)
    ∀ j ∈ Attacked commands :
        Compute new G-code(j)
        Update G-code(j) in G-code$_{Attacked}$
return G-code$_{Attacked}$
Attack-3-function(Attack$_{param}$, G-code$_{original}$) **:**
(Infill to wall structure bonding attack)
while A$_{Loc}$ ∉ Infill-structure: shift A$_{Loc}$
∀ i ∈ A$_{Layers}$:
    ∀ j ∈ A$_{instances}$ **:**
        Calculate new x and y coords, such that
        No change in slope for infill lines or connecting segments
        $|d_{IF_j}|$ ← $|d_{IF_j}|$ - $|A_{Mag}|$
        $|d_{s_j}|$ not modified
        Add IF_j and S_j coords in Attacked-commands list

Algorithm 1: Low-Magnitude Extrudates Bonding Kinetic Attacks (Continued)

```
        ∀ k ∈ Attacked-commands :
            Compute new G-code(k)
            Update G-code(k) in G-code_Attacked
return G-code_Attacked
Attack-4-function(Attack_param , G-code_original) :
(Interlayer bonding attack)
∀ i ∈ A_Layers :
    Identify limits of Infill structure - [IF_0, IF_n]
    Identify IF_st and IF_end  ∋  0 ≤ st < end ≤ n
    if index(i) is even :
        Z_a1 ← Z_i - A_mag
    else:
        Z_a2 ← Z_i + A_mag
    Append Z_a1 to G-code(i_st) position
    Append Z_a2 to G-code(i_end) position
return G-code_Attacked
```

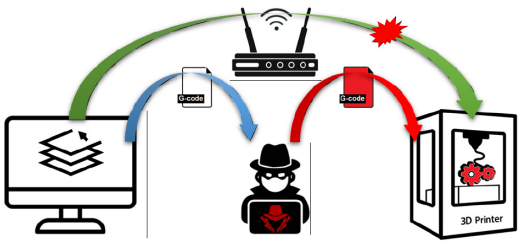

Fig. 7. MiTM attack to manipulate G-code file.

For demonstration, we assume the MiTM attack vector, where the attacker having access to the LAN manipulates the G-code file. To initiate a MiTM attack, the attacker employs ARP poisoning to associate their own MAC address with the IP addresses of both the 3D printer and the control PC. This allows the attacker to intercept and proxy the network traffic between these two entities. Specifically, the attacker sends spoofed ARP messages to the 3D printer, associating the attacker's MAC address with the IP address of the control PC and vice versa. This two-way ARP poisoning ensures that all communication between the control PC and the 3D printer passes through the attacker's machine, as illustrated in Figure 7. Once in control of the communication, the attacker can intercept the G-code commands generated by the slicer software on the control PC. They can then determine the appropriate attack type, modify the G-code instructions accordingly, and transmit the altered instructions over the network to the 3D printer. This manipulation allows the attacker to potentially compromise the integrity of the printed object without detection.

4.2 Attack Implementation

Algorithm 1 outlines the attack process for the four proposed attacks described in Section 3.5. After successfully targeting the communication channel and extracting the G-code file based on adversary input, one of the four attack actions is performed to output a manipulated G-code file. The malicious file is then transferred to the printer while ensuring that the communication between

the control PC and printer is maintained smoothly despite the MiTM attack. The four subroutines corresponding to each attack type are detailed as follows:

4.2.1 Subroutine: Attack 1. In this implementation, we choose the midpoint instruction within the infill section of the targeted layer as the attack location, A_{Loc}. The algorithm checks whether the selected instruction corresponds to connecting segments rather than infill lines and adjusts A_{Loc} if necessary. Based on the attack magnitude, A_{Mag}, the algorithm recalculates the modified x and y coordinates for three consecutive connecting segments. Finally, we apply the newly calculated coordinates to modify the specific G-code commands.

4.2.2 Subroutine: Attack 2. In this attack subroutine, similar to the previous one, the algorithm iterates through targeted layers to identify the connecting segment closest to the attack location. However, unlike the previous approach, only two segments, A_1 and A_2, are altered by a magnitude A_{Mag}. While the adjustment introduces slight changes to the slopes of the infill lines, the slope of the connecting segment remains unchanged. The corresponding G-code commands are then updated with the newly calculated coordinates.

4.2.3 Subroutine: Attack 3. Unlike the previous two attacks, Attack 3 adjusts the length of the infill lines by A_{Mag} to reduce the overlap between the wall structure and the central three pairs of connecting segments within the targeted layers. The algorithm calculates the new coordinates and updates the corresponding instructions in the G-code file.

4.2.4 Subroutine: Attack 4. This attack subroutine alters the infill layer height of two consecutive layers to compromise interlayer bonding. The algorithm analyzes the G-code file to identify the start (IF_{st}) and end (IF_{end}) points of the infill structure within the targeted layers. It then calculates the new layer heights, Z_{a_1} for the preparatory layer (reduced height) and Z_{a_2} for the attacked layer (increased height). The G-code file is modified by injecting the new layer-height commands at the start of the infill structure, followed by a reversion to the original height at its conclusion.

4.3 Experimental Settings

This section details the specification of the FFF printer, printing parameters, and the specimens used for the experiment. The overall dimensions of the tensile bars used for the experiment are $115 \times 19 \times 4.07$ mm ($length \times width \times thickness$), with the central part 6.5 mm in width. Due to the test equipment having a maximum span of 41 mm, a smaller thickness specimen is used for the flexure test with dimensions $76.8 \times 12.7 \times 2.4$ mm. The dimensions, however, comply with the standard's requirement to maintain a span-to-thickness ratio of 16. All the parts were printed with **polylactic acid (PLA)** polymer using an FFF-based printer—Ultimaker 3. The printer is connected over the LAN to the control machine hosting the Windows 10 operating system and running the Cura version 4.10 slicer application. Table 1 presents the printing profile of the slicer software chosen for the experiment based on the printing and attack heuristics.

4.4 Design of Experiment

Each of the attacks is implemented for a range of magnitudes to identify the strength reduction trend using statistical parameters, including mean and standard deviation. Table 2 outlines the design of the experiment. Kinetic manipulations in the first three attacks involve the x and y axes. As attacks 1 and 2 target the bonding between two consecutive infill lines, they are performed on the middle infill lines pair of the internal layers. For Attack 3, which targets the bonding between the infill and wall segment, 3 consecutive infill lines in the central infill region are attacked. Attack 4 involves z-axis manipulation only for the infill structure of the selected internal layers. Initially, the attack magnitudes selected for all attacks are 0.05 mm, 0.1 mm, and 0.2 mm to stay close to

Table 1. Printing Parameters Selected for the Experiment

S/No	Printing parameter	Selected value
1	Layer thickness	0.2 mm
2	Nozzle diameter	0.4 mm
3	Build plate temperature	60°C
4	Nozzle temperature—Layer 1	210°C
5	Nozzle temperature—Layer 2 onwards	205°C
6	Infill pattern	LINE at 45°
7	Infill percentage	100%
8	Infill overlap with walls	20%
9	Number of layers for tensile specimens	20
10	Number of layers for flexure specimens	12
11	Bottom layers	2
12	Top layers	Nil
13	Printing speed for initial layer	20 mm/sec
14	Printing speed for top/bottom layers	45 mm/sec
15	Infill printing speed	70 mm/sec
16	Walls printing speed (outer/inner)	50/55 mm/sec
17	Number of walls in Attack 1, 2, & 4	2
18	Number of walls in Attack 3	4
19	Printing sequence for Attack 1, 2, & 3	Infill first
20	Printing sequence for Attack 4	Walls first

Table 2. Design of Experiment for the Proposed Attacks

	Proposed attacks			
	Attack 1	**Attack 2**	**Attack 3**	**Attack 4**
Attack target: Inter-extrudates bonding between	2 consecutive infill lines across the span	2 consecutive infill lines at one edge	Infill and wall structure	Infill across 2 consecutive layers
Kinetic manipulation axes	x , y	x , y	x , y	z
Attack location	Internal layers, middle infill	Internal layers, middle infill	Internal layers, central infill zone	3 instances/attack in internal layer
Attack instances	**Attack magnitudes (mm)**			
1	0.015	0.025	0.05	0.05
2	0.025	0.05	0.10	0.10
3	0.05	0.10	0.15	0.15
4	0.10	0.20	0.20	0.20
5	0.20	–	–	–

the trueness specifications of FFF printers. Where needed to further examine the trend, additional steps are inserted at appropriate places. Attack 1 is examined up to 0.015 mm magnitude, whereas an extra step of 0.15 mm is inserted for Attack 3 and 4, and it was not felt necessary to examine these attacks below 0.05 mm magnitude. Five samples are printed for each attack instance.

5 Experiment Results

This section presents the performance of the attacked specimens in accordance with the success criteria outlined in Section 3.1.

Table 3. Stealthiness Performance: Impact on Dimensions and Mass of the Attacked Parts

Attack type	Attack magnitude (mm)	Width (mm)			Thickness (mm)			Mass (g)		
		Mean	Std dev	Diff.	Mean	Std dev	Diff.	Mean	Std dev	Diff.
Infill lines spacing (A1)	0	12.877	0.023	0.000	2.555	0.021	0.000	2.876	0.014	0.000
	0.015	12.885	0.028	0.008	2.518	0.012	−0.037	2.881	0.110	0.005
	0.025	12.853	0.009	−0.023	2.510	0.010	−0.045	2.883	0.012	0.007
	0.05	12.843	0.012	−0.033	2.537	0.014	−0.018	2.868	0.016	−0.008
	0.10	12.853	0.021	−0.023	2.543	0.019	−0.012	2.859	0.011	−0.014
	0.20	12.868	0.020	−0.008	2.558	0.011	0.003	2.853	0.018	−0.023
Infill vertices spacing (A2)	0	12.877	0.023	0.00	2.555	0.020	0.000	2.876	0.014	0.000
	0.025	12.885	0.042	0.008	2.590	0.023	0.035	2.869	0.011	−0.007
	0.05	12.907	0.038	0.030	2.588	0.018	0.033	2.874	0.019	−0.002
	0.10	12.898	0.045	0.021	2.580	0.021	0.025	2.886	0.017	0.010
	0.20	12.892	0.041	0.015	2.583	0.012	0.028	2.865	0.012	−0.011
Infill to walls bonding (A3)	0	12.873	0.017	0.000	2.563	0.017	0.000	2.870	0.022	0.000
	0.05	12.902	0.016	0.029	2.563	0.012	0.000	2.895	0.032	0.025
	0.10	12.891	0.017	0.018	2.547	0.005	−0.017	2.896	0.006	0.026
	0.15	12.887	0.028	0.014	2.567	0.012	0.003	2.881	0.021	0.012
	0.20	12.901	0.008	0.028	2.557	0.005	−0.007	2.870	0.022	0.000
Interlayer bonding (A4)	0	12.724	0.043	0.000	2.563	0.017	0.000	2.897	0.023	0.000
	0.05	12.721	0.022	−0.003	2.587	0.025	0.024	2.890	0.019	−0.006
	0.10	12.743	0.017	0.019	2.578	0.007	0.014	2.881	0.023	−0.015
	0.15	12.750	0.027	0.026	2.590	0.022	0.027	2.878	0.016	−0.019
	0.20	12.726	0.030	0.002	2.574	0.016	0.011	2.871	0.022	−0.026

5.1 Attack Stealthiness Against Part Inspection

As the attacks are performed on the internal layers without manipulating the wall structure, no visual impairment or modification is observed in any of the printed parts. The measurements of the printed parts confirm that the statistical difference between the thickness and width of the attacked versus non-attacked samples remains less than 0.1 mm for all cases. Similarly, the deviation in the mass of the printed parts is less than 0.03 g. Table 3 presents the mean and standard deviation of the attacked specimens' mass and dimensions along with their difference from the corresponding non-attacked specimens' measurements.

5.2 Impact of Attacks on Tension and Flexure Strength

In this subsection, the results of tensile and bending tests are presented for the proposed attacks. As the study was performed over a span of a few months using different PLA spools, a set of non-attacked parts is printed for each category except for the tensile specimens for Attack 1 and 2 (being printed through the same spool and settings).

5.2.1 Attack 1 Mechanical Tests Results. The tensile test results and the stress vs strain curves for Attack 1 are presented in Table 4 and Figure 8, respectively. This attack shows up to a 33% reduction in peak stress value at 0.1 mm or higher attack magnitude. The attacked specimens always broke from the point of attack and at a lower strain value. Table 5 presents three-point bending test results showing a 28% reduction in the peak flexure stress for the highest attack magnitude. Flexure stress vs flexure strain curves for Attack 1 are presented in Figure 9.

5.2.2 Attack 2 Mechanical Tests Results. Table 6 and Figure 10 presents the tensile tests results, whereas Table 7 and Figure 11 presents the three-point bending tests results for Attack 2. The maximum tensile stress reduction of 12.4% is observed at 0.2 mm attack magnitude. At all magnitudes,

Table 4. Tensile Test Results for Attack 1: Infill Lines Spacing Attack

	Peak load (N)			Peak stress (MPa)			Strain at break (mm/mm)		
Attack magnitude (mm)	**Average**	**Std dev**	**%age diff**	**Average**	**Std dev**	**%age diff**	**Average**	**Std dev**	**%age diff**
0	936.9	98.8	0.0	35.5	3.4	0.0	0.035	0.003	0.00
0.015	938.1	40.9	0.1	35.5	1.7	0.1	0.034	0.004	−2.86
0.025	919.9	35.7	−1.8	34.4	1.6	−3.1	0.03	0.002	−14.29
0.05	694.7	18.0	−25.8	25.9	0.7	−26.9	0.031	0.004	−11.43
0.1	622.6	34.7	−33.6	23.2	1.4	−34.6	0.026	0.002	−25.71
0.2	624.3	32.6	−33.4	23.3	1.3	−34.3	0.024	0.004	−31.43

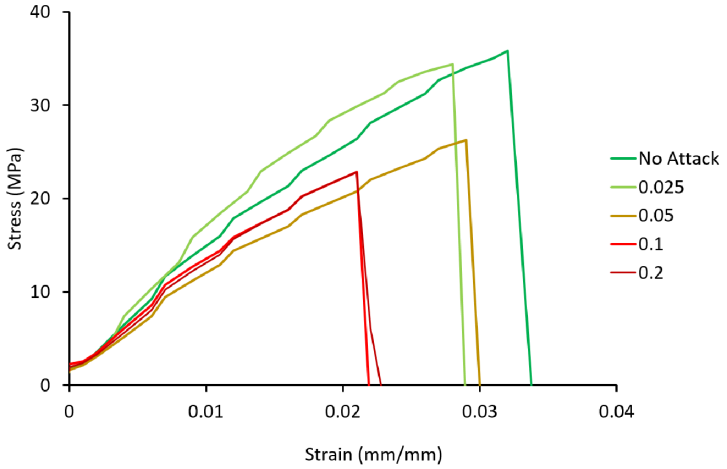

Fig. 8. Attack 1 stress vs strain curves for the tensile tests.

Table 5. Three Point Bending Test Results for Attack 1: Infill Lines Spacing Attack

	Peak load (N)			Peak flexure stress (MPa)		
Attack magnitude (mm)	**Mean**	**Std dev**	**%age diff**	**Mean**	**Std dev**	**%age diff**
0	103.09	4.44	0.00	74.54	3.12	0.00
0.025	101.98	4.64	−1.08	71.82	3.02	−3.64
0.05	99.84	1.44	−3.16	69.61	1.93	−6.61
0.1	77.51	4.96	−24.82	53.28	3.60	−28.52
0.2	74.18	13.08	−28.05	53.23	10.64	−28.58

the attacked specimens broke at a lower strain value. The bending tests show a reduction of 25% in the peak flexure stress value.

5.2.3 Attack 3 Mechanical Tests Results. Table 8 and Figure 12 presents the tensile tests results, whereas Table 9 and Figure 13 presents the three-point bending tests results for Attack 3. Although

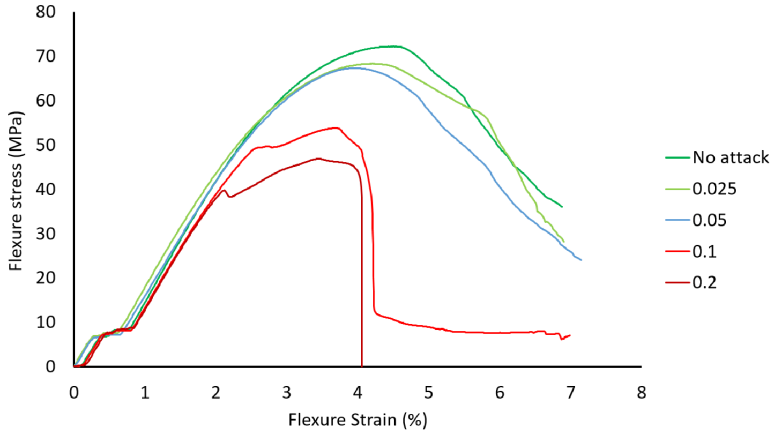

Fig. 9. Attack 1 flexure stress vs strain curves for the three-point bending tests.

Table 6. Tensile Test Results for Attack 2: Infill Vertices Spacing Attack

	Peak load (N)			Peak stress (MPa)			Strain at break (mm/mm)		
Attack magnitude (mm)	**Average**	**Std dev**	**%ag diff**	**Average**	**Std dev**	**%age diff**	**Average**	**Std dev**	**%age diff**
0	1,036.1	42.5	0.0	38.6	1.5	0.0	0.035	0.003	0.0
0.025	1,041.9	59.8	0.6	38.7	2.4	−0.5	0.03	0.003	−13.5
0.05	1,008.7	39.1	−2.6	37.9	1.3	−2.5	0.027	0.001	−21.5
0.1	953.4	44.4	−8.0	35.5	1.8	−8.7	0.026	0.008	−25.0
0.2	916.3	36.5	−11.6	34.1	1.2	−12.4	0.025	0.006	−27.8

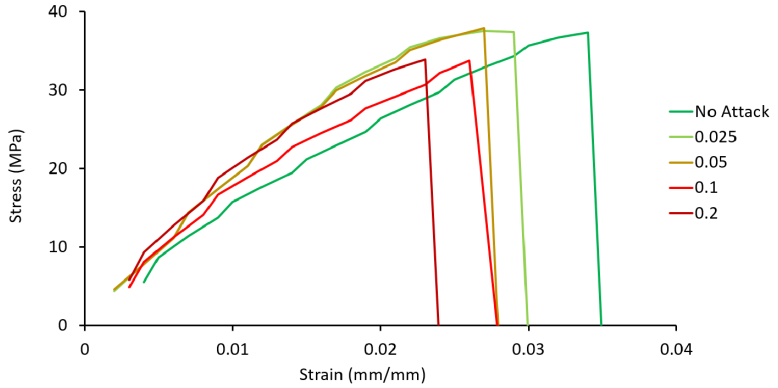

Fig. 10. Attack 2 stress vs strain curves for the tensile tests.

the attacks show a consistent reduction in peak tensile load and stress values, the maximum reduction is only 5.8%, which is not as pronounced as in Attacks 1 and 2. All attacked specimens still broke at a lower strain value. Similarly, the maximum reduction in bending strength is 6.17%. The reason for this low impact is discussed ahead in Section 7.

Table 7. Three Point Bending Test Results for Attack 2: Infill Vertices Spacing Attack

	Peak load (N)			Peak flexure stress (MPa)		
Attack magnitude (mm)	Mean	Std dev	%age diff	Mean	Std dev	%age diff
0	103.09	4.44	0.00	74.54	3.12	0.00
0.025	102.12	3.95	−0.94	73.43	1.92	−1.49
0.05	98.06	1.89	−4.88	70.85	1.17	−4.94
0.1	93.27	2.46	−9.52	67.67	1.06	−9.22
0.2	74.94	2.10	−27.31	55.52	2.52	−25.51

Table 8. Tensile Test Results for Attack 3: Infill to Wall Structure Bonding Attack

	Peak load (N)			Peak stress (MPa)			Strain at break (mm/mm)		
Attack magnitude (mm)	Average	Std dev	%ag diff	Average	Std dev	%age diff	Average	Std dev	%age diff
0	1,195.9	14.5	0.0	42.5	0.8	0.0	0.0305	0.006	0.000
0.025	1,180.2	15.5	−1.3	42.1	0.5	−0.9	0.0293	0.004	−3.780
0.05	1,147.5	21.9	−4.0	41.2	0.9	−3.1	0.0223	0.001	−26.776
0.1	1,130.1	15.2	−5.5	40.5	0.6	−4.6	0.0237	0.002	−22.404
0.2	1,121.9	8.0	−6.2	40.0	0.5	−5.8	0.0254	0.004	−16.721

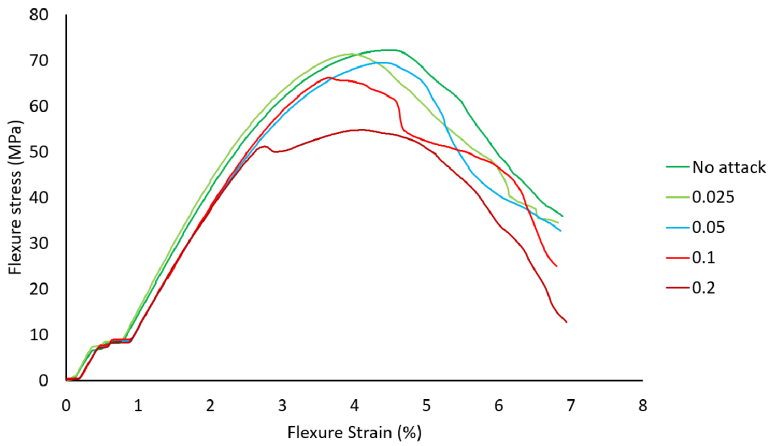

Fig. 11. Attack 2 flexure stress vs strain curves for the three-point bending tests.

5.2.4 Attack 4 Mechanical Tests Results. Table 10 and Figure 14 presents the tensile tests results, whereas Table 11 and Figure 15 presents the three-point bending tests results for Attack 4. Peak tensile stress reduction observed at the highest attack magnitude is 23%. Unlike the other three attacks, these attacked specimens did not break earlier except for the ones attacked with the highest magnitude (0.2 mm). Maximum reduction in the peak flexure stress is recorded as 16.56%.

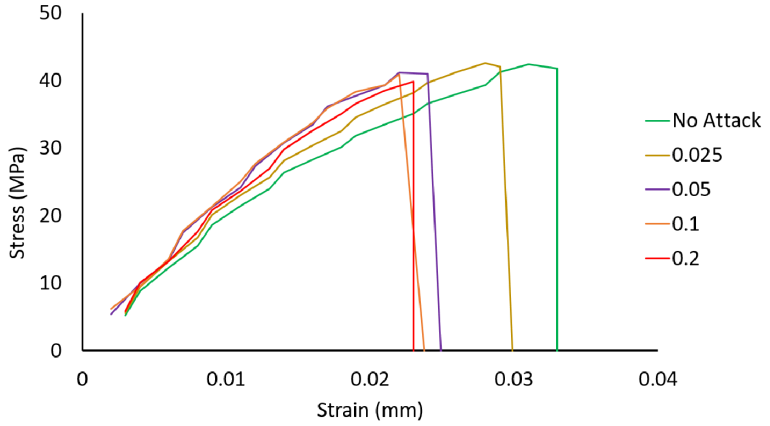

Fig. 12. Attack 3 stress vs strain curves for the tensile tests.

Table 9. Three Point Bending Test Results for Attack 3: Infill to Wall Spacing Attack

	Peak load (N)			Peak flexure stress (MPa)		
Attack magnitude (mm)	**Mean**	**Std dev**	**%age diff**	**Mean**	**Std dev**	**%age diff**
0	107.78	0.62	0.00	78.35	1.67	0.00
0.05	105.26	0.26	−2.72	75.75	2.83	−3.31
0.1	103.03	0.44	−5.12	73.51	1.23	−6.17
0.15	103.48	1.89	−4.63	73.92	1.29	−5.64
0.2	103.24	0.54	−4.91	73.71	0.72	−5.91

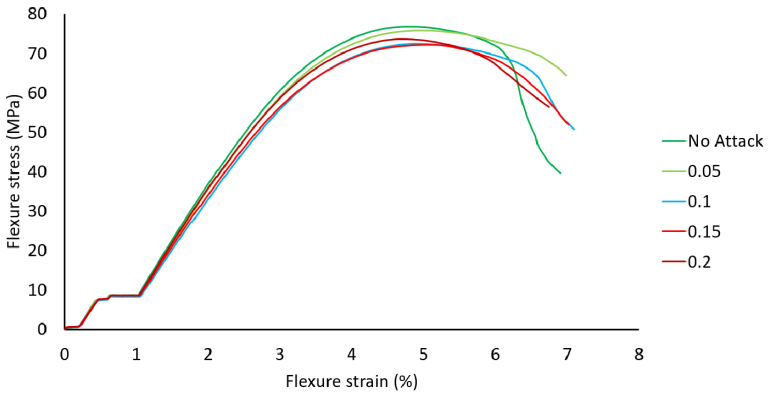

Fig. 13. Attack 3 flexure stress vs strain curves for the three-point bending tests.

Table 10. Tensile Test Results for Attack 4: Interlayer Bonding Attack

	Peak load (N)			Peak stress (MPa)			Strain at break (mm/mm)		
Attack magnitude (mm)	**Average**	**Std dev**	**%ag diff**	**Average**	**Std dev**	**%age diff**	**Average**	**Std dev**	**%age diff**
0	1,345.1	13.49	0.00	50.60	0.163	0.00	0.023	0.000	0.00
0.05	1,269.8	9.88	−5.59	48.53	0.822	−4.08	0.024	0.001	2.86
0.10	1,244.1	21.46	−7.51	46.50	0.852	−8.10	0.022	0.001	−7.14
0.15	1,139.6	4.13	−15.28	43.07	0.665	−14.89	0.022	0.001	−7.14
0.20	1,043.3	8.79	−22.44	38.93	0.471	−23.06	0.020	0.000	−14.29

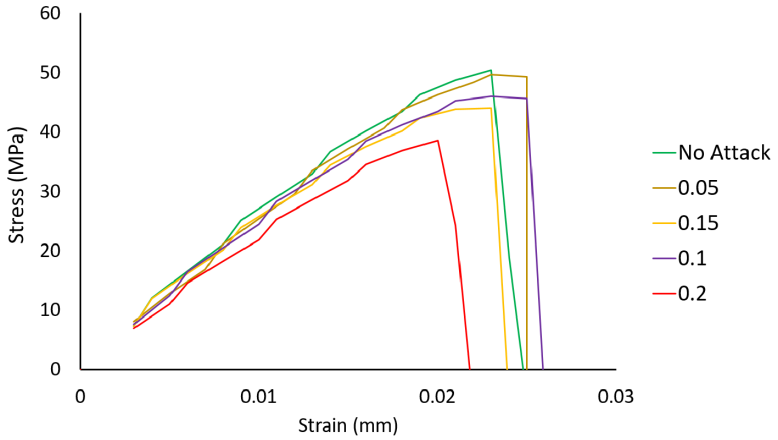

Fig. 14. Attack 4 strain vs strain curves for the tensile tests.

Table 11. Three Point Bending Test Results for Attack 4: Interlayer Bonding Attack

	Peak load (N)			Peak flexure stress (MPa)		
Attack magnitude (mm)	**Mean**	**Std dev**	**%age diff**	**Mean**	**Std dev**	**%age diff**
0	103.81	0.00	0.00	74.99	0.00	0.00
0.05	103.14	0.16	−0.65	73.74	0.60	−1.67
0.10	102.08	2.10	−1.67	73.84	1.33	−1.56
0.15	98.36	2.87	−5.25	69.46	1.18	−7.49
0.2	86.98	1.15	−16.22	63.49	0.75	−16.56

6 Evaluation of Proposed Attacks Against Detection Schemes

6.1 Stealthiness Performance: Deviations in the Printing Process

In addition to the common inspection parameters of the printed parts described in Section 5.1, certain unusual behaviors in the printing process can also disclose the attacks. Table 12 highlights a set of these parameters to examine the attacks' stealthiness against the detection schemes. All the attacks are launched after the user initiates the printing operation from the control machine by

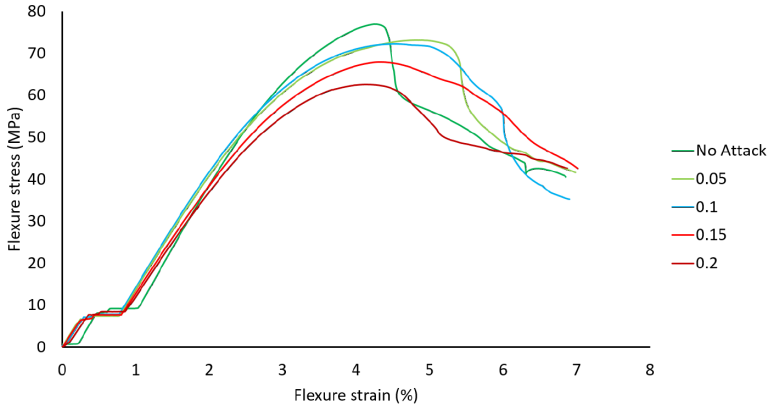

Fig. 15. Attack 4 flexure stress vs strain curves for the three-point bending tests.

Table 12. Stealthiness Performance Against Commonly Monitored Parameters

Attack type	Attack magnitude (mm)	Launch time delay (sec)	Max spatial deviation Linear (mm)	Max spatial deviation Angular (degree)	Per attack command time difference (sec)	Per command filament length difference (mm)	FNR % for detection scheme [34]
Infill lines spacing attack (A1)	0.015	0.25	0.015	0	0.005	None	100
	0.025	0.25	0.025	0	0.005	None	100
	0.05	0.25	0.05	0	0.005	None	100
	0.1	0.25	0.1	0	0.005	None	100
	0.2	0.25	0.2	0	0.005	None	18
Infill vertices spacing attack (A2)	0.025	0.25	0.025	0.154	0.005	None	100
	0.05	0.25	0.05	0.308	0.005	None	100
	0.1	0.25	0.1	0.612	0.005	None	100
	0.2	0.25	0.2	1.211	0.005	None	100
Infill to wall bonding attack (A3)	0.05	0.25	0.05	0	0.005	None	100
	0.1	0.25	0.1	0	0.005	None	100
	0.15	0.25	0.15	0	0.005	None	100
	0.2	0.25	0.2	0	0.005	None	100
Interlayer bonding attack (A4)	0.05	0.25	0.05	0	None	None	40
	0.1	0.25	0.1	0	None	None	0
	0.15	0.25	0.15	0	None	None	0
	0.2	0.25	0.2	0	None	None	0

sending the G-code file. The overall time in finding, recalculating, and modifying the G-codes in the input file is less than 250 ms on a Core i7-8700, 16 GB RAM machine. The spatial linear deviation is equivalent to the attack magnitude (0.2 mm max), and the maximum angular deviation is 1.2° for Attack 2 and 0° for the other three attacks.

The attacks effectively maintain the timing profile integrity on a per-instruction basis. When sampled at 5 ms, no statistical difference is observed in the execution time for Attacks 1 to 3 at the selected printing settings. The printing bed movement in Attack 4 takes from 50 to 150 ms in our printed specimen. Malicious bed movement is incorporated within the infill move command, ensuring no extra time for Attack 4. No attack modifies the filament consumption for any G-code instruction. These values are well below the detection performance of the existing techniques presented in the literature.

6.2 Attacks Evasion Performance Case Study

Our attacks are demonstrated at the cyber-physical boundary and their impact is primarily visible in the physical domain. To assess the attacks' performance, we considered the detection schemes

that utilize physical process information to identify anomalous actions [5, 8, 12, 13, 34, 39]. Among the various schemes, we selected the one with the highest precision against kinetic attacks [34].

6.2.1 Sophos—Spatiotemporal Anomaly Detection Scheme. Sophos [34] investigates spatiotemporal and thermal anomalies immediately after each layer is printed. This approach enables near real-time detection of attack instances, as opposed to relying on the accumulated print profile of the entire object, which may overlook instantaneous malicious modifications. High-resolution heterogeneous sensors monitor critical printing parameters such as the positions of the print head and bed, filament length, and nozzle and bed temperatures. The framework uses the G-code file as a reference for anomaly detection, transforming both G-code instructions and real-time sensor data into space and time-domain vectors for comprehensive comparison. After each layer, Sophos conducts multiple checks, including assessments of layer geometry, timing profiles, G-code command verification, and thermal profiles, to identify any anomalies. For instance, Sophos reverse-engineer a 2D image of each layer using sensor data, filament thickness functions, and motion equations. It then employs a custom image synchronization function to identify discrepancies between the images generated from G-code and those derived from sensor data. To establish the system's performance baseline, the authors printed a series of unaltered objects, reporting zero false negatives for attack magnitudes of 0.3 mm or greater on the X/Y axes. However, they observed a significant increase in both false negatives and false positives for deviations smaller than this threshold.

6.2.2 Attacks Evasion Performance Against Sophos. Our scope is focused on identifying the evasion potential of the proposed attacks. Consequently, we did not include benign cases to determine the false positive rate of the detection technique. Our primary interest lies in assessing the **false negative rate (FNR)** for the proposed attacks at various attack magnitudes.

The detection scheme uses a print-start sequence to initiate the detection engine, therefore the 200 milliseconds delay caused by MiTM prior to the print-start sequence does not trigger any anomaly. For the attack magnitude of 0.2 mm, Attack 1 is successfully detected with only 18% FNR. Once we reduce the magnitude to 0.15 mmm, FNR increased to 54%. For magnitudes of 0.1 mm and less, the scheme could not detect a single instance of Attack 1.

Unlike Attack 1, all instances of Attacks 2 and 3 (200 instances for each attack) successfully bypassed the detection scheme. This evasion is attributed to the smaller spatial deviation caused by Attacks 2 and 3 compared to Attack 1. However, the impact on the attacked object's strength is also less pronounced compared to Attack 1. Attack 4, or the interlayer bonding attack, manipulates the Z-axis kinetics, which is more precise than the XY-axis kinetics. For Attack 4, 60% of instances with a 0.05 mm magnitude were able to evade detection. However, all instances of Attack 4 with magnitudes ranging from 0.1 to 0.2 mm were successfully detected by the scheme. The FNR results are summarized in Table 12.

7 Analysis and Discussion

The results presented in the previous section demonstrate that manipulating extrudate bonding through kinetic variations can significantly weaken the tensile and bending strength of printed parts. These low-magnitude attacks leave a minimal footprint, allowing them to evade most existing attack or anomaly detection techniques. Attacks 1 and 2 fully satisfy the four-point success criteria outlined in Section 3.1. Attack 3 leads to a modest reduction in the tensile and flexural strength of the parts. In Attack 4, although the attack magnitude remains within 0.2 mm, some techniques [34] can reliably detect layer thickness deviations as small as 0.05 mm. However, the effectiveness of layer-thickness detection techniques in the context of auto-leveling in modern printers remains

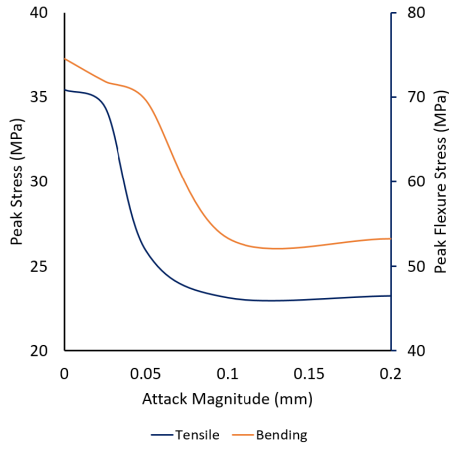

(a) Attack 1: Infill lines spacing attack

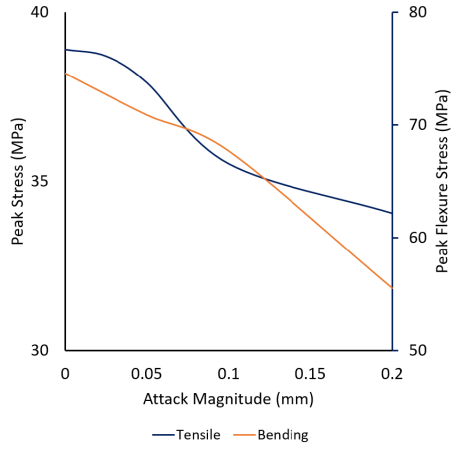

(b) Attack 2: Infill vertices spacing attack

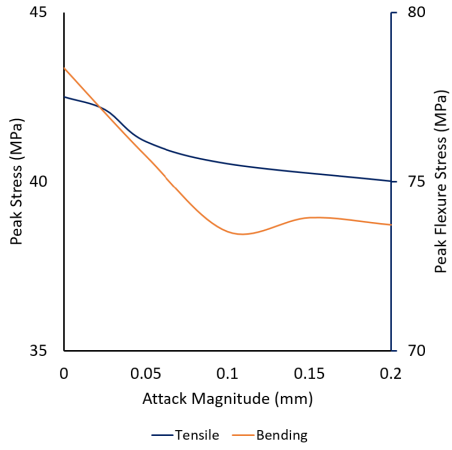

(c) Attack 3: Infill to wall structure bonding attack

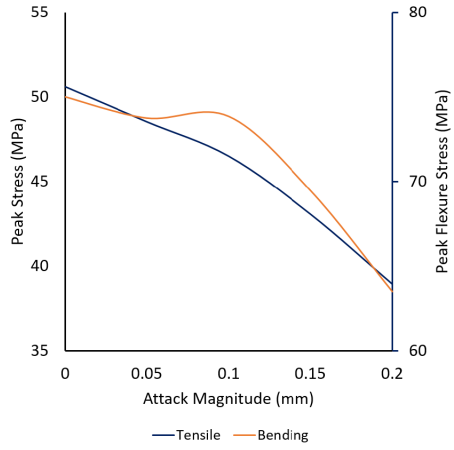

(d) Attack 4: Interlayer bonding attack

Fig. 16. Peak tensile and bending stresses at various attack magnitudes.

untested. The continuous bed's movement caused by auto-leveling could confuse detection schemes, potentially camouflaging the attack with false positives.

Figure 16 shows the peak tensile and flexural strength values plotted against the attack magnitude. Attacks 1 and 2 caused the greatest reduction in tensile and bending strength but are suitable only for solid parts. Attacks 3 and 4 resulted in less strength reduction but are applicable to both solid and non-solid geometries. Since the attacks introduce imperfections along different axes, the direction and type of load influence the choice of attack. The higher reduction in tensile strength in the first two attacks is due to the alignment of imperfections with the direction of the applied load.

For attack magnitudes greater than 0.1 mm in Attack 1, the attacked layers did not contribute to the tensile strength of the specimens. As presented in Figure 16(a), the peak stress value becomes nearly constant after a magnitude of 0.1 mm. The imperfection in Attack 2 is only introduced at one end of the infill lines pair (refer to Figure 4), resulting in a less impact on tensile strength as

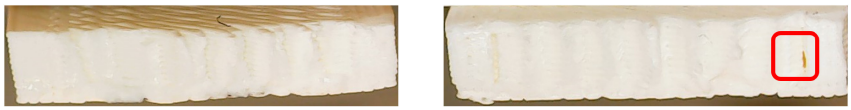

Fig. 17. Attack 3 specimens after bending tests.

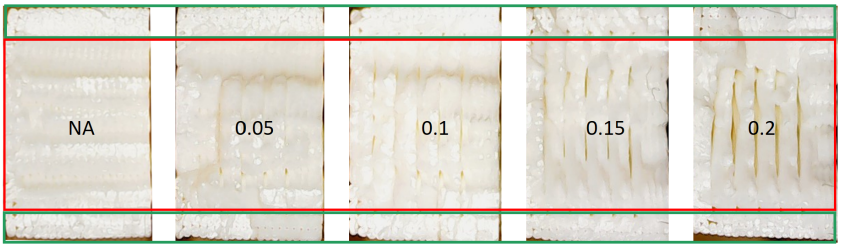

Fig. 18. Interlayer bonding attack (Attack 4) specimens after performing tensile tests.

presented in Figure 7. Unlike Attack 1, the reduction in Attack 2 continues after 0.1 mm but at a slower rate.

Attack 3 causes a minimal reduction in tensile strength. A cross-sectional view of the attacked specimens after destructive tests shows a tiny crack across multiple layers at 0.2 mm magnitude (see Figure 17). The defect introduced in Attack 3 is aligned with the tensile load direction; thus not causing any significant impact. This kind of attack could be effective for parts under compression or shear stress. In Attack 4, which targets interlayer bonding, the tensile strength reduction continues till 0.2 mm and beyond at a nearly linear rate. Figure 18 presenting cross-sectional views of 5 instances of Attack 4 highlights that the outer wall structures on both edges (in green color) do not show notable signs of attack. The weak interlayer bond in the infill structure becomes obvious as the attack magnitude increases from zero to 0.2 mm (from left to right in the figure).

A similar trend with slight differences is observed in the bending test results. As Attacks 1 and 2 are planted at the center of the part, the three-point bending results show considerable strength reduction proportional to the attack magnitude. The zone of steep reduction for Attacks 1 and 2 was shifted by approximately 0.05 mm in comparison to the tensile test results. In Attack 3, a small reduction in bending stress is observed at a lower attack magnitude, as visible in Figure 16(c), but the trend did not continue as the attack magnitude increased to 0.2 mm. As this experiment was restricted to a maximum deviation of 0.2 mm, the study did not investigate the effect at higher magnitudes. In Attack 4, a considerable impact on the bending stress is observed after the attack magnitude is raised from 0.1 towards 0.2 (see Figure 15).

This study focuses on evaluating the performance of the proposed attacks on PLA, selected for its widespread use, consistent mechanical properties, and ease of printing, making it an ideal baseline for this research. The impact of these kinetic attacks may differ when applied to other materials. For instance, ABS, with its greater ductility, and PETG, valued for its balance of printability and durability, could experience varying levels of mechanical degradation under similar attack conditions. Investigating these attacks across different materials offers a promising future research direction to better understand material-specific vulnerabilities and resilience in 3D printing.

8 Attack Countermeasures

This section discusses possible attack avoidance and detection measures against the proposed attacks. As the attacks are launched by hijacking the network connection, cybersecurity measures, including

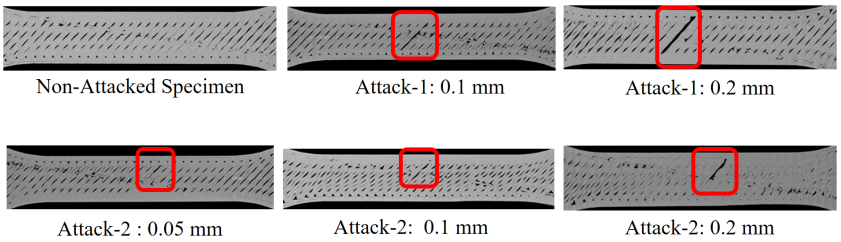

Fig. 19. Micro CT scan results for a few selected specimens.

access control and encrypted communication, DHCP snooping, and dynamic ARP inspection, are effective in avoiding MiTM attacks. These attacks can also be performed by compromising the printer's firmware or by using a kernel module to manipulate the G-code file in the slicer process memory. Tools such as Tracee [25] that detect and report the loading of any kernel module are helpful in detection and investigation. However, it is still challenging to automatically detect these attacks without any ground truth. To ascertain firmware integrity, periodic firmware verification through out-of-band methods may be employed [32].

If an attacker successfully launches these attacks, a forensic readiness framework could assist in investigating the attack traces [29]. Non-destructive tests might also reveal abnormalities in the internal structure. To explore this possibility, the authors conducted x-ray **micro-computed tomography (micro-CT)** tests on selected specimens from Attacks 1 and 2. The scans, performed using a Skyscan 1173 machine with 1.8-second exposure, 0.5-degree rotational steps, and a 20 µm pixel size, show that attack magnitudes of 0.1 mm or higher are easily detectable in Figure 19. However, lower-magnitude attacks do not produce obvious defects. Another challenge with micro-CT scans is the manual examination, which is time-consuming and requires expertise. Yoginath et al. proposed using trained probabilistic models to detect low-magnitude anomalies by analyzing phase-angle deviations in noisy side-channel data [42]. A real-time detection solution based on this approach could help discriminate malicious low-magnitude deviations from noise.

9 Limitations and Future Work

In FFF-based printing, the two primary physical processes that impact the bonding of extrudates and can be influenced by cyberattacks are kinetics and thermodynamics. This study was limited to kinetic manipulations. A promising direction for future research is to examine low-magnitude thermodynamic manipulations, followed by hybrid attacks that leverage the full range of modifiable parameters while remaining below the detection threshold. Additionally, it is crucial to develop cybersecurity solutions capable of detecting malicious deviations and simultaneously estimating their impact on the printed part. Such solutions would enable informed decisions on whether to continue, abort, or mitigate the attack.

This study examined the potential of attacks against the parts printed with PLA filament. An interesting avenue for future research would be to explore the applicability and effectiveness of the proposed and the existing cyberattacks on different printing materials.

This study focused on the common scenario of hijacking network communication to manipulate G-codes in transit. However, two significant cases were not covered. First, some printers use only a USB-serial connection, making typical MiTM attacks inapplicable. Second, certain printers perform slicing internally, requiring only the design file as input over the network. The proposed attacks on sliced files would be ineffective on pre-sliced files. Therefore, alternative methods, such as firmware manipulation or bootloader compromise, may be necessary to execute these attacks.

10 Conclusion

With all the benefits and the promise, the material extrusion printing process also brings up some unique challenges. To use material extrusion for printing critical parts, it is important to establish the integrity of the printing operation. Material extrusion printer's trueness and precision specifications offer an opportunity window to launch attacks that do not deviate the process beyond the tolerance window and still impact the mechanical strength. It is challenging for process monitoring-based techniques to declare such small within-tolerance deviations as anomalies. This study proposes four attacks on FFF printing by manipulating the bonding between two neighboring extrudates from the same and adjacent layers. The attacks were demonstrated by modifying the G-code file through an MiTM attack on the network segment between the control machine and the printer. An experiment was designed to evaluate the impact of the proposed attacks on ASTM D638 Type IV standard tensile bars and ASTM D790 compliant flexure bars using attack deviation magnitudes ranging from 0.015 mm to 0.2 mm. The destructive tests conducted on the attacked specimens confirm that the attacks are capable of reducing tensile and bending strength by up to 28% and 25%, respectively. Manual analysis of Micro CT scans of the attacked specimen shows that higher magnitude attacks can be spotted. However, manual scanning and analysis is not a scalable solution for commercial manufacturing. In addition to the standard cybersecurity tools and methods to avoid MiTM attacks, the authors recommend that researchers investigate the option of real-time CT scanning and automated analysis to ascertain the printing process's integrity. If FFF has to be successful in printing functional parts, a dedicated research effort is required to safeguard it against low-magnitude sabotage attacks on extrudates bonding.

References

[1] Muhammad Ahsan and Muhammad Ali. 2023. LsStk: Lightweight solution to preventing Stack from buffer overflow vulnerability. In *Proceedings of the 17th International Conference on Open Source Systems and Technologies (ICOSST)*, 1–7. DOI: https://doi.org/10.1109/ICOSST60641.2023.10414205

[2] Muhammad Ahsan, Eunice Pak, Kate Jackson, Muhammad Haris Rais, Barry Najarro-Blancas, Nastassja Lewinski, and Irfan Ahmed. 2024. BioSaFe: Bioprinting security framework for detecting sabotage attacks on printability and cell viability. In *Proceedings of the 40th Annual Computer Security Applications Conference (ACSAC '24)*. IEEE.

[3] Muhammad Ahsan, Muhammad Haris Rais, and Irfan Ahmed. 2023. Sok: Side channel monitoring for additive manufacturing-bridging cybersecurity and quality assurance communities. In *Proceedings of the IEEE 8th European Symposium on Security and Privacy (EuroS & P)*. IEEE, 1160–1178.

[4] Hina Alam, Muhammad Shaharyar Yaqub, and Ibrahim Nadir. 2022. Detecting IoT attacks using multi-layer data through machine learning. In *Proceedings of the 2nd International Conference on Distributed Computing and High Performance Computing (DCHPC)*, 52–59. DOI: https://doi.org/10.1109/DCHPC55044.2022.9732117

[5] S. Belikovetsky, Y. A. Solewicz, M. Yampolskiy, J. Toh, and Y. Elovici. 2019. Digital audio signature for 3D printing integrity. *IEEE Transactions on Information Forensics and Security* 14, 5 (2019), 1127–1141.

[6] Sofia Belikovetsky, Mark Yampolskiy, Jinghui Toh, Jacob Gatlin, and Yuval Elovici. 2017. dr0wned–Cyber-physical attack with additive manufacturing. In *Proceedings of the 11th USENIX Workshop on Offensive Technologies (WOOT 17)*. USENIX Association. Retrieved from https://www.usenix.org/conference/woot17/workshop-program/presentation/belikovetsky

[7] Javaid Butt and Raghunath Bhaskar. 2020. Investigating the effects of annealing on the mechanical properties of FFF-printed thermoplastics. *Journal of Manufacturing and Materials Processing* 4, 2 (2020). DOI: https://doi.org/10.3390/jmmp4020038

[8] S. R. Chhetri, A. Canedo, and M. A. Al Faruque. 2016. KCAD: Kinetic cyber-attack detection method for cyber-physical additive manufacturing systems. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 1–8.

[9] Ugur M. Dilberoglu, Bahar Gharehpapagh, Ulas Yaman, and Melik Dolen. 2017. The role of additive manufacturing in the era of industry 4.0. *Procedia Manufacturing* 11 (2017), 545–554. DOI: https://doi.org/10.1016/j.promfg.2017.07.148

[10] Yanzhou Fu, Austin Downey, Lang Yuan, Avery Pratt, and Yunusa Balogun. 2021. In situ monitoring for fused filament fabrication process: A review. *Additive Manufacturing* 38 (2021), 101749.

[11] Xia Gao, Shunxin Qi, Xiao Kuang, Yunlan Su, Jing Li, and Dujin Wang. 2021. Fused filament fabrication of polymer materials: A review of interlayer bond. *Additive Manufacturing* 37 (2021), 101658. DOI: https://doi.org/10.1016/j.addma.2020.101658

[12] Yang Gao, Borui Li, Wei Wang, Wenyao Xu, Chi Zhou, and Zhanpeng Jin. 2018. Watching and safeguarding your 3D printer: Online process monitoring against cyber-physical attacks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (Sep. 2018), Article 108, 27 pages. DOI: https://doi.org/10.1145/3264918

[13] J. Gatlin, S. Belikovetsky, S. B. Moore, Y. Solewicz, Y. Elovici, and M. Yampolskiy. 2019. Detecting sabotage attacks in additive manufacturing using actuator power signatures. *IEEE Access* 7 (2019), 133421–133432.

[14] Ian Gibson, David Rosen, Brent Stucker, and Mahyar Khorasani. 2021. *Material Extrusion*. Springer International Publishing, Cham, 171–201. DOI: https://doi.org/10.1007/978-3-030-56127-7_6

[15] Xiao Zi Hang. 2016. Three Demos of Attacking Arduino and Reprap 3d Printers, Code to Keynote at XCon2013. Retrieved from https://github.com/secmobi/attack-arduino-and-reprap

[16] J. Hanssen. 2013. *Fortus 360mc/400mc Accuracy Study*. Technical Report. Stratasys, Eden Prairie.

[17] Bilal Imran, Bilal Afzal, Ali Hammad Akbar, Muhammad Ahsan, and Ghalib A. Shah. 2019. MISA: Minimalist implementation of oneM2M security architecture for constrained IoT devices. In *Proceedings of the IEEE Global Communications Conference (GLOBECOM)*, 1–6. DOI: https://doi.org/10.1109/GLOBECOM38437.2019.9013863

[18] Bilal Imran, Muhammad Ahsan, Ali Hammad Akbar, and Ghalib Asadullah Shah. 2024. D4GW: DTLS for gateway multiplexed application to secure MQTT(SN)-based pub/sub architecture. *Internet of Things* 26 (2024), 101172. DOI: https://doi.org/10.1016/j.iot.2024.101172

[19] Soo-Yeon Kim, Yoo-Seok Shin, Hwi-Dong Jung, Chung-Ju Hwang, Hyoung-Seon Baik, and Jung-Yul Cha. 2018. Precision and trueness of dental models manufactured with different 3-dimensional printing techniques. *American Journal of Orthodontics and Dentofacial Orthopedics* 153, 1 (2018), 144–153. DOI: https://doi.org/10.1016/j.ajodo.2017.05.025

[20] Elizabeth Kurkowski, Alyxandra Van Stockum, Joel Dawson, Curtis Taylor, Tricia Schulz, and Sujeet Shenoi. 2022. Manipulation of g-code toolpath files in 3D printers: Attacks and mitigations. In *Proceedings of the International Conference on Critical Infrastructure Protection XVI*. Jason Staggs and Sujeet Shenoi (Eds.), Springer Nature Switzerland, Cham, 155–174.

[21] Matthew McCormack, Sanjay Chandrasekaran, Guyue Liu, Tianlong Yu, Sandra DeVincent Wolf, and Vyas Sekar. 2020. Security analysis of networked 3D printers. In *Proceedings of the International Conference on IEEE Security and Privacy Workshops (SPW)*, 118–125. DOI: https://doi.org/10.1109/SPW50608.2020.00035

[22] Matthew McCormack, Sanjay Chandrasekaran, Guyue Liu, Tianlong Yu, Sandra DeVincent Wolf, and Vyas Sekar. 2020. Security analysis of networked 3d printers. In *Proceedings of the IEEE Security and Privacy Workshops (SPW)*. IEEE, 118–125.

[23] Samuel Moore, Phillip Armstrong, Todd McDonald, and Mark Yampolskiy. 2016. Vulnerability analysis of desktop 3D printer software. In *Proceedings of the International Conference on Resilience Week (RWS)*. IEEE, 46–51.

[24] Samuel Bennett Moore, William Bradley Glisson, and Mark Yampolskiy. 2017. Implications of malicious 3D printer firmware. In *Proceedings of Hawaii International Conference on System Sciences*, 1–10. DOI: https://doi.org/10.24251/HICSS.2017.735

[25] Assaf Morag and Itamar Maouda. 2021. Understanding the evolving threat landscape – APT techniques in a container environment. *Network Security* 2021, 12 (2021), 13–17. DOI: https://doi.org/10.1016/S1353-4858(21)00145-8

[26] Bilal Msallem, Neha Sharma, Shuaishuai Cao, Florian S. Halbeisen, Hans-Florian Zeilhofer, and Florian M. Thieringer. 2020. Evaluation of the dimensional accuracy of 3D-printed anatomical mandibular models using FFF, SLA, SLS, MJ, and BJ printing technology. *Journal of Clinical Medicine* 9, 3 (2020). DOI: https://doi.org/10.3390/jcm9030817

[27] Ahsan Muhammad, Bilal Afzal, Bilal Imran, Asim Tanwir, Ali Hammad Akbar, and Ghalib Shah. 2019. oneM2M Architecture based secure MQTT binding in Mbed OS. In *Proceedings of the IEEE European Symposium on Security and Privacy Workshops (EuroS & PW)*, 48–56. DOI: https://doi.org/10.1109/EuroSPW.2019.00012

[28] Hammond Pearce, Kaushik Yanamandra, Nikhil Gupta, and Ramesh Karri. 2022. FLAW3D: A trojan-based cyber attack on the physical outcomes of additive manufacturing. *IEEE/ASME Transactions on Mechatronics* 27, 6 (2022), 5361–5370.

[29] Muhammad Haris Rais, Muhammad Ahsan, and Irfan Ahmed. 2023. Fromepp: Digital forensic readiness framework for material extrusion based 3d printing process. *Forensic Science International: Digital Investigation* 44 (2023), 301510.

[30] Muhammad Haris Rais, Muhammad Ahsan, and Irfan Ahmed. 2024. SOK: 3D printer firmware attacks on fused filament fabrication. In *Proceedings of the 18th USENIX WOOT Conference on Offensive Technologies*. USENIX Association.

[31] Muhammad Haris Rais, Muhammad Ahsan, Vaibhav Sharma, Radhika Barua, Rob Prins, and Irfan Ahmed. 2022. Low-magnitude infill structure manipulation attacks on fff-based 3D printers. In *Proceedings of the International Conference on Critical Infrastructure Protection XVI*. Springer, 205–232.

[32] Muhammad Haris Rais, Rima Asmar Awad, Juan Lopez, and Irfan Ahmed. 2021. JTAG-based PLC memory acquisition framework for industrial control systems. *Forensic Science International: Digital Investigation* 37 (2021), 301196. DOI: https://doi.org/10.1016/j.fsidi.2021.301196

[33] Muhammad Haris Rais, Ye Li, and Irfan Ahmed. 2021. Dynamic-thermal and Localized Filament-kinetic Attacks on Fused Filament Fabrication based 3D Printing Process. *Additive Manufacturing* (2021), 102200. DOI: https://doi.org/10.1016/j.addma.2021.102200

[34] Muhammad Haris Rais, Ye Li, and Irfan Ahmed. 2021c. Spatiotemporal G-code modeling for secure FDM-based 3D printing. In *Proceedings of the ACM/IEEE 12th International Conference on Cyber-Physical Systems (ICCPS '21)*. ACM, New York, NY, 177–186. DOI: https://doi.org/10.1145/3450267.3450545

[35] L. Sturm, C. Williams, J. Camelio, J. White, and R. Parker. 2014. Cyber-physical vunerabilities in additive manufacturing systems. *Context* 7, 8 (2014), 951–963.

[36] Logan D Sturm, Christopher B Williams, Jamie A Camelio, Jules White, and Robert 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 (2017), 154–164.

[37] Alyxandra Van Stockum, Elizabeth Kurkowski, Tiffany Potok, Curtis Taylor, Joel Dawson, Mason Rice, and Sujeet Shenoi. 2022. Attack-defense modeling of material extrusion additive manufacturing systems. In *Proceedings of the International Conference on Critical Infrastructure Protection*. Springer, 121–153.

[38] Mingtao Wu, Zhengyi Song, and Young B Moon. 2019. Detecting cyber-physical attacks in CyberManufacturing systems with machine learning methods. *Journal of Intelligent Manufacturing* 30, 3 (2019), 1111–1123.

[39] Mingtao Wu, Heguang Zhou, Longwang Lin, Bruno Silva, Zhengyi Song, Jackie Cheung, and Young Moon. 2017. Detecting attacks in cybermanufacturing systems: Additive manufacturing example. *MATEC Web of Conferences* 108 (2017), 06005. DOI: https://doi.org/10.1051/matecconf/201710806005

[40] Mark Yampolskiy, Wayne E. King, Jacob Gatlin, Sofia Belikovetsky, Adam Brown, Anthony Skjellum, and Yuval Elovici. 2018. Security of additive manufacturing: Attack taxonomy and survey. *Additive Manufacturing* 21 (2018), 431–457.

[41] Muhammad Shaharyar Yaqub, Haroon Mahmood, Ibrahim Nadir, and Ghalib Asadullah Shah. 2022. An ensemble approach for IoT firmware strength analysis using STRIDE threat modeling and reverse engineering. In *Proceedings of the 24th International Multitopic Conference (INMIC)*, 1–6. DOI: https://doi.org/10.1109/INMIC56986.2022.9972941

[42] Srikanth Yoginath, Michael Iannacone, Varisara Tansakul, Ali Passian, Rob Jordan, Joel Asiamah, M. Nance Ericson, Gavin Long, and Joel A. Dawson. 2022. Stealthy Cyber anomaly detection on large noisy multi-material 3D printer datasets using probabilistic models. In *Proceedings of the ACM CCS Workshop on Additive Manufacturing (3D Printing) Security (AMSec'22)*. ACM, New York, NY, 25–38. DOI: https://doi.org/10.1145/3560833.3563564

[43] Steven Eric Zeltmann, Nikhil Gupta, Nektarios Georgios Tsoutsos, Michail Maniatakos, Jeyavijayan Rajendran, and Ramesh Karri. 2016. Manufacturing and security challenges in 3D printing. *JOM* 68, 7 (Jul. 2016), 1872–1881. DOI: https://doi.org/10.1007/s11837-016-1937-7

Received 20 June 2023; revised 30 August 2024; accepted 26 October 2024