Chapter 5

ATTACK-DEFENSE MODELING OF MATERIAL EXTRUSION ADDITIVE MANUFACTURING SYSTEMS

Alyxandra Van Stockum, Elizabeth Kurkowski, Tiffany Potok, Curtis Taylor, Joel Dawson, Mason Rice and Sujeet Shenoi

Abstract The use of additive manufacturing in the critical infrastructure makes it an attractive target for cyber attacks. However, research on additive manufacturing threats has tended to focus on specific vulnerabilities and specific attacks against specific systems. The narrow scope hinders the understanding of the attack vectors that constitute the attack surfaces as well as the various targets and impacts of attacks. This results in vulnerabilities, potential attacks and countermeasures being overlooked during security analyses.

This research addresses the limitations by focusing on material extrusion, the most common additive manufacturing process. A material extrusion workflow (process chain) that comprehensively covers the design, slicing and printing phases is specified. Analysis of the workflow in conjunction with attack and defense frameworks yields attack-defense models for the three material extrusion phases. The attack-defense models, which specify the attack vectors, attack vector vulnerabilities and countermeasures, attack surfaces, system targets, target vulnerabilities and vulnerability countermeasures, and attacks and attack impacts, directly support risk identification, risk assessment and analysis, and risk mitigation and planning.

Three material extrusion printers ranging from hobbyist to industrial systems are used as case studies. Four attacks on the printers during the design, slicing and printing phases are described, including vulnerability identification, exploit development and countermeasures. The case studies demonstrate the effectiveness of attack-defense modeling and its ability to clarify and bolster the cyber security and risk management postures of material extrusion additive manufacturing environments.

Keywords: Additive manufacturing, material extrusion, attack-defense modeling

Published by Springer Nature Switzerland AG 2022
J. Staggs and S. Shenoi (Eds.): Critical Infrastructure Protection XVI, IFIP AICT 666, pp. 121–153, 2022.
https://doi.org/10.1007/978-3-031-20137-0_5

1. Introduction

Additive manufacturing is a multi-step process for building physical objects (parts) from computer-aided designs [24]. Unlike traditional subtractive manufacturing that removes material to create parts, additive manufacturing applies material layer by layer to build parts. Additive manufacturing combines manufacturing automation and custom part creation in ways that subtractive manufacturing cannot accomplish [10].

Additive manufacturing is a key component of Industry 4.0 – the fourth industrial revolution [6]. Industry 4.0 is the digital transformation of manufacturing and production industries characterized by the intelligent networking of machines that bridges the physical and digital worlds via cyber-physical systems that define and implement the manufacturing steps for flexible and customizable part production. The digital transformation supports autonomous decision-making and real-time monitoring of assets and processes. Additive manufacturing enables new capabilities in product design, prototyping, remote control, predictive maintenance, system monitoring and more.

Additive manufacturing is a multibillion-dollar industry [13]. Many critical infrastructure sector industries rely on additive manufacturing for mission-critical parts. The incorporation of additive manufacturing systems and their products in the critical infrastructure makes them attractive targets for hackers, criminal entities and nation-state actors.

In general, there are two types of additive manufacturing threats. The first are threats that use additive manufacturing for malicious purposes – concealing illicit objects such as drugs or explosives in printed parts, and creating objects such as untraceable "ghost guns" and spoofed biometrics of fingerprints and facial features [9]. The second are threats against additive manufacturing – intellectual property theft, part sabotage and additive manufacturing environment sabotage [24]. This research focuses on the threats against additive manufacturing, which are more serious in the context of the critical infrastructure.

Several researchers have investigated threats against additive manufacturing. However, the research efforts have primarily examined specific vulnerabilities and specific attacks against specific additive manufacturing systems [3, 7, 25]. Also, the research primarily focuses on firmware and stereolithography (STL) design file manipulations [5]. The research is interesting and important – it provides valuable insights into threats and their mitigation, and stimulates efforts at securing additive manufacturing systems. However, the deficiency is that the research efforts do not adopt holistic perspectives of additive manufacturing systems, let alone families of additive manufacturing systems corresponding to the seven standard additive manufacturing processes [10].

The narrow focus is problematic. The consideration of a specific additive manufacturing system instead of an additive manufacturing process hinders the overall understanding of the attack vectors that constitute the attack surface as well as the various targets and impacts of attacks. The lack of comprehension and comprehensiveness can result in vulnerabilities, potential attacks and countermeasures being overlooked during security analyses, negatively impacting risk management efforts.

This research attempts to address the limitations by focusing on the most common additive manufacturing process – material extrusion, also called fused deposition modeling or fused filament fabrication [12]. The material extrusion process involves heating material and depositing it on a print bed via an extruder layer by layer according to G-code toolpath instructions. The research comprehensively models the material extrusion workflow (process chain) over three additive manufacturing phases: (i) design, (ii) slicing and (iii) printing. The fourth phase, post-processing, is not considered because an analysis of the material extrusion process reveals that the overwhelming majority of cyber threats target the earlier design, slicing and printing phases.

The material extrusion workflow facilitates the specification of attack-defense models for complex material extrusion additive manufacturing systems. An attack-defense model is created for each phase by specifying the original attack surface and implemented attack vector countermeasures to establish the current attack surface. Next, the system targets that can be accessed using the current attack surface are identified. Following this, the material extrusion workflow and the MITRE ATT&CK Knowledge Base [14] are employed to identify vulnerabilities in the targets and potential attacks that exploit the vulnerabilities. Next, countermeasures based on the MITRE D3FEND Knowledge Graph [15] are identified to combat the attacks. Attacks without adequate countermeasures would be successful and their potential negative impacts are specified. The attack-defense model directly supports three key risk management steps, risk identification, risk assessment and analysis, and risk mitigation and planning [16].

Three material extrusion printers are used as case studies in this research. The first is a material extrusion printer with a price point of $25,000 that is used in industry. The second is a fused filament fabrication printer priced at $4,000 that is commonly used in laboratory environments. The third is a $300 fused filament fabrication printer primarily used by educators and hobbyists.

Four real attacks on the material extrusion printers are described in detail. The first is a printer-independent, design phase attack that causes part sabotage. The second is a man-in-the-middle attack that targets the first printer during the slicing phase. The third is a G-code toolpath file modification attack that targets the second printer during the slicing phase. The fourth is a malware implant attack that targets the third printer during the printing phase. The case studies demonstrate the effectiveness of attack-defense modeling and its ability to help understand and bolster the cyber security postures and risk management of material extrusion additive manufacturing environments.

2. Additive Manufacturing Workflow

The ability to rapidly design and create complex parts with intricate internal structures have led to dramatic increases in the use of additive manufacturing by industries across the critical infrastructure sectors. Additive manufacturing offers environmental, socioeconomic and technical advantages compared with traditional manufacturing [24]. The advantages include speed, accuracy, efficiency and cost savings. Additive manufacturing also results in less wasted material compared with traditional manufacturing. Parts can be printed on-site and on-demand without the added financial and temporal costs of off-site production.

Design files used for additive manufacturing can be shared to allow for reliable repeatability, enabling the printing of precisely the same parts by any capable printer. A design file can be used to print a part with identical properties (shape, size, weight and internal structures) anywhere in the world. Large warehouses of additive manufacturing printers, known as “print farms,” are used to increase the number of print jobs completed simultaneously to further improve efficiency [24].

Figure 1 presents a generic additive manufacturing workflow (process chain). The workflow comprises four phases: (i) design, (ii) slicing, (iii) printing and (iv) post-processing:

- During the design phase, a 3D design of the desired part, including its shape, size, weight and other intricate details, is created using computer-aided design (CAD) software. Parts are designed for a range of uses from hobbyist toys and medical prosthetics to mission-critical components and weapons. The design details of the parts are saved in stereolithography (STL) design files. Part design files are often archived in online databases, enabling users to upload and download designs for dissemination and printing by compatible printers, respectively.

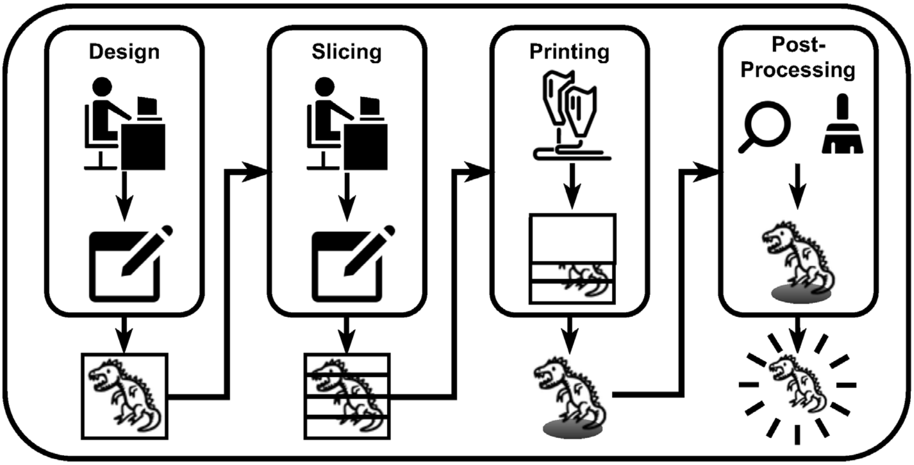

Figure 1. Additive manufacturing workflow.

- During the slicing phase, an STL design file is processed by a slicer, a type of computer-aided manufacturing (CAM) software. The slicing software divides the STL design file into segments of geometric code (G-code). Each G-code segment conveys the toolpath instructions for printing a slice of the part. Segments of the G-code file are sent directly to a printer or the entire file may be stored on removable storage media for subsequent input to a printer.

- During the printing phase, the printer firmware executes the G-code toolpath instructions to control actuator movements. The printer builds the part by depositing printing material (filaments) layer by layer according to the instructions that determine characteristics such as extruder motion, material temperature, thickness and distribution speed.

- During the post-processing phase, quality control and part finalization steps are performed, for example, to improve part strength and obtain the desired part finish. These steps are heavily dependent on the printer technology, material types and printed parts.

3. Additive Manufacturing Threats

This section discusses the two principal types of threats involving additive manufacturing systems: (i) threats that leverage additive manufacturing for malicious purposes and (ii) threats that target additive manufacturing environments.

3.1 Threats Leveraging Additive Manufacturing

Additive manufacturing can be used for nefarious purposes such as concealing illicit objects [9]. This is accomplished by pausing the printing process, inserting an illicit object inside the unfinished part and continuing the print job to hide the illicit object. Example illicit objects include explosives, illegal drugs and espionage devices.

Untraceable weapons such as "ghost guns" can be printed without serial numbers and other identifying information [9]. Digital part files for handguns and assault rifles have been available on the Internet for almost a decade. Accessories can be printed for illegal modifications to weapons. The perpetrator of the October 2019 synagogue shooting in Halle, Germany used improvised guns that incorporated 3D-printed components [4].

A novel feature of additive manufacturing is the ease with which parts can be reverse engineered to create digital part files for producing counterfeit parts. Additionally, modifications can be introduced in the reverse-engineered part files to produce hazardous items.

Biometric authentication devices scan human features such as fingerprints, handprints, retinas and faces. Additive manufacturing can be used to print high-quality spoofed fingerprints, handprints and facial features that defeat biometric authentication [9].

3.2 Threats Against Additive Manufacturing

The primary threats against additive manufacturing are intellectual property theft, part sabotage and additive manufacturing environment sabotage [24].

Researchers have theorized attacks that compromise the intellectual property of 3D-printed parts [1, 2, 8, 21]. One approach is to steal a digital part file from a control device that interfaces with a 3D printer. Another is to steal a part file directly from a printer. Yet another approach is to use a man-in-the-middle attack to steal a part file during its transfer from a control device to a printer over a network. Additionally, it is possible to scan a part and create a part file to replicate the part at will.

Sabotage attacks may target printed objects as well as print environments. Zeltmann et al. [25] discuss the potential risks and impacts of embedded defects and orientation changes on part strength. Moore et al. [17] analyzed a variety of open-source 3D printer software products.

They employed static and dynamic code analyses to reveal vulnerabilities such as buffer overflows and unencrypted communications that could be used to compromise printed parts. Additionally, they discovered weaknesses that could be exploited to manipulate G-code in toolpath files to sabotage parts.

Belikovetsky et al. [3] leveraged a phishing attack to install a backdoor on a control device. The backdoor enabled compromises of STL design files that resulted in weakened objects being printed. This attack was subsequently confirmed by Sturm et al. [22] who used malware to modify STL design files, leading to the premature failure of printed objects.

Moore et al. [18] implanted malicious code in 3D printer firmware. The modified firmware ignored incoming print commands, substituted malicious print commands and manipulated printer feed rates. The research amply demonstrated the negative impacts that malicious firmware can have on printed parts as well as on print environments.

As early as 2013, Xiao [23] demonstrated the malicious modification of a print environment. The firmware in a RepRap Prusa desktop 3D printer was changed to make the printer believe that the extruder temperature was twice as high as the actual temperature.

Pearce et al. [19] installed Trojan bootloaders in more than 100 Marlin-compatible commercial 3D printers to modify their print environments and compromise printed part integrity. The bootloaders scanned the firmware for certain byte patterns in the G-code and triggered manipulations that reduced printer extrusion rates and reordered G-code commands.

Most research in additive manufacturing has investigated weaknesses and avenues for attacks against design files and firmware. The narrow body of research involving real attacks focuses on STL design file manipulations and firmware modifications. In contrast, this research, in addition to demonstrating working attacks, presents an additive manufacturing attack-defense model that supports the discovery and exploitation of vulnerabilities in diverse material extrusion printers as well as the articulation of appropriate countermeasures.

4. Material Extrusion Additive Manufacturing

The additive manufacturing workflow differs based on printing technology, print material and other characteristics. Several types of additive manufacturing technologies have been developed, each with specific use cases, benefits and challenges. This research focuses on material extrusion additive manufacturing.

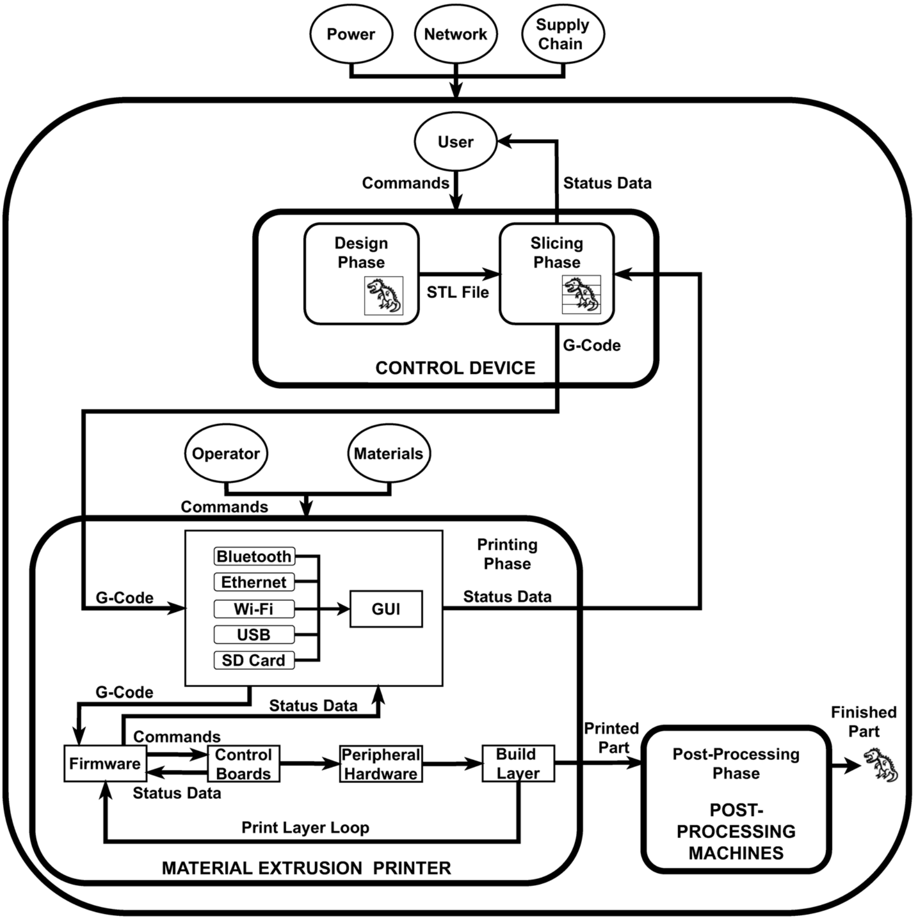

Figure 2. Material extrusion additive manufacturing workflow.

Material extrusion, also called fused deposition modeling or fused filament fabrication, is the most common additive manufacturing process [12]. The process involves heating print material and using an extruder to deposit it on a print bed layer by layer. Material is deposited along three dimensions according to the G-code instructions. Material extrusion is primarily used for printing prototypes, household items, toys, games and products with rough surface finishes.

Figure 2 shows the material extrusion additive manufacturing workflow. It comprises the four additive manufacturing phases: design, slicing, printing and post-processing. However, to provide background information and support the creation of the attack-defense model specified later, details about the four phases are only provided for material extrusion additive manufacturing.

The control device in Figure 2 is responsible for the design and slicing phases of material extrusion. The design phase inputs include electric power, network communications, supply chain components and user commands to the control device. The principal design phase output is the STL design file, which is transmitted to the slicing phase for processing by the slicer.

The slicing phase inputs include electric power, network communications, supply chain components and control device user commands, as well as the STL design file input from the design phase. Since the slicing software acts as an interface between the control device and printer, it receives print status data inputs from the printer during the printing phase. The slicing phase also outputs status data to the control device user who interacts with the slicing software.

The material extrusion printer is responsible for the printing phase. The printing phase inputs include electric power, network communications, supply chain components, printer operator commands and extruder materials. The G-code file, a key printer input, is transmitted by the slicing software remotely via Bluetooth, Ethernet or Wi-Fi, or manually by a printer operator via a USB device or SD card. The printer also receives status data from the firmware as the part is printed.

During the printing phase, the printer firmware processes the G-code file. The firmware communicates G-code toolpath instructions to the control boards, which control the peripheral hardware that prints the part layer by layer (in a loop) until all the G-code instructions are executed. The firmware sends status data as necessary to the printer.

The printing phase outputs the printed part to the post-processing phase, which may have multiple automated/manual sub-phases depending on the part and its desired properties. The post-processing phase receives inputs such as electric power, network communications, supply chain components, materials and technician/operator commands. The output of the post-processing phase, indeed the ultimate product of the material extrusion additive manufacturing workflow, is the finished part.

Post-processing operations are highly specific to the print materials and parts. Additionally, an analysis of the material extrusion process conducted in this research revealed that the overwhelming majority of cyber threats target the earlier design, slicing and printing phases. Therefore, the post-processing phase is considered to be out of scope in this research and is not described in detail in the material extrusion workflow in Figure 2.

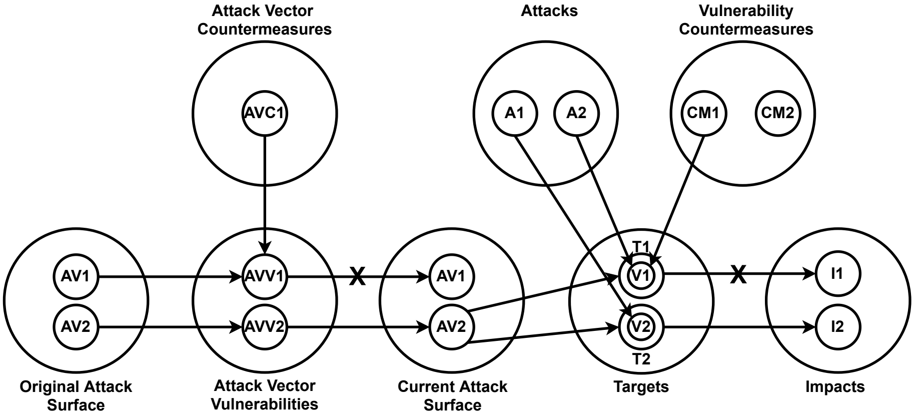

Figure 3. Attack-defense model components.

5. Attack-Defense Modeling

This section describes an attack-defense model of the security environment of a complex cyber-physical system. The model specifies key components such as attack vectors, attack vector vulnerabilities and countermeasures, attack surfaces, system targets, target vulnerabilities and vulnerability countermeasures, and attacks and attack impacts. The model directly supports three key risk management steps, risk identification, risk assessment and analysis, and risk mitigation and planning (the remaining two steps are risk allocation and risk monitoring and control) [16].

Figure 3 shows the components of an attack-defense model. An attack vector gives an adversary cyber or physical access to one or more targets in the system of interest. The collection of possible attack vectors comprises the original attack surface of the system ([AV1, AV2]).

An attack vector (AV2) is effective when it exploits an attack vector vulnerability (AVV2) to achieve the desired access. However, if an attack vector countermeasure (AVC1) is implemented to combat an attack vector vulnerability (AVV1), the associated attack vector (AV1) is ineffective. The collection of effective attack vectors comprises the current attack surface of the system ([AV2]).

An attacker can leverage effective attack vectors in the current attack surface to access targets in the system ([T1, T2]). Having gained access to a target, the attacker proceeds to launch an attack that exploits a vulnerability in the target. If appropriate countermeasures that address the target vulnerability are implemented, the attack is unsuccessful; otherwise, the attack is successful and causes negative impacts.

In Figure 3, attack A2 that exploits vulnerability V1 in target T1 is unsuccessful because vulnerability countermeasure CM1 is implemented for vulnerability V1. However, attack A1 that exploits vulnerability V2 in target T2 is successful because no countermeasures are implemented for vulnerability V2, resulting in impact I2.

The attack-defense model of a cyber-physical system is created by specifying the original attack surface and implemented attack vector countermeasures to obtain the current attack surface. Next, the system targets that can be accessed using the current attack surface are identified. Following this, a cyber-physical system workflow as in Figure 2 and the MITRE ATT&CK Knowledge Base [14] are employed to identify vulnerabilities in the targets and possible attacks. Countermeasures based on the MITRE D3FEND Knowledge Graph [15] are then identified to combat the attacks. Attacks without adequate countermeasures would be successful and their potential negative impacts on the system are specified.

Attack-defense modeling effectively conveys the security environment of a complex cyber-physical process such as material extrusion manufacturing. It clearly specifies the attack vectors that provide access to targets and the attack vector countermeasures that combat the vulnerabilities exploited by attack vectors to reduce the overall attack surface. Having identified the targets reachable by attacks, it clarifies the target vulnerabilities that could be exploited and demands that countermeasures be considered to address the vulnerabilities, defeating the attacks and reducing or eliminating the negative impacts.

6. Material Extrusion Attack-Defense Model

This section specifies a general attack-defense model for material extrusion additive manufacturing systems. The overall model includes separate models for the design, slicing and printing phases. Each model comprises the current attack surface (set of attack vectors), targets, target vulnerabilities, attacks, vulnerability countermeasures and attack impacts. Note that the current attack surface includes all the attack vectors because it is assumed that no attack vector countermeasures are implemented.

A graphical representation of an attack-defense model with circles and arrows as shown in Figure 3 offers clarity. However, in the case of the attack-defense models for the design, slicing and printing phases, the graphical representations are cumbersome because there are large numbers of vulnerabilities, attacks and vulnerability countermeasures.

Alternatively, attack-defense models may be presented as tables with columns: attack vectors, targets, target vulnerabilities, attacks, vulnerability countermeasures and impacts. The tables simplify the presentation while providing details of individual vulnerabilities, attacks and vulnerability countermeasures. A table may not provide a comprehensive description of all the vulnerabilities, attacks and vulnerability countermeasures, but it does provide significant examples to understand the security environment, including the gaps that must be filled by adding new rows to the table. Additionally, the tables are readily implemented in an automated system for presenting the security environment of an material extrusion additive manufacturing system and evaluating vulnerability countermeasures and attack impacts for various targets.

Another significant advantage of the tabular representation compared with its graphical counterpart is its ability to express one-to-many, many-to-one and many-to-many relationships involving target vulnerabilities, attacks and vulnerability countermeasures. One example is a single attack that can exploit multiple vulnerabilities on multiple targets. Another is a single target vulnerability that can be exploited by multiple attacks. Yet another example is a single vulnerability countermeasure that can be applied to address multiple target vulnerabilities.

6.1 Design Phase Attack-Defense Model

Table 1 shows an attack-defense model table created for the design phase of the material extrusion workflow. The attack vectors in the attack-defense table correspond to the four design phase inputs in the workflow, power supply, network, supply chain and user. The targets, control device, design software and STL design file, correspond to the three design phase components whose compromise can impact an STL design file, which is the output of the design phase.

A target reachable by an attack vector may have vulnerabilities. An attack exploits one or more vulnerabilities. A vulnerability countermeasure addresses one or more vulnerabilities and combats the associated attacks.

As seen in Table 1, the three principal impacts of attacks on the design phase are intellectual property theft, part sabotage and print environment sabotage. Note that power supply attacks only result in part sabotage because they prevent the STL design file from being created. In contrast, the network, supply chain and user attack vectors may ultimately result in intellectual property theft, part sabotage and print environment sabotage. Intellectual property theft occurs when an STL design file is exfiltrated. Part sabotage occurs when the 3D surface geo-

Table 1. Design phase attack-defense table.

Attack Vectors	**Targets**	**Vulnerabilities**	**Attacks**	**Vulnerability Countermeasures**	**Impacts**
Power supply	Control device	Unprotected power supply	Power shut off	Backup power supply	PS
			Power surge	Power surge protection	PS
Network	Control device	Memory access	File modification	Access control	IPT, PS, PES
		Root access	File theft	Access control	IPT
		Open ports	File theft	Port security	IPT
	Design software	Software access	Malware implant	Integrity checking	IPT, PS, PES
	STL design file	No STL file integrity checking	STL file modification	Integrity checking	IPT, PS, PES
			STL file replacement	File hashing	IPT, PS, PES
		Memory access	STL file modification	Access control	IPT, PS, PES
Supply chain	Control device	Operating system access	Malware implant	Integrity checking	IPT, PS, PES
		Firmware access	Malware implant	Integrity checking	IPT, PS, PES
		Network access	Malware implant	Integrity checking	IPT, PS, PES
		Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
	Design software	Software access	Malware implant	Integrity checking	IPT, PS, PES
	STL design file	Vendor USB drive	Malware implant	USB port security	IPT, PS, PES
			Malicious STL file	USB port security	IPT, PS, PES
User	Control device	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
			Erroneous use	User training	IPT, PS, PES
			Memory modification	Quality control	IPT, PS, PES
			File theft	Access control	IPT
	Design software	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Erroneous STL file	User training	IPT, PS, PES
			Malicious STL file	Access control	IPT, PS, PES
				Quality control	IPT, PS, PES
	STL design file	Physical access	Erroneous STL file	User training	IPT, PS, PES
			Malicious STL file	Access control	IPT, PS, PES
				Quality control	IPT, PS, PES

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

metry encoded in an STL design file is manipulated. Print environment sabotage occurs (for example) when malware is incorporated in an STL design file to target slicing software, causing it to incorporate malicious G-code instructions that impact the print environment.

6.2 Slicing Phase Attack-Defense Model

Table 2 shows an attack-defense model table created for the slicing phase of the material extrusion workflow. The attack vectors in the attack-defense table correspond to the four slicing phase inputs in the material extrusion workflow, power supply, network, supply chain and user. The targets, control device, slicing software and G-code file, cor-

Table 2. Slicing phase attack-defense table.

Attack Vectors	Targets	Vulnerabilities	Attacks	Vulnerability Countermeasures	Impacts
Power supply	Control device	Unprotected power supply	Power shut off	Backup power supply	PS, PES
			Power surge	Power surge protection	PS, PES
Network	Control device	Memory access	File modification	Access control	IPT, PS, PES
		Root access	File theft	Access control	IPT
		Open ports	File theft	Port security	IPT
	Slicing software	Software access	Malware implant	Integrity checking	IPT, PS, PES
		No printer authentication	Man-in-the-middle	Printer authentication	IPT, PS, PES
			ARP spoofing	Network authentication	IPT, PS, PES
		Print queue access	Print queue modification	Print queue access control	IPT, PS, PES
	G-code file	No G-code file integrity checking	G-code file modification	Integrity checking	IPT, PS, PES
			G-code file replacement	File hashing	IPT, PS, PES
		Memory access	G-code file modification	Access control	IPT, PS, PES
Supply chain	Control device	Operating system access	Malware implant	Integrity checking	IPT, PS, PES
		Firmware access	Malware implant	Integrity checking	IPT, PS, PES
		Network access	Malware implant	Integrity checking	IPT, PS, PES
		Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
	Slicing software	Software access	Malware implant	Integrity checking	IPT, PS, PES
	G-code file	Vendor USB drive	Malware implant	USB port security	IPT, PS, PES
			Malicious G-code file	USB port security	IPT, PS, PES
User	Control device	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
			Erroneous use	User training	IPT, PS, PES
			Memory modification	Quality control	IPT, PS, PES
			File theft	Access control	IPT
	Slicing software	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Erroneous G-code file	User training	IPT, PS, PES
			Malicious G-code file	Access control	IPT, PS, PES
				Quality control	IPT, PS, PES
	G-code file	Physical access	Erroneous G-code file	User training	IPT, PS, PES
			Malicious G-code file	Access control	IPT, PS, PES
				Quality control	IPT, PS, PES

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

respond to the three slicing phase components whose compromise can impact the G-code file, which is the output of the slicing phase.

As shown in Table 2, the three principal impacts of attacks on the slicing phase are intellectual property theft, part sabotage and print environment sabotage. Power supply attacks result in part sabotage and print environment sabotage due to the dependence of the slicing software on the control device. Attacks leveraging the access provided by the network, supply chain and user attack vectors result in intellectual property theft, part sabotage and print environment sabotage. Intellectual property theft occurs when a G-code or control device file is exfiltrated. Part sabotage occurs when a G-code file is modified to alter the toolpath, which modifies the printed part. Print environment sabotage occurs when the environment is disturbed by modifying a G-code file or by directly interacting with the printer.

Attacks during the slicing phase that modify a G-code toolpath file are a concern because G-code determines the toolpath and print environment variables such as temperature and fan speed. Alterations to a G-code file can result in part sabotage and print environment sabotage regardless of the intent of the alteration. The direct connection between the slicing software and a printer provides an avenue for accessing the print environment. Attacks against the direct connection between the slicing software and a printer can result in the exploitation of several vulnerabilities.

6.3 Printing Phase Attack-Defense Model

Table 3 shows the attack-defense model table for the printing phase of the material extrusion workflow. The attack vectors in the attack-defense table correspond to the four design phase inputs in the material extrusion workflow, power supply, network, supply chain and operator. As seen in the table, the targets vary based on the attack vector. The potential targets include the control device, printer, printer firmware, peripheral hardware, print layer, control boards and printer material. Each target represents a printing phase component whose compromise can impact the printing process and print environment, which collectively produce the final printed part.

Table 3 shows the three principal impacts of attacks on the printing phase, intellectual property theft, part sabotage and print environment sabotage. The most concerning impacts of successful attacks against the printing phase are part sabotage and print environment sabotage. Attacks against the power supply can be used to target vulnerabilities in the control device and printer. The impacts of successful attacks against the power supply are part sabotage and print environment sabotage.

The network attack vector may provide an attacker with access to targets such as a printer, printer firmware, peripheral hardware and print layer. Communications between slicing software and a printer are commonly unencrypted, and therefore, subject to eavesdropping, interference and malicious modification of G-code and status data in transit unless strong access controls are implemented.

The supply chain attack vector may enable an attacker to access a printer, printer firmware, control boards, peripheral hardware, print layer and printer material. Physical access to a printer via the supply chain provides opportunities to implant malware and parasitic devices. A malware implant may involve malicious modifications to printer

Table 3. Printing phase attack-defense table.

Attack Vectors	**Targets**	**Vulnerabilities**	**Attacks**	**Vulnerability Countermeasures**	**Impacts**
Power supply	Control device	Unprotected power supply	Power shut off	Backup power supply	PS, PES
			Power surge	Power surge protection	PS, PES
	Printer	Unprotected power supply	Power shut off	Backup power supply	PS, PES
			Power surge	Power surge protection	PS, PES
Network	Printer	No control device authentication	HTTP packet spoofing	Control device authentication	IPT, PS, PES
	Printer firmware	No integrity checking	Firmware implant	Integrity checking	IPT, PS, PES
		Remote update access	Firmware implant	Integrity checking	IPT, PS, PES
		Firmware access	Firmware implant	Integrity checking	IPT, PS, PES
	Peripheral hardware	No access control	Data modification	Access control	IPT, PS, PES
	Print layer	No access control	G-code layer theft	Access control	IPT
			G-code modification	Access control	IPT, PS, PES
Supply chain	Printer	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
	Printer firmware	No integrity checking	Firmware implant	Integrity checking	IPT, PS, PES
		Firmware access	Firmware implant	Integrity checking	IPT, PS, PES
	Control boards	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
	Peripheral hardware	Physical access	Faulty hardware implant	Physical inspection	IPT, PS, PES
	Printer material	Physical access	Faulty material	Quality control	IPT, PS, PES
Operator	Printer	Physical access	Malware implant	Integrity checking	IPT, PS, PES
			Parasitic device implant	Physical inspection	IPT, PS, PES
	Printer firmware	Firmware access	Firmware implant	Integrity checking	IPT, PS, PES
		Physical access	Erroneous use	Operator training	IPT, PS, PES
			Malicious use	Access control	IPT, PS, PES
	Control boards	Physical access	Erroneous use	Operator training	IPT, PS, PES
			Malicious use	Access control	IPT, PS, PES
	Peripheral hardware	Physical access	Erroneous use	Operator training	IPT, PS, PES
			Malicious use	Access control	IPT, PS, PES
	Print layer	Physical access	Erroneous use	Operator training	IPT, PS, PES
			Malicious use	Access control	IPT, PS, PES
	Printer material	Physical access	Erroneous use	Operator training	IPT, PS, PES
			Malicious use	Access control	IPT, PS, PES

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

firmware that could alter printer functionality, thereby sabotaging print jobs and the print environment.

The operator attack vector enables an attacker to access a printer, printer firmware, control boards, peripheral hardware, print layer and printer material. Operators often have unfettered access to the targets, providing opportunities to erroneously or maliciously interfere with printed parts and the print environment.

7. Material Extrusion Case Studies

This section describes the three material extrusion printers used as case studies to demonstrate the effectiveness of the attack-defense model and help understand the cyber security and risk management postures of material extrusion additive manufacturing environments. For security reasons, certain details about the printers and their environments are obfuscated.

7.1 Printer Annamieke

Printer Annamieke is a proprietary material extrusion printer. The printer facilitates efficient and durable printing with plastic and metallicized-plastic materials. An Annamieke printer has a unique device name and serial number, neither of which can be changed.

Printer Annamieke is typically used in industry because of its size and $25,000 price. The printer comes equipped with proprietary software for slicing STL design files to produce G-code toolpath files, and for interfacing between the control device and printer.

A vulnerability in the network discovery process utilized by the proprietary interface software and printer was exploited to obtain a man-in-the-middle position. Attacks that exploit the vulnerability result in intellectual property theft, part sabotage and print environment sabotage.

7.2 Printer Beatrijs

Printer Beatrijs is an industrial fused filament fabrication printer. It uses a dual extruder and a partially-enclosed environment to print parts using a variety of materials, including plastics, wood and stainless steel. The printer costs approximately $4,000 and is commonly used in laboratory environments.

Printer Beatrijs is equipped with open-source slicing and printer interface software that allows for reliable and persistent access. A vulnerability discovered in the open-source slicing and interface software enables the unauthorized modification of G-code toolpath files in control device memory [11]. Modifications to the G-code toolpath file during the slicing phase, before it is sent to the printer, can result in part sabotage and print environment sabotage.

7.3 Printer Cathelijne

Printer Cathelijne is a fused filament fabrication printer with a single extruder capable of printing with plastic or wood filament. The printer has an open design and comes with a removable print bed, built-in filament tray, patent extruder, touch screen and multiple communications modes, including USB cable, USB drive, Ethernet and Wi-Fi.

Printer Cathelijne is primarily used by educators and hobbyists due to its low $300 price and ease of use. It has proprietary slicing and interface software and printer firmware.

Printer Cathelijne is vulnerable to several attacks, including firmware modification, remote code execution and malware implants.

8. Material Extrusion Attacks

Attack-defense models for the design, slicing and printing phases were created for the three printers in the case study. The attack-defense models comprise the attack vectors, targets, target vulnerabilities, attacks, vulnerability countermeasures and attack impacts. Since the models were developed from a common process workflow, the attack vectors, targets and attack impacts are common to all three printers. However, differences exist in the target vulnerabilities between printers due to differences in printer designs, features and implementations. As a result, the target vulnerabilities, attacks and vulnerability countermeasures in the attack-defense models vary from printer to printer.

This section describes four exemplar attacks in the printer attack-defense models. The first exemplar attack is a design phase attack on a control device that is printer-independent. The second and third exemplar attacks, which focus on the slicing phase, are unique to printers Annamieke and Beatrijs, respectively. The fourth exemplar attack targets the printing phase of printer Cathelijne.

8.1 Design Phase Attack

During the design phase, a 3D rendering of a part is created using computer-aided design software running on a control device that is independent of the eventual printer. The 3D rendering of the part is saved on the control device as an STL design file. Vulnerabilities in the control device that creates and/or stores the STL design file can enable attacks on the STL design file, which is the output of the design phase.

A classic attack is to modify an STL design file to sabotage the resulting printed parts. The attack, first demonstrated by Belikovetsky et al. [3], leveraged general infiltration methods to target a control device hosting an STL design file. Access to the control device (target) was gained using a phishing attack (network attack vector) and the data alteration attack exploited a ZIP file vulnerability (target vulnerability). Specifically, the STL design file was modified to introduce a void in

Table 4. STL design file attack.

Attack Vectors	Targets	Vulnerabilities	Attacks	Vulnerability Countermeasures	Impacts
. . .	. . .	. . .	. . .	. . .	. . .
Network	. . .	. . .	. . .	. . .	. . .
	STL design file	. . .	. . .	. . .	. . .
		Memory access	STL file modification	Access control	IPT, PS, PES
		. . .	. . .	. . .	. . .
	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

the part as it was printed (attack), causing a time-delayed part failure (attack impact).

Attacks on an STL design file during the design phase can be executed independently of a printer. Table 4 shows a portion of the design phase attack-defense model corresponding to the STL design file attack. The impacts of the STL design file attack include part sabotage as well as intellectual property theft (theft of the STL design file) and print environment sabotage (malware implant in the STL design file). Table 4 also shows that an access control countermeasure can address the vulnerability and counter the STL design file attack, eliminating the negative impacts.

8.2 Slicing Phase Attacks

The slicing software hosted on a control device transforms an STL design file to a G-code toolpath file for eventual printing. The slicing software may also act as an interface between the control device and a compatible printer. Penetration tests revealed that the slicing software systems designed for the Annamieke and Beatrijs printers had vulnerabilities that could be exploited during the slicing phase to cause intellectual property theft, part sabotage and print environment sabotage.

Printer Annamieke Man-in-the-Middle Attack. A control device may connect to an Annamieke material extrusion printer via a Wi-Fi or Ethernet link, or directly via a cable. However, dynamic or static connections via a Wi-Fi network are most common. If the network is configured for the Dynamic Host Configuration Protocol (DHCP), an IP address is automatically assigned to the printer. Otherwise, a user

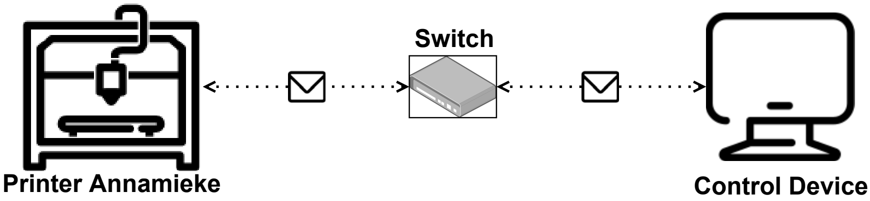

Figure 4. Legitimate connection between the slicing software and printer Annamieke.

may manually enter a static IP address in the printer user interface and enter the same IP address in the proprietary printer interface software.

The control device executes slicing/interface software developed for an Annamieke printer. The software searches the network for a compatible printer and establishes a connection if one is discovered. The software then slices the STL design file to create a printer-compatible G-code toolpath file. The software may be used to view, resize or place a 3D rendering on the print bed. Additionally, the software provides data about the connected printer, including its name, material status, print status, print history and current print job data.

The exemplar attack developed for the Annamieke printer slicing phase leverages local network connectivity as the attack vector to target the slicing software. Analysis of the material extrusion workflow using the MITRE ATT&CK Knowledge Base led to the discovery of a vulnerability in how the slicing software establishes a connection with the Annamieke printer. Specifically, the software and printer Annamieke use plaintext HTTP communications without authentication to establish their connection. This vulnerability is exploited to obtain a man-in-the-middle position before or after the connection between the slicing software and printer Annamieke is established.

During its execution, the slicing software spawns a network process that sends Simple Service Discovery Protocol (SSDP) multicast messages in the local network looking for compatible printers. Upon receiving the message, printer Annamieke sends a plaintext HTTP response containing its printer name, serial number and IP address. The software stores the data received from printer Annamieke and proceeds to establish a connection as shown in Figure 4. After the connection is established, printer Annamieke sends the slicing software status reports about the printer material status and extruder location and temperature. Additionally, the slicing software sends user commands to and requests status data from printer Annamieke.

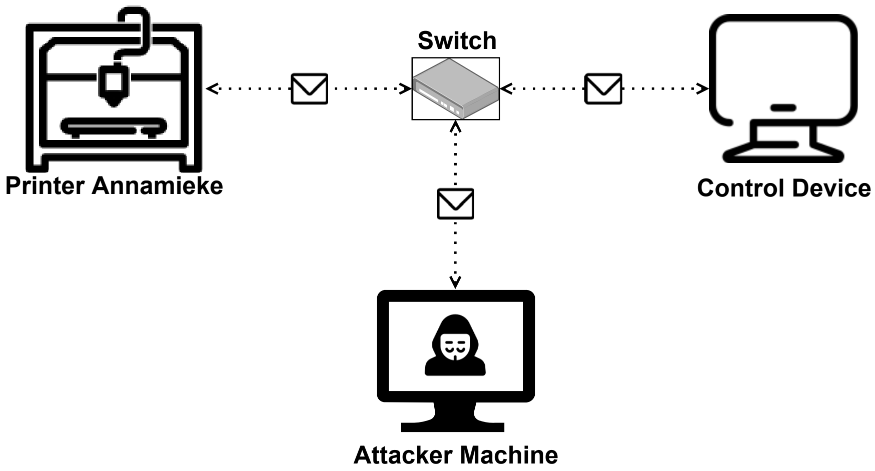

Figure 5. Printer Annamieke man-in-the-middle attack position.

An attacker can leverage the lack of encryption and authentication to assume a man-in-the-middle position between the slicing software and printer Annamieke. This is accomplished by actively monitoring the local network traffic from an attacker-controlled machine for SSDP multicast messages sent by the slicing software to find a compatible printer. Upon detecting an SSDP message and the response from printer Annamieke, a spoofed response is created by the attacker claiming to be printer Annamieke (using the unique identifiers in the Annamieke response packet), but replacing the legitimate IP address with the IP address of the attacker-controlled machine. Figure 5 shows the attacker-controlled machine after it has assumed a man-in-the-middle position during initial session establishment.

Note that the attacker can also assume a man-in-the-middle position after a legitimate connection is established between the slicing software and printer Annamieke. This is because the slicing software allows dynamic updates to IP addresses and only requires the printer name, unique identification number and IP address to update the connection, all of which can be captured from network traffic. The attacker then hijacks the legitimate connection between the slicing software and printer Annamieke by sending a packet to replicate an Annamieke IP address update. Finally, the attacker maintains persistence by establishing network traffic forwarding rules on the attacker-controlled machine to ensure that all communications are forwarded through the attacker-controlled machine.

Table 5. Printer Annamieke man-in-the-middle network attack.

Attack Vectors	Targets	Vulnerabilities	Attacks	Vulnerability Countermeasures	Impacts
. . .	. . .	. . .	. . .	. . .	. . .
Network	. . .	. . .	. . .	. . .	. . .
	Slicing software	. . .	. . .	. . .	. . .
		No printer authentication	Man-in-the-middle	Printer authentication	IPT, PS, PES
		. . .	. . .	. . .	. . .
	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

Table 5 shows a portion of the slicing phase attack-defense model corresponding to the printer Annamieke man-in-the-middle network attack. Intellectual property theft is perpetrated by copying the G-code toolpath file from the man-in-the-middle position. Part sabotage and print environment sabotage are accomplished by modifying the G-code file during its transmission. Table 5 also shows that a device authentication countermeasure can counter the printer Annamieke man-in-the-middle network attack, eliminating the negative impacts.

Printer Beatrijs G-Code File Modification Attack. Printer Beatrijs uses open-source slicing/interface software that is employed by many other additive manufacturing printers. Analysis of the material extrusion workflow using the MITRE ATT&CK Knowledge Base led to the discovery of a vulnerability in the open-source slicing software [11].

Specifically, after the G-code is generated by the slicing software, but before it is saved on the control device hosting the slicing software, the entire G-code toolpath file is stored unencrypted as ASCII characters in the heap memory of the control device. Root access to the control device enables the ASCII representation of the G-code in heap memory to be modified while the user views the 3D rendering of the G-code using the slicing software.

A tool was created to locate and extract the ASCII G-code in heap memory, and reconstruct the G-code toolpath layers in ascending order by layer number [11]. The tool also facilitates surreptitious alterations of the G-code such as excluding infill from certain layers and reducing the extruder temperature when certain layers are printed. When the user saves the G-code toolpath file to the control device, the modified

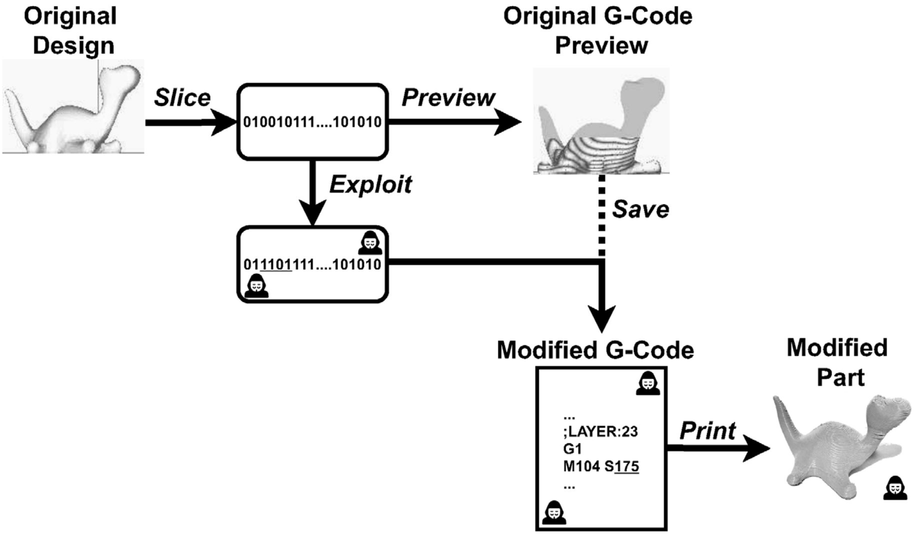

Figure 6. Printer Beatrijs G-code file modification attack workflow.

G-code in heap memory is saved instead of the original version. Figure 6 shows the printer Beatrijs G-code file modification attack workflow.

Table 6. Printer Beatrijs G-code file modification attack.

Attack Vectors	**Targets**	**Vulnerabilities**	**Attacks**	**Vulnerability Countermeasures**	**Impacts**
. . .	. . .	. . .	. . .	. . .	. . .
Network	. . .	. . .	. . .	. . .	. . .
	G-code file	. . .	. . .	. . .	. . .
		Memory access	G-code file modification	Access control	IPT, PS, PES
		. . .	. . .	. . .	. . .
	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

Table 6 shows a portion of the slicing phase attack-defense model corresponding to the G-code file modification attack against the Beatrijs printer. Intellectual property theft is perpetrated by copying the ASCII version of the G-code toolpath file from heap memory. Part sabotage and print environment sabotage are accomplished by modifying the G-code file in heap memory.

Experiments revealed that G-code modifications that cause infill to be excluded from certain layers and the extruder temperature to be reduced while printing certain layers have significant ramifications [11].

In the infill exclusion experiments, excluding infill from just five of the 127 total layers in printed plastic cylinders yielded a 10.6% reduction in the average failure force under compression. Excluding infill from 25 of the 127 layers yielded a 19.9% reduction in the average failure force under compression. In both instances, the mass reductions were negligible (within standard error) and no perceptible differences were visible between the original and modified cylinders.

In the temperature reduction experiments, G-code was modified to reduce the extruder temperature slightly (from the normal 198°C to the new 190°C) when just seven centrally-located layers of the 530 total layers of plastic parts were printed. No perceptible differences were visible between the original and modified parts. However, the average breaking force under tensile testing dropped by 14% for the modified parts.

Table 6 shows that a G-code file hashing vulnerability countermeasure can counter the printer Beatrijs G-code file modification attack, eliminating the negative impacts.

8.3 Printing Phase Attack

The control board of a Cathelijne printer has a debug port for analyzing printer activity and errors. The supply chain attack vector enables physical access to a Cathelijne printer control board (target). Physical access to the control board is a critical vulnerability that enables malware to be implanted using the debug port. Note that physical access could also be leveraged to attack the printer after it is operational at the print facility. Table 7 shows a portion of the printing phase attack-defense model corresponding to the printer Cathelijne malware implant attack.

The printer Cathelijne malware implant attack requires a debug port connection. Jumper cable connections are made between the RX, TX and ground pins of the debug port and an FTDI Basic chip connected to a laptop running a Windows operating system. The FTDI Basic chip converts the serial communications from the debug port to the USB protocol, enabling them to be monitored via a `PuTTY` application that offers a terminal emulator, serial connection and network file transfers to the Windows laptop.

Table 7. Printer Cathelijne malware implant attack.

Attack Vectors	**Targets**	**Vulnerabilities**	**Attacks**	**Vulnerability Countermeasures**	**Impacts**
. . .	. . .	. . .	. . .	. . .	. . .
Supply chain	. . .	. . .	. . .	. . .	. . .
	Printer	. . .	. . .	. . .	. . .
		Physical access	Malware implant	Integrity checking	IPT, PS, PES
		. . .	. . .	. . .	. . .
	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .

IPT: Intellectual property theft, PS: Part sabotage, PES: Print environment sabotage

The next steps in the attack are to interrupt the boot process and modify the boot settings to execute a shell instead of the Linux operating system. After the `PuTTY` terminal displays the communications from the debug port, the printer boot process is interrupted by depressing the escape key repeatedly. Following this, the new boot environment variables are set by issuing the following two commands in sequence:

-
  ```
  setenv bootargs ‘noinitrd root=/dev/mmcblk0p2
  rootfstype=ext4 init=/bin/sh/ rootwait console=ttyS0,
  115200n8’
  ```

-
  ```
  saveenv
  ```

The `printenv` command is executed to confirm that the environment variables have been set.

Printer Cathelijne is restarted after confirming that the environment variable has been changed to execute a shell at bootup. When the terminal prompts for a username and password, the default credentials provided in the Cathelijne printer manual are entered, enabling access to the printer filesystem. The filesystem access enables any file to be moved to a USB drive plugged into the printer. In this case, the `/etc/shadow` file is moved to the USB drive and a forensic tool is used to decrypt the file to obtain root credentials.

Next, a USB drive is used to implant malware in printer Cathelijne. This is accomplished by navigating to the `/media` directory and copying the malware file to the `/bin/obfuscated` directory. Leveraging root access via the operating system shell, a startup script is modified to connect the printer to the Wi-Fi network and execute the malware as

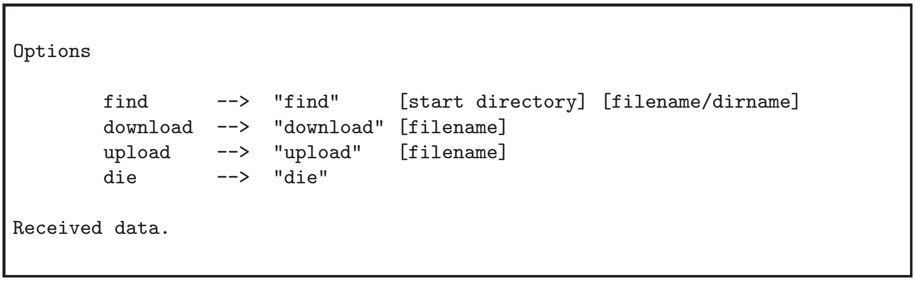

```
Options

        find      -->  "find"     [start directory] [filename/dirname]
        download  -->  "download" [filename]
        upload    -->  "upload"   [filename]
        die       -->  "die"

Received data.
```

Figure 7. Malware command terminal options.

a persistent background process. Printer Cathelijne is then rebooted to launch the malware. The malware executes whenever printer Cathelijne boots up.

The executing malware establishes a client-server connection between the printer and a remote attacker-controlled device. When the attacker device executes the client code, a user interface with malware command options is presented. Figure 7 shows the four malware command options, `find`, `download`, `upload` and `die`.

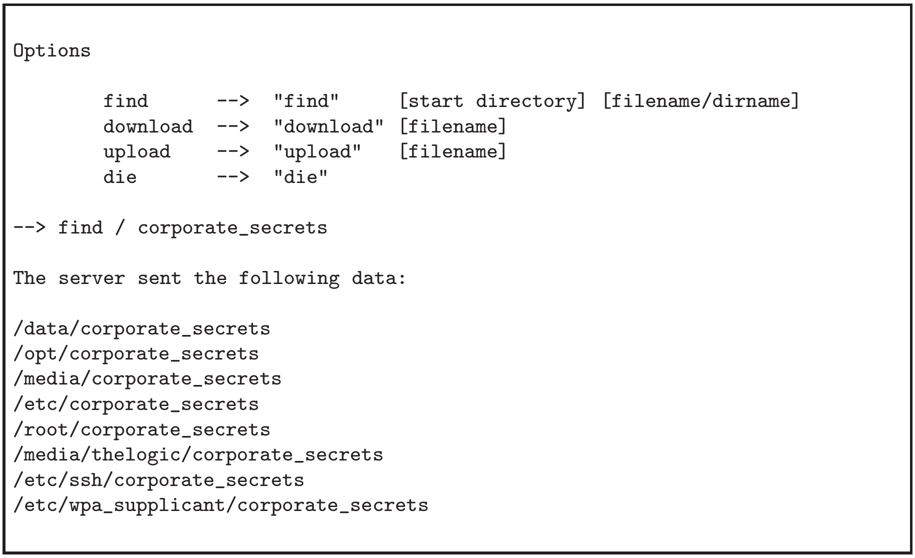

```
Options

        find      -->  "find"     [start directory] [filename/dirname]
        download  -->  "download" [filename]
        upload    -->  "upload"   [filename]
        die       -->  "die"

--> find / corporate_secrets

The server sent the following data:

/data/corporate_secrets
/opt/corporate_secrets
/media/corporate_secrets
/etc/corporate_secrets
/root/corporate_secrets
/media/thelogic/corporate_secrets
/etc/ssh/corporate_secrets
/etc/wpa_supplicant/corporate_secrets
```

Figure 8. Malware `find` command execution results.

The `find` command searches through directories for filenames. Figure 8 shows the execution results of a `find` command that searches the Cathelijne printer filesystem for directory names and/or filenames con-

```
Options

        find      -->  "find"     [start directory] [filename/dirname]
        download  -->  "download" [filename]
        upload    -->  "upload"   [filename]
        die       -->  "die"

--> download /data/corporate_secrets as stolen_data

Received data.
```

Figure 9. Malware `download` command execution results.

taining `corporate_secrets`. Options are provided to prune directory paths in the file-search tree to shorten the search time.

Intellectual property theft is perpetrated using the `download` command to transfer files from the printer to the remote attacker-controlled device. Figure 9 shows the downloading of the `corporate_secrets` file discovered using the `find` command. The downloaded file is given the name `stolen_data`.

The `upload` command enables files to be moved to the printer filesystem. The files may include G-code files to print sabotaged parts, firmware files to sabotage the print environment and malware files with sophisticated functionality.

```
Options

        find      -->  "find"     [start directory] [filename/dirname]
        download  -->  "download" [filename]
        upload    -->  "upload"   [filename]
        die       -->  "die"

--> upload /etc/altered_firmware_file

File written to path on server.
```

Figure 10. Malware `upload` command execution results.

Figure 10 shows `/etc/altered_firmware_file` being uploaded from the attacker-controlled device to the working directory of the printer. It was observed that an uploaded file overwrites an existing file with the same name in the printer directory. This feature can be exploited to overwrite the Wi-Fi configuration files in the `/etc/wpa_suplicant`

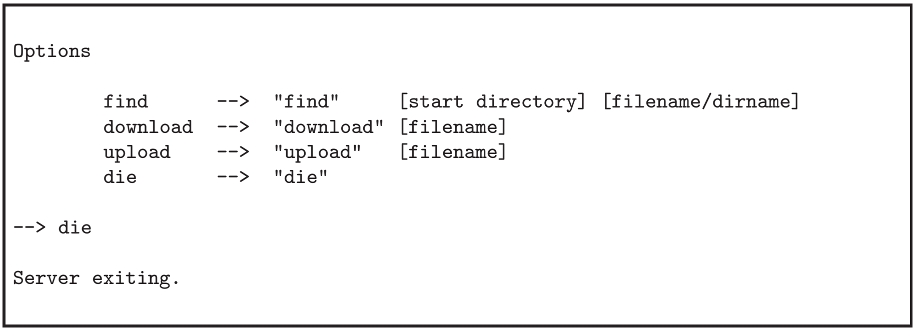

```
Options

        find      -->  "find"     [start directory] [filename/dirname]
        download  -->  "download" [filename]
        upload    -->  "upload"   [filename]
        die       -->  "die"

--> die

Server exiting.
```

Figure 11. Malware `die` command execution results.

directory or any other system configuration files. As a result, any number of file manipulations and malware updates could be performed to alter the physical, storage and network behavior of the printer.

The `die` command halts malware execution until the printer is rebooted. Figure 11 shows the `die` command execution results. The command to halt execution enables the malware to remain dormant for an extended period of time to prevent the discovery of an open network port on the Cathelijne printer. The malware is reactivated automatically when the printer is rebooted.

The malware can be deployed on any Linux kernel running on an ARM or x86 architecture, which enables it to target a variety of printers. The case study demonstrates how a supply chain attack vector and physical access vulnerability can enable malware that causes intellectual property theft, part sabotage and print environment sabotage to be implanted.

Table 7 shows that integrity checking can address the physical access vulnerability and counter the printer Cathelijne malware implant attack, eliminating the negative impacts.

9. Discussion

Material extrusion additive manufacturing is a complex cyber-physical process system. Attempting to secure the process system in a robust and (ideally) comprehensive manner requires a holistic perspective provided by a workflow that describes the operational phases, their systems and subsystems, and inputs and outputs. In the case of material extrusion additive manufacturing, separate workflows were created for its three principal phases, design, slicing and printing. Based on their workflows, separate attack-defense models were constructed for the three material extrusion phases.

Each attack-defense model comprises a set of attack vectors, targets, target vulnerabilities, attacks, vulnerability countermeasures and attack impacts. The specification of the attack surface is the first step in developing an attack-defense model. The attack surface is the collection of attack vectors that provide cyber or physical access to targets. The attack vectors and targets are clearly discernible in the process workflow. At this juncture, the vulnerabilities exploited by the attack vectors must be identified and the countermeasures that would address the vulnerabilities and combat the associated attack vectors must be specified. Attack vectors for which no countermeasures are implemented constitute the current attack surface, which provides insights into the accessible targets and types of access.

An attack framework is employed to identify target vulnerabilities and devise potential attacks that exploit the vulnerabilities. Simultaneously, a defense framework is used to identify countermeasures that combat the attacks by addressing the vulnerabilities they exploit. Attacks without adequate countermeasures would be successful and their potential negative impacts on the system are specified.

Attack-defense models are often represented graphically, but the graphical models developed for the design, slicing and printing phases were cumbersome. Alternative representations of the attack-defense models as tables with attack vectors, targets, target vulnerabilities, attacks, vulnerability countermeasures and impacts columns proved to be superior. The tables simplify the presentation while providing details about vulnerabilities, attacks and vulnerability countermeasures.

Considerable effort was invested in creating the attack-defense model tables for the design, slicing and printing phases of material extrusion additive manufacturing. The tables are large and detailed, but they are certainly not comprehensive specifications of the vulnerabilities, attacks and vulnerability countermeasures. What is important is that they provide adequate examples to understand the security environments and the gaps in the security analysis that must be filled by adding new rows to the tables.

Finally, the three attack-defense model tables provided deep insights that contributed immensely to the vulnerability discovery, exploit development and countermeasure identification efforts in this research on material extrusion additive manufacturing. Vulnerability discovery, exploit development and countermeasure identification are essential to security analyses of cyber-physical systems. In this light, a construct, such as the attack-defense model, that advances vulnerability discovery, exploit development and countermeasure identification, has considerable value.

10. Conclusions

Additive manufacturing systems, which produce mission-critical parts used in the critical infrastructure, are exposed to cyber threats that perpetrate intellectual property theft, part sabotage and print environment sabotage. Research on additive manufacturing threats has tended to focus on specific vulnerabilities and specific attacks against specific systems. The narrow scope hinders the overall understanding of the attack surfaces and targets, causing vulnerabilities, potential attacks and countermeasures being overlooked during security analyses.

This research addresses the limitations in the context of material extrusion additive manufacturing, the most common additive manufacturing process. A material extrusion workflow that comprehensively covers the design, slicing and printing phases is specified. Analysis of the workflow in conjunction with attack and defense frameworks (MITRE ATT&CK Knowledge Base and MITRE D3FEND Knowledge Graph) yield detailed attack-defense models for the design, slicing and printing phases of material extrusion systems. The attack-defense models, which specify the attack vectors, attack vector vulnerabilities and countermeasures, attack surfaces, system targets, target vulnerabilities and vulnerability countermeasures, and attacks and attack impacts, directly support risk identification, risk assessment and analysis, and risk mitigation and planning. Although the attack-defense models are very detailed, they do not specify all the target vulnerabilities, attacks and vulnerability countermeasures. However, they provide adequate examples to understand the threat environment and security posture, and the gaps that must be filled to make the models more comprehensive.

The case studies involving three material extrusion printers ranging from a $300 hobbyist device to a $25,000 industrial system demonstrate the effectiveness of attack-defense modeling at advancing vulnerability discovery, exploit development and countermeasure identification as well as its ability to clarify and bolster the cyber security and risk management postures of material extrusion additive manufacturing environments.

Future research will focus on vulnerability discovery, exploit development and countermeasure identification for a larger subset of additive manufacturing systems. It will also develop workflows and attack-defense models for the remaining six standard additive manufacturing processes.

Acknowledgement

This research was supported by the National Science Foundation under Grant no. DGE 1501177 and by UT-Battelle under Contract no. DE-AC05-00OR22725 with the U.S. Department of Energy.

References

[1] M. Al Faruque, S. Chhetri, A. Canedo and J. Wan, Acoustic side-channel attacks on additive manufacturing systems, *Proceedings of the Seventh ACM/IEEE International Conference on Cyber-Physical Systems*, 2016.

[2] M. Al Faruque, S. Chhetri, S. Faezi and A. Canedo, Forensics of Thermal Side Channels in Additive Manufacturing Systems, CECS Technical Report #16-01, Center for Embedded and Cyber-Physical Systems, University of California, Irvine, Irvine, California, 2016.

[3] S. Belikovetsky, M. Yampolskiy, J. Toh, J. Gatlin and Y. Elovici, dr0wned – Cyber-physical attack with additive manufacturing, presented at the *Eleventh USENIX Workshop on Offensive Technologies*, 2017.

[4] M. Berger, The attack on a German synagogue highlights the threat posed by do-it-yourself guns, *The Washington Post*, October 11, 2019.

[5] S. Bridges, K. Keiser, N. Sissom and S. Graves, Cyber security for additive manufacturing, *Proceedings of the Tenth Annual Cyber and Information Security Research Conference*, article no. 14, 2015.

[6] A. Damani, The fundamentals and impact of Industry 4.0, *Forbes*, June 24, 2020.

[7] Q. Do, B. Martini and K. Choo, A data exfiltration and remote exploitation attack on consumer 3D printers, *IEEE Transactions on Information Forensics and Security*, vol. 11(10), pp. 2174–2186, 2016.

[8] A. Hojjati, A. Adhikari, K. Struckmann, E. Chou, T. Nguyen, K. Madan, M. Winslett, C. Gunter and W. King, Leave your phone at the door: Side channels that reveal factory floor secrets, *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pp. 883–894, 2016.

[9] Homeland Security Advisory Council, Final Report of the Emerging Technologies Subcommittee 3D-Printing, U.S. Department of Homeland Security, Washington, DC, 2020.

[10] Hybrid Manufacturing Technologies, Seven Families of Additive Manufacturing (According to ASTM F2792 Standards), McKinney, Texas (www.additivemanufacturing.media/cdn/cms/7_families_print_version.pdf), 2021.

[11] E. Kurkowski, A. Van Stockum, J. Dawson, C. Taylor, T. Schulz and S. Shenoi, Manipulation of G-code toolpath files in 3D printers: Attacks and mitigations, in *Critical Infrastructure Protection XVI*, J. Staggs and S. Shenoi (Eds.), Springer, Cham, Switzerland, pp. 155–174, 2022.

[12] Manufactur3D Magazine, The seven types of additive manufacturing technologies, Thane, India (manufactur3dmag.com/7-types-additive-manufacturing-technologies), April 6, 2018.

[13] T. McCue, Additive manufacturing industry grows to almost $12 billion in 2019, *Forbes*, May 8, 2020.

[14] MITRE Corporation, ATT&CK for Industrial Control Systems, Bedford, Massachusetts (collaborate.mitre.org/attackics/index.php/Main_Page), 2021.

[15] MITRE Corporation, D3FEND: A Knowledge Graph of Cybersecurity Countermeasures, Bedford, Massachusetts (d3fend.mitre.org), 2021.

[16] K. Molenaar, S. Anderson and C. Schexnayder, Guidebook on Risk Analysis Tools and Management Practices to Control Transportation Project Costs, NCHRP Report 658, The National Academies Press, Washington, DC, 2010.

[17] S. Moore, P. Armstrong, T. McDonald and M. Yampolskiy, Vulnerability analysis of desktop 3D printer software, *Proceedings of the 2016 Resilience Week*, pp. 46–51, 2016.

[18] S. Moore, W. Glisson and M. Yampolskiy, Implications of malicious 3D printer firmware, *Proceedings of the Fiftieth Hawaii International Conference on System Sciences*, 2017.

[19] H. Pearce, K. Yanamandra, N. Gupta and R. Karri, FLAW3D: A Trojan-Based Cyber Attack on the Physical Outcomes of Additive Manufacturing, arXiv: 2104.09562 (arxiv.org/abs/2104.09562), 2021.

[20] N. Shevchenko, B. Frye and C. Woody, Threat Modeling for Cyber-Physical System-of-Systems: Methods Evaluation, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania (resources.sei.cmu.edu/library/asset-view.cfm?assetid=526365), 2018.

[21] C. Song, F. Ling, Z. Ba, K. Ren, C. Zhou and W. Xu, My smartphone knows what you print: Exploring smartphone-based side-channel attacks against 3D printers, *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pp. 895–907, 2016.

[22] L. Sturm, C. Williams, J. Camelio, J. White and R. Parker, Cyber-physical vulnerabilities in additive manufacturing systems: A case study attack on the .STL file with human subjects, *Journal of Manufacturing Systems*, vol. 44(1), pp. 154–164, 2017.

[23] C. Xiao, Security attack on 3D printing, presented at the *xFocus Security Conference* (`www.claudxiao.net/Attack3DPrinting-Claud-en.pdf`), 2013.

[24] M. Yampolskiy, W. King, J. Gatlin, S. Belikovetsky, A. Brown, A. Skejellum and Y. Elovici, Security of additive manufacturing: Attack taxonomy and survey, *Additive Manufacturing*, vol. 21, pp. 431–457, 2018.

[25] S. Zeltmann, N. Gupta, N. Tsoutsos, M. Maniatakos, J. Rajendran and R. Karri, Manufacturing and security challenges in 3D printing, *Journal of the Minerals, Metals and Materials Society*, vol. 68(7), pp. 1872–1881, 2016.